PosePal

Authors: Sruti Srinidhi, Ankita Kalkar, Youssef Kallel

Department of Electrical and Computer Engineering, Carnegie Mellon University

***Abstract*—**Many beginners find yoga to be difficult to do with the physically demanding poses and constant readjustments required. PosePal aims to alleviate these issues by providing an easy, effective, and accessible way for any self-starter to teach themselves yoga. PosePal is a smart mirror that has 3 major components: Pose Detection, Pose Correction, and Mirror UI. It provides real-time audio and visual feedback on whether the user is matching the specified pose. Using the BlazePose CNN model for Pose Detection and a custom-built node computation library for Pose Correction, PosePal presents a novel and effective yoga learning experience.

***Index Terms*—BlazePose, Computer Vision, Machine Learning, MediaPipe, Pose Correction, Pose Detection, Real-time Feedback, Smart Mirror, Yoga**

I. Introduction

Yoga has been linked to tremendous benefits not just physically, but for one's overall wellbeing. It improves flexibility, helps with stress relief, boosts immunity, improves mental health, and so on [1]. With there being a plethora of benefits linked to doing yoga, more and more people are attempting to practice yoga. In fact, the number of yoga practitioners in the US has increased by over 50% in the past 5 years with 1 in 3 Americans having at least tried yoga [4]. Unfortunately, many beginners find yoga to be difficult citing reasons like yoga postures being physically demanding or the deep breathing technique being unusual. Many beginners may get frustrated when trying to adjust their body all the time and having to maintain a fine balance between strength and flexibility while focusing on mindfulness [2].

Enter PosePal: an easy-to-use, accessible, and cheap solution to help beginners learn the physically demanding yoga poses, so they can simply focus on reaping the bounty of mental benefits. PosePal is a smart mirror that aids in pose correction by providing real time audio and visual feedback on whether the user is matching the specified form. Our primary focus is to help beginners learn yoga poses and improve their proficiency. The visual feedback will be provided on the mirror by highlighting specific errors in the user's form directly onto their reflection. Mirrors along with the verbal instructions will especially help yoga practitioners who learn best using visual information. One yoga instructor cites, "This external, reliable, and unbiased reference point can keep you on track and enhance your understanding of alignment principles and relationships—like the tilting and tucking required to find a 'neutral' pelvis in Tadasana (Mountain Pose)" [5]. The final product will also include an interactive mat that further aids in guiding users to help them correct their pose. Furthermore, sections of the mat will light up with different colors according to where users put their limbs to visually indicate whether the user is in the correct section of the mat.

The differentiating factor of our project in comparison to past work, is that it provides real-time feedback, clear and actionable pose correction instructions, and it includes a hardware component to better integrate the physical world around the user. PosePal also considers both the front and side view to ensure no matter what the pose performed is, it can provide an accurate audio cue. Many past ECE capstone yoga corrective projects in the past employed the use of an error calculation without rendering any suggestions for improvement. They also did not have real-time feedback and did not integrate the mirror component to further augment the learning experience. Other competing technologies that exist include smart mirrors such as the Lululemon Mirror that fails to consider the side view as the mirror just processes information from the front and is not as cost effective (can range up to $2200) making it not as accessible to some users[6].

The primary users or target group of PosePal are beginner to intermediate level yoga practitioners that are looking for a way to learn yoga on their own in the comfort of their own homes. From a societal perspective, demand for an in-home solution to learning yoga has grown exponentially over the pandemic, with 8 in 10 yoga practitioners preferring to practice yoga at home. Users find it more convenient, cheaper, and often prefer the privacy of their own home as they are trying to learn a new skill [3]. PosePal provides an easy, effective, and accessible way for any self-starter to teach themselves yoga and embark on a journey to be happier and healthier.

II. Use-Case Requirements

For PosePal to provide a valid, alternative solution to at-home yoga, we defined the use-case requirements to emphasize providing live feedback, accuracy, accessibility, and usability.

A. Live Feedback

In order for PosePal to be effective in real-time when users are practicing yoga, it needs to be able to provide quick, fast feedback live. We determined that the latency should be less than 100 ms from when PosePal first estimates the user's pose to when it renders the errors on the screen. Latency delays within this threshold will be imperceivable to the eyes of users and enable the real-time experience they expect to justify PosePal's use case as a replacement for in-person yoga classes.

B. Accuracy

To effectively replace the need for more expensive and inconvenient yoga classes with human instructors, PosePal needs to have reliable corrective suggestions for an enjoyable, positive learning experience that users can trust. Thus, our accuracy benchmark is 90% or higher, which we deemed possible as the accuracy for BlazePose, our pose estimation component, was stated to be approximately 90% as well.

C. Accessibility

The PosePal system needs to be easy to interact with such that the fitness experience is not disrupted and everyone able to participate in a traditional yoga class can use PosePal. Users should not need to interrupt the flow of their yoga routine to interact with the mirror in the middle of an exercise. With our audio suggestions and visual highlighting of errors, PosePal provides a hands-free approach to learning yoga. We hope to achieve a 90% or higher user satisfaction rate on surveys evaluating accessibility of PosePal. Accessibility also includes making sure the Pose Detection and Correction component of PosePal will provide the same suggestions to users making the same error regardless of height, skin color, gender, and etc.

D. Usability

To ensure PosePal is a reliable and effective alternative, it should be able to perform the same for users of all different heights and sizes. This means that it should be large enough such that any user can see themselves in the mirror comfortably and not be concerned about fitting within the mirror size constraints. Similar to the previous requirement, we are aiming for an average of 90% or higher on user feedback for usability.

III. Architecture and/or Principle of Operation

PosePal's principle of operation is essentially a Pose Detection and Correction feedback loop to aid the user in performing physically-demanding yoga poses. The general sequence a user will follow when they interact with PosePal will be as follows:

1. User clicks "Start Yoga Routine"
2. A new yoga pose is shown
3. Wait 5 seconds for user to maintain pose
4. Show any errors in the pose on the user's reflection with feedback on how to fix it
5. Wait until the user holds 80% accurate pose for 5 seconds
6. Go back to step 2 if there are more poses in the routine

The Mirror UI is detailed in Figure 1 and provides a representation of what the user sees displayed on the smart mirror. At the top of the mirror, the pose name is displayed below the first camera which is used to capture the front view of the user. Not detailed in the diagram is the secondary camera that will be used to capture live feed of side poses. The user's pose is initially highlighted yellow and will turn green once the correct pose is matched within some predetermined error threshold. At the bottom of the screen, a timer will be displayed to help the user keep track of how long each pose should be held along with the side view camera feed so the user has better insight on how to adjust their body without constantly toggling between examining side and front views.

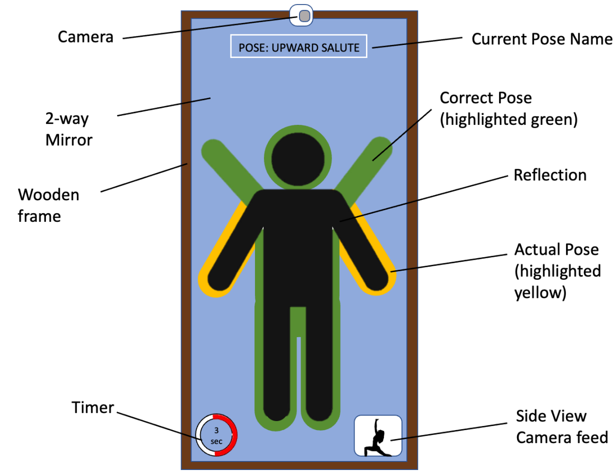

Figure 1: The proposed Mirror UI display

In terms of the Software, PosePal primarily consists of the Pose Detection, Pose Correction, and Mirror UI components (see Appendix Figure 1 for full Software Systems diagram). For the Pose Detection, both the front and side camera feed is inputted into the BlazePose model which is a Pose Detection model that outputs the (x, y, z) coordinates, visibility, and depth of 33 skeleton key points. This is then inputted into the Pose Correction component where we calculate key joint angles and node distances which we compare with reference angles we obtained from a yoga instructor. If the accuracy is higher than 80%, we output the corresponding positive audio cue and send the node positions, angles, and accuracy computation to be rendered on the mirror. If the accuracy is lower than 80%, the corresponding suggestions along with the node positions, angles, and accuracy is sent. For rendering this information on the mirror, the Magic Mirror library will be used and will render the various components detailed in the Mirror UI section above. If the time allotted for that pose is up, then a new pose will be displayed and the Pose Detection component is called again.

The overall Hardware system is outlined in Figure 2. PosePal consists of two equally-sized monitors which will act as our LCD screen that renders the Mirror UI components. On top of the monitors, a two-way mirror will be overlaid that allows the user to see both their reflection along with the Mirror UI components juxtaposed on each other. The Pose Detection and Correction computation will be performed on a laptop which takes the camera feed as input and outputs the rendering to the LCD screen. To combine all the components, a wooden frame will be used to encompass both monitors, mirror, and camera mounted at the top.

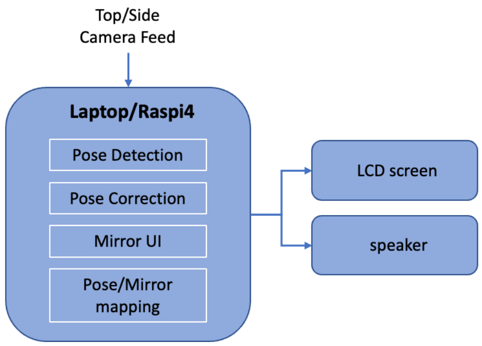

Figure 2: Hardware System and Components

IV. Design Requirements

The designed solution for this use case of making the yoga learning experience more convenient, accessible, and personalized should satisfy the following design requirements to ensure the actual implementation on the engineering side addresses the original requirements outlined for this use case. Each of the design requirements addresses one of the critical use case requirements outlined in the previous section.

The first major design requirement addresses the use case requirement of ensuring the final PosePal system operates in real-time and provides a live feedback and learning experience system. Research for this requirement found that humans take on average ~100ms to process and recognize images visually, implying the need to have an overall latency of the complete pose correction to pose correction and suggestion rendering loop below this 100ms threshold [X]. Given that the benchmark inference time for the out-of-the-box BlazePose model running on a 2019 Macbook Pro has a third quartile time of 38 FPS or approximately 26ms, the rest of the closed loop system (pose correction and rendering) has 100 - 26 = 74ms to run until corrective suggestions are relayed to the user visually and auditorily in order to meet the 100ms threshold use case requirement. Satisfying this requirement from the engineering side will ensure that PosePal truly operates in real-time and can act as a viable substitute for in-person classes.

The next major design requirement satisfies the accuracy use case requirement, where based on feedback, a threshold of 80% accuracy would be sufficient for users to have a fulfilling learning experience without noticeable gaps in quality or inaccuracies in corrective suggestions. Based on published benchmark performance metrics for the BlazePose model deployed as our pose estimation component, PosePal has a 93% accuracy rate based on absolute 2D Euclidean error between inference pose estimates and ground truth examples. Thus, in order to achieve an overall lower bound accuracy greater than 80%, our implemented pose correction algorithm must have an average accuracy of 86% again by the absolute 2D Euclidean error measured against ground truth labeled examples.

In addition, another crucial design requirement is to address the use case requirement of the accessibility of the system. To effectively satisfy the original motivation of this project, PosePal needs to enable ease-of-use in that the user does not have to disrupt the flow of their yoga learning session and workout to use it. To satisfy this from the engineering side, we have a design requirement that PosePal should implement a "one-click" system where the user only needs to physically interact with the mirror once at the start of their session to begin the workout. We also require that user feedback for the corrective suggestions are delivered both visually and auditorily to allow users to adjust even when the mirror isn't easily visible for some poses which require a downward head angle (i.e. downward dog). Achieving this makes sure that PosePal most effectively simulates in-person instruction where users can receive auditory feedback from the instructor as well as visuals to correct their technique.

Finally, the last major design requirement addresses the need to have a usable system for people in that the device allows users to see themselves properly in the mirror, where this property holds for people of diverse heights and sizes as well. From an engineering perspective to accomplish this, we identified that mirror lengths need to be at least half the height of the person to give an entire reflection of their full body [7]. Thus, to accommodate the majority of possible users, we have a required minimum mirror height of 45 inches which accommodates people of a height of 90 inches (or 7'6" in feet) which handles the vast majority of possible users easily. We also require that the material used for the mirror itself is sufficiently reflective and gives the users a clear view of their reflection and body pose. These two design requirements ensure that the system accomplishes its fundamental design goals as a smart mirror based device.

V. Design Trade Studies

Pose Estimation

Pose estimation is a relatively well-studied and established topic in the fields of computer vision and machine learning. However, as is frequently common whenever considering the deployment of machine learning models in almost any domain in a production or user-ready environment, the computational cost, speed, and overall efficiency while maintaining sufficiently high accuracy is the primary distinguishing factor between different existing options. For the purpose of this project, the overall set of considered options were narrowed down to deploying the OpenPose model, BlazePose, or a custom solution trained by us on a yoga-specific pose dataset. Ultimately BlazePose was the most appropriate approach for the outlined use case due to its ability to provide truly real-time inferences from a video feed of users and its high accuracy and ease of deployment and use. In order to make this decision, a thorough tradeoff study was completed comparing these three options against several critical metrics: ability to provide real-time inference, accuracy, and ease of deployment and integration with constrained resources.

A. Real-time Inference

All three major options under consideration do offer relatively real-time inference capabilities, however, the conditions under which this fast inference is possible for each model is distinct. For both the OpenPose and custom model options, performance of on-device inference time on CPU is not as optimized as BlazePose which was designed specifically for this purpose, and instead, real-time inference capabilities are only feasible and consistently reliable when deployed with GPU support. In order to simplify the development process and avoid straining the limited resources allocated for this project for the purchase of materials, BlazePose is the most apt option of the three for this metric.

B. Accuracy and Quality of Inference

Again, accuracy for both BlazePose and OpenPose is high, with OpenPose having slightly higher pose estimation accuracy and more detailed information from the inference as an entire meshgrid point-cloud of the user's pose. The third option of a custom model was not actually implemented so specific benchmarking of its performance is unknown, but assuming it would use OpenPose as a baseline that would be fine tuned to yoga poses, it would have similar accuracy to OpenPose. While OpenPose gives more detailed information and has marginally higher accuracy overall in comparison to BlazePose, the additional information it extracts was deemed to be unnecessary for the error signals PosePal would need to detect and correct on based on discussions with a yoga expert instructor who helped guide and advise some of the fundamental principles on how to perform pose correction in an automated way.

C. Deployment Method and Dependencies

BlazePose is again the lightest weight solution out of the three for this use case - it is very simple to deploy and only requires installation as a lightweight python library module to be run on-device using CPU. In comparison, OpenPose and a custom model solution would require GPU-enabled deployment in order to produce inferences in near real-time and maintain their high benchmark accuracies.

Model Option	**Real-time**	**Accuracy**	**Deployment Method**
BlazePose	Yes (26.3 ms)	93%	On-device CPU
OpenPose	Delayed (10396 ms w/ CPU)	77.9%	GPU-enabled
Custom Model	Delayed (~10396 ms w/ CPU expected)	~77.9% expected	Likely GPU-enabled

Based on these metrics and the different tradeoffs, for the given use case BlazePose is the best solution for the pose estimation component of the project because of its optimizations to provide real-time, sufficiently accurate inferences while running on CPU.

VI. System Implementation

Our main software subsystems are the pose detection system, the pose correction system, and error rendering system. Our hardware subsystem is the smart mirror, and our firmware subsystem is our mirror camera system. The software subsystem workflow is shown in Figure 6 and the overall workflow is shown in Figure 5.

A. Smart Mirror Subsystem

The smart mirror system is built up of three main hardware components – the monitor, the two way mirror and the wooden frame. It also has two webcams – one on top of the mirror and one on the left to provide a side view of the yoga pose. Details of each one are discussed below. These are assembled together as shown in Figure 3.

Figure 3: Smart Mirror Components and Assembly

1) Monitor

As specified by the design requirements, we want a mirror that is at least 45 inches tall. Since we were unable to find monitors of that length within our budget, we decided to use two monitors that are vertically stacked above each other. Each monitor is 24in x 14in in dimension, so the total height is 48in while the width is 14in. To ensure that the experience is as close to a single monitor as possible, we decided to use seamless monitors so that there is a very small bag in between the two screens, making it look like one large monitor. The monitor being used is a Koorui 27-inch monitor.

2) Two-way Mirror

The goal of the two way mirror is to allow the person to see their own reflection while still being able to see what is rendered on the monitor through the mirror. This helps has the same principle as the mirror film in terms of reflection and transparency. We expect it to not distort the reflection as much.

3) Wooden Frame

The wooden frame is built to hold the monitors together with the mirror, as well as create the effect of a mirror. The wooden frame is going to be built with plywood that will be 48in x 14in for its internal dimensions with 2 inches of width for the frame. The depth of the frame depends on whether we use the acrylic mirror or the film, as we will have to account for the thickness of the acrylic if we end up using that.

provide an overlay for human reflection. For the two way mirror, we are testing two possible solutions and going to choose the best option:

a) Mirror Film: This is a film that is commonly placed on windows to block out sunlight from entering the house, like tint. We have tested out the solution, and it is transparent enough to see the monitor's display while still being reflective enough to see the person. The only concern we have with this solution is that since it is a film that is pasted onto the monitors, there tend to be bubbles and slight blemishes, which distorts the reflection slightly. While it is not very distorted to the point where a person cannot see themselves, it is more than ideal and so we plan to keep this idea as a back up and try the next one.

b) Acrylic mirror: This is an acrylic two-way mirror which is ¼ inch thick so does not have to be pasted on the monitor.

B. Pose Detection Subsystem

The pose detection algorithm uses the Blaze Pose model by Google which uses Google's media pipe solution [8]. It provides 33 different body nodes that correspond to different joints and can detect it in real time. Blaze Pose provides the x, y, and z coordinates as well as the visibility of each node. We decided to use this pre-built model over creating our own because finding a dataset was challenging and we decided we had other major challenges to tackle and would rather spend time on that over building our own pose detection algorithm.

The output of the Blaze Pose pose detection algorithm is shown below in Figure 4

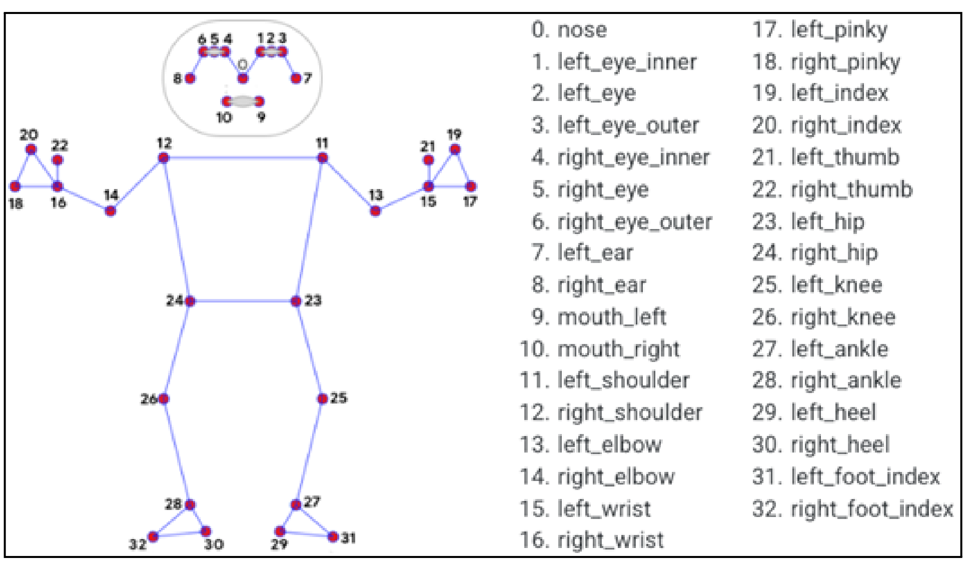

Figure 4: Blaze Pose detection outputs

C. Pose Correction Subsystem

The pose correction system takes the detected pose as input and compares it against a reference yoga pose that we received from a yoga expert. After extensive research, we found that body angles, distances between nodes as well as node alignment are the key factors for pose comparison and correction in yoga. We analyzed several different pictures of the yoga expert doing poses, as well as online research on expected angles and joint mobility, to come up with the expected body angle, alignment and distances for each of the 6 poses we chose. We then built up a system which checks the detected pose's features and compares it to these expected values, with a buffer which we also got from our research and tuning with multiple users.

D. Error Rendering Subsystem

The error rendering subsystem takes the pose inaccuracies that are generated by the pose correction subsystem and renders it on the monitors that are part of the smart mirror. Using python's tkinter library, yellow highlights are drawn over the inaccurate parts of the pose, while green is drawn at the area where the body should also be. The rest of the background is black so the two-way mirror is as reflective as possible.

We are using google text to speech (GTTS) to provide audio cues on how to correct the pose. We will either pre-record a set of audio cues and just play the appropriate one, or we will generate the cues in real time, based on the latency of the audio generation.

E. Mirror Camera Subsystem

The mirror and camera need to be integrated with each other so that what is rendered on the monitor matches with where the reflection shows up on the mirror. This is especially important for the error overlays. This is our biggest technical challenge right now, but we plan to calibrate where a person who stands in front of the mirror would show up in the front camera. Based on this, we can understand horizontally and vertically where the person's reflection will show up in the mirror, so we can render the overlays on the monitor accordingly.

VII. Testing, Verification, and validation

Described below are the tests we intend to conduct to validate our design requirements and ultimately satisfy our use case requirements.

A. Pose Correction Accuracy Test

We plan to run our pose detection and correction algorithm on both static images (from a yoga dataset as well as sample images we got from the yoga expert) to test our algorithm. We then plan to use our pose correction algorithm in front of the yoga expert and allow them to evaluate whether the errors detected are accurate. Based on the number of static images that have inaccurate error correction and the number of times the yoga expert points out an error in our correction algorithm, we calculate the accuracy of the pose correction system. Our goal is to have an 80% total accuracy.

B. Real Time Latency Test

To test the latency of the system, we measure the system time from when a camera frame is captured, to when the error is rendered on the monitors and audio feedback begins to play. Our goal is to have it under 100ms.

C. Accessibility Test

This test is to determine the ease of use for a user. We will have 10 completely new users test out the PosePal solution. We will record the number of times they needed to go back and restart the yoga routine, how many tries it took them to understand the audio and visual feedback on error correction, and how many questions they asked. Our goal is to have a 90% satisfaction rate which means <=2 questions per user and <=5 sec to understand the feedback.

D. Usability Test

The main test here is to see if the size of the mirror and clarity of the mirror is enough for use as a smart mirror. We will test on 20 people and ask them if they can see their whole body clearly. We will get their feedback on a scale of 1-10 and hope to get 75% clarity for the mirror.

VIII. Project Management

This section describes the team's project schedule, breakdown of responsibilities, bill of materials as well as risk mitigation plans.

A. Schedule

Given the large number of different aspects of our project, we decided to break down the tasks into detailed sub tasks and schedule each section of it. This helps us stay on top of all the tasks instead of just looking at overarching milestones. The detailed gantt chart that shows this is viewable in table 2.

B. Team Member Responsibilities

Our project is relatively serial in nature – i.e. many of the tasks in the project are dependent on other projects. Thus, we work together for most of the tasks. In addition, there is a lot of discussion and collaboration for each part of the project. However, each team member is assigned to a section of the project that they take point on and are the main person responsible for it.

- **Youssef**: Pose detection and comparison
- **Sruti**: Camera-mirror integration and error rendering
- **Ankita**: Hardware construction and mirror interface

C. Bill of Materials and Budget

The primary cost for this project is the two-way acrylic mirror. The monitors were received from ECE's stock. While the price of the monitors is reflected in the table, its cost did not go out of our budget. The rest of the items were relatively cheaper and small. The cost of wood is still unknown as it depends on what our final mirror will look like and is so left out of the table. The bill of materials and budget for our project can be found in Table 1.

Item	Manufacturer	Quantity	Cost	Total
27-inch seamless Monitor	Koorui	2	$140.99	$281.89
Mirror Film	KESPEN	1	$11.99	$11.99
2-way Mirror Acrylic 48x24inch	SPEEDYORDERS	1	$266.50	$266.50
1080p Webcam	EMEET	2	$29.99	$59.98
			Total Cost:	$620.26
			Money out of Budget:	$338.37

Table 1: Bill of materials

D. Risk Mitigation Plans

One of the risks for the project was understanding the different metrics to judge and correct yoga poses as it can be quite arbitrary. However, after significant research and consulting a yoga expert, we have now built a strong pose correction algorithm. The second risk is to integrate the camera feel with the mirror. Our plan is to spend 2 weeks working on this and if it does not make progress then we will replace using the mirror with a camera feed.

IX. Related Work

There is a lot of related work for our project. They are mainly divided into two main sections - smart mirrors, and workout mirrors.

The smart mirror solutions talk about various different interactions that can happen with a mirror, especially real time overlays and custom widgets [9].

There are also several workout mirrors, like the lululemon mirror which corrects form and acts like a personal trainer [10,11].

X. Summary

Overall, our system can be broken down primarily into the Pose Detection and Pose Correction component. We are attempting to provide a robust and efficient solution to enable us to fulfill our stated use-case requirements of providing live-feedback, creating an accurate algorithm, ensuring the system is accessible to a diverse set of users, and making sure PosePal is a usable and effective smart mirror. Using a relatively simple hardware interface of the LCD screen with the mirror, we hope to mimic a mirror with the augmented ability of pose correction for physically-demanding yoga poses.

Currently we have finished most of our key design choices and will modify accordingly as we continue the building and testing process. We have finished an initial prototype of the mirror using a mirror sheet stuck to the monitor and tested the visibility of both the reflection and LCD output.

XI. Glossary of acronyms

CNN – Convolutional Neural Network
CPU – Central Processing Unit
GTTS – Google Text to Speech
ML – Machine Learning
RPi – Raspberry Pi

XII. References

[1] Ezrin, Sarah. "16 Science-Based Benefits of Yoga." *Healthline*, Healthline Media, 14 Dec. 2021, https://www.healthline.com/nutrition/13-benefits-of-yoga#TOC_TITLE_HDR_4.

[2] "Here's Why Yoga Is so Hard for Beginners." *Yoga Kali*, 3 Oct. 2022, https://yogakali.com/yoga/yoga-faq/heres-why-yoga-is-so-hard-for-beginners/.

[3] "The Global Yoga Survey 2021: How and Why People Practice Yoga." *DoYou*, 12 Apr. 2022, https://www.doyou.com/yoga-survey/.

[4] Andre, Louie. "47 Compelling Yoga Statistics: 2023 Data on Industry Growth & Effects on Health." *Financesonline.com*, FinancesOnline.com, 9 Jan. 2023, https://financesonline.com/yoga-statistics/.

[5] *Mirrors in Yoga: Are They Helpful or Distracting? - Yoga Journal*. https://www.yogajournal.com/poses/anatomy/hips/mirrors-yoga-helpful-distracting/.
[6] Lang, Adam. "Lululemon Mirror: Pros and Cons: Driven Fit Personal Training." *Driven Fit*, 13 Nov. 2020, https://www.drivenfit.com/lululemon-mirror-pros-and-cons.
[7] Shortest mirror to see full length of your body. (n.d.). Retrieved March 3, 2023, from https://mammothmemory.net/physics/mirrors/flat-mirrors/shortest-mirror-to-see-full-length-of-your-body.html
[8] "Blaze Pose." *Mediapipe*, https://google.github.io/mediapipe/solutions/pose.html .
[9] "Himirror Slide Smart Makeup Mirror Provides Personalized Skin Tips & Also Streams Content." *Gadget Flow*, Gadget Flow, 4 July 2022, https://thegadgetflow.com/portfolio/himirror-slide-smart-makeup-mirror-provides-personalized-skin-tips-streams-content /.
[10] "Echelon Reflect Smart Fitness Mirror Includes Both Live & On-Demand Fitness Classes." *Gadget Flow*, Gadget Flow, 27 May 2022, https://thegadgetflow.com/portfolio/echelon-reflect-smart-fitness-mirror-includes-thousands-live-on-demand-fitness-classes/ .
[11] "Smart Fitness Mirror: Lululemon Studio Mirror." *Lululemon Studio*, https://www.mirror.co/shop/mirror .

XIII. Appendix

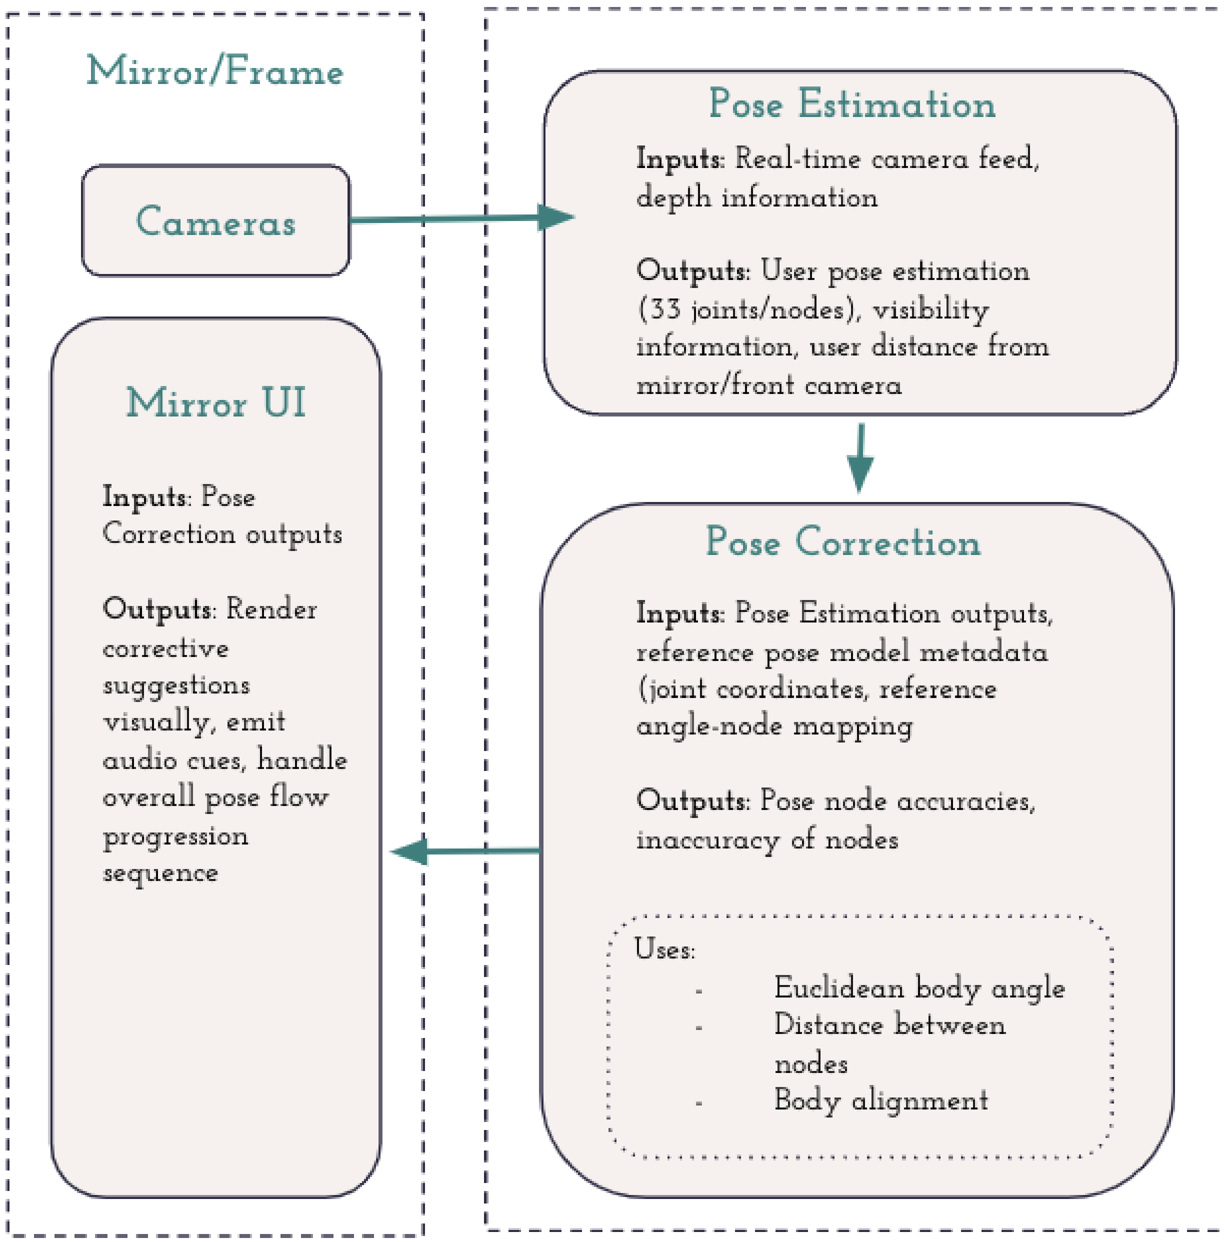

Figure 5: Overall Subsystem Interaction

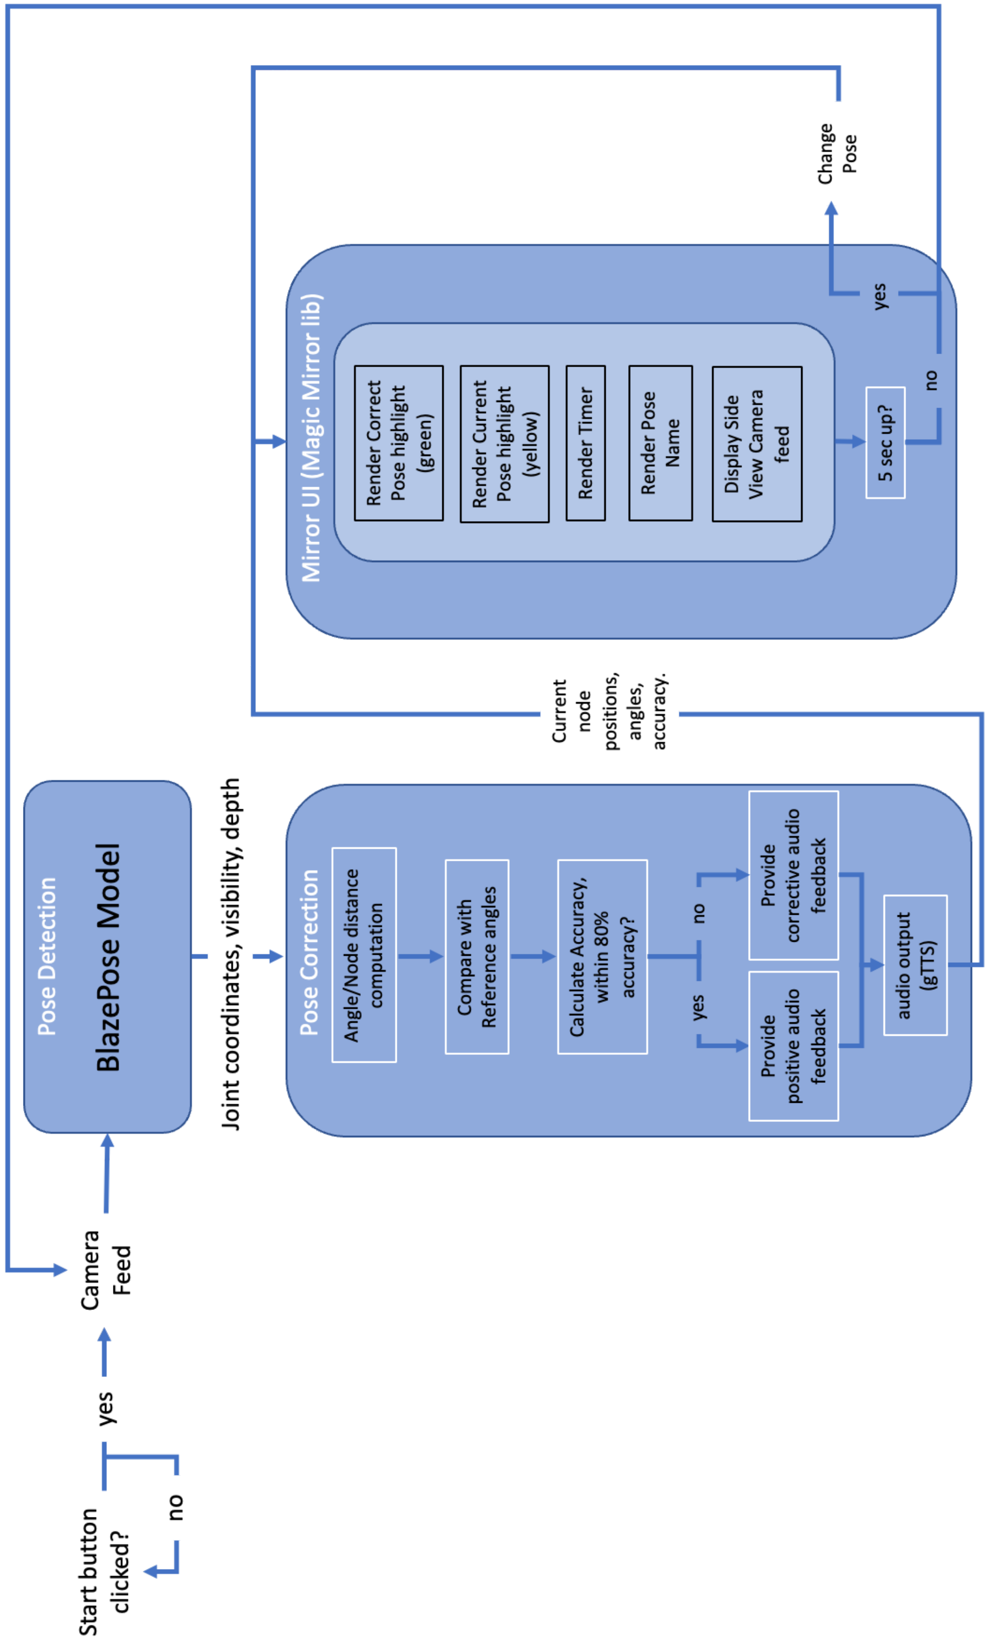

Figure 6: Software Implementation Solution

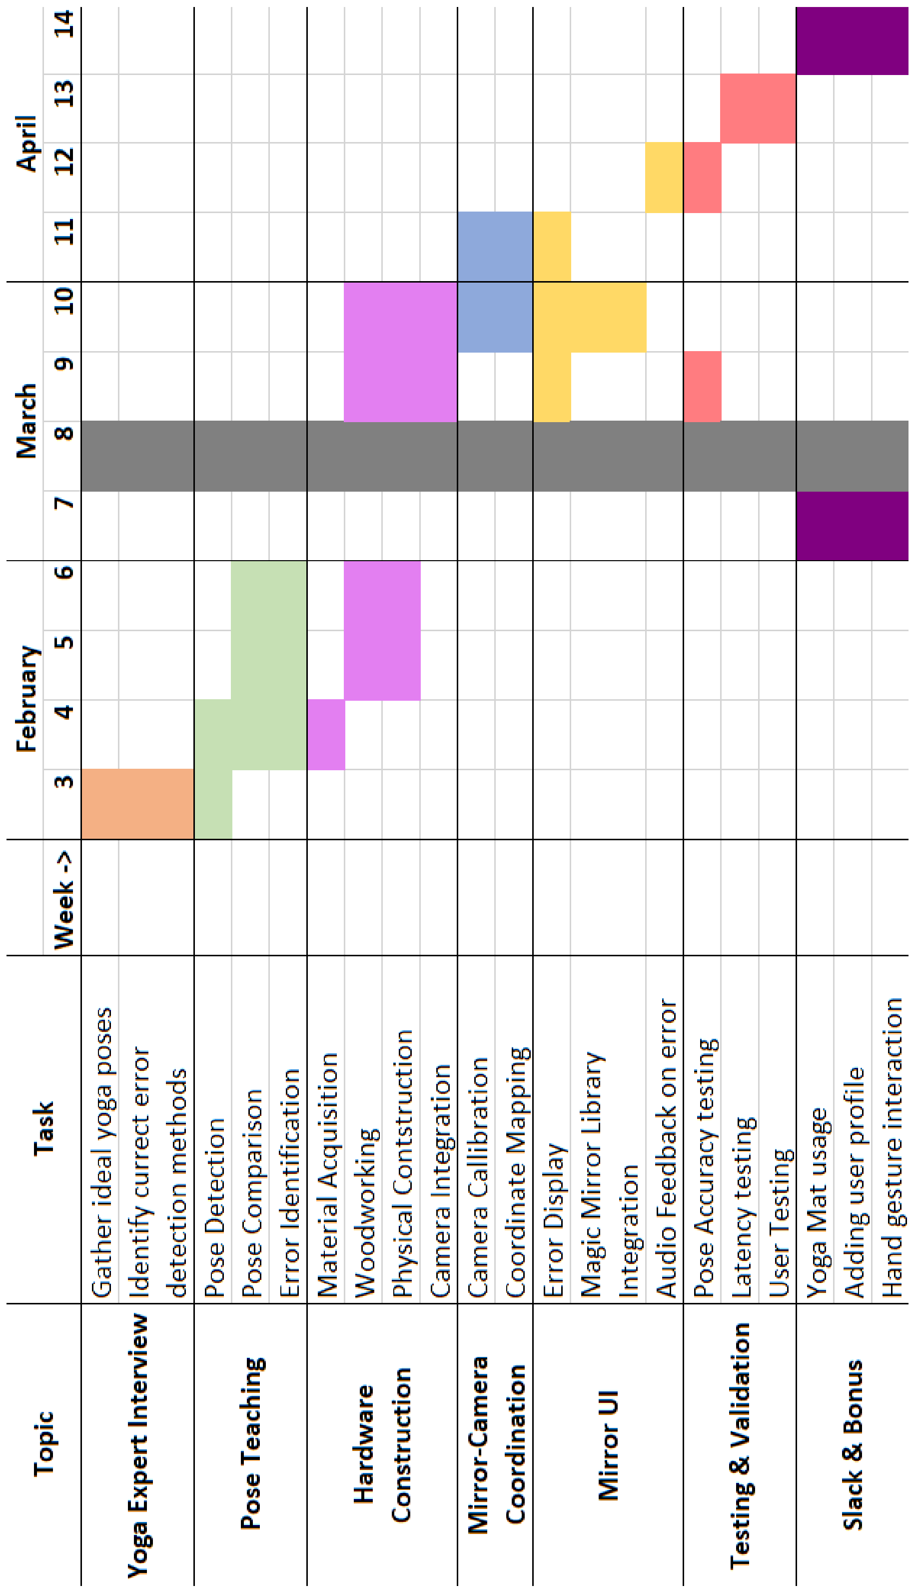

Table 2: Gantt chart