

DEADLIFT ASSISTANT PLUS

By

Deniz Yildirim

Erik Jacobson

Sang Baek Han

Final Report for ECE 445, Senior Design, Spring 2020

TA: Weihang Liang

8 May 2020

Project No. 7

Abstract

The Deadlift Assistant Plus provides real-time form feedback for weight trainers exercising a deadlift - amateurs and professionals alike. Our team has outlined a possible solution using computer vision algorithms and a novel wristband recognition system to improve over the existing design by providing the user with a lightweight, portable and reliable system to improve the form of their lifts. Upon recognition of poor form that could damage the user, our system will provide feedback in the form of a warning beep. Similarly, the wristband has the capability to detect and track user heart rate, sounding a separate beep if exceeding a dangerous rate. While direct implementation of the idea lays outside the scope of this project due to environmental constraints, discussed below is a plan meant to guide future implementations to a successful end product.

Contents

1. Motivation.....	1
1.1 Problem Statement.....	1
1.2 Solution	1
1.3 High-Level Requirements	1
1.4 Visual Aid.....	2
1.5 Block Diagram.....	2
2 Implementation.....	3
2.1 Input Subsystem	3
2.1.1 Heart Rate Sensor	3
2.1.2 Camera	3
2.2 Control Subsystem	4
2.2.1 Microcontroller	4
2.2.2 Single-board Computer	4
2.2.3 Bluetooth.....	4
2.3 Output Subsystem	4
2.3.1 Speaker.....	4
2.3.2 Liquid Crystal Display	4
2.4 Power Supply Subsystem	5
2.4.1 Voltage Regulators.....	5
2.4.2 Batteries	5
2.5 Schematics and Board Layout	6
2.6 Bill of Materials.....	8
2.7 Software	8
2.8 Tolerance Analysis	11
3. Conclusion.....	13
3.1 Accomplishments.....	13
3.2 Uncertainties	13
3.3 Ethical considerations	14
3.4 Future work	14
References	15

1. Motivation

1.1 Problem Statement

The deadlift is infamous for its risk when done with poor form because of the heavy weight involved [1]. Normally lifters will start with 1.5 times their bodyweight for the deadlift barbell; this is approximately 315 pounds [2]. Considering this weight, the deadlift poses potential harm to any lifter.

The study cited shows that professional weight trainers get back injuries from deadlifting 31% of the time [3]. Compared to the other popular weight exercises, the percentage is twice as high. This is what happens to professional weight trainers, and it is obvious that those who are not professional in weight training will get more injuries due to their poor form albeit with lighter weights.

The deadlift is one of the top three most popular weight training exercises along with the squat and the bench press. Yet, it is considered the most difficult amongst the three of them due to the weight involved in it. Many people tend to do the deadlift in a poor form to successfully complete the attempt and end up hurting their backs or spines. They do not realize that they are doing the deadlift in the wrong way because they tend to avoid having an expensive lesson from a personal trainer.

1.2 Solution

Our solution focuses on preventing injuries while doing the deadlift by alerting the users whenever they are doing the deadlift in a poor form. We aim to make this cheaper than taking the lesson from a personal trainer. A video camera with computer-vision records the deadlift performance from the user and provides instant feedback on their forms by making a sound alert if the deadlift form is poor. The user wears a color wristband which can be recognized by the computer-vision to recognize when the deadlift is in progress. The wristband also measures the heart rate of the user and warns the user if the heart rate goes too high with a different sound alert. In addition to the alert, the user is also able to get the overall feedback on how their form got better or worse over time, as well as their heartbeat data. This helps users track their deadlift training and motivates them to improve their form.

1.3 High-Level Requirements

- The system must recognize form that poses potential harm to the user and emit an alert sound in less than 0.5 second with a 95% accuracy.
- The camera device must be able to make the alert sound when the heart rate of the user reaches 180 beat per minute.
- Both wristband and camera must be able to run for at least 1 hour continuously to satisfy the regular workout time.

1.4 Visual Aid

Figure 1 below shows good and poor deadlift forms. The back needs to be straightened and in line with the neck. If the back is curved inward or outwards, it puts the user in risk. The first two forms shown in Figure 1 are the poor forms because the head is not straight, and the back is curved. The proper form is shown on the most right, where the neck and back are both straight.



Figure 1 Visual Aid of Good VS Bad Deadlifting Form

1.5 Block Diagram

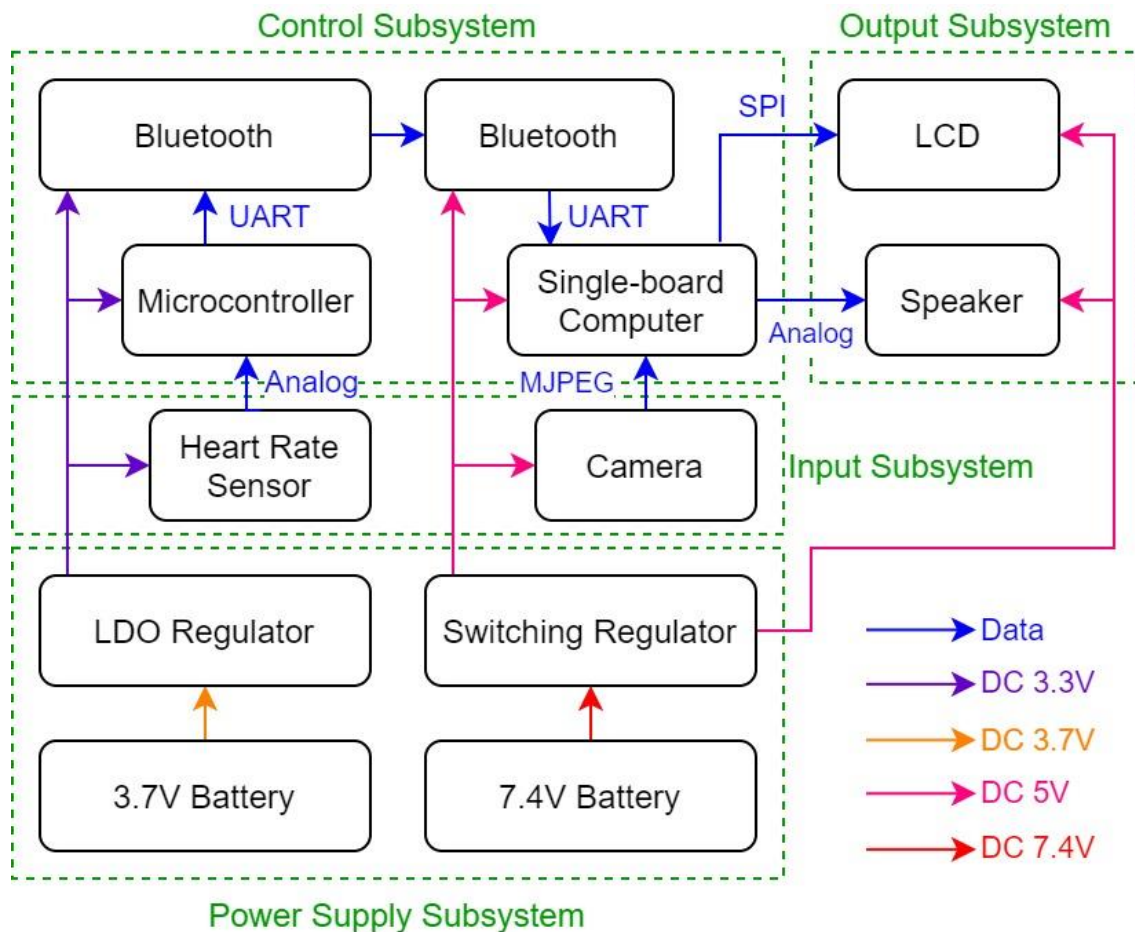


Figure 2 Block Diagram of Input Subsystem, Output Subsystem, Control Subsystem, and Power Supply Subsystem

2 Implementation

Our entire design is simply comprised of two devices: the wristband device and the camera device. The wristband device is worn by the user, and it measures the heart rate of the user while doing the deadlift. The heart rate data is sent to the camera device through Bluetooth. Figure 2 in the previous page shows the block diagram of both wristband device and camera device. The five components on the left are the components of the wristband device and the rest of components on the right are the components of the camera device.

The camera device is a smart camera with a tripod. This is set next to where the user does the deadlift and it captures the motion of deadlift attempt. The camera lens captures the motion in video, and the single-board computer inside the camera processes it through computer-vision. If the deadlift form is good then there is no feedback, but if the deadlift form is bad then there is instant feedback from the camera with a unique alert sound through the speaker. The unique color of the wristband acts as the indicator for the computer-vision to understand that the deadlift attempt has started.

The heart rate data sent from the wristband is processed by the single-board computer, and if the heart rate is too high it alerts the user. For our product, we set it at 180 bpm (beats per minute), which is the average maximum heart rate at 100 percent exertion [4]. Once the heart rate is above 180 bpm, the unique alert sound is made through the speaker. This is different from the sound for the deadlift form feedback.

The overall summary of successful deadlift attempts and heart rate data is shown through the LCD (liquid crystal display) at the end of the exercise.

2.1 Input Subsystem

Each device has its own way of receiving the data. The wristband device receives the heart rate data from the user through the heart rate sensor. The camera device receives the motion of the user through the camera lens.

2.1.1 Heart Rate Sensor

The heart rate sensor monitors the heart rate of the user and warns the user if the heart rate is too high. The heart rate data is sent through the microcontroller via Bluetooth to the single-board computer. This is powered by dc (direct current) 3.3V. We use the hardware module “Pulse Sensor Amped”, which is made to communicate easily with existing Atmel libraries [5].

2.1.2 Camera

The camera lens captures the video of the deadlift attempt of the user. The video is recorded and analyzed on the single-board computer. We use the supported Raspberry Pi Camera V2 for easy compatibility with our single-board computer.

2.2 Control Subsystem

The video recorded through the camera lens is analyzed by the single-board computer through the computer-vision. The computer-vision algorithm uses OpenPose to detect the motion of the body movement [6]. After the analysis, the subsystem sounds an alert if the deadlift form is poor. The analysis is done in real-time. The Bluetooth module is used to communicate between the wristband and the camera so the heart rate data can be received by the single-board computer.

2.2.1 Microcontroller

The microcontroller processes the heart rate data sent from the heart rate sensor and sends it to the single-board computer via Bluetooth. The heart rate sensor itself cannot really parse the heart rate data, so the microcontroller needs to assist in order to convert it into digital data. We used the ATmega328 for this project.

2.2.2 Single-board Computer

The single-board computer processes the images with computer-vision so the deadlift form can be analyzed after receiving the video through the camera lens. It is also responsible for measuring heart rate. The board then calculates the workout summary and displays it on the LCD. For this project, we use a Raspberry Pi 4.

2.2.3 Bluetooth

The Bluetooth module communicates between the wristband and camera devices to send the heart rate data. This acts as the wireless communication between the microcontroller and the single-board computer. We use the HC06 as the Bluetooth module for the microcontroller and use the built-in Bluetooth module for the single-board computer.

2.3 Output Subsystem

There are two alert sounds for the speaker output. The first alert sound is for determining whether the deadlift form is good or bad. If the deadlift form is determined to be bad, the speaker sounds a primary alert. The second alert sound is for the heart rate. If the heart rate is above 180bpm, the speaker alerts a secondary sound.

The LCD displays the final summary of proper deadlift attempts as well as heart rate data. Each attempt will be stacked accordingly in the single-board computer memory. The result is calculated in percentage and displayed on the LCD.

2.3.1 Speaker

The speaker outputs the alert sound when the bad deadlift form is detected or when the heart rate is too high. These two unique alert sounds are different so the user can distinguish them. For this project, we connected Adafruit stereo speakers to the Raspberry Pi audio jack.

2.3.2 Liquid Crystal Display

The LCD displays the entire workout summary report. The result shows the total number of deadlifts attempts since the device started as well as the total number of successful deadlift attempts. It also shows heart rate progression. We use an Adafruit LCD to display the result.

2.4 Power Supply Subsystem

The wristband device is powered by a low-dropout voltage regulator that converts 3.7V DC to 3.3V DC. One 3.7V Li-ion (Lithium-ion) battery is connected to the low-dropout voltage regulator. The camera device is powered by a switching regulator that converts 7.4V DC to 5V DC. Two 3.7V Li-ion batteries are connected to the switching regulator.

2.4.1 Voltage Regulators

An LDO (low-dropout regulator) voltage regulator converts 3.7V DC to 3.3V DC. An LDO voltage regulator is used because the dropout voltage needs to be small and the device on the wristband needs to be small. The circuit size is comparably small as the number of components attached to LDO voltage regulator is smaller than that of the switching regulator. We use LP2985 LDO voltage regulator for the wristband device.

A switching voltage regulator converts 7.4V DC to 5V DC. Instead of using an LDO voltage regulator, we use the switching voltage regulator to conserve battery life. The single-board computer and the camera consume a lot of battery power and the switching voltage regulator helps make this circuit more efficient. The audio output gets a slight bit of noise due to the switching property of the regulator, but the quality of the audio is not very important for this product. We use TPS82150 from Texas Instruments for the camera device.

2.4.2 Batteries

We use one 3.7V Li-ion battery to power the wristband device. 3.7V DC is converted to 3.3V DC so the entire wristband device can be powered. The battery is small so that the user cannot feel much weight from it while doing the deadlift. The battery we use for the wristband device is a 40mAh Li-ion battery and its size is 22mm long and 11.5mm wide.

We use two 3.7V Li-ion batteries to power the camera device. 7.4V DC is converted to 5V DC so that the entire camera device can be powered. The batteries we use for the camera device are two 1Ah Li-ion batteries.

2.5 Schematics and Board Layout

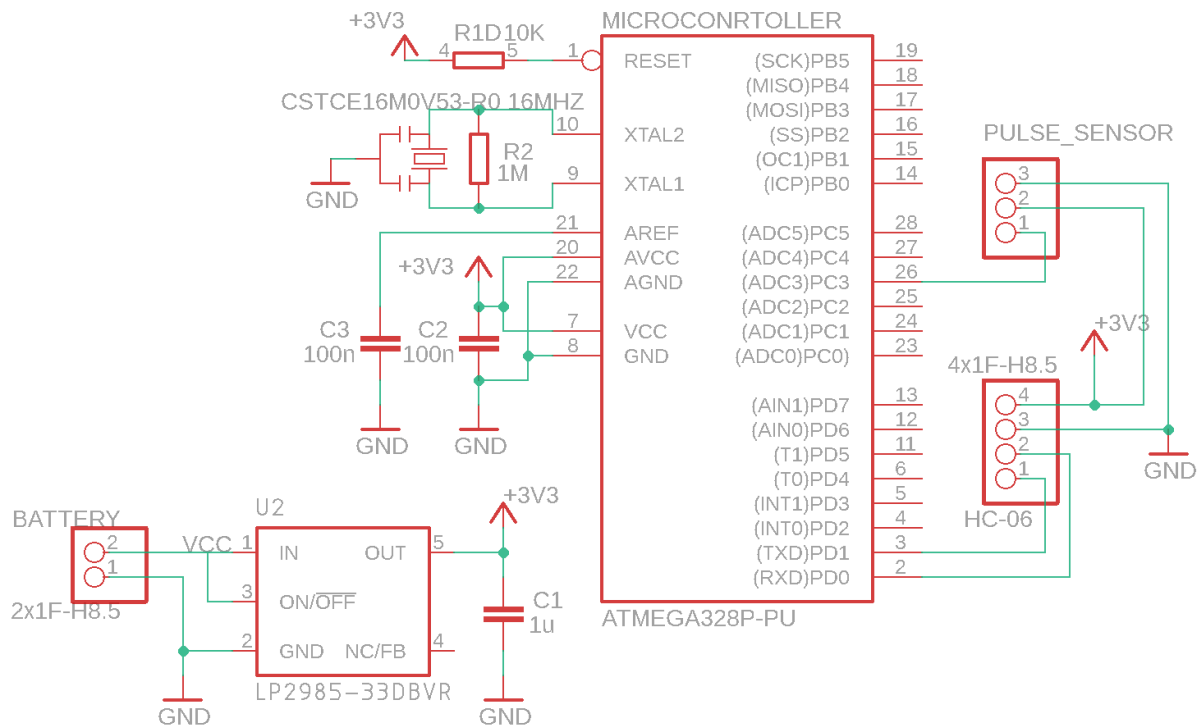


Figure 3 Schematic of LDO Regulator and Microcontroller for Wristband device




Figure 4 PCB (Printed Circuit Board) Layout of LDO Regulator and Microcontroller for Wristband device

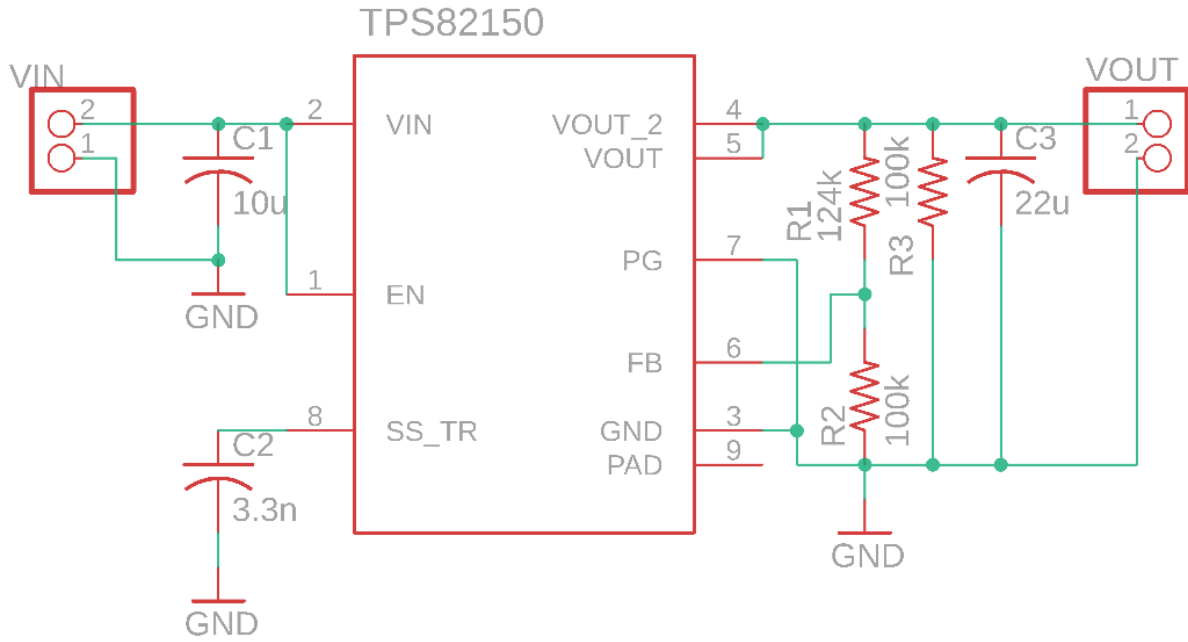


Figure 5 Schematic of Switching Regulator for Camera device




Figure 6 PCB Layout of Switching Regulator for Camera device

2.6 Bill of Materials

Table 1 Parts Costs

Parts	Quantity	Unit Price	Cost
Pulse Sensor Amped	1	\$25.00	\$25.00
Microchip ATMEGA328 Microcontroller	1	\$2.08	\$2.08
Olímex HC-06 Bluetooth Module	1	\$2.99	\$2.99
Texas Instruments LP2985 3.3V Low-dropout Regulator	1	\$0.11	\$0.11
SparkFun Polymer Lithium Ion Battery 40mAh	1	\$4.50	\$4.50
Sony IMX219 Raspberry Pi Camera Module	1	\$29.96	\$29.96
Raspberry Pi 4 Model B	1	\$35.00	\$35.00
Adafruit 1669 Speaker	1	\$7.50	\$7.50
Adafruit 3627 LCD Display	1	\$2.95	\$2.95
Texas Instruments TPS82150 1A Switching Regulator	1	\$1.20	\$1.20
SparkFun Polymer Lithium Ion Battery 1Ah	2	\$9.95	\$19.90
Wristband	1	\$3.00	\$3.00
Total Cost:			\$134.19

2.7 Software

There are two different sides to the software implementation of our system – the software involved on our single-board computer, and the software implemented on our microcontroller. For the former, we need to use OpenCV and Python OpenPose libraries to create a basic skeletal identifier for our user [7]. This needs to run relatively fast, which is difficult as the Raspberry Pi has relatively low processing power for machine learning applications. This is important as OpenPose builds upon TensorFlow, a machine learning or a neural network library, when doing image recognition and classification. Additionally, we use blob detection to find the wristband and track the trajectory of the barbell.

The only benchmark we were able to find is a similar project that uses a Raspberry Pi to classify objects in an image – and while this is not an exact replica of our project, it is a useful insight into the capability of our program. This program runs at approximately 0.9 FPS (frames per second) [7]. However, this program can recognize multiple different objects over a pretrained set of data.

Our skeleton program essentially needs to do a fraction of this. Furthermore, the resolution size of the image can be lowered to make for faster analysis. For the exact benchmark of what should be acceptable, 3-4 FPS is achievable and more than suitable. The algorithm will use wristband detection to only calculate the skeleton when the user has started a lift. Using a simple motion algorithm while tracking the wristband, we can wait until the user begins a lift, record 3-4 FPS of video in the duration of the lift, and then process any frames we could not after the lift is over. This makes a small delay in our system, similar to the delay in the original project, yet would easily achieve a high enough FPS. The OpenCV tutorial written by Gupta is an informative article about where we based our skeletal detection algorithm [6]. We then created a separate neural network and used this processed skeletal data to detect good or bad form based on a set of data we created, implemented with a machine learning library like Keras or TensorFlow. This secondary network runs very quickly compared to OpenPose. It

more than likely better classifies form than the joint-angle based and hardcoded approach done by the prior group.

In addition to skeleton detection and analysis, we track the position of the wristbands to figure out where/when the lift begins and ends, and how fast the lift was. Using trigonometry mentioned in the tolerance analysis, we find the exact location of the wristband at any given timeframe which lets us track the velocity and path of the entire lift. For a healthy lift, the barbell, and thus the wristband, should follow a straight path up. Additionally, the initial acceleration should be as big as possible. The path of the wristbands and their acceleration are added to the data from the skeleton detection to be evaluated, and at the end the algorithm calculates a confidence level of proper or non-proper form.

The second aspect of our program is the software on our microcontroller. This, in order to interface easily with our heartrate monitor, uses the built in Arduino software package. Furthermore, there is plenty of support for this as well as the Bluetooth module – this is a straightforward bit of code. This also needed to communicate with the Raspberry Pi, so some small bit of code is implemented on the single-board computer in order to properly send and receive signals.

Figure 7 shown on the next page is a software flowchart. As showing communication between the two devices in the pseudocode would only obfuscate the process, and we have already discussed communication responsibilities, this is intentionally omitted from the diagram.

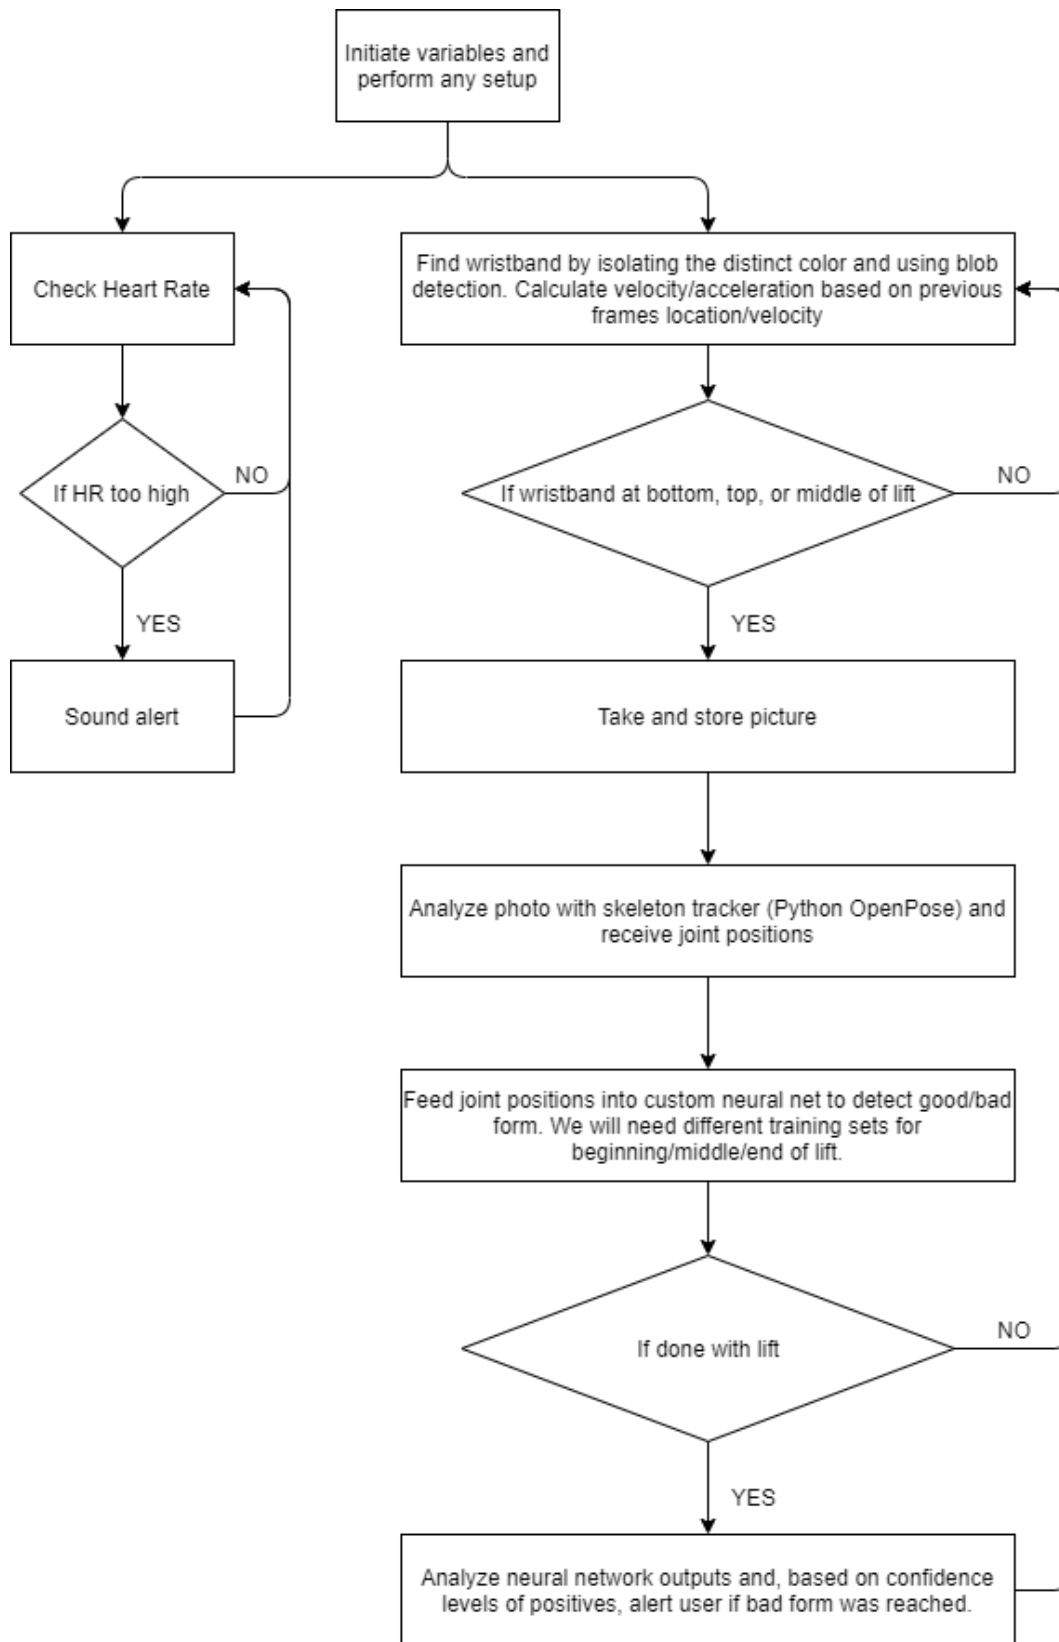


Figure 7 Software Flowchart of Deadlift form detection and Heart rate detection

2.8 Tolerance Analysis

The hardest part of this project is without doubt the computer vision algorithm that detects the wristband and determines the position. The software also detects how bent the person's back gets while lifting. A noisy detection that results in false positives or false negatives can be very dangerous, so we need to figure out everything needed to get accurate detection. Different lighting conditions and different positions of the person related to the camera makes it difficult to make confident decisions.

We check both the pose of the skeleton to make sure the back of the lifter is straight, along with the velocity and trajectory of wristband to make sure the barbell moves in a straight trajectory perpendicular to the ground. For this to work properly, both wristbands as well as the core parts of the lifter's body must be visible to the camera. In particular, the barbells can obstruct the vision of the camera. To prevent this the camera should be high enough that the wristbands are visible during the entire lifting process.

The position of the camera is very important because, even though determining back straightness is relatively independent of camera position, to find the position and velocity of the wristband we must know the exact position of the camera in relation to the wristband.

If the camera is placed at the side of the lifter, then a straight path up would still look like a straight line when viewed from a camera that is placed slightly higher than the position of the wristband. This means that the height of the camera does not interfere with our straight path detection. Determining the velocity, though, requires us to know the position of the wristband.

Because the size of the wristband is known, if detected properly, one can use perspective to figure out how far away the wristband is from the camera. First, we need to find the focal length of the camera - which can either be already known or discovered by using triangle similarity with the equation (1) shown below [8].

$$Focal\ Length = \frac{Pixels \times Distance}{Width_{Known}} \quad (1)$$

Once we know the focal length, we can use the same formula to find the distance of the wristband from the camera if we know the size of the wristband.

After finding this distance, we can find the starting position of the wristband. One can imagine a right triangle between the camera, the ground right under the camera, and the starting position of the wristband. From simple trigonometry we know that at the beginning of the lift we have the equation (2) below.

$$Distance_{Camera} = \sqrt{(Height_{Camera} - Height_{Barbell})^2 + (Distance_{Ground})^2} \quad (2)$$

From this equation we could find $Distance_{Ground}$ which is the distance parallel to the ground. $Distance_{Camera}$ is the distance from the wristband to the camera tripod, which is the hypotenuse distance

between the height difference and the distance parallel to the ground. $Height_{Camera}$ is the height of camera tripod and $Height_{Barbell}$ is the distance from the wristband and the ground.

Assuming the lift is perpendicular to the ground, doing this right before the lift would give us three corners of a trapezoid whose corners are the camera, the feet of the camera, the starting position of the wristband, and the current position of the wristband. With this information we would know, at any point, how high up the wristband is from the ground. Thus, we can find the velocity.

Additionally, it might be difficult to design a wristband that is big and unique enough for the algorithm to detect while also making sure it is comfortable. An ergonomic design is important because, if it is uncomfortable, people would not use it. It could also distract people while lifting and worsen their form instead of improving it.

3. Conclusion

3.1 Accomplishments

As far as direct improvements over the original project go, we begin by quoting the prior team's Final paper - "First and foremost, the Barbell Subsystem ended up weighing roughly 2 lbs because it was attached to the barbell via a large metal clamp. It was contained in a plastic box with 6in. x 8 in. x 3.5 in. dimensions. We would like to reduce the size and weight of this device as much as possible by designing a more compact PCB, casing, and attachment to reduce its impact on the deadlift workout as much as possible" [9]. Our project fixes the issue of a bulky subsystem by reducing the weight and size of the mechanism to detect the start and end of the lift. This is done through a wristband system rather than a box mounted on the end of a weight set, which might in fact limit the number of weights a user can safely place on the end of the bar.

Furthermore, due to the new sound alert system, the user does not have to bend their neck to look at the indicator LEDs. We thought that this might negatively impact the experience by making it less natural. We also considered that the LEDs might be blocked by particularly large weights, requiring the user to set down the bar to observe their results. This issue was shown in the previous project's demonstration video [10].

The one aspect that the original system did not have was a heart rate sensor. While exact numbers in scientific papers are difficult to find, we found a source that suggests max heart rate of a user should be 220 less their age [4]. We have programmed an alert to prevent breaking this limit, assuming an age of 40 for a conservative limit. Our wristband can also track and store heart rate data, which might replace the function of a FitBit or other wearable workout device as the user might not want to wear two wristbands. The heart rate sensor is a secondary priority to form detection, but it does add some functionality to the system.

3.2 Uncertainties

Due to the limitations of a physical implementation, it is difficult to measure the effectivity of our implementation over the prior one. For example, our secondary neural net would theoretically work better than a hardcoded, angle-driven recognition system. Similarly, the wristband recognition system might not work as consistently as the hardware-based system that recognized the beginning and the end of the lift. However, these uncertainties are inherent with the lack of a physical prototype and are not directly linked to shortcomings in our design.

It is impossible for our detection system to be 100% accurate and false negatives of poor form may result in user injuries. Furthermore, our system cannot work on every body type, size, etc. The accuracy of OpenPose is well documented, however this might not carry over to our form detection neural net. False positives, while not as dangerous, could cause the user to change their already proper form, resulting in worsening their form.

As also referenced in our future work, we are unsure if the wristband will be detectable in all configurations of weights. For example, a particularly bulky weight could obscure the wristband from the vision of camera.

3.3 Ethical considerations

The IEEE Code of Ethics #9 states that “avoid injuring others, their property, reputation, or employment by false or malicious action” [11]. In order to satisfy this Code of Ethics, this product aims to reduce the risk of injury and it is very important that it does not do the opposite by causing an injury. No matter how perfect a computer vision algorithm is, it is prone to errors. If the user trusts the product completely, the user might injure themselves. We should make it clear to the user that they must pay attention to their form and proceed with care regardless of what the product says.

The other concern is the Li-ion batteries. We are using total of three Li-ion batteries: one on the wristband and two on the camera. Two Li-ion batteries on the camera is not a huge issue, but the Li-ion battery on the wristband needs careful attention. Li-ion battery is known for its potential risk of explosion, and we will prevent this by making sure there is no shorting of the circuit and the circuit has the protection to prevent over charging [12].

As we have chosen to design our product to be private use only, this also raise the concern with the IEEE Code of Ethics #1 where it states that “disclose promptly factors that might endanger the public or the environment” [11]. We are in a COVID-19 pandemic, and the main modes of transmission for COVID-19 is liquid transmission. Respiratory transmission, aerosol transmission, and contact transmission falls under this category [13]. When the user wears the wristband and exercises the deadlift, the user will eventually sweat. It is highly recommended that this product is never shared with other people, and we will make sure to warn this to the user.

3.4 Future work

The location of wristband might be unrecognized by the camera if the barbell obscures its view. To fix this, we recommend replacing a wristband to an armband to change to position or changing the camera angle, both of which would require slightly different approaches to the software.

It would be much better to have this as a single device rather than two devices. This can be done by either implementing a smartphone app for the camera device or removing the wristband device and allowing the camera device to recognize the start of deadlift exercise without it.

References

- [1] J. Snape, "How To Deadlift," Coach, 18-Dec-2019. [Online]. Available: <https://www.coachmag.co.uk/barbell-exercises/3725/how-to-deadlift>. [Accessed: 03-Apr-2020].
- [2] John, "Physical Living: Health-First Fitness," Physical Living Health First Fitness, 23-Mar-2017. [Online]. Available: <https://physicalliving.com/how-much-should-i-be-able-to-deadlift/>. [Accessed: 03-Apr-2020].
- [3] V. Bengtsson, L. Berglund, and U. Aasa, "Narrative review of injuries in powerlifting with special reference to their association to the squat, bench press and deadlift," *BMJ open sport & exercise medicine*, 17-Jul-2018. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6059276/>. [Accessed: 03-Apr-2020].
- [4] M. MacGill, "Heart rate: What is a normal heart rate?" *Medical News Today*, 15-Nov-2017. [Online]. Available: <https://www.medicalnewstoday.com/articles/235710>. [Accessed: 03-Apr-2020].
- [5] "Heartbeats in Your Project, Lickety-Split," World Famous Electronics LLC. [Online]. Available: <https://pulsesensor.com/>. [Accessed: 17-Apr-2020].
- [6] V. Gupta, "Home," Learn OpenCV, 29-May-2018. [Online]. Available: <https://www.learnopencv.com/deep-learning-based-human-pose-estimation-using-opencv-cpp-python/>. [Accessed: 03-Apr-2020].
- [7] A. Rosebrock, "Raspberry Pi: Deep Learning Object Detection with OpenCV," *PyImageSearch*, 16-Oct-2017. [Online]. Available: www.pyimagesearch.com/2017/10/16/raspberry-pi-deep-learning-object-detection-with-opencv/. [Accessed: 17-Apr-2020].
- [8] Adrian Rosebrock, "Find distance from camera to object/marker using Python and OpenCV," *PyImageSearch*, 19-Jan-2015. [Online]. Available: <https://www.pyimagesearch.com/2015/01/19/find-distance-camera-objectmarker-using-python-opencv/>. [Accessed: 17-Apr-2020].
- [9] J. Mufuta, N. Egiebor, and S. Luo, "Deadlift Assistant," *ECE 445 Senior Design*, 10-Dec-2019. [Online]. Available: <https://courses.engr.illinois.edu/ece445/getfile.asp?id=16352>. [Accessed: 08-May-2020].
- [10] J. Mufuta, N. Egiebor, and S. Luo, "Deadlift Assistant Demonstration," YouTube, 12-Dec-2019. [Online]. Available: <https://www.youtube.com/embed/BWjsyt4PYm0>. [Accessed: 08-May-2020].
- [11] IEEE.org, "IEEE Code of Ethics," 2020. [Online]. Available: <http://www.ieee.org/about/corporate/governance/p7-8.html>. [Accessed: 03-Apr-2020].
- [12] "Safe Practice for Lead Acid and Lithium Batteries," University of Illinois ECE445 Course Staff, 13-Apr-2016. [Online]. Available:

<https://courses.engr.illinois.edu/ece445/documents/GeneralBatterySafety.pdf>. [Accessed: 03-Apr-2020].

- [13] T. Ghose, "How are people being infected with COVID-19." LiveScience, 07-Apr-2020. [Online]. Available: <https://www.livescience.com/how-covid-19-spreads-transmission-routes.html>. [Accessed: 17-Apr-2020].