Vitruvian - Personal Training App

Inter-Departmental Senior Design

29 April 2020

Team 03
eseseniordesign201920team03@gmail.com

Chloe Dietz (SSE)	dietzc@seas.upenn.edu
Ajmain Hossain (SSE)	ajmain@seas.upenn.edu
Halil Can Memoglu (CIS)	memoglu@seas.upenn.edu
Abraham Milhem (SSE)	milhem@sas.upenn.edu

Advisor

Pratik Chaudhari pratikac@seas.upenn.edu

I. Table of Contents

II. Executive Summary

2.1 Executive Summary

While it is recommended to work out and widely accepted that working out increases one's health, it is often difficult to know how to start. Many beginners turn to social media fitness influencers, e-book programs, online videos, and/or personal trainers to start their fitness journey. While these resources can all instruct a user on what exercises to perform, they often do not accurately show a user how to perform the exercises and the proper technique to use to activate the correct muscles. The problem with using improper form is that there is an increased risk of injury as well as less than optimal progress. Our solution, Vitruvian, is an accessible, low cost product targeting individuals who are starting or who are already well into their fitness journey to provide them with technique feedback in order to reduce the risk of injury. The project consists of a front end user interface as well as a backend processing algorithm. The phone's camera is used to record a user working out and the backend computer vision algorithm uses the recording to analyze an individual's skeletomuscular structure, limb orientation, and joint positioning. The backend software evaluates an individual's technique by determining joint angles. Currently, exercises have explicit rules which provide ranges for acceptable and effective movement for that exercise. The software compares a user's angles with the safe rule angles and provides technique correction feedback to the user in easy to understand recommendations. The product is accessible and has an easy to use interface for individuals with a smartphone.

While much of the fall semester was spent iterating through various problem and solution pairs, the team was able to make great technical strides in the spring semester. After deciding to use a smartphone app solution instead of the original drone solution, the team created a figma mockup of the front end user interface and simultaneously started to build out the backend product. Two of the team members spent time learning the ins and outs of react native and expo and were able to build the skeleton structure for the app. As far as the backend software, the team determined that OpenPose would still be a sufficient platform to build off of. The team began by converting the Python-based model bindings to JavaScript bindings in order to make the model accessible on most devices and also added movement specific expected limits to overcome 2D vision limitations. The team chose a basic chest fly to be the first exercise to create rules for. The team had planned on integrating the front end and back end after coming back from spring break; however, our progress was halted due to COVID-19. Although we were not able to present a complete product, we are extremely proud of the progress we made this semester and how far we have come since the fall semester.

2.2 Business Plan

2.2.1 Value proposition

Vitruvian is a phone application that makes physical training healthier and more effective.Vitruvian allows for users to achieve their fitness goals in an accessible and cost-effective way. Most importantly, it allows for safer workouts with real-time corrective feedback.

2.2.2 Stakeholders

End consumers: It is particularly important to not give faulty advice to people in any medical context. While what we are proposing to do is a conceptually straightforward application, we will do our best to ensure accuracy. That being said, we must make it clear that our software's feedback, like the advice of many personal trainers, is strictly non-medical.

2.2.3 Market research

There are 55 million health club members in the US. This includes those with memberships to gyms, studios, and other similar facilities. Each person in this group has already demonstrated substantial desire and willingness to pay for physical fitness and health. Given US statistics, we would strongly expect more than 50 million in this group to have access to smartphones.

2.2.4 Competition

Vitruvian is a disruptive product, therefore it is unique and does not have any direct competition. There are numerous apps that suggest and help track exercises over time. Since this is not our value proposition, they are not direct competition. However, they are conceptually close enough that user confusion may lead them to believe that those apps are in fact competition.

Physical trainers may be considered competition, however our price points and offerings differ significantly. Our primary motivation can be boiled down to making the optimal physical trainer accessible to every individual. Currently, physical trainers offer feedback for training and nutrition at high price points. They require for users to be physically present with them at all times as well, which can be an inconvenience in many cases. We believe Vitruvian provides a more convenient, more economical, and overall more widely accessible solution.

During our research, we found there are at least two apps that take advantage of computer vision to detect and analyze exercise form. In our research, we found them to be emulating social networks with a focus on the exercise as social-proof-of-work. While an intriguing idea, we found their exercise offerings to be limited and focus to be misplaced. Therefore, they are competition to us; however, our value proposition is significantly more focused on the exercise itself.

III. Overview of Project

While exercise and physical activity are commonly recommended by doctors and are known to enhance the health of an individual, there is always a risk associated with any form of physical activity. The most common cause of exercise related injury is improper technique. Technique is the way in which an individual performs an exercise to target a specific muscle group. The effects of improper technique include injuries such as strains or tears as well as slower than expected progress.

Individuals who struggle with improper form often turn to workout videos that they can follow along with or personal trainers who can correct their form in real time. While these solutions can improve form for some individuals, neither are a perfect solution for everyone. Workout videos do have professional trainers leading workouts; however, they often do not mention proper technique or what muscle group is being activated and so a user could be performing an exercise improperly without knowing. The limitation of personal trainers is that not everyone can afford to have a personal trainer or even has the time to meet a personal trainer at a gym. Because of the limitations to current solutions, there is an even greater need for Vitruvian, which is both cost effective and gives real time technique feedback.

Vitruvian is a low cost solution that uses a phone camera to record a user while working out. The quantitative information such as joint and limb positioning is collected from the computer vision algorithm run on the phone camera video and compared against proper form to give the user position correcting feedback after each set. Feedback will be received through a smartphone application, which will visually display the proper form technique overlayed on top of the users last set thus showing the clear difference and form improvements to be made.

IV. Technical Description

4.1 Specifications and requirements

Vitruvian has two goals: to make healthy and effective physical exercise more accessible to as many people as possible, and to collect pose estimation data to create a valuable database for future human-centric computer vision projects. Therefore, Vitruvian needs to conform to the following specifications:

Functionality	Primarily, the software needs to be able to observe the user's movements while exercising in order to recognize, measure, and correct the movements. Thus, the project must be able to observe the user through the full exercise and give feedback during and after the exercise. Depending on user feedback (which will function akin to labelling), the collected data and inferred poses must be usable to extend the dataset and retrain the model.

Cost	Given that higher cost alternatives such as personal trainers exist, and that the project's goal is to increase accessibility, the monetary cost to each individual user must be very low.
Interface	The interface must be nonintrusive and easy-to-use. Given that the project aims to increase accessibility, using personal phones as the interface makes the most sense.
Regulation	To avoid regulatory problems, the project’s feedback needs to conform to a non-medical advisory standard.
Ergonomics	Given the choice of interface, the ergonomics of the project mainly concerns dynamic sensor placement such that the sensors can effectively observe exercises. Stationary sensor placement will not be a concern of ergonomics. **Phones**: Since possible phone placement will vary greatly on a case-by-case basis, we need to make sure that the software works effectively with many different phone positions. Ultimately, we will direct the user through the interface for phone placement, and the baseline positions we determine must work particularly effectively. **Wearables**: Unlike phone placement, we can specify wearable placement relative to the body. Since wearables will be on the body during physical exercise, they must be light and avoid limiting the user’s movements.
Accessibility	The monetary cost, cognitive burden of using the solution, and interface complexity must be minimized.

4.2 Iterations and alternative solutions

We went through numerous alternative solutions to solve the problem of monetary and epistemic inaccesibility of healthy and effective physical exercise. The primary alternatives are as follows:

Physical Trainer Drone - By using a small autonomous drone to dynamically observe the user while they exercise, the possible problems that can stem from poor camera placement can be solved. In addition, the mobile camera allows for the project to also be used for outdoors sports and non-stationary exercises. The user's phone would still be the primary interface except for the camera.

Facility-installed Physical Trainer Software - Instead of creating a phone application that every user would use individually, we can create a project that would be used by a facility (such as a gym) or a sports team. In this solution, a small number of stationary sensors (be it 2d or 3d cameras) would observe multiple people simultaneously, the software would track them through their movements and create performance/correction reports, and this analysis would be used as a part of the facility's offering or as an addition to the existing expertise in a sports team.

Physical Training Regimen Creator - This solution would not only observe and correct exercise movements, but also create a dynamic regimen based on user input and ongoing movement observations. Software that recommend and track regimens already exist, and the challenge with this solution would be to meaningfully utilize the software's observations in order to modify the regimen. In effect, this would mean a very close analogue of the thinking process of a physical trainer.

Wearable sensors as observation tools - In this solution, we would use wearable sensors with accelerometers in order to create an accurate 3d reconstruction of the user's movements. The sensors would be cheap and possibly more accurate than 2d camera inference. A possible combination would be to use this in conjunction with camera observation.

3d cameras as observation tools - In this solution, we would use a 3d camera like a Kinect to observe the user's movements. Given the depth data that accompanies the 2d image in such a setup, the software's reconstruction of movements is expected to be more accurate than using 2d images.

3d cameras as 2d model training tools - In this solution, which could be coupled with most of the others, we would use a 3d camera to further train the machine learning model that the 2d cameras would later use. This is expected to increase the accuracy and consistency of the final product.

4.3 Extrinsic considerations that influenced the design

As noted in our specifications and requirements, we carefully evaluated extrinsic considerations such as high economic accessibility, high impact on improving physical health of the populace, and low environmental impact while determining our project's design. The first two have been discussed at length, being more pertinent to Vitruvian's possible designs. Environmental impact was mainly a factor when evaluating whether drones' impact was justified by their benefits. We found that it was not, and that was a strong factor in going forward with a software-centric design.

4.4 Technical description and approach

We desire a solution that is highly accessible to the end user and is able to make sufficient observations with low cognitive overhead to the user. That is to say, we aim for the solution to be non intrusive so that the effort to use it would not dissuade its use. We analyzed our possible solutions with this and the above specifications in mind.

We found that the use of a small drone as the observation tool was an insufficient solution for this specific problem. Low battery life, intrusive use, and the additional cost all make that solution a poor choice if the only goal is to make healthy and effective exercise more accessible. Like a drone, a 3d camera is also a poor fit for a specialized solution given its additional cost and use burdens.

Regimen creation and recommendation apps already exist. Unless we can meaningfully use the observation and reconstruction data to modify the recommended regimen, any solution we deliver will be unoriginal (which matters for the course) and likely less effective (which matters for the project).

A group solution for businesses or teams remains the most intriguing alternative. Ultimately, limiting the usage to physical spaces or organizations is counterintuitive given the teams goals for accessibility.

The use of wearable sensors would require many of them to be attached to specific parts of the body at the start of every training period. We determined that this would go against the specifications outlined in *Accessibility* and, to a lower extent, *Cost*. Therefore, we considered using resources that users would already have access to.

With our need for high accessibility and low cost, using the personal phone as the main driver of the project is a natural conclusion. Almost everyone has phones. Given the project's goals, we would expect an even higher percentage of the targeted demographic to have phones, if not all of them. The onboard camera and increasingly high processing power of these devices would also mean that we can entirely contain the hardware requirements of the project within a single, highly accessible device.

Finally, with our advisor Professor Chaudari's guidance, we realized that we have two ways to tackle even the purely software-based project design. In particular we could create a system where explicit rules were used to check the user's movements, and we could collect pose data as a time series derived from professionals' recorded movements for future comparison.

4.5 Status of the project and test results

We have completed the back-end work of Vitruvian with the following capabilities:

- Video feed can be converted to error-corrected and rescaled skeleton vectors. Error correction is based on maintaining a time-series based threshold to discard artefact-level errors caused by the deep learning model's out-of-scale errors given very similar data. In practical terms, this reduced the amount of "jitter" in the detected pose that is not reflective of the real world, where the user may be standing still.
- Skeleton pose data can be converted to joint angle data based on the required joints by the explicit or implicit rules. The required joints are implicitly measured against their idiosyncratic "base" vectors; for instance obtaining the vertical upper-arm angle requires measuring that vector against a new vector created from the shoulder and hips.
- The joint angle data can be joined into a time series for duration-sensitive rule evaluation.
- Users can create explicit rules by determining the "required" joints and ranges of movement for those joints. These custom rules would be equivalent to our existing ones.
- Both the joint angle time series and frame-by-frame joint angle data can be used to evaluate movement correctness based on explicit rules' range-of-movement thresholds.

- Footage can be converted to time-series data, and multiple time-series data can be clustered to create implicit rules based on those ranges of motion.

Testing-wise, we determined a +/-3 degree precision with near-perfect accuracy (accurate being in that range) at a range of 2 meters (~6.5 feet) with a waist-high camera placement (+/- 30 cm ~ 1 foot) for orthogonal angle detection using a high-definition phone camera (1080p, 60fps). Real-time processing through a mid-range 2018-model smartphone ran at a rate of 10fps. Deferred processing (using saved footage) matched the footage framerate. We found that the camera quality is a more important factor than camera placement height, which was *not* our expectation to begin with. In particular, a common webcam with <480p resolution suffered from a significant amount of jitter until we implemented our small-error correction algorithm. We filmed exercises from as low as the ground level and the system did not suffer in single-person environments. However, the confidence level for the user's body using such perspectives was lower and people in the distance were more likely to interfere in non-standard placement height setups. Finally, we found that accuracy tended to suffer when observing poses that are "less likely" in the general sense. This was difficult to quantify as a general rule. However, specific cases were easily definable: for example, raising arms straight up (fully vertically) created a figure that Vitruvian found hard to detect at angles greater than 75 degrees vertical. Crossing arms fully *behind* the head caused even more inaccuracy. While exploring this oddity, we found that poses that could better be explained as "contortionist-like" were difficult to detect. This general limitation was one of the primary technical motivations to retrain the model through detected poses and user input.

In addition, we have done front-end design, built a working web app (**Figure 1)**, and partly-built a React-based mobile app for Vitruvian. Our designs for implementation are in **Figure 2**.

Capabilities we do not have due to impacted development:

- Implicit rules cannot be automatically generated due to an inability to align movement-specific substate times. Example substates would be "resting", "extension", and "transition"; that is the general components of the routine that constitutes a movement. As it is, I have to align angle data "frames" with these times by hand. This can be done automatically and I know how, but we did not have the opportunity to do so.
- A finished mobile app. We decided on a Javascript stack to be able to have an instantly-usable application. We achieved that with the web app, but we were not able to finish the native mobile application to go with it.
- User-generated time-series cannot currently be fed back into the dataset to retrain the model. The data can be saved for future use, but that use is not yet implemented.

4.6 Conclusion

We aimed to create a low-cost, highly accessible solution to solve the problem of monetary and epistemic inaccesibility of healthy and effective physical exercise. Advances in neural networks and the capabilities of smartphones allowed us to seamlessly take advantage of devices that our users are familiar with and have access to. Using the smartphone as the sensor and interface, we created a solution that fits our specifications and sufficiently reduces the problem to a statistical fit problem with multiple useful avenues of approach. **Figure 3** describes how the project works from a higher technical point of view.

V. Self-learning

5.1 Areas of learning

The main subject areas involved in our project are computer vision programming, phone app development, and working with new software. For implementing the limb algorithm program, we had to have a better understanding of computer vision by reading research papers within the field. Some group members had prior experience in this field which made it easier to use computer vision tools such as OpenCV. We also had to learn how to use different software in Linux such as Tensorflow, Python, and Anaconda; although it took some time to learn, we were able to develop enough of an understanding in order to implement the limb detection correctly. We used a Python-based model with JavaScript bindings and Node to enable the limb detection to work on any platform, which took some effort to understand these new tools. On the phone app development side, Abe and Chloe spent a lot of time learning Figma and React in order to create a presentable application on the mobile front.

5.2 Useful classes and prior knowledge

In terms of classes beneficial for our project, Chloe and Abe’s experience in their Foundations of Data Science course helped streamline our limb detection program. Hal’s background in different computer science courses has helped us figure out how to plan out our system design through phone apps, machine learning, and programming. Ajmain’s prior research in the GRASP Lab taught him about computer vision, machine learning and drones, which helped get the limb detection program running with Hal’s programming help. Having experience in Tensorflow helped us integrate different subject fields of our project into one solution.

VI. Ethical and Professional Responsibilities

6.1 Global/Societal context

Our project is geared towards promoting fitness and a healthy lifestyle. We believe that most individuals understand the importance of staying healthy and how exercise is vital to that goal, but there are many factors that inhibit individuals from exercising. These factors include the cost of going to the gym, lack of knowledge, and need for instruction. Our project hopes to reconcile these factors by making exercise accessible for anyone; the only thing you will need is a phone. Within a societal and global context, our project can work anywhere and for anyone as long as they have a phone that can meet the requirements to run our app. We hope that this will get more people interested in getting into exercise, which will make our society healthier. In terms of environment, there are not too many implications related to our project, but using our exercise app indoors will reduce transportation costs and related emissions from driving to the gym everyday. At a more current level, the widespread COVID-19 has made it difficult for people to go to the gym. Our app can help people exercise appropriately indoors and is especially useful for beginners who would like to use their newly found free time proactively. By helping new people exercise, we can limit the amount of injuries in workouts through our specially devised fitness app.

6.2 Potential ethical issues

The main ethical issues we face are related to data and privacy-related concerns. Our app records videos of people performing exercises with their faces as well as keeping any personal information needed to register accounts or other similar concepts. In order to resolve these issues, we plan on encrypting personal information. In regards to videos, we do not plan on storing any videos after their original use because we only need the videos for immediate processing to give real feedback to users during exercise routines. In addition, we are not professionals so we have to ensure that we do the necessary research to provide correct and honest feedback for workouts. Although we are not directly liable for injuries, it is our ethical duty to prevent people from hurting themselves while using our app. Aligned with our data privacy standards, we must abide by client confidentiality and not share any personal information with others, even for business reasons.

VII. Meetings

Throughout the spring semester specifically, the team properly utilized the teaching assistant staff and professors. The team first met with Sid and Jan at the beginning of the semester to discuss fall semester progress and milestones for the spring semester. However, this quickly turned into weekly meetings where the team was able to discuss different business paths and present our weekly technical progress. This greatly helped in keeping the team focused and on track as well helping to continually push the team to do more than previously thought capable. In addition to these weekly meetings with the professors, the team met with

Matt (TA) to discuss the front end user interface. He was able to steer us in the right direction and suggested the use of react native and expo.

In addition to advisor and teaching staff meetings, the team met every Sunday night prior to spring break for at least an hour in order to discuss the progress we made, needed to make, and important decisions in our pivoting process. We also met at least one more time during the week in order to practice upcoming presentations or work on any deliverables due for that week. After spring break, we had video calls at least twice a week in preparation for the final demo day.

VIII. Proposed Schedule With Milestones

The spring semester milestones, as shown in detail in **Figure 4** were meant to build on our work from the previous semester in order to create a fully fleshed-out Vitruvian for our Senior Design project. We realized that we had a lot of work ahead of ourselves from the very beginning, so we made sure to work closely with our instructors and our advisor, Pratik Chaudhari, to ensure that we would be able to properly complete these milestones in our limited time.

The completed spring semester milestones were the following:

- By **January 27**, we completed our research in optimal workout form and safety, through reading various medical journals and consulting with professional trainers and physical therapists. Our research primarily focused on the first exercise we decided to implement, the upright chest fly, however we compiled sufficient research to understand the basic safety requirements (e.g. upright back, straight torso) for other exercises we hoped to implement.
- By **February 2**, we were able to fully implement the limb detection data analysis, and the ability for the software to respond to user positioning and provide feedback. We also were able to validate our model by comparing actual limb angles vs. the detected angles. We continued to refine this in coordination with our advisor and instructors.
- We decided that Vitruvian should have one fully implemented exercise and a user interface by our Midterm Demo Day on **February 26.** Therefore, we worked on the front-end and back-end in parallel. While Abraham and Chloe developed the front-end in react-native, Aj and Hal improved on the back-end and implemented detection and feedback for the upright chest fly exercise. We were able to complete this in time for the Midterm Demo Day.

There were a few milestones that we, unfortunately, were not able to complete due to the COVID-19 crisis.

- While this was an important milestone for our project, we were unable to fully integrate the front-end and back-end components of Vitruvian. This was supposed to be completed by **March 23**. The current state of Vitruvian does include a fully functional front-end and back-end, they are just not in the form of a mobile app.

- We hoped to add functionality for squats, push-ups, and sit-ups by **March 27**. This would not have been very difficult to implement, since Vitruvian already had the capacity for limb detection, analysis, and feedback in its back-end, however we were not able to complete it in time.
- Since we were not able to create a fully integrated app, we also were not able to beta-test it. We had planned to refine our app based on beta-testing from our classmates by **April 1**, however due to the unforeseen circumstances, we were unable to do it.

IX. Discussion of Teamwork

Our product is composed of a front end user friendly phone application and a back end computer vision software. After moving away from a solution that included a drone, the team shifted milestones and divided up the work according to each team member's past experiences and skills. As mentioned previously, Hal's previous computer science coursework was essential in developing the back end software. Chloe and Abe's knowledge from the Foundations Data Science course better equipped the team to tackle the limb detection software needed to analyze an individual's skeletomuscular structure. Ajmain's prior research in the GRASP Lab provided the team with the foundation to implement different computer vision techniques.

The team divided the project tasks into two groups - the front end user interface and the back end software analysis - to better divide tasks and utilize each individual's knowledge and skills to their fullest extent. While the team did split the work, the team met weekly to discuss the improvement made on each side as well as explain to the other team members how this was technically done and in doing so, all team members have gained an understanding of a topic they otherwise would not know about.

Due to Hal and Aj's previous computer science coursework and experience, they spent the semester working on the back end software. Aj devoted much of his time to learning the OpenPose library and was able to get the limb detection software to run with his computer camera. Hal used Node and C++-to-Javascript bindings to create the CV application, built the web app, wrote the processing software and relevant algorithms, and created the Figma mockups. This allowed us to run the program from our phone and multiple computers. While Aj researched safe angles and movement for a proper chest fly, Hal was able to implement those angles as rules in the program for a user's angles to be checked against. On the other hand, Chloe and Abe tackled the front end user interface. They began by creating mockups of the home page and exercise pages within figma. They then spent time learning react native and expo in order to create the skeleton of the phone app. The team successfully implemented the chest fly exercise and had a skeleton app that showed feedback. The team would have liked to add a number of other exercises and rules as well as integrate the front end and back end components.

X. Budget and Justification

Our original budget was $0 and remained $0 throughout the semester. Vitruvian is primarily a software-based product and we had no costs in programming the actual app. Additionally, in order to test it, we were able to use the cameras on our laptops and phones, so we did not need to purchase any additional hardware either.

XI. Standards and Compliance

Regarding the phone application development, we would need a registered app developer license and abide by the privacy policy. Our app is still in development but we still comply with these guidelines. Our other main standards are related to the technology side. The main standard we comply with is IEEE P7002: Data Privacy which is related to maintaining our data securely and privately. We do not store any private videos or footage of people exercising because we perform real-time feedback for the users. In addition, we do not ask for any personal information that would be at risk for security concerns. The next standards we comply with is IEEE P2675: DevOps: Building Reliable and Secure Systems Including App Build and 12207-2017:Systems and Software Engineering - Software Life Processes which are related to the technological development of our project. Everything related to the project from the limb detection and app is coded so these standards are important to follow. We ensured our code is secure through rigorous testing and followed the normal software engineering cycle for developing our code by mapping out the different interfaces of the project. The final standard we comply with is ISO 42: Photography which is related to the privacy of all the images and videos used in our project. We do not store any videos as said earlier so there are no privacy concerns regarding that. The only things we record are the data from exercise workouts in order to see how a user is progressing.

XII. Work Done Since Last Semester

Since last semester, we have built upon our OpenPose limb detection software in order to create a fully-functional back-end and front-end. Last semester, we did not have a full product, we simply had a bare-bones software that could detect limb positioning based on image data from a webcam. While this limb detection serves as the foundation for Vitruvian in its current state, this semester we have built upon it in order to accurately calculate body limb positioning and relative angles. We have conducted research on specific exercises and integrated the research into our software in order to provide the user with real-time feedback on his or her correct and incorrect movements. Our back-end can now accurately detect and correct the upright chest fly exercise. We have also been able to create a fully functional front-end for the product in react-native, for use on Android phones. In short, we have been able to take the software that we developed last semester, build it up, and turn it into a real, marketable product that can serve as a virtual training app for people everywhere.

XIII. Discussion and Conclusion

This year, we have been able to come together as a team and develop our personal training platform, Vitruvian. While we began the year on a very different path, with ideas ranging from an artificial intelligence coffee selection platform to a high-frequency mosquito-repelling device, throughout our ideation and iteration process we were able to converge on a useful, interesting, and fun project. As mentioned throughout this report, we have been able to develop a functional product that accurately detects and corrects user workouts, providing useful feedback to users in real time. While we were not able to fully integrate our front-end with our back-end due to unforeseen circumstances, we are proud that we were able to develop Vitruvian as best we could.

We have learned a tremendous amount about the design process as a whole. From the very beginning, it was a challenge for us to narrow down our proposed ideas to something that would be feasible and useful as a Senior Design project. We iterated through several different designs, with our advisor and instructors guiding us every step of the way. While we had initially wanted to utilize drone hardware, we quickly realized it was not the optimal solution to the problem we were trying to solve. We pivoted our design to be usable on a phone or computer camera since it seemed like a much more realistic solution for users, and ultimately it improved our project significantly. This whole process of iterating and adjusting based on feedback taught us valuable lessons about properly designing

This project also taught us to be able to plan effectively to meet requirements within a short period of time. Since we spent so much time iterating through different designs, the majority of our implementation of our project had to be done in the second semester. This meant that we needed to flesh out our back-end analysis and begin our user interface from scratch within the span of a few months. This was definitely challenging, but we were able to utilize our respective strengths and proper planning to effectively meet the requirements for our project. With guidance from our advisor and instructors, we were able to learn how to properly split up work and avoid bottlenecks to producing the final product. This is an essential skill that we all will certainly use in our engineering careers after Senior Design.

Overall this project was a period of significant learning and growth. This was a unique opportunity for us to utilize the skills we have gained as engineers throughout our time at Penn and work to solve a complex, interesting problem of our choosing. Working to meet deadlines as a team is never easy, but we believe it was an important experience and an excellent capstone to our time at Penn. We are excited to take the skills that we have gained from this experience to our careers after Penn, and we are grateful to our instructors and our advisor for all of their help along the way.

XIV. Appendices

Figure 1: Vitruvian back-end web application
A screenshot of the web app demonstrating Vitruvian’s internal processes

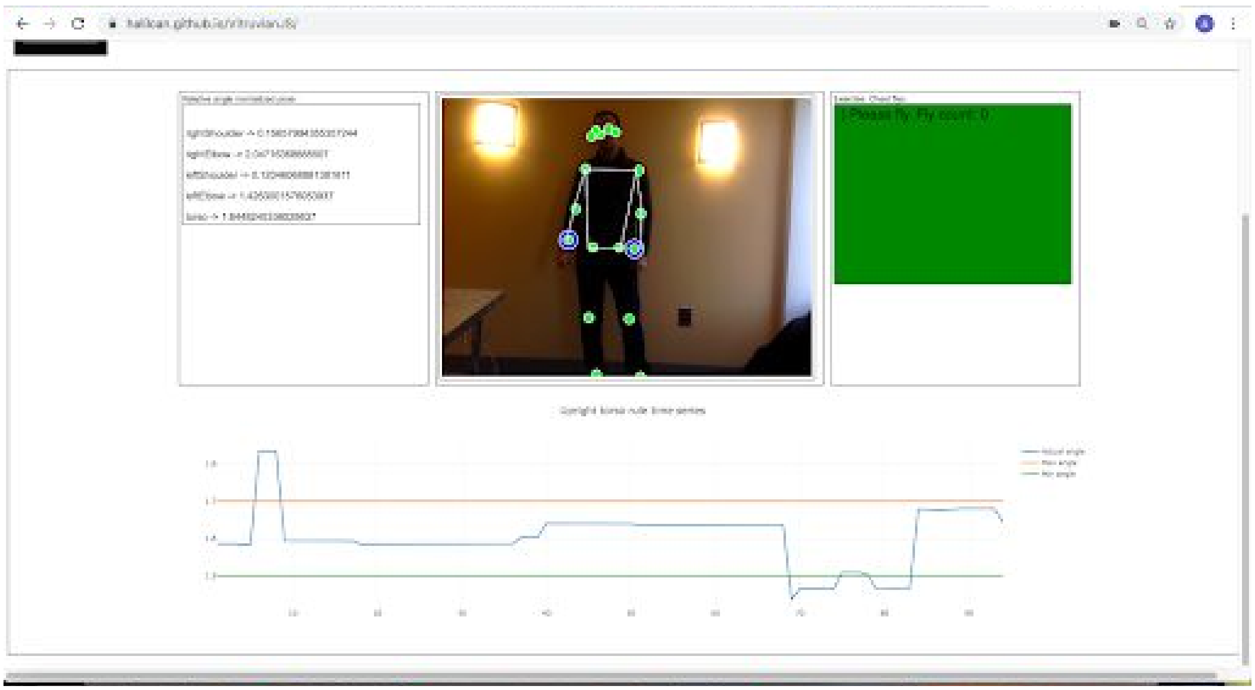

Figure 2: Figma Wireframe Designs

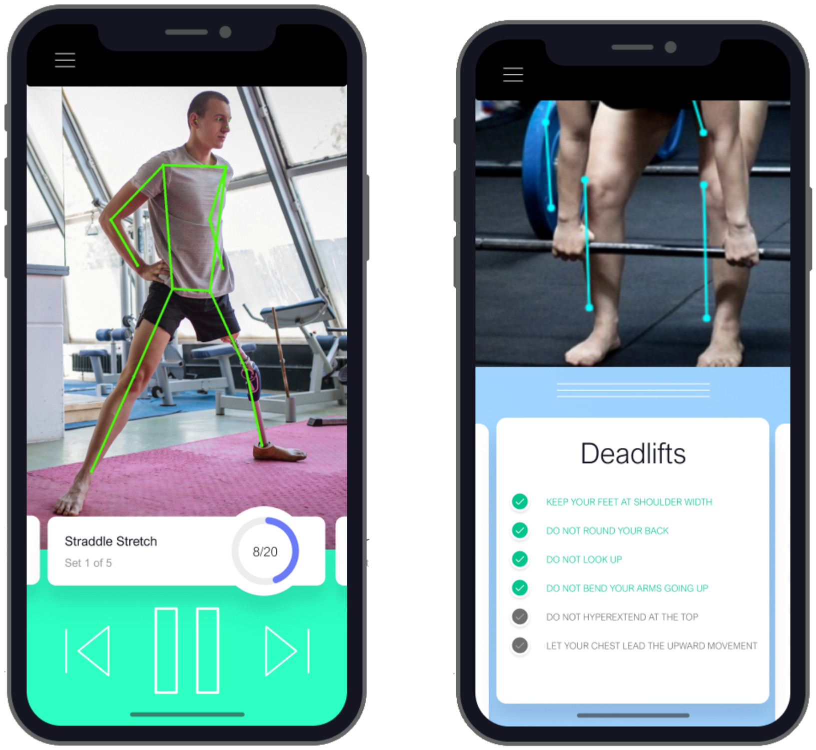

Figure 3: Visual Explanation of How Vitruvian Works

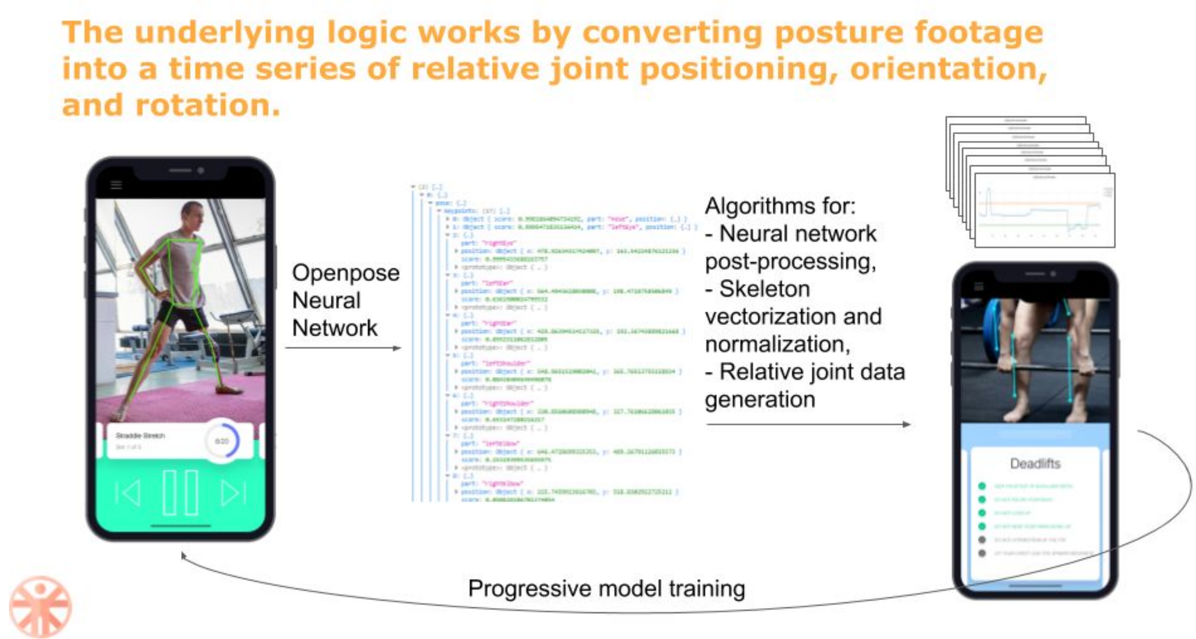

Figure 4. Spring Semester Milestones

Description of our milestones for the Spring. The green represents the milestones that we were able to complete, while the red represents the milestones we could not fully complete due to the COVID-19 crisis.

Spring Semester Milestones

Task	Chloe	Ajmain	Abraham	Halil
Establish understanding of optimal workout practices	1/27	1/27	1/27	
Validate and implement data analysis/response from limb recognition	2/4	2/4	2/4	2/4
Develop Phone Application with intuitive user interface	2/26		2/26	
Implement Chest-Fly Exercise		2/26		2/26
Integrate Front-end with Backend	3/23	3/23	3/23	3/23
Add squat/push-up/sit-up exercises		3/27		3/27
Test prototype/refine if any bugs appear	4/1	4/1	4/1	4/1