Spring 2022

Enforcing a More Complex Notion of Proper Posture with Pose Estimation

Marco Alessandro Scanni
Bard College

Enforcing a More Complex Notion of Proper Posture with Pose Estimation

A Senior Project Submitted to

The Division of Science, Math, and Computing

of Bard College

by

Marco Scanni

Annandale-on-Hudson, New York

May, 2022

Abstract

Practicing poor posture for many hours throughout the day can lead to a myriad of both physical and mental morbidities. People tend to put their postures in jeopardy for multiple hours per day when using the computer. Pose estimation has the ability to track human positions in real-time with high accuracy and performance. The goal of this thesis is to provide evidence for what proper posture is, and enforce this proper posture through the use of a pose estimation model, namely BlazePose.

Contents

Dedication

I would like to dedicate this work to Alyssa, Duke and Joj for their endless love and support.

Acknowledgements

I would like to express gratitude to my Senior Project advisor, Professor Robert McGrail, for guiding me throughout the process of building this project and being a reliable and knowledgeable mentor and educator. I would also like to thank my siblings Giordana, Giuliana, and Antonio for remaining my strongest support system throughout my four years at Bard College.

1 Introduction

A complex notion of proper posture is seldom discussed in everyday life, however the implications of poor posture, which may also present itself as seemingly “good” posture, can be very diffuse and harmful. The second chapter evaluates in detail some of the overlooked, yet crucial biomechanics of proper posture, including: what anatomical structures are involved, how those structures are supposed to function, the maladies associated with their dysfunction, and how to restore proper function.

The literature shows that habits are the main culprit of poor posture, and that certain cues should be used to help restore proper posture. However, it may be difficult to understand the purpose of these cues when considering them, alone. Thus, the underlying biomechanics must be properly understood first, to not only understand how to execute the postural cues correctly, but also to understand why they are so paramount to our health. Moreover, intricately defining proper posture aims to help the user understand why it is important to enforce that notion not only throughout everyday life, but in the context of this thesis, when using the computer. The goal of this thesis is to try and connect this sophisticated definition with real-time enforcement through pose estimation.

Pose estimation takes an image, video, or live feed as input and gives as output the location of a person or an object within the frame. In the case of human pose estimation, this usually involves providing keypoints, which are landmarks that identify specific points on the person’s body, such as their eyes, neck, shoulders, knees, and so on. Enforcing proper postural cues show promise through utilizing the heights of and angles between specific keypoints. To help achieve this, OpenCV and NumPy are used alongside the Media Pipe framework, which gives access to their pose estimation solution, BlazePose. BlazePose is an extremely fast and lightweight pose estimation solution that has had success previously in being utilized for the purpose of health and fitness applications, and relies on single-person pose estimation to maintain its speed. These factors make the framework a fitting tool to be used to build the bridge between a complex notion of proper posture and real-time enforcement when on the computer.

2 Biomedical Background

2.1 How is "Posture" Defined?

Posture is a controversial topic in the medical community, and practitioners of all kinds; sports therapists, chiropractors, neurologists, and spine surgeons all may have different views on what it really means to have a healthy posture. Basic postural cues commonly prescribed by physical therapists and the like will be examined briefly in this chapter, and their validity in promoting overall functional health to the individual will be questioned.

The aim of this chapter is to intricately define "good posture;" more specifically, to define a posture that conforms with the neurological and vascular systems of the human body using scientifically backed, biomechanical logic. In terms of the structures that we tend to address when it comes to maintaining adequate posture while on the computer, or sitting in general, we will most closely examine the shoulders (more specifically the scapulae and their positioning) and the neck. These structures tend to show promise when identifying dysfunction based solely on visual signs, and therefore will also be highly useful in monitoring with a webcam.

2.2 Anatomical Explanations & Potential Postural Dysfunctions: The Scapula & Cervical Spine

The two structures that we will most closely examine are the scapulae (shoulder blades) and the cervical spine (neck). These two bodily structures are often discussed when regarding posture, as the terms "rounded shoulders" and Forward Head Posture (FHP) have become common diagnoses amongst practitioners and even everyday people. Research has shown that rounded shoulders is a result of faulty positioning and movement of the shoulder blades (scapular dyskinesis), and that FHP and rounded shoulders are highly correlated. In a study regarding sixteen subjects, who all worked at least 4 hours a day on the computer, and half of which reported some form of daily neck and shoulder pain, Szeto states that:

> *Symptomatic subjects also tended to have more protracted acromions compared with asymptomatic subjects and showed greater movement excursions in the head segment and the acromion. All subjects demonstrated an approximately 10% increase in forward head posture from their relaxed sitting postures when working with the computer display, but there were no significant changes in posture as a result of time-at-work* (Szeto et al., 2002).

It should be noted that the term “acromion” refers to the top and outer-most edge of the scapula. Moreover, Forward Head Posture (FHP) has been shown to be a compensatory mechanism of the body’s central nervous system as it attempts to maintain balance, given its altered center of gravity due to poor posture (Lee, 2016). Now, a more intricate look into the anatomy associated with these structures will be reviewed.

2.2.1 The Scapula: Anatomy

The importance of proper scapular positioning in posture is seldom mentioned, however it is paramount that we must understand why, from a biomechanical perspective, that we need to vanquish rounded shoulders. It can bring a plethora of associated maladies, especially over time with prolonged improper positioning. To make the following sections more comprehensible, some of the important anatomical structures surrounding and associated with the scapula will be reviewed, which will in turn help explain why the scapula itself plays a vital role in ensuring our posture does not conflict with some of the body's essential musculoskeletal, neurological, and vascular functions.

The scapula and clavicle (collar bone) are closely connected, as they make up the "shoulder girdle" (Figure 1). The upper trapezius muscle attaches to these two bony structures as well, and its job is to stabilize and allow for proper movement and positioning of the scapula, specifically. When the upper trapezius muscle becomes inhibited and functions improperly, (in the case of posture, through poor habits and lack of awareness), the muscle will inhibit, which will inevitably facilitate poor posture.

Figure 1

The Shoulder Girdle

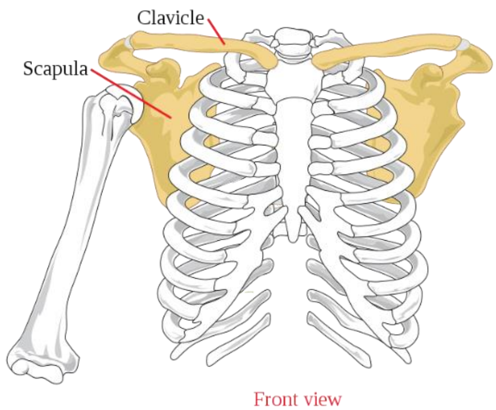

At the same time, lack of proper activation of the upper trapezius will cause the scapula to drop, shrinking the interval between the clavicle and the first rib, known as the costoclavicular space

(CCS). Within this space exists the neurovascular bundle made up of the brachial plexus and the subclavian artery and vein (Figure 2).

Figure 2

The Neurovascular Bundle Residing within the CCS

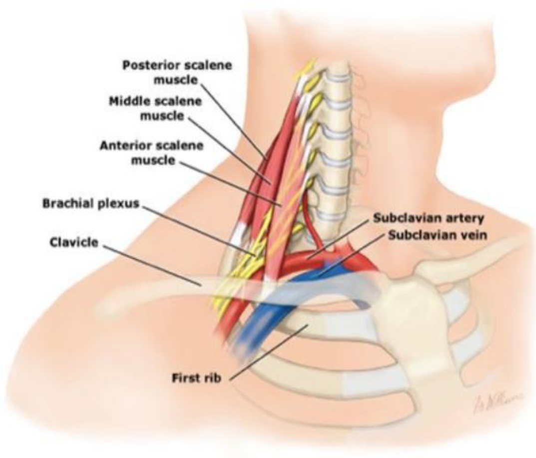

The brachial plexus is a network of nerves that runs into the arm, parts of the chest, neck and back, and supplies the arm with sensory and movement signals into the arms and hands. The brachial plexus, when faced with either muscular or osseous (bony) compression can be responsible for a myriad of symptoms such as weakness, tingling, pain in the fingers, wrists, arms, shoulders, neck, and between the shoulder blades. The vasculature that exists within the costoclavicular space, made up of the subclavian artery and vein, brings blood to and from the arm and shoulder, respectively. When this vasculature becomes compressed, that flow of blood has potential to be interrupted, also leading to a variety of different complications that will be discussed further in the next section.

A brief overview of the scapula and its surrounding anatomy shows that scapular movement and positioning is quite a bit more involved and important than one would expect, as poor scapular movement and positioning will lead to issues regarding compression within the costoclavicular space of vital anatomical structures. In the next subsection, what is considered a “poor” resting position of the scapula, the potential maladies associated with such positioning, and how we can alter our posture in order to avoid such maladies will be further explored.

2.2.2 The Scapula - Postural Dysfunctions: Causes, Symptoms, and Solutions

As mentioned prior, proper resting position and movement of the scapula is important for maintaining the functionality of the upper trapezius muscle, as well as the neurovascular bundle which sits within the costoclavicular space, namely the brachial plexus and the subclavian artery and vein. Although seldom mentioned when regarding posture, science has shown that muscular dysfunction within that region, namely of the upper trapezius, as well as compression of one or more of the structures within the costoclavicular space is not rare, and in fact quite common among those with scapular dyskinesis, more specifically and especially rounded shoulders.

Both neurological and vascular symptoms can manifest if the structures that sit within the costoclavicular space become compressed between the clavicle and first rib. This can occur through poor posture, when the scapula becomes anteriorly tilted and downwardly rotated, as in the case with “rounded shoulders.” However, compression within the CCS can also occur because of iatrogenic postural cues, such as “shoulders back and down.” Although this cue is commonly prescribed by various practitioners, it is in fact dangerous. It is a common but incorrect notion to pull the shoulders “back and down” for proper posture, as it puts the aforementioned neurovascular bundle at risk for bony compression (Watson et al., 2010). This cue will shove the clavicle back and downward, also narrowing the CCS, and thus also leaving the residing neurovascular bundle vulnerable to osseous compression (Figure 3).

Figure 3

Visualization of the Iatrogenic Cue of "Shoulders Back and Down"

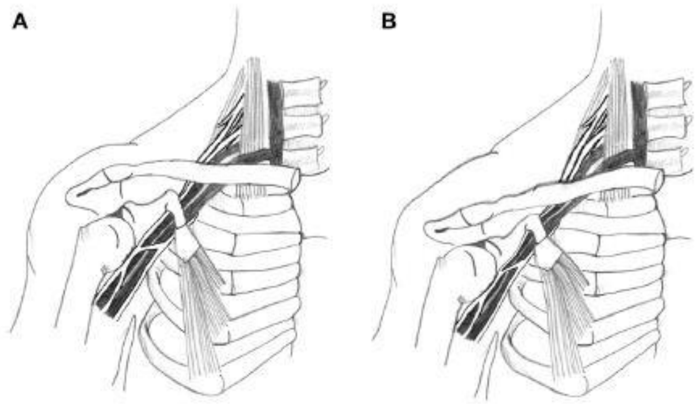

Note. 'A' represents the neurovascular bundle uncompressed by the clavicle. 'B' represents the neurovascular bundle becoming crushed by the clavicle after executing the iatrogenic cue.

In fact, although "shoulders back and down" is a popular postural cue for the scapula, it has instead been used in clinical settings as a provocative measure to *prove* osseous compression within the CCS, (i.e. it is meant to provoke symptoms and thus confirm CCS compression). This measure is commonly called "Halstead's Test", or the "military posture" stress test (Hixson et al., 2017).

If the underlying structures become compressed, this can cause an extremely wide array of symptoms, since the dysfunction can have both neurological and vascular implications such as chest, neck and shoulder pain, positional ischemia or venous insufficiency (Archie & Rigberg, 2017), which may contribute to headaches, migraines (Sina et al., 2017), intracranial hypertension (Flanagan, 2015), and more. A vast amount of research has been done on the costoclavicular space, and the symptoms can be very diffuse. To ameliorate posturally-induced CCS compression, Larsen states:

> *Thus, postural CCS compression may be ameliorated by raising the shoulders slightly, and staying there. This will also upwardly rotate the scapula, which is important for SD and SIS treatment. To correct scapular slouching, ask the patient to lift their acromion until the clavicle is elevated and superior scapular angle is approximately leveled with the T2 vertebra, and the scapula is in mild upwards rotation. The patient must be educated with regards to the etiology of costoclavicular space compression syndrome, so that they become sufficiently motivated to maintain their newly acquired posture. Further, he or she must learn to maintain adequate scapular height during glenohumeral articulation* (Larsen et al., 2020).

Addressing posturally-induced CCS compression is in essence quite simple: slightly raise the shoulders to increase the interval between the first rib and clavicle and hold that position.

2.2.3 The Cervical Spine – Anatomy

The cervical spine, or the neck, has many implications regarding posture. There are seven vertebrae of the spine that make up the cervical spine (neck), aptly named C1-C7. The neck plays an important role in ensuring that the head is properly situated. Similarly with the scapula, there are many neurological and vascular structures within the neck that are rarely mentioned when regarding proper posture, however they absolutely have the potential to be compromised through poor cervical posture, and thus impair their functionality and cause a wide variety of symptoms. The neck has many different muscular, neurological, and vascular structures that can become compromised due to poor posture and cause a wide variety of symptoms such as general neck pain due to improperly functioning musculature, as well as many symptoms of vascular nature.

Regarding musculature, one common group of muscles that become inhibited due to improper use, and thus weakness, is the suboccipital muscle group (Figure 4). The suboccipital muscle group allows us to extend the occiput (look up without collapsing at the upper neck). This muscle group is also extremely important to our vision, proprioception, and balance. Studies have shown that these muscles have extremely high spindle fiber densities:

> *The SOM has a high density of muscle spindle per gram. Muscle spindles per gram are found here between 98 (Rectus capitis posterior major - RCPma) and 242 (M. obliquus inferior), which is immense compared to the already well-structured hand muscles (M. opponens pollicis - 17 spindles / gram) [9] … The high number of muscle spindles of the SOM seems to be a prerequisite both for the function as a receptor, as well as an effector and for the interaction with various equilibrium and orientation systems* (Reinhardt, 2018).

Among other things, spindle fibers are highly important in mammals for regulating posture and overall movement control (Liu et al., 2003). This relation is also proven quite literally by the intimate connection between the eyes and the suboccipital muscles; anyone can test for themselves and put their fingertips on their suboccipital muscles (just at the base of the skull), move their eyes around, and feel their suboccipital muscles in action, attempting to ensure that the head is synchronized with these eye movements. Lack of this synchronization because of weak suboccipital muscles has been connected to instances of vertigo and motion sickness, such as in a car or on a boat, and strengthening of said muscles has been shown to ameliorate these symptoms.

Figure 4

The Suboccipital Muscle Group

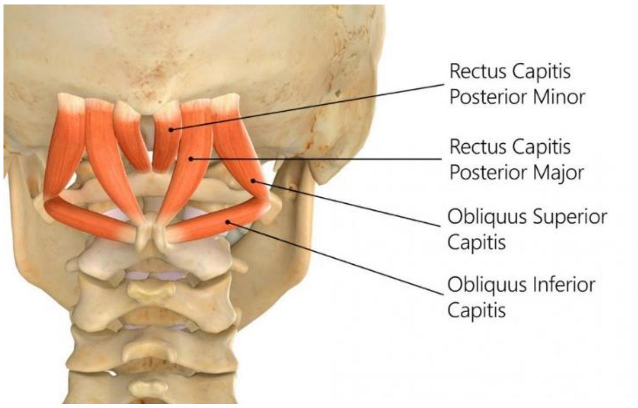

Another muscle group that lies within the neck are the deep neck flexors, namely the longus colli and longus capitis muscles (Figure 5). These muscles help to stabilize the C1, so as not to allow it to glide forwards, as well as preventing hyperextension of the cervical spine, or "hinging." The longus capitis extends all the way down from the skull to the C6 cervical vertebra, and the longus colli extends all the way from the C1 (atlas) all the way down to the T3 vertebra of the thoracic spine. Since they run throughout the entire neck, these two muscles essentially ensure that there is proper tensegrity throughout. People who hyperextend or "hinge" in the neck habitually will always find that these deep neck flexors are weak at the cervical level in which the hinge occurs.

Figure 5

The Deep Neck Flexors: The Longus Colli & Longus Capitis

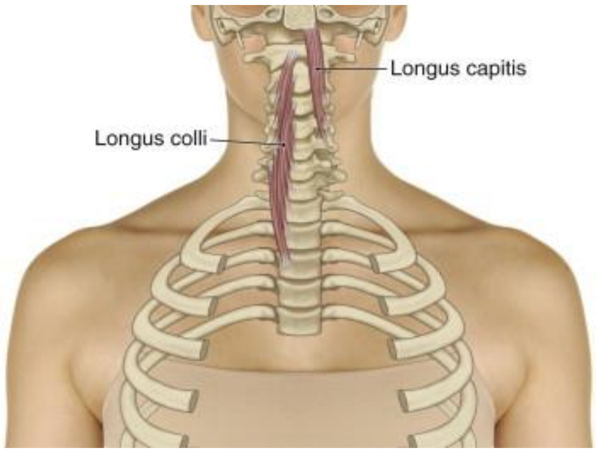

A relatively new discovery regarding these two muscle groups, the suboccipital group and the deep neck flexors, is the notion that vascular issues can arise due to inhibition of these muscles in tandem. Running just between the styloid process of the skull and the C1 vertebra (atlas bone) is the internal jugular vein (IJV), which circulates blood and cerebral spinal fluid (CSF) from the brain. This area, often referred to as the atlantostyloid interval (ASI) has the potential to narrow because of faulty upper cervical posture, which could cause various problems of vascular nature. Improper use of the suboccipital and deep neck flexor musculature will promote the narrowing of this region, potentially causing either the transverse process of the C1 to compress the IJV on its own, or for the IJV to become compressed between the transverse process of the C1 and the styloid process of the skull. Posturally-induced maladies related to these muscle groups will be explored in the next section.

2.2.4 The Cervical Spine - Postural Dysfunctions: Causes, Symptoms, and Solutions

As mentioned in the previous section, the suboccipital muscles as well as the deep neck flexors play an important role in ensuring proper cervical posture. One poor habit that many people tend to practice when on the computer as well as in daily life is hyperextension, or the cervical hinge. A cervical hinge occurs when one or more vertebrae collapses, causing one vertebra to flex, and one to hyperextend. This is a very common phenomenon, and, if it occurs at the lower cervical region for long enough, a hump may manifest at the base of the neck, named the “Dowager’s hump” or a “buffalo hump” (Figure 6). Cervical hinging can occur at any level of the neck, however. Now, the common upper cervical hinge will be examined more in-depth.

Figure 6

The “Dowager’s Hump,” as a Result of Lower Cervical Musculoskeletal Dysfunction

One of the associated maladies regarding dysfunction of the suboccipital muscle group is general musculoskeletal pain, which can manifest in the form of neck pain or headaches up and over the head, as well as in and around the eyes (Fernández-De-Las-Peñas et al., 2006). However, as mentioned before, there are more serious implications regarding dysfunction in the upper cervical region, the suboccipital muscles and the deep neck flexors, namely vascular issues. As mentioned in the previous section, if these two muscle groups become inhibited, in our case, through posture, the person will exhibit a "hinging" at the upper cervical level (Figure 7), wherein the C1 will come forward and potentially compress or even occlude one or both internal jugular veins, either on its own, or in tandem with the styloid process of the skull (Larsen, 2018). If the IJV becomes compressed or occluded in this way, the person may experience varying degrees of outlet insufficiency, a situation where the head cannot drain a regular amount of blood and CSF from the brain. With regards to this situation occurring due to poor posture, it most commonly manifests in the form of headaches, migraines, vision issues, concentration issues, fatigue, tinnitus among others.

Figure 7

Neck "Hinge" (Cervical Hyperextension)

The solution to cervical hyperextension is to become “long” in the neck, pulling the back of the head up toward the ceiling (Osar, 2012). However, because of the nature of how the styloid process of the skull is situated, there is still potential for IJV compression within the ASI, particularly if the styloid process is congenitally longer than average or posteriorly angulated (Li et al., 2019). Thus, the solution is to remain “long” in the neck while also slightly extending the atlanto-occipital junction. In other terms, regarding musculature, this combination of movements essentially means to adequately utilize both the suboccipital muscles and the deep neck flexors.

2.2.5 The Scapula and the Cervical Spine: Putting it All Together

As we have explored, the scapula and cervical spine have many implications regarding posture-associated maladies. These biomechanical complications can manifest in ways other than posture. For instance, CCS compression caused by faulty scapular mechanics may originate in cervical whiplash (Schenardi, 2005), clavicular fracture (Ishimaru et al., 2012), among other things. However, proper habits will, at the least, ensure that these maladies do not manifest through poor posture, and that the corresponding stabilizing muscles of the upper back and neck are properly used and remain strong and healthy.

Regarding scapular positioning, it must be made sure that adequate resting height of the shoulders has been achieved, and that they are not anteriorly nor posteriorly tilted but rather that they go “straight up.” A sound cue to ensure this is to raise the acromion rather than the “shoulder.” It has also been noted that the scapula should be raised “at least one centimeter higher than maximal depression” to decompress the costoclavicular space and adequately activate the upper trapezius muscle in posture.

Regarding the neck, one should aim to be "long in the neck" as previously mentioned, and to slightly extend the atlantooccipital junction (extend at the level just below the base of your skull, bringing the chin slightly and gently up). These two cues will ensure that the adequate posture stabilizing muscles are being properly utilized, namely the longus capitis, longus colli, and suboccipital group. This will ensure that the vertebrae of the cervical spine be at the least amount of risk for musculoskeletal pain, disk herniations, and vascular compression at the atlantostyloid interval.

General postural habits have been proven to worsen when on the computer. Practicing the aforementioned solutions will help combat this lack of postural attention and ensure that we are preserving our some of our body's most vital musculoskeletal, neurological and vascular systems. However, it is obvious that these findings regarding proper habits are a bit more complex than the average person would be aware of, if asked about what it means to have "proper posture." This thesis is meant to not only uncover and understand these more complex notions of proper posture, but also to see if it is viable to use pose estimation to enforce these notions when using the computer. Because the nature of resolving these postural dysfunctions is rooted mostly in habitual practices of the user, application of these resolutions through real-time pose estimation shows promise, and some applications have already had success on enforcing a much more general sense of posture (Chen, 2019). Next, whether this more complex notion of proper posture can be reinforced with pose estimation will be explored.

3 Pose Estimation

Pose estimation refers to capturing the positioning of bodies within images, videos, and live feeds, and has had a massive surge in use beginning in the late 2000s. It has been applied to various spheres such as human-computer interaction (Weidenbacher et al., 2006), detection of driver's attentiveness in driver assistance systems (Murphy-Chutorian et al., 2007), multi-person pedestrian tracking (Andriluka et al., 2010), person re-identification (Qian et al., 2018), animating images (Siarohin et al., 2019), various forms of physical fitness applications, etc. This chapter will serve as an introduction of the basics of different types of pose estimation and how they work on a high level, as well as an overview of just a few real-world applications of human pose estimation.

3.1 Uses for Pose Estimation

Human pose estimation aims to locate the human body parts and build human body representations through input data, i.e. images, videos, and live feeds. More recent implementations involve deep learning to obtain accurate, reliable results for human poses. Deep learning has essentially become the norm for human pose estimation because of its reliability. Deep learning as a basis for pose estimation and how it compares to other categories of pose estimation will be explored later in the chapter. As an introduction to the various complex uses of modern-day pose estimation, a few of the use cases for deep-learning-based human pose estimation will now be discussed.

3.1.1 Person Re-Identification

Person re-identification refers to properly identifying a person across various camera views. This problem can be challenging for a multitude of reasons. First, there are many factors that may vary independent of the person's identity, namely, the point-of-view of the camera, the person's body configuration (i.e. their pose) at the time of the captured frame, lighting, and the potential for partial occlusion.

Person re-identification methods usually involve a deep learning approach of combining identity-sensitive information (facial recognition, height, shape, clothing etc.) with a large, annotated dataset of a person's images. With this approach, features that are not discriminative of a person's identity, such as their pose, lighting, occlusion and so on are compensated for through brute-force data collection and annotation until the model has a chance to effectively learn what features are discriminative of each person and which are identity-insensitive. This approach lacks scalability in larger camera networks, though, as sufficient data must be collected and annotated from each camera view, and in an expansive network, such as at a mall or airport, this is too time consuming and difficult, and thus becomes infeasible. This specific deep learning approach also lacks generalizability, as it becomes difficult to redeploy a preexisting re-identification model to a brand-new network, (e.g. from one airport to another), as view points and poses tend to vary across these networks. Consequently, more data needs to be collected and annotated in order to revise the preexisting model.

To overcome the complications that previous person re-identification methods encounter, pose estimation has been used to separate the discriminative characteristics of a person from the non-discriminative ones. More specifically, pose estimation has been used to remove the influence of pose on a person's appearance. This is done through using any person's image along with a desired pose as input. The model would then synthesize a new image of the same person, but with their pose in the original image replaced with the desired pose. Qian et al. use this approach with 8 different generic poses as their "desired pose" input, effectively expanding their training dataset by 8 times per every new identity. This method thus not only requires less data collection and annotation than generic approaches to person re-identification, but it also handles the problem of the complementary features associated with identity-sensitive and camera view-sensitive factors of a particular person. For instance, high heels may influence a person's pose, which would generate an instance of interconnectedness between a discriminative characteristic (the type, color, and texture of a person's shoes), with an identity-insensitive, but view-sensitive characteristic (their pose). This is an example of a case that the generic person re-identification approach would have to overcome through more data, as it attempts to discriminate which aspects are discriminatory. On the other hand, the pose estimation approach can effectively and efficiently overcome this case without such measures, and instead learn the complementary features that stem from the original and pose-normalized images.

This pose-invariant method of person re-identification allows models to decouple pose, the proposed main culprit for preventing learning of identity-sensitive, camera view-insensitive features, from the other non-discriminatory characteristics of a particular person. Normalizing poses in this manner not only improves the generic model by eliminating body configuration from the pool of non-discriminative characteristics, but it also solves the problems of scalability and generalization that the generic approaches encountered.

3.1.2 Safety in Driver Assistance Systems

Pose estimation has been used to further the utility of driver assistance technologies. Consider modern technology in cars that alert the driver, when parking or when in reverse, when the vehicle is within close proximity to an object or person. The car uses sensors to detect the distance between the car and the object/person, and once it is close enough, the car may play an audible alert to indicate that the car is dangerously close to said object/person.

An issue that may arise from this technology is that it is either always in operation, even if the driver does not need it, or it is never in operation. For instance, when driving behind a separate vehicle, the car may alert the driver if its sensors indicate that it may be travelling at a speed that could lead to a collision with the vehicle in front of them. Despite the driver's level of attentiveness, (i.e., even if the driver is in full control of the situation and is certain they will not rear end the car in front of them), their car will still alert them that there is potential for a collision given a certain set of variables (speed, distance from car in front etc.).

Human pose estimation has been used as a means of solving this problem. More specifically, research has been done to see whether the orientation of the driver's head could be used as another variable to help decide whether the car should utilize these assistance systems:

> *Instead, if the automobile were able to recognize when a driver had not noticed a possible hazard, it could trigger an alert only when these dangerous situations arise. This fusion of interior and exterior observations [of the driver's vehicle] comprises an important new paradigm in advanced vehicular safety" (*Murphy-Chutorian et al., 2006).

The paper proposes that the orientation of the head is a valid indicator of field of view and attention of the driver. Although more work needs to be done, combining human pose estimation with car safety spreads the focus from solely obstacles outside of the vehicle towards all important aspects of the environment, namely the surroundings of the vehicle in combination with the driver themselves, paving the way for more advanced technologies in the realm of not only driver assistance systems, but self-driving vehicles as well.

3.2 Types of Pose Estimation

3.2.1 Model-Based and Deep Learning-Based Methods

The two main methods of estimating poses are through models and deep learning. Model-based methods focus solely on the relationship between body parts themselves, as opposed to precise locality of each individual keypoint. To represent these relationships, an explicit hand-crafted model is created to estimate the human poses (Andriluka et al., 2019). This method of pose estimation disregards specific keypoints, and thus it is most useful for whole-person detection within stiff boundaries of freedom (Figure 8).

Figure 8

Model-Based Human Pose Estimation

The deep-learning method of pose estimation involves learning the mapping from an inputted image/video to joint coordinates by explicit models, similar to the model-based approach. However, the difference lies in how these explicit models are derived. Contrary to the model-based approach, the models in the deep-learning approach implicitly capture the relations between body parts from a large number of annotated samples. Therefore, the deep-learning method involves a lot of training data and has a much higher computational cost when compared to the more traditional model-based approach, but also severely outperforms the traditional method.

3.2.2 Top-Down & Bottom-Up Methods

Among the deep learning-based methods of pose estimation, there are two different, two-step approaches to detecting and tracking poses: top-down and bottom-up. The top-down method of pose estimation involves first determining the area that a person takes up within the frame using a detector. Once this area is determined, it goes through a data-processing stage in which that area is turned into a cropped image of the person, and then the pose estimator, which tracks the keypoints, is ran on top of that image. This approach may suffer from runtime issues, as the approach is proportional to the number of people within the frame: for each detection of a person within the frame, they must be cropped and then another single-person pose tracker needs to be ran (Cao et al., 2016). This issue, of course, is trivial in the context of using a top-down approach for single-person pose estimation.

Bottom-up methods, then, are highly sought after in the case of multi-person pose estimation, as there is much greater potential for separation of multi-person tracking and runtime complexity. The most common bottom-up approach uses Part-Affinity Fields (PAFs) to detect the location and orientation of limbs within the frame. First, confidence maps are used to estimate the likelihood of the positioning of a certain body part within the frame. At the same time, body parts are then associated with one another using PAFs to form limbs. Finally, limbs are correctly matched such that each limb belongs to one skeleton until a full skeleton is formed.

4 BlazePose: A Media Pipe Pose Estimation Solution

Media Pipe is a library with various machine learning solutions, one of which is BlazePose. BlazePose is a pose estimation model that has a multitude of uses, such as augmented reality, sign language recognition, full-body gesture control, and quantification of physical exercises. This wide variety allows for various cases within the realm of physical health and fitness, like recognizing yoga poses and counting exercise repetitions (Bazarevsky et al., 2020). We will be implementing our posture enforcement tool using BlazePose.

Because of BlazePose's high speed and accessibility, it is highly viable in the sphere of real-time fitness and physical health applications. The speed and low computational power stems from its reliance on assuming that there is just one person in the inputted frame, allowing it to use the previously mentioned "top-down" method of deep-learning human pose estimation (it first detects the bounds of a person within the frame, then subsequently estimates their keypoints). This assumption allows for a lower network capacity, allowing BlazePose to preserve a quality of pose estimation not unlike more popular, but more computation-heavy frameworks, such as OpenPose. To reiterate, BlazePose is able to achieve the same quality of pose estimation with reliable speed and less computational power than OpenPose because the latter has the capacity for multi-person pose estimation; BlazePose does not.

4.1 BlazePose's Pipeline for Detection

The BlazePose framework utilizes a unique pipeline to consistently detect the human body, splitting it into two parts. First, the framework uses a detector to track the face properly. Then, once the face is recognized, essential virtual keypoints are predicted to detect the center, size, and rotation of the entire area that the human occupies within the frame; one landmark at the midpoint of the hips, a circle circumscribing the whole person, and one landmark that extends through the midline of the body, respectively. BlazePose uses these essential elements; a bounding box of the face, along with the two predicted virtual keypoints of alignment as a radius of a circumscribing circle, to consistently and accurately predict the three aforementioned characteristics of a human within the frame.

Regarding the face detector alone, the pipeline flows as follows: a detector first locates and crops a region-of-interest (ROI) within the given frame, in essence, creating a bounding box around the face. In the case of videos and live feeds, the face detector is run only on the first frame, and then the ROI of the face is derived from the location of the previous frame's keypoints. In other words, the area that the face resides in within the frame only needs to be detected once. Then, the estimated keypoints of the face are used to maintain the ROI for the rest of the video/live feed (Figure 9). This method of detecting and forming the bounding box only once, and then subsequently using it as a basis for predicting future ROIs is highly efficient, as it allows the model to dedicate most of its capacity towards accuracy of keypoint predictions, as opposed to data augmentation. BlazePose utilizes a face-recognition model, BlazeFace, a previously trained face detector, as a proxy for this specific step in the human pose estimation pipeline.

Figure 9

Detector-Tracker Pipeline

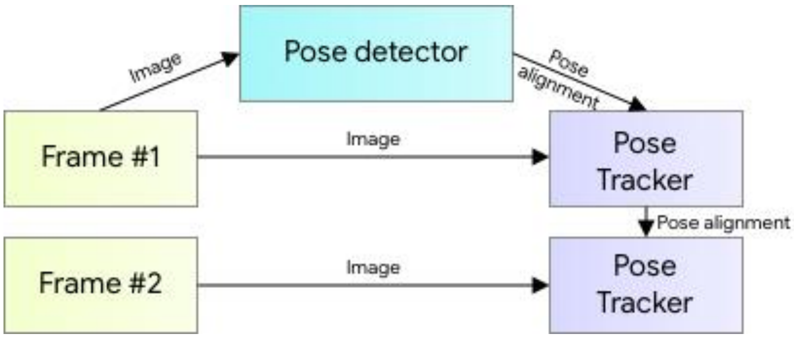

Next, to detect the rest of the body, BlazePose predicts the alignment, rotation, and size of the human. It infers that a human's head should be visible if their torso is visible. The reason this assumption is considered valid and useful is because the most significant indicator to the model about the position of a human's torso is their face. The face is mostly rigid, highly recognizable, has small variations in appearance and high-contrast features (i.e. it is much simpler to detect that a face is indeed a face, as opposed to, say, defining a type of limb), making it an excellent marker for detecting a human. In the case of posture estimation, specifically, this is also a very fitting assumption. In simplest terms: in the eyes of the program, once a head is found, a body will follow.

From this assumption, additional parameters are predicted: The middle point between the hips, the radius of a circle that would circumscribe the person, and the angle between this hip midpoint and the midpoint of the shoulders, which would extend out to complete the radius of the person. The aforementioned face bounding box, along with these three additional parameters, would complete the ROI of the entire person within the frame. These size, center, and rotation landmarks can be best visualized by Davinci's "Vitruvian Man," which was the inspiration for this method of determining the ROI (Figure 10).

Figure 10

"Vitruvian Man" Inspired Bounding Box

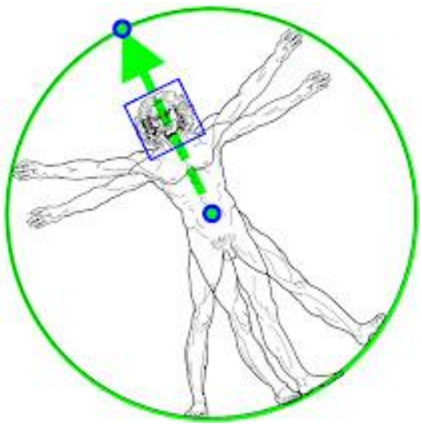

This is but a brief overview of the detector that BlazePose uses to consistently predict the size and orientation of a human within the frame. This thesis uses the technologies associated with pose estimation as a "black box." In other words, the focus is primarily on providing an input (a live video feed via webcam) and receiving output (keypoint locations to be used for the sake of posture enforcement). The inner workings of how we receive these keypoints as output are not a primary concern for this thesis, and thus will not be discussed in further detail. Now, enforcing our novel approach to "good" posture using BlazePose will be discussed.

5 Posture Estimation

Posture estimation can be considered a subset area of research that utilizes human pose estimation to monitor or enforce posture. This approach to exploring how users can use human pose estimation to their benefit for posture has been proven useful in various instances, from evaluating general slouched posture in the classroom (K. Chen, 2019) to improving exercise form when weightlifting (S. Chen & Yang, 2020). However, considering the relatively young age of the literature surrounding our new, more complex approach to proper posture, a posture correction tool to enforce it has yet to be seen. Therefore, it will be attempted here.

5.1 The Approach

As mentioned in Chapter 2, there are a few postural cues that are often overlooked and habitually ignored, especially while on the computer. The goal of this chapter is to evaluate whether each of these postural cues can be reliably enforced in real-time, by making use of the heights and angles between the keypoints that the BlazePose framework provides with a standard webcam.

Regarding the specific biomechanical structures and patterns that have been explored previously in this thesis, all of them share a specific set of characteristics that make them suitable to be enforced in the manner that has been explored within this project:

- All postural dysfunctions can be corrected through adequate habitual awareness and understanding of the postural cues.
- All postural dysfunctions can be reliably evaluated through visual hints.
- All postural dysfunctions follow a pattern of being related to the upper body.

With these characteristics in mind, the complex notion of proper posture that has been explored previously in this thesis will be enforced with keypoints. More specifically, the correct postural cues to facilitate proper posture, and thus avoid overall atrophy of certain musculature and compression of neurovascular bundles, will be enforced through the use of heights and the angles between the keypoints relevant to the areas of the scapula and cervical spine.

5.1.1 Item 1: Adequate Habitual Awareness and Understanding of Postural Cues

The first characteristic is an important one to ponder, as it changes the way most practitioners would view posture rehabilitation. Rather than relying on exercises, the literature shows that lack of proper habits and awareness is the largest component in postural dysfunction. Therefore, there must be some degree of reliance on the user to understand what good posture entails. For this reason, the program relies on the user to understand the proper postural cues adequately, and set the thresholds for proper and improper positioning themselves within the webcam's frame. This approach's validity is furthered when considering the 2D plane that BlazePose is utilized on, meaning that if the user decides to sit elsewhere in the frame, they are able to recalibrate their thresholds to accommodate for the two-dimensional space.

5.1.2 Item 2: Evaluation through Visual Hints

Regarding visual hints, the literature shows us that identifying these postural dysfunctions is not necessarily easy, but it is simple. For scapular positioning, recent literature has shown that levelling the superior angle of the scapula at least as high as the T2 thoracic vertebrae is considered an adequate height. This is a difficult notion to enforce through monitoring the upper body with a two-dimensional camera frame. However, numerous case studies have shown that this more nuanced cue can be achieved by putting the user into maximal depression of the scapula (pushing the shoulders all the way down), and then raising them at least 1 centimeter from that point. This is very reliably enforced through visual hints, as it becomes a matter of taking a snapshot of the height of the scapula when in maximal depression, and then ensuring that it remains a bit above that threshold.

Regarding the avoidance of a cervical hinge, the user must be cued to bring the back of their neck up towards the ceiling. Studies have shown that a cervical hinge can be detected based upon the positioning of the transverse processes of the C1 vertebra in relation to the bimastoid line(Larsen, 2018). In simpler terms, if a vertebra in the neck is forward in relation to the part of the skull that sits just below the ear, then that part of the neck is hinging. Although we do not have the ability to evaluate the positioning of the vertebrae within the neck externally, we can rely on the notion that once the hinge is vanquished, the eyes will become level with the horizon (Osar, 2012). Therefore, our approach is similar to the scapular correction approach. To enforce the user to get out of the hinge by bringing the back of their neck up toward the ceiling in real-time, they must execute the cue correctly, and then a snapshot of the angle between the eye and the ear will be taken. The angle should be roughly 90 degrees. A significant deviation over that (eyes falling below the ears) would indicate an overly flexed neck, and a significant deviation under that (eyes hanging above the ears) would indicate that the user is hinging (Figure 11).

Figure 11

Keypoint Relations of Neck Flexing and Hinging

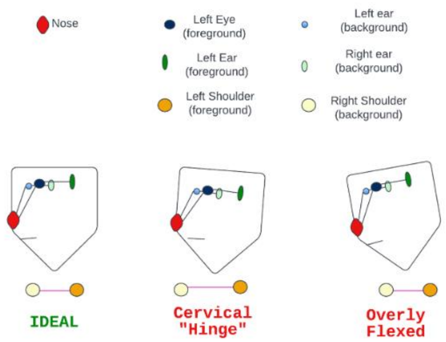

Lastly, to avoid compression of the IJV by the styloid process of the skull, the user needs to gently raise their chin without going back into a hinging position, thus increasing the atlantostyloid interval. This can be reliably enforced by simply vanquishing the neck hinge first, and then ensuring that the previously measured angle between the eye and the ear remains relatively stable while the chin comes up.

5.1.3 Item 3: Relation to the Upper Body

All of the explored postural dysfunctions and their corresponding corrective cues are related to the upper body. This allows for a webcam to be used adjacent to the computer monitor that the user is working on, as nothing below the upper body needs to be detected. This means that no extra special webcam setup is required to properly use the program. On top of this, BlazePose is lightweight, and can run on most newer generation smartphones, and so this characteristic allows for a potential smartphone port in the future to be used as an alternative to an external webcam.

5.2 Combining Biomechanics with Pose Estimation

5.2.1 Scapular Resting Height & Positioning

As mentioned in Chapter 2, the scapula has a particular proper resting height and positioning that must be practiced to maintain the health of the associated musculoskeletal and neurovascular structures. To alleviate any maladies associated with these structures, it must be ensured that the scapula is lifted at least one centimeter higher than maximal depression. It must also be ensured that the shoulders go “straight up,” and that elevation is not compensated for by tilting, most commonly anteriorly (forwardly). Combatting this compensation is as simple as cueing the user to elevate the acromion (the outer edge of the scapula), as opposed to just the shoulder.

The shoulder keypoints should be considered when addressing the scapula, since they are the closest indicators of acromial height. Since it has been shown that at least one centimeter higher than maximal depression is ideal, the user is given a countdown and cued to put themselves in maximal depression of the scapula (Figure 12). This height will be recorded by the program, and then the user will be cued when the height of their scapula is high enough to meet this threshold (Figure 13).

Figure 12

Scapular Positioning: Calibration Step

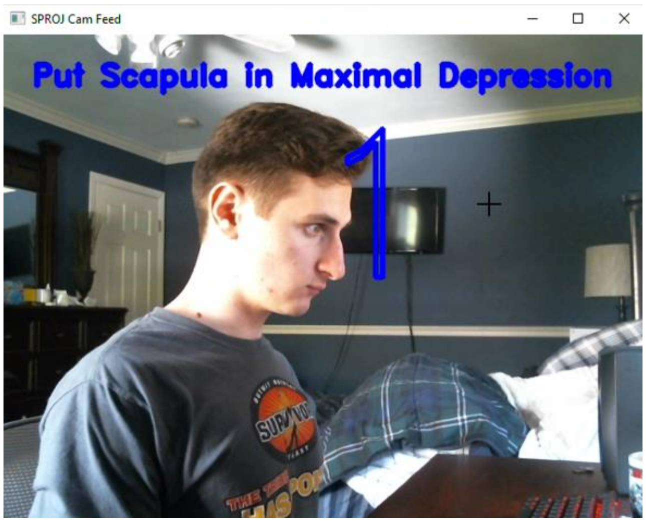

Figure 13

Scapular Positioning: Adequate Measurement

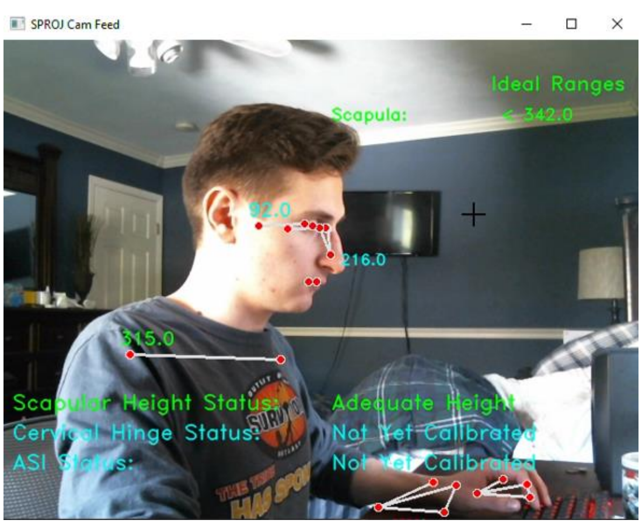

5.2.2 Vanquishing Hinge-Neck Posture

After the scapula has been properly situated, taking care of hinge-neck posture is the next step in attaining adequate posture. As mentioned prior, promoting the user to eliminate their hinge-neck posture first involves getting the user to execute the proper postural cues, and then taking a snapshot of the angle most relevant to the ideal positioning of the neck. In this instance, this involves taking the angle between the eye and ear, with a completely level measurement being 90 degrees. The user is given a countdown and is cued to raise the back of their neck toward the ceiling (Figure 14). If the cue has been executed properly, a level measurement should be achieved (Figure 15). Slight deviations of this measurement do not indicate any severe improper neck posture, however a deviation of larger than 10 degrees from the calibrated measurement would indicate for certain a reappearance of neck hinging or an overly flexed neck, depending on the direction of the deviation (Figure 16, 17).

Figure 14

Neck Positioning (Hinge): Calibration Step

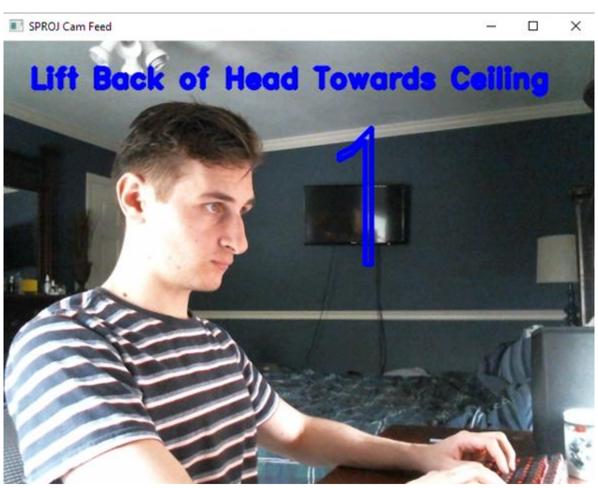

Figure 15

Neck Positioning: Adequate Positioning for Hinge Avoidance

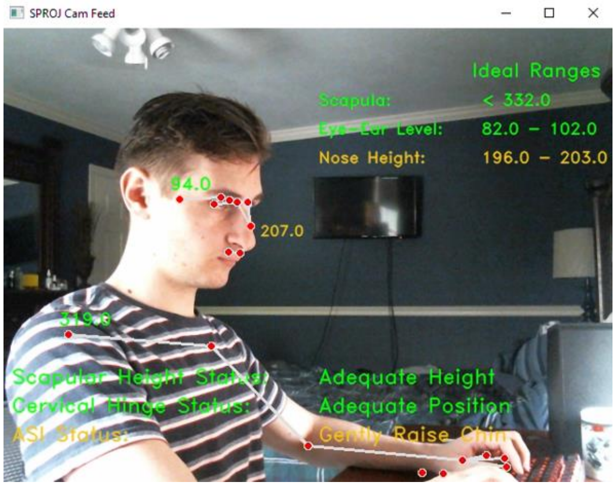

Note. The ASI has not yet been accounted for in this instance. As per the cue, once the hinge has been vanquished, the chin must be gently raised to account for the ASI.

Figure 16

Neck Positioning: Hinging

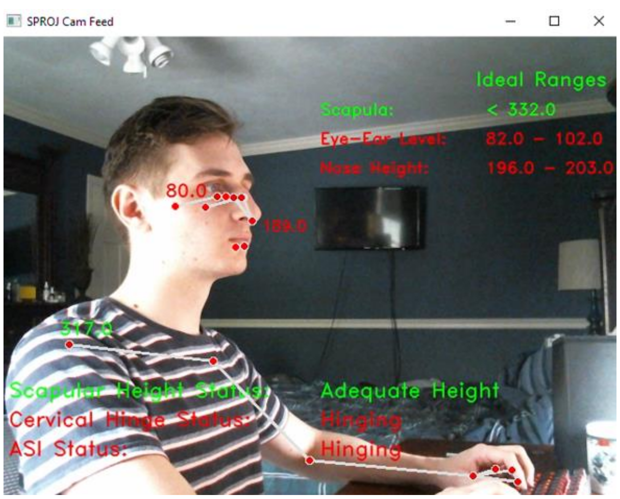

Figure 17

Neck Positioning: Overly Flexing (Tucking the Chin)

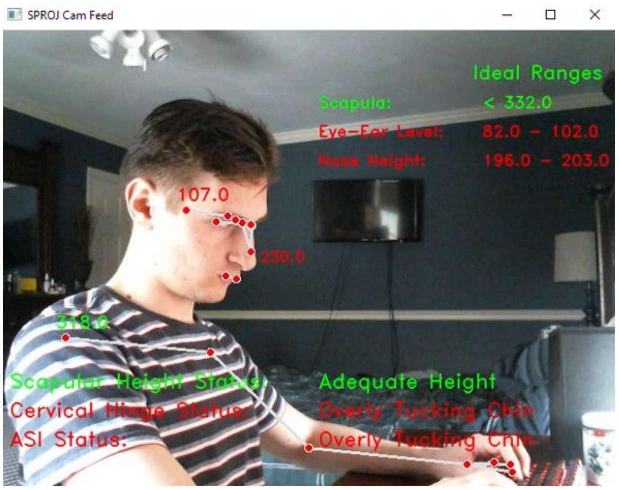

5.2.3 Ensuring an Adequate Atlantostyloid Interval

Ensuring that the ASI has an adequate interval so as not to compress the IJV between the transverse process of the C1 and the styloid process of the skull, involves gently raising the chin without hyperextending the neck (reintroducing a cervical hinge). Biomechanically speaking, this task is the most difficult to understand and execute properly, as it involves extending the atlanto-occipital junction without “cheating” through extending other levels of the cervical spine. In simplest terms, properly executing the cue entails adequately utilizing the small neck muscles that reside just below the base of the skull (the suboccipital muscle group), only.

In terms of the program, a correct execution of this postural cue would, assuming the previous cue was executed correctly as well, be realized through the level measurement between the eye and ear remaining within its targeted threshold, and the chin being gently raised (Figure 18). To reiterate, this is what the cue *looks like*, but not necessarily all that needs to be done to ensure that the cue is executed properly. In essence, the head should be extended while the neck remains nearly still. For a highly nuanced cue, the program does a fine job at detecting whether it is done properly. However, proper proprioception of the responsible musculature, namely the suboccipital muscle group, is much more reliable for executing this cue properly.

Figure 18

Neck Positioning: Adequate ASI & Adequate Positioning for Hinge Avoidance

5.2.4 All Cues in Tandem

The program is able to detect when the user is currently executing all three cues adequately and can also detect when this is not the case. Each cue has a stage, color, and ideal range associated with them. The ideal range is generated during the calibration process of each cue. Each cue's associated stage indicates whether the cue has yet to be calibrated by the user, or if it has, whether the user is currently within the ideal threshold of the related measurement. The stage of each cue indicates which side of the threshold the user is lacking in (e.g., in the case of cervical hinge, whether the user is overly flexing or overly extending their neck), if at all, and gives real-time cues to promote the user to readjust themselves. The color associated with each cue serves a similar role to its corresponding stage, with different colors indicating proper or improper execution of said cue, and if the latter, to what degree (Figure 19-22). The program works exceptionally well under the assumptions that the user has adequate knowledge of the biomechanical background associated with each cue, as well as how each postural cue should be executed. This new notion of proper posture creates a circumstance in which the user must be properly educated on the logic related to the movement and positioning of the associated anatomical structures, as previous notions of proper posture usually entail basic cues that can be executed without much understanding on the biomechanics behind them (e.g., put the shoulders back and down, tuck the chin down etc.).

Figure 19

Adequate Scapular Height but Slightly Hinging at the Neck

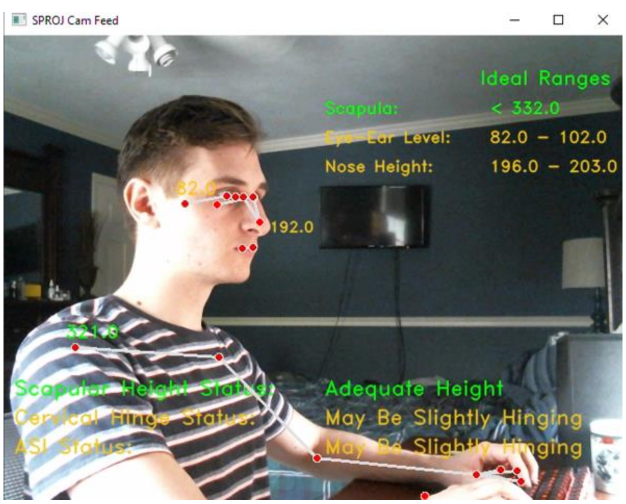

Figure 20

Scapula too Low, but Adequate Neck Positioning for Hinge Avoidance

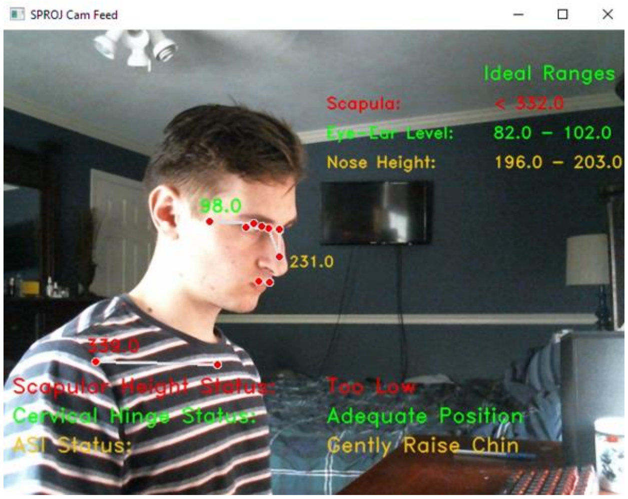

Figure 21

Scapula too Low, Hinging at the Neck

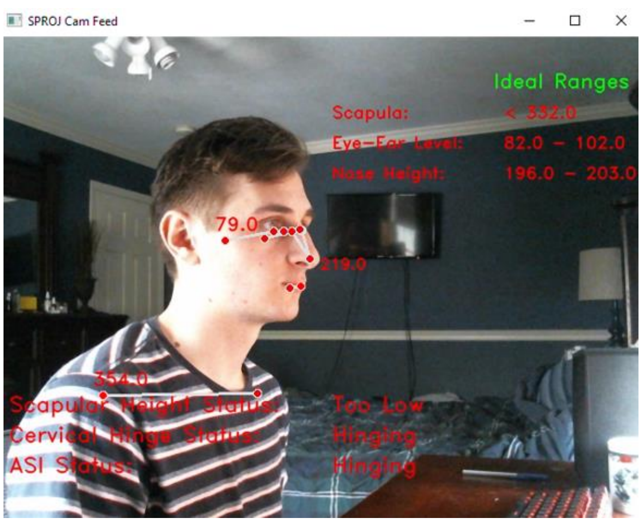

Figure 22

Scapula too Low, Overly Flexing Neck (Tucking the Chin)

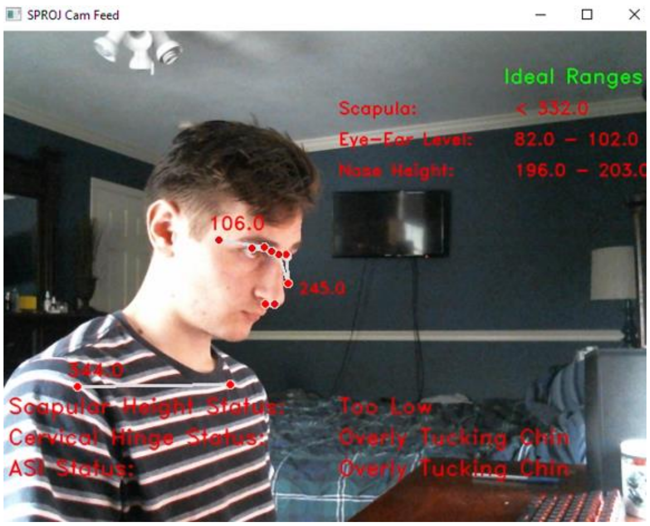

6 Conclusion

The medical literature shows that proper posture entails a much greater understanding of our anatomy than one would initially assume. Viewing proper posture as a way of adequately conforming to the neurological and vascular structures of the human body is a much more reliable, albeit more complicated method of ensuring that our posture does not negatively affect our health, as opposed to solely relying on simple visual markers, such as using the presence or absence of a slouch as an indicator of the quality of one's posture. However, detecting potentially dangerous postural habits can be done through visual cues. Utilizing a pose estimation framework has proven to be an effective way of doing so.

BlazePose is lightweight and fast, making it a very useful tool for implementing a real-time body tracker of any kind. Utilizing BlazePose to capture the relevant keypoints of the user in real-time to use as markers for specific cues is a valid approach to enforcing good posture, as much of the logic behind these more nuanced cues involve heights and angles of and between certain anatomical structures. For instance, maintaining proper scapular height is an important aspect of posture to ensure that the underlying neurovascular bundle does not become compressed between the clavicle and first rib. In essence, it is a simple solution to a complex problem; lift the scapula up at least 1 cm higher than maximal depression to avoid osseous compression. This is a problem that can be reliably tackled by noting the height of the scapula when in maximal depression, and then ensuring that the scapula rests just above that height.

The end result is a program that can reliably enforce a more complex notion of proper posture, but under the assumption that the user has adequate knowledge of the biomechanics associated with this notion. The angles and heights of specific keypoints within the frame can enforce the user to practice proper posture, but the user must be able to complete the calibration process with a comprehensive understanding of what the calibration step and correct execution of how the postural cues should look and feel. Overall, the program serves as a reliable baseline for taking a step in a new direction of reinforcing a seemingly unexplored, more complex notion of proper posture through pose estimation.

6.1 Future Work

The future of this project sees paths towards making the program an overall more useful tool. The program in its current state serves more as an educational tool than a practical one. It, in tandem with proper biomechanical understandings, serves to help the user learn about the relationships between the correct postural cues as well as the overall reasoning and practicality of the cues themselves. Transitioning this posture enforcement tool away from the educational and more toward the practical would involve implementing methods of allowing the user to utilize the program without having to constantly monitor it themselves. In its current state, the user is notified that they are falling out of proper posture through specific visual markers, such as statements and color changes. This means that the user cannot reliably remain in good posture unless they are monitoring the augmented camera feed themselves. An example of transitioning away from this approach would be to utilize sound cues to allow the user to be mindful of their posture while not having to have the posture enforcement tool window in focus.

Taking steps towards a more streamlined, easy-to-use application would also involve using a method of training the application to recognize proper posture on its own, or with minimal calibration. Ideally, the program should be autonomous and not require calibration to the extent that the current one does. Currently, the frequency of recalibration has the potential to be high in our scenario, where keypoint locations are measured in pixels, and these measurements are relied on within a 2D frame. This scenario leaves potential for recalibration every time the webcam moves, or the user adjusts themselves. Streamlining the calibration aspect of the program could help every-day users utilize the program more efficiently, especially if they find the biomechanics explored in this thesis too difficult to understand. However, considering the complexity of the biomechanical background in contrast to how most every-day users understand “good” posture, enforcing a notion of proper posture with a complexity of this magnitude will always require at least some form of education.

Lastly, the pose estimation framework used in the program, BlazePose, is lightweight enough to be run on the latest smartphones. Thus, a smartphone port of the application could greatly heighten the usability of the posture reinforcement tool, as it currently relies on a lateral view of the user, and thus an external webcam is required. Using a smartphone as a standalone device to run the program would solve an issue that many potential users may face, in that they may not have access to an external webcam. A smartphone port would also open the door to a more extensive approach to posture enforcement, involving the entire body, and not just the upper body, as the smartphone would not have the limitation of being connected to the user’s computer. Thus, more flexible camera viewpoints could be achieved.

Bibliography

Andriluka, M., Roth, S., & Schiele, B. (2010). *Monocular 3D Pose Estimation and Tracking by Detection*.

Andriluka, M., Roth, S., & Schiele, B. (2019). *Pictorial Structures Revisited: People Detection and Articulated Pose Estimation*.

Archie, M., & Rigberg, D. (2017). Vascular TOS—Creating a Protocol and Sticking to It. *Diagnostics*, *7*(2), 34. https://doi.org/10.3390/diagnostics7020034

Bazarevsky, V., Grishchenko, I., Raveendran, K., Zhu, T., Zhang, F., & Grundmann, M. (2020). *BlazePose: On-device Real-time Body Pose tracking*.

Cao, Z., Simon, T., Wei, S.-E., & Sheikh, Y. (2016). *Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields*. http://arxiv.org/abs/1611.08050

Chen, K. (2019). Sitting Posture Recognition Based on OpenPose. *IOP Conference Series: Materials Science and Engineering*, *677*(3). https://doi.org/10.1088/1757-899X/677/3/032057

Chen, S., & Yang, R. R. (2020). *Pose Trainer: Correcting Exercise Posture using Pose Estimation*. http://arxiv.org/abs/2006.11718

Fernández-De-Las-Peñas, C., Alonso-Blanco, C., Cuadrado, M. L., Gerwin, R. D., & Pareja, J. A. (2006). Trigger points in the suboccipital muscles and forward head posture in tension-type headache. *Headache*, *46*(3), 454–460. https://doi.org/10.1111/j.1526-4610.2006.00288.x

Flanagan, M. F. (2015). The role of the craniocervical junction in craniospinal hydrodynamics and neurodegenerative conditions. In *Neurology Research International* (Vol. 2015). Hindawi Publishing Corporation. https://doi.org/10.1155/2015/794829

Hixson, K. M., Horris, H. B., McLeod, T. C. V., & Bacon, C. E. W. (2017). The diagnostic accuracy of clinical diagnostic tests for thoracic outlet syndrome. *Journal of Sport Rehabilitation*, *26*(5), 459–465. https://doi.org/10.1123/jsr.2016-0051

Ishimaru, D., Ogawa, H., Sumi, H., & Shimizu, K. (2012). Late Thoracic Outlet Syndrome after Clavicle Fractures in Patients with Multiple Trauma: A Pitfall of Conservative Treatment. *Open Journal of Orthopedics*, *02*(03), 90–93. https://doi.org/10.4236/ojo.2012.23018

Larsen, K. (2018). Occult intracranial hypertension as a sequela of biomechanical internal jugular vein stenosis: A case report CASE REPORT ANAESTHESIA, PAIN & INTENSIVE CARE. In *PAIN & INTENSIVE CARE* (Vol. 22, Issue 2). www.apicareonline.com

Larsen, K., Galluccio, F. C., & Chand, S. K. (2020). Does thoracic outlet syndrome cause cerebrovascular hyperperfusion? Diagnostic markers for occult craniovascular congestion.

Anaesthesia, Pain and Intensive Care, 24(1), 69–86. https://doi.org/10.35975/apic.v24i1.1230

Lee, J.-H. (2016). *Effects of forward head posture on static and dynamic balance control*.

Li, M., Sun, Y., Chan, C. C., Fan, C., Ji, X., & Meng, R. (2019). Internal jugular vein stenosis associated with elongated styloid process: Five case reports and literature review. In *BMC Neurology* (Vol. 19, Issue 1). BioMed Central Ltd. https://doi.org/10.1186/s12883-019-1344-0

Liu, J.-X., Thornell, L.-E., & Pedrosa-Domellöf, F. (2003). Muscle Spindles in the Deep Muscles of the Human Neck: A Morphological and Immunocytochemical Study. In *The Journal of Histochemistry & Cytochemistry* (Vol. 51, Issue 2). http://www.jhc.org

Murphy-Chutorian, E., Doshi, A., & Trivedi, M. M. (2006). *Head Pose Estimation for Driver Assistance Systems: A Robust Algorithm and Experimental Evaluation*.

Osar, E. (2012). *Corrective Exercise Solutions to Common Hip and Shoulder Dysfunction* (1st Edition). Lotus Pub.

Qian, X., Fu, Y., Xiang, T., Wang, W., Qiu, J., Wu, Y., Jiang, Y.-G., & Xue, X. (2017). *Pose-Normalized Image Generation for Person Re-identification*. http://arxiv.org/abs/1712.02225

Reinhardt, R. (2018). Suboccipital Musculature - Morphology, Functions and Variants - An Update. *Journal of Complementary Medicine & Alternative Healthcare, 7*(5). https://doi.org/10.19080/jcmah.2018.07.555721

Schenardi, C. (2005). Whiplash injury. TOS and double crush syndrome. Forensic medical aspects. In *Acta Neurochir* (Vol. 92). Springer-Verlag. http://www.tifaq.com/

Siarohin, A., Lathuilière, S., Tulyakov, S., Ricci, E., & Sebe, N. (2019). *First Order Motion Model for Image Animation*. https://aliaksandrsiarohin.github.io/first-order-model-website/

Sina, F., Razmeh, S., Habibzadeh, N., Zavari, A., & Nabovvati, M. (2017). Migraine headache in patients with idiopathic intracranial hypertension. *Neurology International, 9*(3), 69–70. https://doi.org/10.4081/or.2017.7280

Szeto, G. P. Y., Straker, L., & Raine, S. (2002). A field comparison of neck and shoulder postures in symptomatic and asymptomatic office workers. In *Applied Ergonomics* (Vol. 33).

Watson, L. A., Pizzari, T., & Balster, S. (2010). Thoracic outlet syndrome Part 2: Conservative management of thoracic outlet. *Manual Therapy, 15*(4), 305–314. https://doi.org/10.1016/j.math.2010.03.002

Weidenbacher, U., Layher, G., Bayerl, P., & Neumann, H. (2006). *LNAI 4021 - Detection of Head Pose and Gaze Direction for Human-Computer Interaction*.

Appendix

Full code can be found at https://github.com/marco-scanni/SPROJ-Code.

Below are most relevant blocks and methods for fetching and displaying measurements and giving correct cues.

IDs of specific keypoints returned by BlazePose :

```
keypoint_IDs = {
    'nose':0,
    'r_eye':3,
    'r_ear':8,
    'r_shoulder':12
}
```

Height, Angle Calculations between Keypoints (coordinates → height) (coordinates → angle):

```
def calculate_height_of(a):
    ycoord = a[1]
    height = np.rint(ycoord*480)
    return height

def calculate_angle(a,b,c):
    a = np.array(a)
    b = np.array(b)
    c = np.array(c)

    radians = np.arctan2(c[1]-b[1], c[0]-b[0]) - np.arctan2(a[1]-b[1], a[0]-b[0])
    angle = np.abs(radians*180.0/np.pi)
```

```
    return angle
```

Get coordinates of keypoints (keypoint name → keypoint coordinates):

```
def get_coordinates_of(keypoint: str):
    return[keypoints[keypoint_IDs[keypoint]].x,keypoints[keypoint_IDs[keypoint]].y]
```

Example of Calibration Step, Threshold Ranges and Threshold Display

```
  capkey = cv2.waitKey(10)
            if capkey == ord('n'):
                prev = time.time()

                while timer >=0:
                    ret,image = cap.read()

                    cv2.putText(image, "Lift Back of Head Towards Ceiling",
                                (30,50),
                                font,1, (255,0, 0), 4, cv2.LINE_AA)

                    cv2.putText(image, str(timer),
                                (320,240),
                                font,7, (255,0, 0), 4, cv2.LINE_AA)
                    cv2.imshow('SPROJ Cam Feed', image)
                    cv2.waitKey(1)

                    cur = time.time()
```

```
                    if cur-prev >=1:
                        prev = cur
                        timer = timer-1

                else:
                    ret,image = cap.read()
                    cv2.imshow('SPROJ Cam Feed', image)

                    results = pose.process(image)
                    keypoints = results.pose_landmarks.landmark
                    #capture coordinates after timer

                    r_eye = get_coordinates_of('r_eye')
                    r_ear = get_coordinates_of('r_ear')
                    nose  = get_coordinates_of('nose')

                    cv2.waitKey(1000)
                    timer = int(3) # reset timer
                    eye_ear_level_angle = 
np.rint(calculate_angle(r_eye,r_ear,[r_ear[0],0]))
                    chin_down_nose_height = calculate_height_of(nose)




            #hinge threshold ranges
            if eye_ear_level_angle != 0:

                #Header
                cv2.putText(image,'Eye-Ear Level:',
                              (330,110),
```

```
                              font, 0.55,color_hinge, 1, cv2.LINE_AA
                          )
                #Ideal Range
                cv2.putText(image,str(eye_ear_level_angle - 12) + ' - ' +
str(eye_ear_level_angle + 8),
                             (500,110),
                              font, 0.55,color_hinge, 1, cv2.LINE_AA
                          )

                if  eye_ear_level_angle > 110 or eye_ear_level_angle < 70:
                    stage_hinge = "Bad Calibration, Try Again"
                    color_hinge = blue

                elif  r_eye_ear_angle > eye_ear_level_angle + 8:
                    stage_hinge = "Overly Tucking Chin"
                    color_hinge = red
                #hinging is granted 4 degrees of freedom because a slight change is
inevitable when
                #raising chin to increase ASI
                elif r_eye_ear_angle < eye_ear_level_angle - 12:
                    stage_hinge = "Hinging"
                    color_hinge = red

                elif stage_asi == "May Be Slightly Hinging" and stage_hinge !=
"Hinging":
                    stage_hinge = stage_asi
                    color_hinge = orange

                elif (r_eye_ear_angle <= eye_ear_level_angle + 8) and
(r_eye_ear_angle >= eye_ear_level_angle - 12):
                    stage_hinge = "Adequate Position"
                    color_hinge = green
```

```
        else:
            stage_hinge = "Not Yet Calibrated"
            color_hinge = blue



            # Hinge threshold display
        cv2.putText(image,"Cervical Hinge Status:",
                    (10,400),
                     font, 0.7,(color_hinge), 1, cv2.LINE_AA
                 )

        cv2.putText(image,stage_hinge,
                    (330,400),
                     font, 0.7,(color_hinge), 1, cv2.LINE_AA
                 )
```