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.

The differentiating factor of our project in comparison to past work, is that it provides real-time feedback, clear and actionable pose correction instructions. 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 is not as cost effective (can range up to $2200) making it not as accessible to some users[5].

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 80% or higher.

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. 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.

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 presses enter to begin the flow
2. The current (or first) yoga pose is displayed to the user
3. Visually rendered error detection and audio feedback is given to help the user correct their pose
4. Wait 8 seconds for user to maintain pose once correct
5. Once 8 seconds have passed consecutively with a correct pose, repeat from step 2 if there are more poses, otherwise finish

A high-level overview of the final product and 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 camera is positioned to face the user and generate the pose estimation needed for pose correction and mapping the user's coordinates to their reflection in the mirror. On the UI itself, a reference image for the current pose is displayed in the top left corner of the mirror/screen to help guide the user on the expected positioning for their current pose in the workflow. Once the user's entire body is within the frame of the camera, the relevant nodes on their body that are processed and evaluated for the current pose's correctness are initially red. These node groups turn green once the correct pose for that component is matched within some predetermined and tuned error threshold, using feedback and discussions from a yoga instructor as guidance. Audio feedback is given simultaneously for richer detail and insight into how to correct the user pose and to know when it is correct, signaling to the user to hold their pose until the next one. This audio feedback also enables the user to get post adjustments even when the mirror is not visible for poses like downward dog.

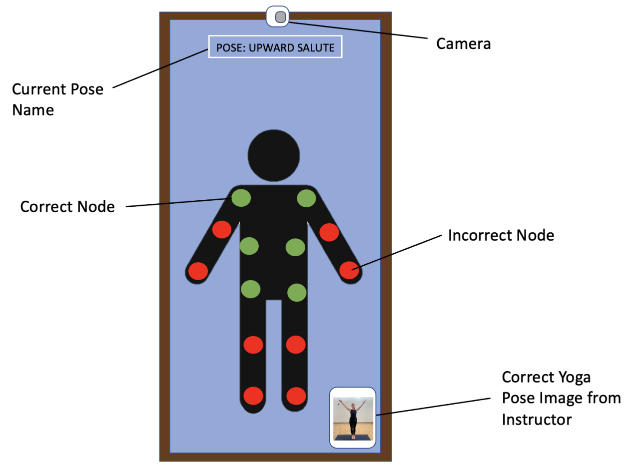

Figure 1: The Mirror UI display

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 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 for each of the evaluated node groups for a given pose is within a certain threshold of satisfiability, 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 below the threshold for any nodes, the corresponding suggestions along with the node positions, angles, and accuracy is sent. A simple, lightweight yet informative PosePal UI library using openCV utilities will render the various components detailed in the Mirror UI section above. Google's text-to-speech library (gTTS) was used to generate the actual audio feedback emitted to the user at different points in time throughout the PosePal workflow. This cycle repeats for each pose as the user progresses after holding a correct pose for a configurable amount of time (currently set to 8 seconds based on yoga instructor feedback).

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. (*Note: Unfortunately, due to unforeseen events, one of these monitors was broken when moved by a non-team member, so the interim demo edition of PosePal features only one functional monitor.) 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.

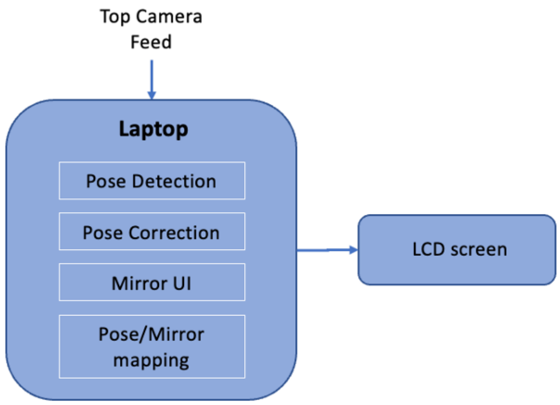

Figure 2: Hardware System and Components

IV. Design Requirements

Our solution has 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 live feedback. The requirement is that PosePal operates in real-time and implying the need to have an overall latency of the complete pose correction to pose correction and suggestion rendering loop below this 100ms threshold. 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.

The next major design requirement is based on the accuracy use case requirement, where 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 and based on yoga instructor evaluation of the system.

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 with diverse heights and sizes. 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 [6]. 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 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

After testing both options, we decided to go with the Acrylic mirror as it proved to warp the user's reflection less and proved to be less flimsy resulting in a more durable and robust design. Furthermore, the reduced warpage in the user's reflection helped our Pose detection system better identify the key nodes and render the error corrections in the correct locations on the user's reflection.

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. We decided to add padding along both sides of the monitor to hold them tightly in place along with a thinner strip of plywood in the back to support the monitors from the back. We made sure to leave space for the HDMI wires and monitor powering cables by laser cutting a small cut-out in the back. Additionally, for extra support we decided to attack mirror brackets to hold the mirror up tightly against the monitors to prevent any warping effects and create a more seamless design.

B. Pose Detection Subsystem

The pose detection algorithm uses the Blaze Pose model by Google which uses Google's media pipe solution [7]. 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

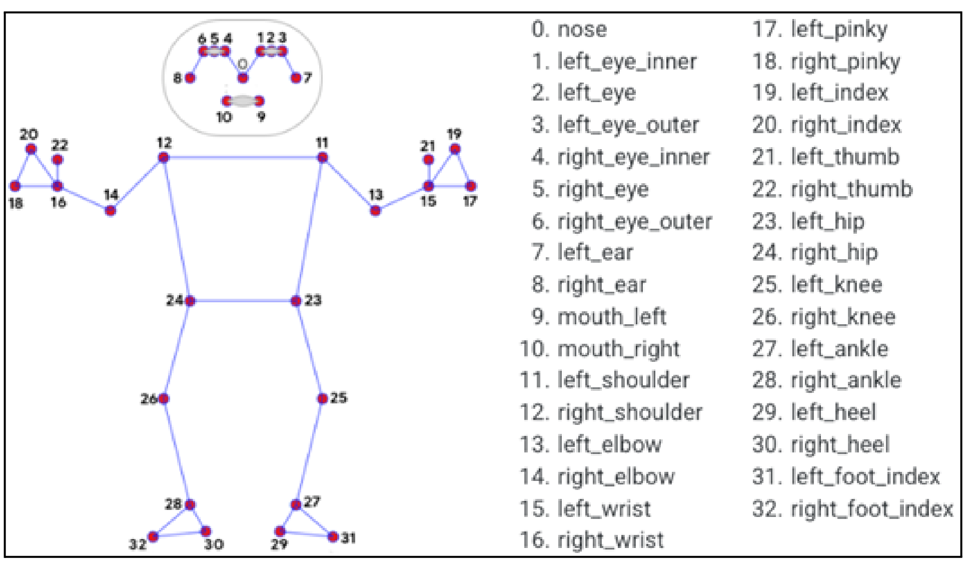

Figure 4: Blaze Pose detection outputs

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 opencv, red 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. After writing the corresponding audio cues for each pose correction, we used GTTS to generate the mp3 files beforehand which helped reduce the latency of our system. Then, we use the playsound module to play the mp3 files in a queue ordered from most to least important pose corrections.

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 on the user's reflection. This was one of our biggest challenges of our project, and we tried several different methods to map the person's position from the camera feed onto the monitors. Ultimately, we decided to include a scaling factor and focused on resizing the window to get an accurate mapping.

VII. Testing, Verification, and validation

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

A. *Pose Correction Accuracy Test*

To test the overall accuracy of our custom pose detection and correction system, we conducted experiments with the yoga expert instructor. They used our completed system and worked through the entire set of poses comprising the overall PosePal flow. The instructor evaluated whether the initial pose instructions were sufficiently clear and whether the detected errors were accurate and appropriately strict for each given pose. They also evaluated whether the visual and auditory corrective suggestions were appropriate for the error and if they accurately aided in reaching the correct desired pose. These results were collected through a thorough survey that the instructor completed immediately after using the system. Our initial goal was to achieve an overall **accuracy of 80%**, and the actual **outcome was 87.2%** (computed as a product of the accuracy of the pose estimation model and the reported accuracy of the pose correction component).

B. *Real Time Latency Test*

To test the latency of the system, we conducted 500 experiments measuring the system time from when a camera frame is captured, to when the error is rendered on the monitors and audio feedback is played. Our goal was to have an **average time under 100ms**. The actual statistics finally measured were an **average of 67.1 ms** and standard deviation of **5.78ms** putting the system latency performance comfortably within our target range.

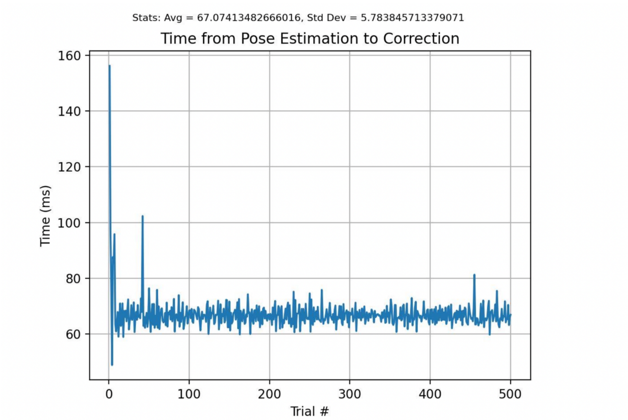

The end-to-end latency of the system is on average 67ms.

C. *Accessibility Test*

This test is to determine the ease of use for a user. We had 10 completely new users test out the PosePal solution. We recorded 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. The survey results can be seen in the appendix. Our goal was to have a 90% satisfaction rate which means <=2 questions per user and <=5 sec to understand the feedback. The recorded outcomes from the user feedback on aggregate was that the average effectiveness of the visual and auditory corrective feedback was found to be **9 out of 10**

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 asked the same 10 people from the previous test if they can see their whole body clearly in the mirror and whether PosePal as a system was easy to use and effective with little difficulty in use. From the aggregated feedback collected through the survey, we had **80%** of those surveyed saying the tool was both easy to use and that they would be inclined to use it again.

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 helped us stay on top of all the tasks instead of just looking at overarching milestones. While some sections took longer than we initially anticipated, we had enough buffer time so we were not pressed for time overall. The detailed gantt chart of what our schedule ended up being 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**: Error rendering and Pose Detection
- **Sruti**: Camera-mirror integration and Pose Correction
- **Ankita**: Hardware construction and audio interface

C. Bill of Materials and Budget

A major cost for this project is the two-way acrylic mirror. We also used multiple monitors as they were not very stable and broke easily. We also spent some money on buying wood and woodworking supplies to build our mirror. We had two webcams and small hardware components that were small parts of our budget as well. The bill of materials and budget for our project can be found in Table 1.

additional work could include adding a social component

Item	Quantity	Cost	Total
Monitor	1	130	130
Monitor	1	86	86
Wood Pieces	1	47.14	47.14
2-way Mirror acrylic	1	266	266
Usb-c hub	1	40	40
Mirror Brackets	1	6.59	6.59
2-way Mirror sheet	1	12	12
Camera	2	39.59	79.18
		Total Cost:	666.91

Table 1: Bill of materials

D. Risk Management

One of the major risks we realized once we began working with our hardware is how fragile it was and how the monitors would constantly break. The only fix for this was to buy a new monitor which is very expensive and also takes a lot of effort to reinstall into our system. Our plan was that if one of the monitors broke close to the demo then we would just render our solution on one monitor and not scale it across two monitors to scale for height. While this will not have the complete intended functionality of our project, it will showcase most aspects of it.

Our second risk that we managed was the woodworking. No one on the team really had much experience with woodworking. While we did make some mistakes with our initial design, we managed to work with more experienced members and redesign our woodwork given the progress we had already made and this helped us bring that aspect of our project together.

IX. Ethical Issues

In terms of public welfare, one ethical issue that could arise is if the smart mirror is not accessible or affordable. Another issue is that the smart mirror, if used as a replacement for in-person yoga classes, could contribute to social isolation and potentially lack of community support. This could lead to the opposite effect on user's mental health and overall well-being. To address this, we made sure to use friendly and encouraging audio cues to motivate users through using a friendly context. Furthermore, where users can see how many poses they hit correctly that day on the first try and compare that with friends.

In terms of public safety, a potential ethical issue that could arise is if the smart mirror is not calibrated properly and provides incorrect feedback. The repercussions of this is that the user could get injured, defeating the purpose of our project. Furthermore, if the data is not secured, it could get hacked leading to privacy concerns. To address this, we made sure to create a system that does not store any user data and make sure the instructions we provide the user have undergone tremendous user testing and are very succinct, accurate, and easy-to-understand.

In terms of public health, our project aims to provide an easier way to get active and learn a new skill. Practicing yoga has been linked to tremendous benefits not just physically, but mentally as well, curbing symptoms of anxiety and depression and overall improving one's mood after practicing yoga for even as little as 30 minutes. However, there are some issues. For example, the mirror may not be suitable to all individuals such as those with medical conditions and other such physical limitations. The mirror is currently not geared to accommodate a slower pace or more relaxed poses. To address this, future work could focus on adding different modes such as easy, intermediate, and hard with parameters tunes more loosely for the easy level and more finely for the hard mode.

X. 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 [8].

There are also several workout mirrors, like the lululemon mirror which corrects form and acts like a personal trainer [9,10]. The way that these differ from our work is that we have on body feedback for a person to know what their error is, and our cost is lower.

XI. 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.

PosePal performed well in most of the tests and it surpassed all of the stated design requirements. The final product has a latency of 67.07 ms on average and an overall Pose Correction accuracy of 87% which surpassed our previous requirements of 100 ms for latency and 80% for Pose Correction accuracy. Lastly, for our accessibility requirement, we were able to produce a system that can be used by people of all different heights, body types, skill levels, and so on. Most of the feedback we received after conducting user testing was positive and users found the audio cues in conjunction with the visual error renderings to be extremely helpful when learning a new yoga pose.

XII. Glossary of Acronyms

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

XIII. 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] 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.

[6] 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

[7] "Blaze Pose." *Mediapipe*, https://google.github.io/mediapipe/solutions/pose.html .

[8] "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 /.

[9] "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/ .

[10] "Smart Fitness Mirror: Lululemon Studio Mirror." *Lululemon Studio*, https://www.mirror.co/shop/mirror

XIV. Appendix

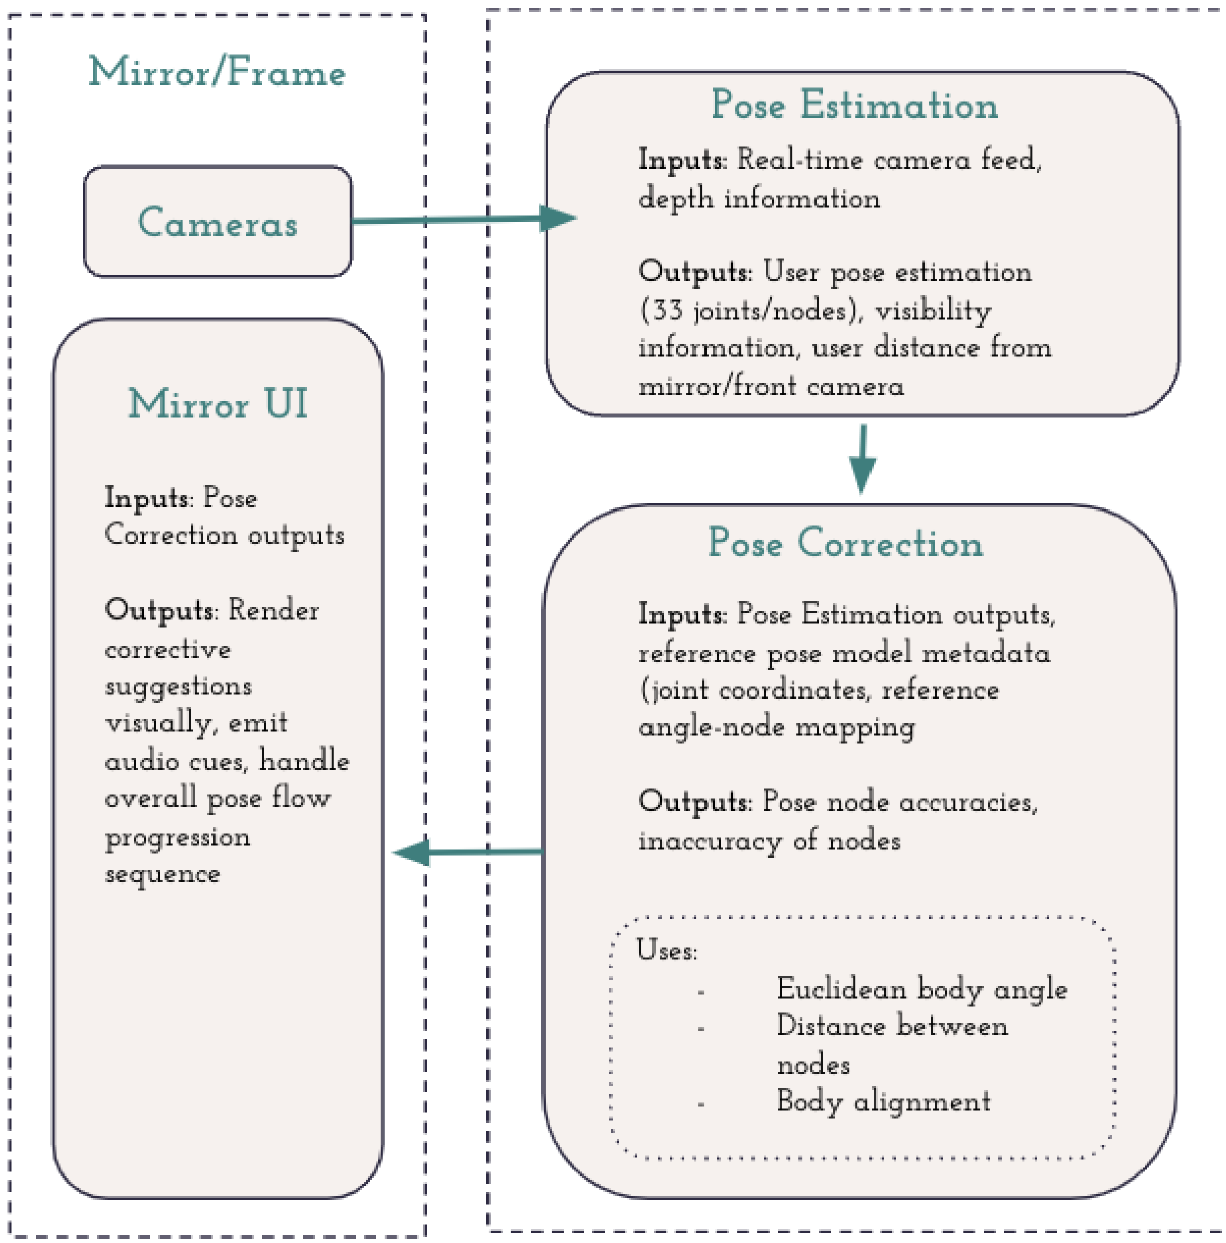

Figure 5: Overall Subsystem Interaction

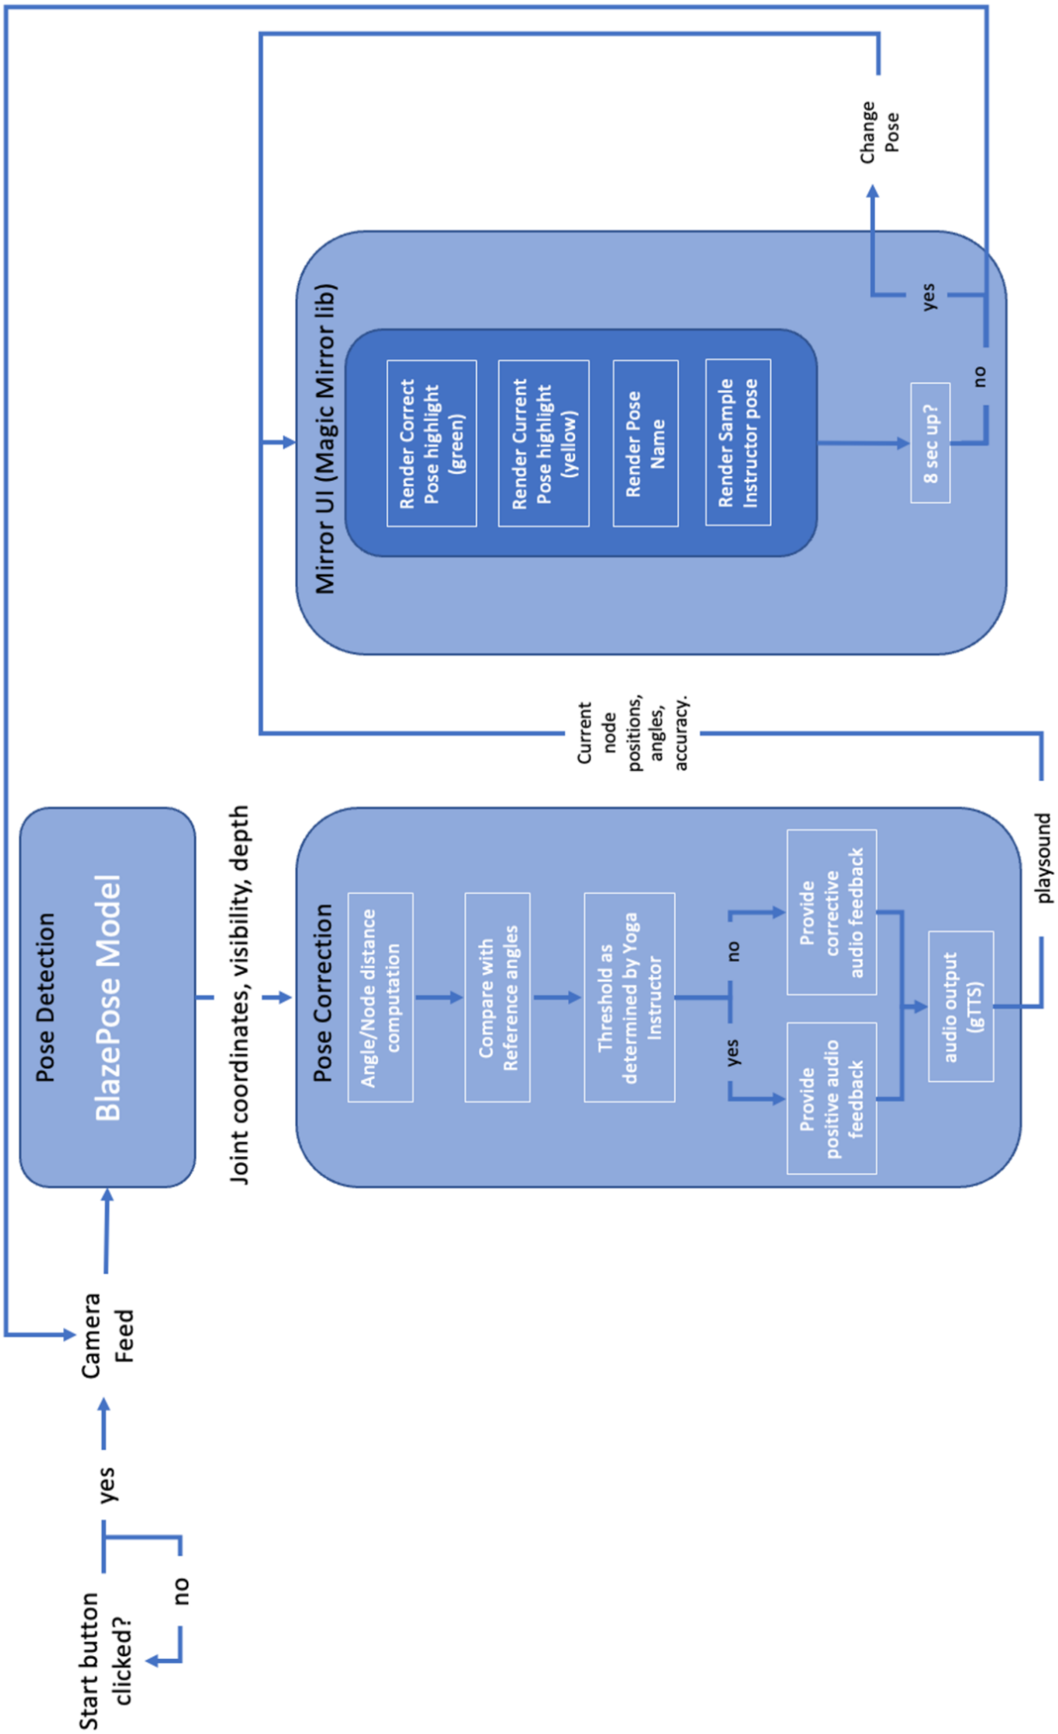

Figure 6: Software Implementation Solution

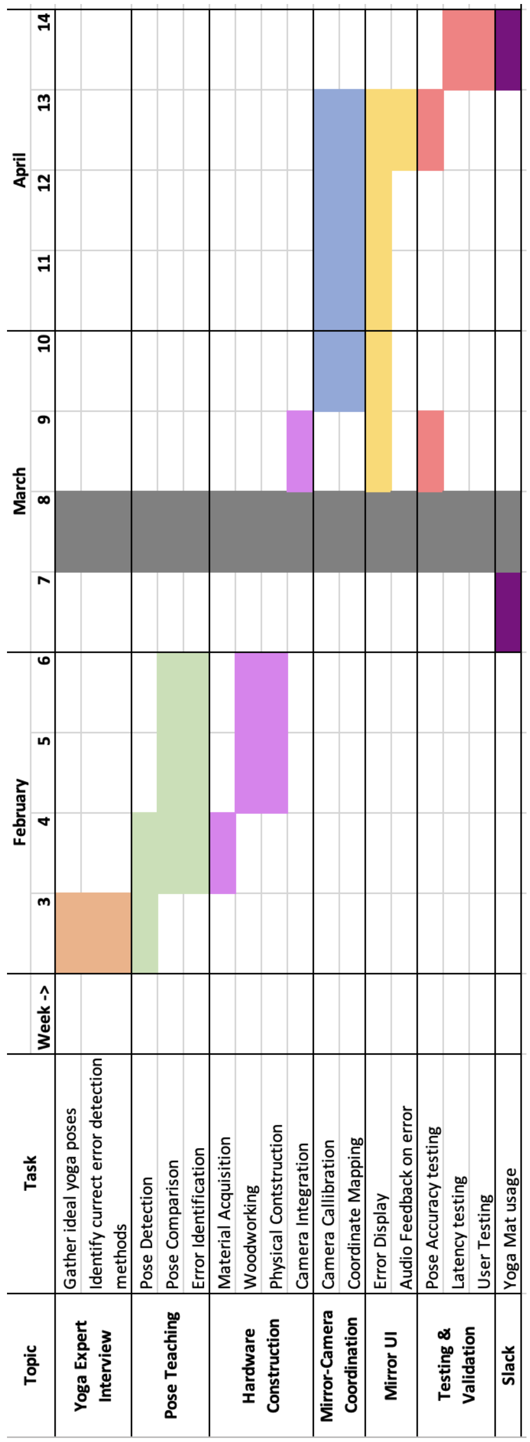

Table 2: Gantt chart

Survey Results

On a scale of 1-10, how easy did you find it to do yoga with PosePal?
10 responses

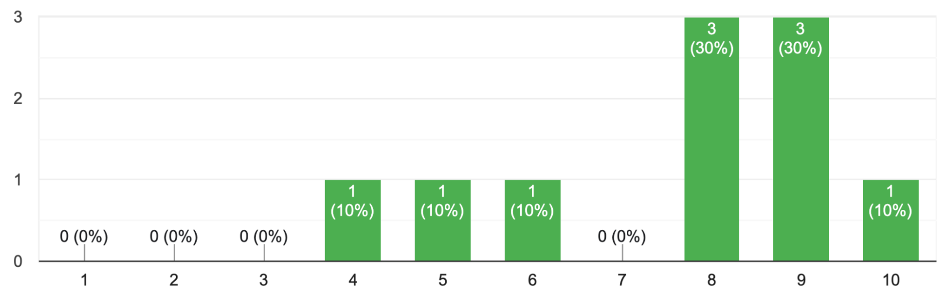

Are the audio instructions clear and easy to understand/follow?
10 responses

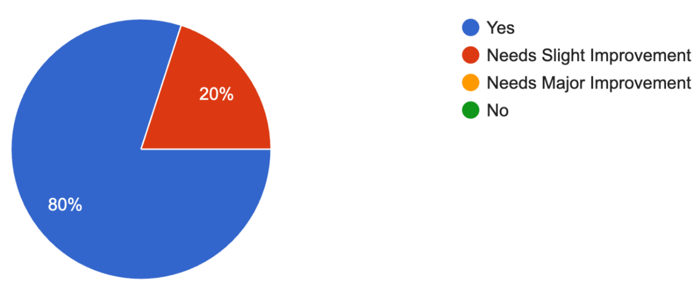

Does PosePal provide audio feedback in a non-judgemental and supportive manner?
10 responses

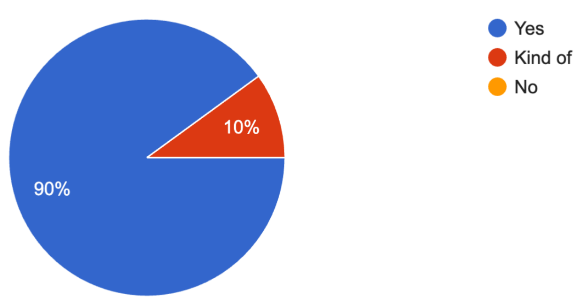

Do you find the audio cues to be a helpful supplement to the visual cues?
10 responses

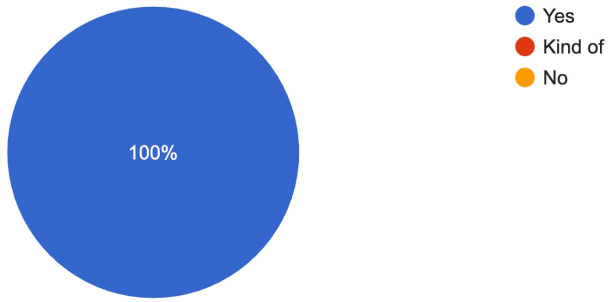

On a scale of 1-10, how effective did you find the visual highlighting of the nodes green and red to be when learning a new pose?
10 responses

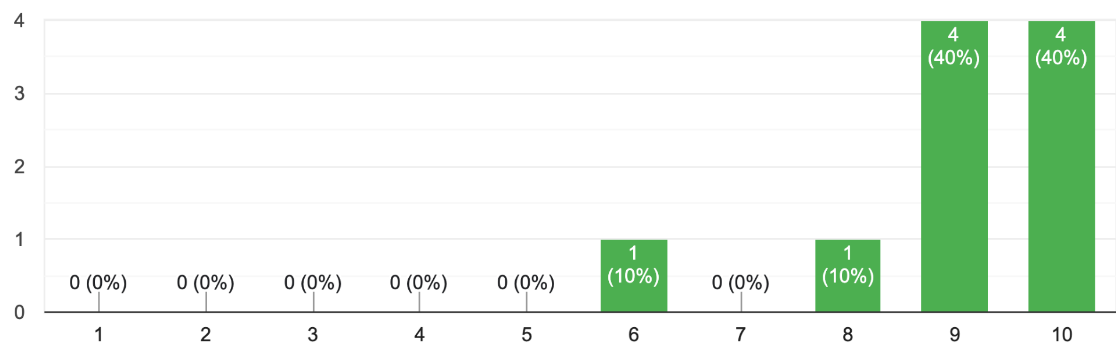

Would you find this product easy to set-up and use as an at-home tool to learn yoga?
10 responses

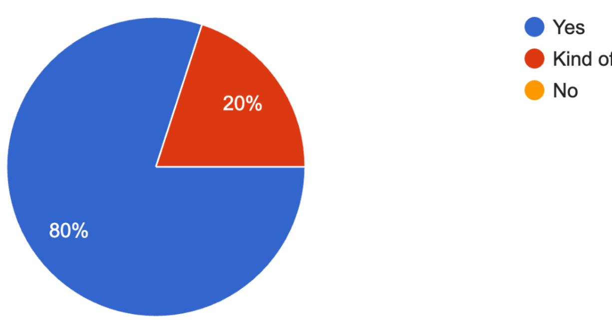

Do you think PosePal respects user's privacy and security given that it doesn't store or save any video footage?

10 responses

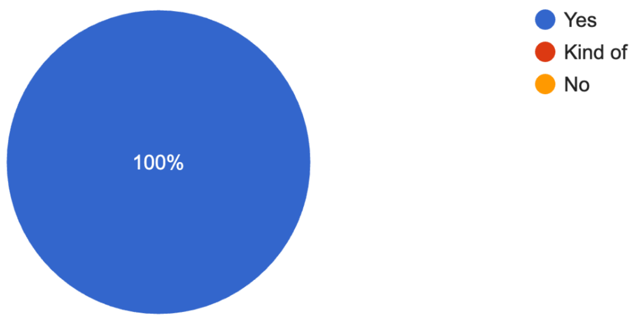

PosePal currently costs $670 which is around $330 less than the leading Smart Mirror brand. Yoga classes also can currently cost up to $75/session. ...you think PosePal is cost effective and affordable?

10 responses

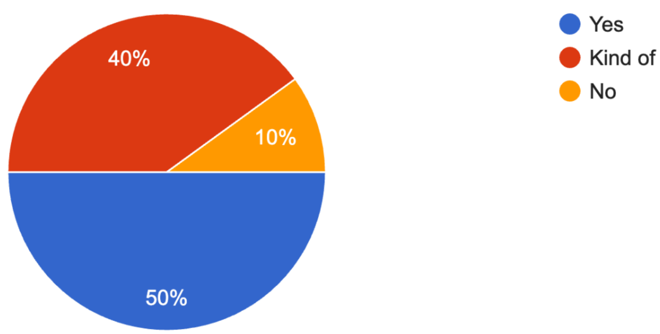

Is PosePal designed with inclusivity in mind? Do you think PosePal is effective for diverse groups of users, including people of different ages, genders, cultures, and body shapes?

10 responses

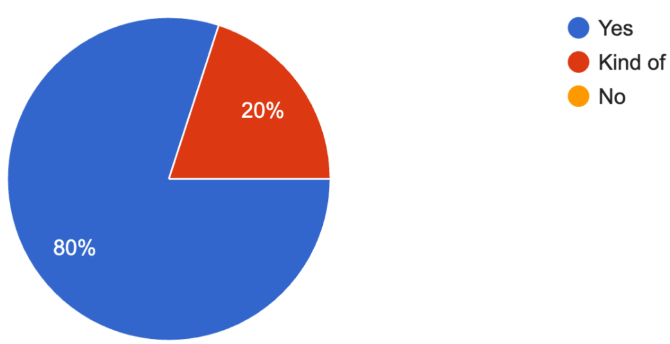

Can PosePal be customized to meet individual needs and preferences?

10 responses

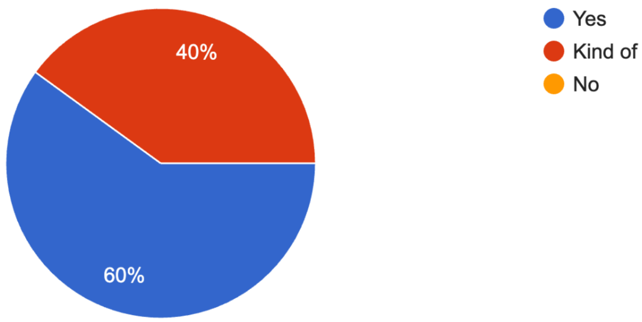

On a scale of 1-10, how would you rate PosePal's effectiveness as a tool to learn yoga?

10 responses

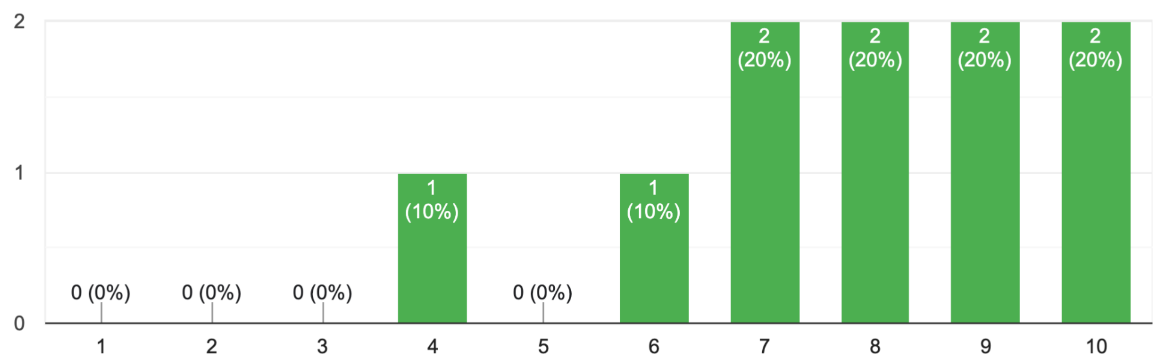

Would you use PosePal again?

10 responses