Facial Matching Display Mirror w/ Motion Sensor
ECE 445 Design Document - Spring 2026

Project Team #97
Keenan Peris, Krish Sahni, Connor Tan

Professor: Yang Zhao
TA: Argyrios Gerogiannis

1. Introduction

1.1 Problem

Modern museums, classrooms, and science outreach programs face a persistent challenge: effectively engaging visitors with science concepts in a way that is interactive, memorable, and personalized. Static displays and traditional informational panels often fail to capture attention for long periods, particularly for younger audiences who are accustomed to interactive digital experiences. As a result, many STEM outreach tools struggle to communicate scientific ideas in a way that encourages curiosity and sustained engagement.

Our project proposes an interactive "scientist mirror" display that transforms a regular mirror into an educational interface. When a person stands in front of the mirror, a camera and computer vision system analyze the individual and trigger the display to overlay information about a scientist that visually resembles them. This interaction encourages users to imagine themselves as scientists and learn about notable figures in STEM fields. The display remains mirror-like when no user is present, creating a natural and intuitive interaction experience.

This solution improves upon traditional museum displays by combining computer vision, embedded systems, and interactive displays to create a personalized educational interaction. Unlike static posters or informational screens, the mirror provides an engaging and surprising experience that connects the user directly with the scientific narrative being presented. The system therefore functions both as an outreach tool and as a demonstration of modern AI-driven interactive technology.

1.2 Visual Aid
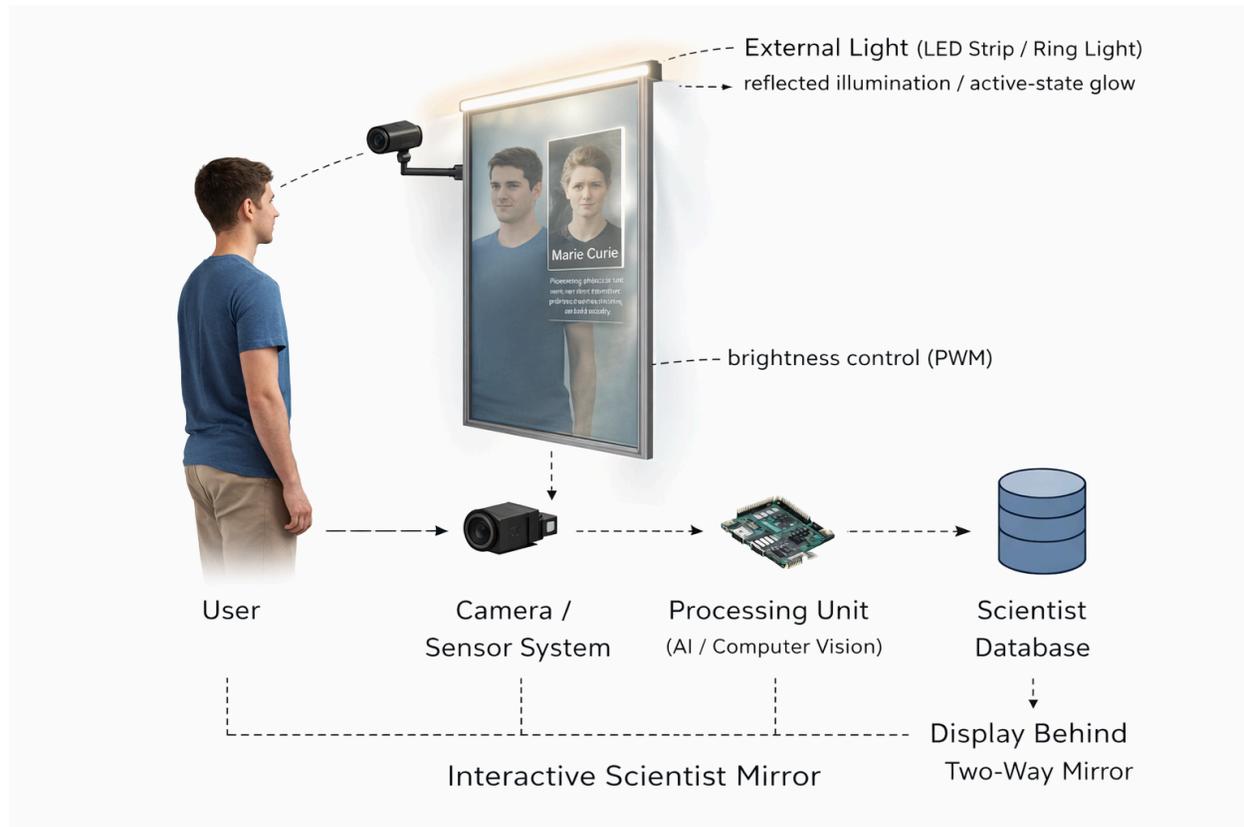
1.3 High-Level Requirements

To consider the project successful, the mirror system must satisfy the following requirements:

1. The system must detect the presence of a user standing in front of the mirror within 2 seconds and activate the display overlay.
2. The system must identify or match a user to a scientist profile from a predefined dataset and display the associated information on the screen.
3. The mirror must function as a reflective surface when inactive and seamlessly transition to an interactive display when a user is detected.
4. The system must present readable visual information (scientist name, image, and short description) on the display while maintaining the mirror illusion.
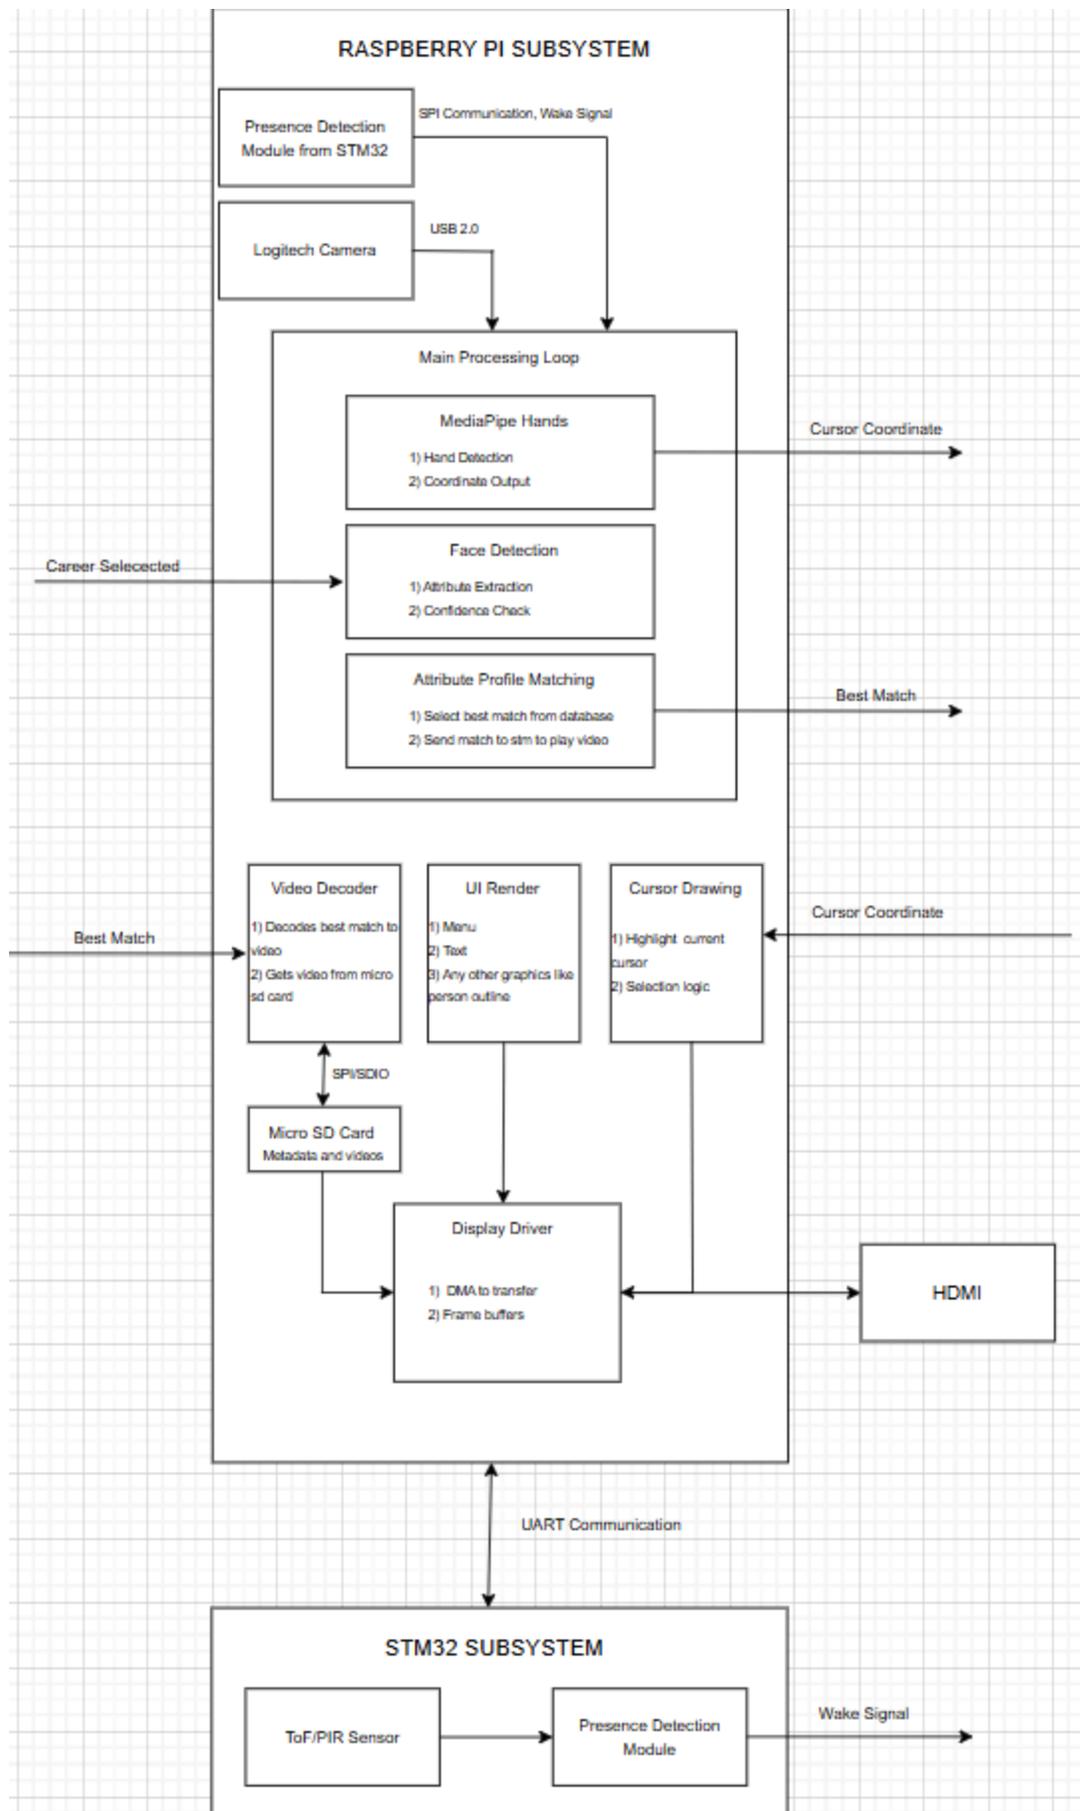
2. Design

2.1 Overall System Architecture

The Smart Mirror system is implemented as a distributed embedded architecture consisting of two coordinated processing platforms: a Raspberry Pi 4 (8GB) and an STM32F401 microcontroller. This partitioned design separates computationally intensive perception algorithms from deterministic embedded control logic, improving system robustness and ensuring predictable real-time behavior.
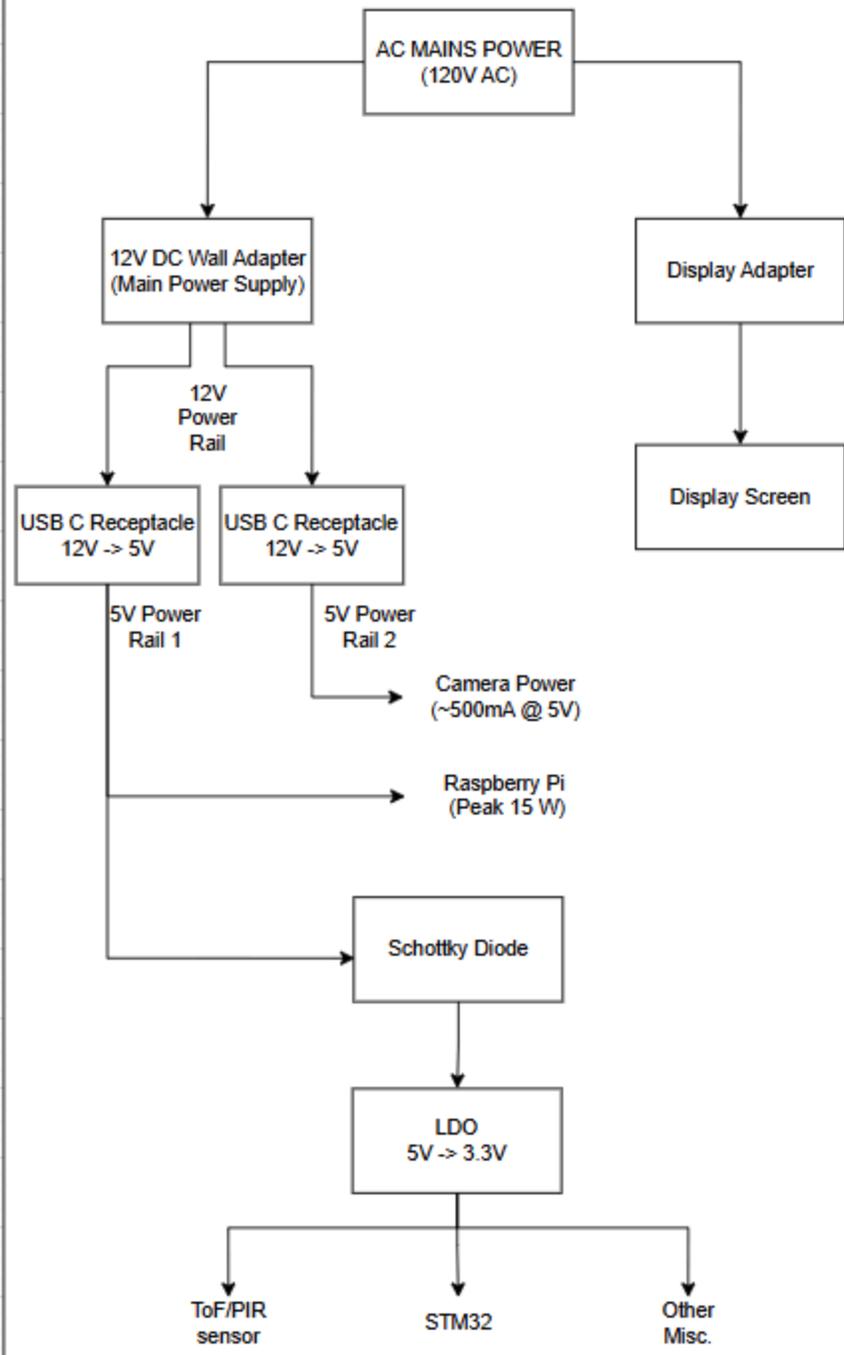
The Raspberry Pi operates within a Linux environment and executes high-level computer vision algorithms, including hand tracking, face detection, confidence evaluation, and scientist profile selection. The STM32F401 functions as a real-time embedded controller responsible for presence detection, finite state machine (FSM) execution, peripheral interfacing, and system-level timing.

Communication between the two processors is implemented using a 3.3 V UART interface. Only compact command packets are transmitted, including cursor coordinates, selection confirmations, and profile identification values. No image data is exchanged between processors. This design choice minimizes communication latency and prevents high-bandwidth transfers from affecting deterministic system behavior.

By isolating perception tasks from hardware control, the system maintains consistent UI responsiveness even during temporary CPU load spikes on the Raspberry Pi.

RASPBERRY PI SUBSYSTEM

Presence Detection Module from STM32

SPI Communication, Wake Signal

Logitech Camera

USB 2.0

Main Processing Loop

MediaPipe Hands
1) Hand Detection
2) Coordinate Output

Cursor Coordinate

Face Detection
1) Attribute Extraction
2) Confidence Check

Career Selecected

Attribute Profile Matching
1) Select best match from database
2) Send match to stm to play video

Best Match

Video Decoder
1) Decodes best match to video
2) Gets video from micro sd card

UI Render
1) Menu
2) Text
3) Any other graphics like person outline

Cursor Drawing
1) Highlight current cursor
2) Selection logic

Cursor Coordinate

Best Match

SPI/SDIO

Micro SD Card
Metadata and videos

Display Driver
1) DMA to transfer
2) Frame buffers

HDMI

UART Communication

STM32 SUBSYSTEM

ToF/PIR Sensor

Presence Detection Module

Wake Signal

POWER DIAGRAM

AC MAINS POWER
(120V AC)

12V DC Wall Adapter
(Main Power Supply)

Display Adapter

12V
Power
Rail

USB C Receptacle
12V -> 5V

USB C Receptacle
12V -> 5V

Display Screen

5V Power
Rail 1

5V Power
Rail 2

Camera Power
(~500mA @ 5V)

Raspberry Pi
(Peak 15 W)

Schottky Diode

LDO
5V -> 3.3V

ToF/PIR
sensor

STM32

Other
Misc.

2.2 Power Management Subsystem

The STM32 subsystem is powered via a USB-C receptacle configured as a downstream-facing device (UFP). Dual 5.1 kΩ pull-down resistors are connected to CC1 and CC2 to correctly advertise device mode and ensure proper USB-C negotiation.

The 5 V VBUS input passes through a Schottky diode for reverse polarity protection before being regulated to 3.3 V using an AP2112K-3.3 low-dropout linear regulator. The regulator supports up to 600 mA output current, providing sufficient margin for the STM32, ToF sensor, USB interface, and supporting peripherals.

Input and output capacitors of 4.7 µF are placed according to the regulator datasheet to ensure stability. Additional local decoupling capacitors are distributed throughout the board to suppress high-frequency noise.

Worst-case LDO power dissipation is calculated as:

$P = (V_{in} - V_{out}) \times I$

Assuming a conservative peak current of 250 mA:

$P = (5\ V - 3.3\ V) \times 0.25\ A = 0.425\ W$

This value remains within safe thermal limits for the selected regulator.

Voltage test points are included to allow oscilloscope-based validation of ripple and load response during system integration.

2.3 STM32F401 Embedded Control Subsystem

The STM32F401 serves as the deterministic control core of the system. It manages presence detection, system state transitions, peripheral communication, and hardware-level interfacing.

Clocking and Stability

The microcontroller uses a 25 MHz external high-speed crystal oscillator (HSE) with matched load capacitors to provide a stable clock reference. The internal PLL multiplies the external clock to generate the required system frequency for peripheral operation and real-time control.

Each VDD pin includes a dedicated 100 nF ceramic decoupling capacitor to minimize supply noise. A 2.2 µF capacitor is connected to the VCAP pin in accordance with STMicroelectronics

guidelines to stabilize the internal regulator. Proper decoupling ensures reliable operation under high-frequency switching conditions.

Reset and Boot Configuration

The NRST pin includes a 10 kΩ pull-up resistor and a 100 nF capacitor to ensure stable power-on reset behavior. A manual reset pushbutton is provided for debugging and recovery. BOOT0 is pulled low by default to ensure boot from internal flash memory during normal operation.

Brown-out reset protection is enabled at 2.9 V to prevent flash corruption during undervoltage conditions. An independent watchdog timer (IWDG) is implemented to automatically reset the system in the event of firmware lockup.

2.4 Presence Detection Subsystem

Presence detection is performed using a Time-of-Flight (ToF) distance sensor connected to the STM32 via an I2C interface.

Distance is calculated using the relationship:

$$d = (c \times t) / 2$$

where d is the measured distance, c is the speed of light, and t is the measured round-trip time of the infrared pulse.

The I2C bus includes 5.1 kΩ pull-up resistors to 3.3 V. Additional GPIO lines are used for sensor shutdown (XSHUT) and interrupt signaling.

The STM32 polls the ToF sensor at approximately 15 Hz. A valid interaction zone is defined as a measured distance between 1.5 m and 3.5 m, ensuring proper camera framing and user comfort.

To reduce false activations due to transient motion or measurement noise, the measured distance must remain within this range continuously for 500 ms before presence is confirmed. Total activation latency is therefore approximately 505 ms, satisfying the requirement for automatic activation within one second.

2.5 Raspberry Pi Perception and Computer Vision Subsystem

The Raspberry Pi 4 performs all high-level perception tasks within a Linux-based environment. A Logitech C920 USB camera captures video at 1280 × 720 resolution and 30 frames per second. The camera's 78-degree horizontal field of view ensures consistent framing within the defined interaction zone.

To reduce computational load, frames are downscaled to 640 × 480 prior to processing. This reduces pixel count by approximately 56 percent while maintaining sufficient resolution for reliable face detection and hand tracking. Frame acquisition and preprocessing latency is approximately 20 to 30 milliseconds per frame under nominal conditions.

The perception pipeline consists of three stages: hand tracking, face detection and validation, and profile selection.

Hand tracking is implemented using MediaPipe-based landmark detection. For each frame, 21 hand landmarks are extracted when visible. The index fingertip coordinate is used as a cursor reference and proportionally mapped from camera space to display space using scaled coordinate transformation.

Cursor updates are transmitted to the STM32 at approximately 30 Hz via UART. Selection is implemented using a dwell-time threshold of 1.5 seconds to prevent accidental activation.

Face detection is performed using optimized classifiers suitable for real-time inference on the Raspberry Pi. To improve reliability and reduce false positives, a composite confidence score is calculated:

$$C = 0.4P_d + 0.3(A_{face} / A_{frame}) + 0.3 \bullet V_{laplacian}$$

where P_d represents detection probability, A_{face} is the bounding box area, A_{frame} is the total frame area, and $V_{laplacian}$ is a Laplacian-based sharpness metric. A threshold of C greater than 0.65 is required before proceeding to profile selection.

The system does not perform biometric identification. Instead, after a STEM category is selected, the Raspberry Pi filters a predefined dataset of scientist profiles and selects a representative match based on metadata constraints. The selected 8-bit profile ID is transmitted to the STM32.

End-to-end perception latency, including capture, inference, communication, and rendering, remains below 65 milliseconds, ensuring a responsive and natural interaction experience.

2.6 System Control and Finite State Machine

The STM32 executes a deterministic finite state machine within a main control loop operating at approximately 1 kHz. The FSM governs overall system behavior and ensures structured transitions between operating modes.

Primary states include:

1. Idle (Mirror Mode)

2. UI Selection Mode

3. Face Validation Mode

4. Video Playback Mode

5. Timeout Return to Idle

State transitions are triggered by validated presence detection, UART messages from the Raspberry Pi, face confidence threshold confirmation, and inactivity timeouts.

If no presence is detected for 10 seconds, the system automatically returns to Idle mode to reduce unnecessary processing and power consumption.

Because the STM32 operates without a preemptive operating system, state transitions occur without scheduling variability, ensuring sub-millisecond responsiveness to internal and external events.

2.7 Communication and Protection Features

Communication between subsystems is limited to low-bandwidth command packets over UART, minimizing timing interference and reducing integration complexity.

The USB interface includes ESD protection using a USB transient suppression device to protect against electrostatic discharge events. Reverse polarity protection is implemented on the 5 V input rail.

These safeguards enhance long-term system reliability and ensure safe operation in public interactive environments.

2.8 System-Level Integration Rationale

The Smart Mirror architecture combines Linux-based artificial intelligence processing with microcontroller-based deterministic control. This hybrid design enables sophisticated perception capabilities while maintaining predictable timing for hardware operations.

By isolating perception algorithms, minimizing inter-processor communication bandwidth, and incorporating hardware-level protection features, the system achieves a balance between computational flexibility and embedded reliability.

This structured architecture ensures that the Smart Mirror meets quantitative requirements for activation latency, detection reliability, UI responsiveness, and electrical stability while remaining scalable for future expansion.

2.9 STM PCB
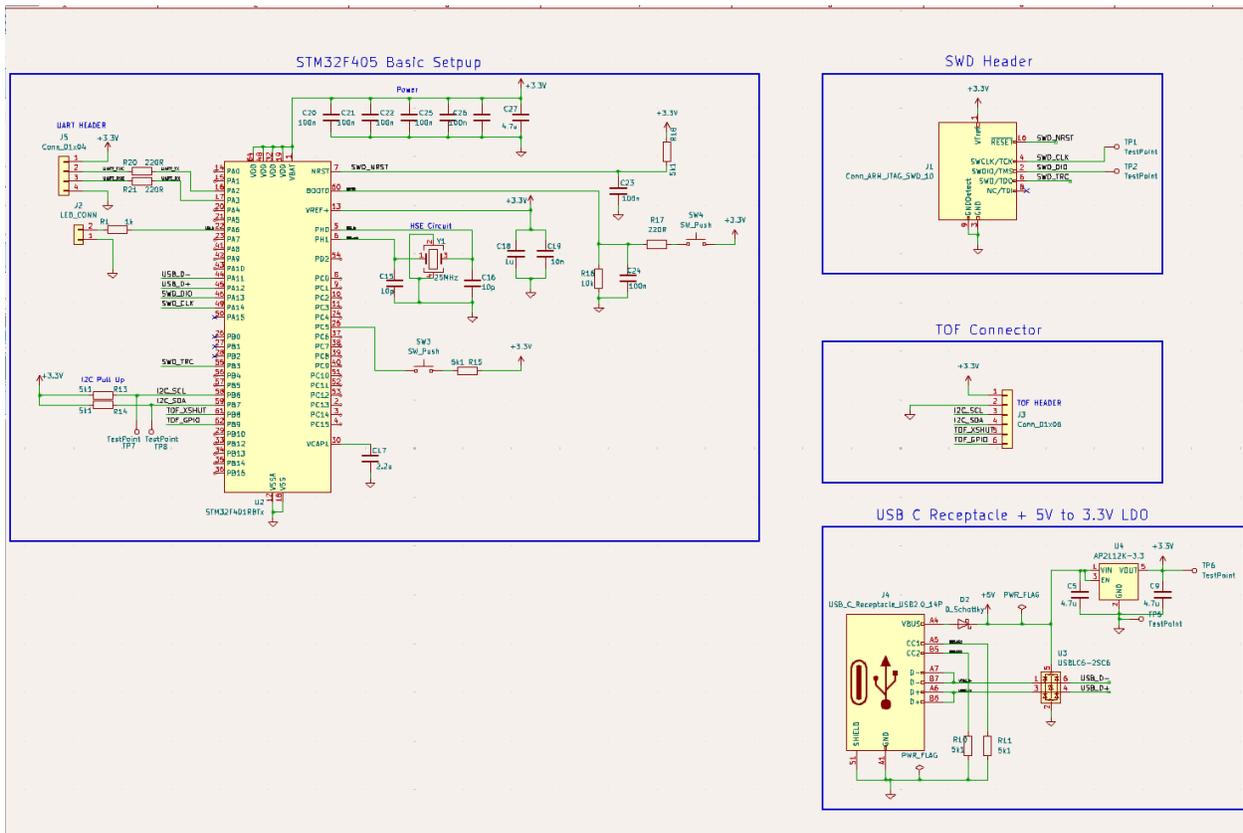
3. Societal Impact, Engineering Standards, Ethics, and Safety Considerations

3.1 Societal Impact

The Smart Mirror project promotes STEM awareness through an interactive, touchless educational interface. By presenting curated profiles of scientists and engineers, the system contributes positively to public welfare by encouraging interest in science and technology careers. The design emphasizes inclusive representation across disciplines and backgrounds in order to reduce stereotypes and broaden participation in STEM fields. Because the system is built using commercially available, low-cost embedded components, it can be replicated in schools, museums, and community environments. Environmentally, the system uses energy-efficient electronics and reduces reliance on printed materials. Additionally, all image processing is performed locally on the device, avoiding external data transmission and supporting a privacy-conscious and globally responsible design approach.

3.2 Engineering Standards

The project was developed in alignment with professional engineering standards and best practices. System testing and verification procedures reflect principles found in IEEE standards for hardware and software validation. The embedded processing and computer vision components are guided by accepted software engineering practices emphasizing reliability, transparency, and maintainability. Although this prototype is not intended for commercial distribution, the electrical design follows general safety principles consistent with UL and IEC standards for low-voltage information technology equipment. All components operate within manufacturer-specified voltage and current ratings, and voltage regulation circuitry was implemented according to published datasheets to ensure safe operation.

3.3 Ethical Considerations

Ethical considerations were central to the architecture of the Smart Mirror. Because the system uses a camera for presence detection, gesture tracking, and facial feature matching, privacy was treated as a primary concern. All image processing occurs locally on the Raspberry Pi, and no images or identifying data are stored or transmitted. The system does not perform biometric identification, facial recognition, or identity verification. Instead, it uses confidence-based detection metrics to determine whether a valid face is present before displaying educational content. The interface clearly indicates when the camera is active, ensuring transparency and informed user interaction.

Another ethical consideration is performance consistency across users. Computer vision algorithms may behave differently under varying lighting conditions, skin tones, or facial features. To mitigate this concern, widely validated detection models were used and tested under varied environmental conditions. The system does not assign traits, intelligence, or personal characteristics to users; it only verifies face presence using measurable image quality and detection confidence metrics. By limiting functionality to non-identifying, educational outputs, the project avoids reinforcing stereotypes or making evaluative judgments.

These decisions align with IEEE and ACM ethical principles requiring engineers to minimize harm, respect privacy, ensure fairness, and maintain transparency in system capabilities and limitations.

3.4 Electrical and Mechanical Safety

Electrical safety was addressed by operating the system entirely at low-voltage DC levels. A 12V input supply is regulated to 5V and 3.3V for embedded subsystems using properly rated regulators. All components function within specified voltage and current limits, and protective circuit elements such as decoupling capacitors and polarity protection were included. Wiring is insulated, connections are secured, and exposed conductors are minimized to reduce the risk of short circuits or accidental contact. Voltage and current levels were verified prior to full system integration.

Thermal management was considered for both the Raspberry Pi and display assembly. Heat sinks were installed on high-power components, and ventilation spacing was maintained behind the display to prevent overheating. Continuous operation testing confirmed stable temperatures within manufacturer-recommended operating ranges.

Mechanical safety concerns primarily involve the 32-inch display and acrylic mirror panel. The enclosure was reinforced with stable mounting hardware to prevent tipping or detachment. Acrylic edges were smoothed to reduce injury risk, and strain relief was implemented on cables to prevent mechanical stress at connection points.

3.5 Mitigation Procedures and Ongoing Considerations

All assembly and testing were conducted with power disconnected during circuit modifications. Voltage levels were verified with a multimeter prior to subsystem connection. The enclosure was tested for structural stability before demonstration. Since the original proposal, additional attention was given to user perception of surveillance and mechanical stability; these concerns were addressed through clearer operational indicators and reinforced mounting structures.

The project does not involve high voltage, hazardous chemicals, lithium battery packs, or moving mechanical actuators, which reduces overall system risk. Through careful electrical design, local image processing, and mechanical reinforcement, the Smart Mirror sufficiently protects both users and developers while delivering an educational and socially beneficial system.

4. Cost and Schedule

4.1 Cost Analysis

This section presents a complete cost analysis of the Smart Mirror project, including labor estimates, parts costs, and the total projected system cost. While the prototype hardware cost remains within the allocated course budget, a full engineering cost assessment includes estimated professional labor.

4.1.1 Labor Costs

To estimate labor cost, we assume a reasonable starting salary for an Electrical or Computer Engineering graduate from the University of Illinois. Based on industry averages, an entry-level salary of approximately $85,000 per year corresponds to roughly $40 per hour (assuming 40 hours per week over 52 weeks).

Per course guidelines, labor cost is calculated as:

Hourly Rate × 2.5 × Total Hours Worked

Each team member contributed approximately 120 hours over the semester.

Partner 1:
 $40/hour × 2.5 × 120 hours = $12,000

Partner 2:
 $40/hour × 2.5 × 120 hours = $12,000

Total Labor Cost:
 $12,000 + $12,000 = $24,000

Although this labor cost is not part of the prototype budget, it reflects the true engineering value of the project.

4.1.2 Parts Cost

The following table summarizes the non-standard components required for the project. Development boards supplied by the department are listed with zero cost where applicable.

Table 4.1 – Parts List

VL53L3CX ToF Distance Sensor (Breakout)
 Manufacturer: DFRobot
 Part Number: SEN0378
 Quantity: 1
 Unit Cost: $8.90
 Total: $8.90

LM2596 DC-DC Buck Converter Module
 Manufacturer: Generic (TI design reference)
 Part Number: LM2596
 Quantity: 2
 Unit Cost: $7.99
 Total: $15.98

MCP1700-3302E 3.3V Linear Regulator
 Manufacturer: Microchip
 Part Number: MCP1700-3302E
 Quantity: 5
 Unit Cost: $0.60
 Total: $3.00

AP2112K-3.3 LDO
 Manufacturer: Diodes Incorporated
 Part Number: AP2112K-3.3TRG1
 Quantity: 5
 Unit Cost: $0.22
 Total: $1.10

USBLC6-2SC6 ESD Protection
 Manufacturer: STMicroelectronics
 Part Number: USBLC6-2SC6
 Quantity: 5
 Unit Cost: $0.36
 Total: $1.80

USB-C Receptacle
 Manufacturer: GCT
 Part Number: USB4105-GF-A
 Quantity: 3
 Unit Cost: $0.78
 Total: $2.34

25 MHz Crystal
 Manufacturer: KYOCERA AVX
 Part Number: CX2016SA25000DFPLG1
 Quantity: 1
 Unit Cost: $0.54
 Total: $0.54

10-Pin SMD Header
 Manufacturer: Samtec
 Part Number: FTSH-105-01-L-DV-007-K
 Quantity: 1
 Unit Cost: $2.68
 Total: $2.68

PCB Components Subtotal: $36.34

Additional System Components:

Raspberry Pi 4 (4GB)
 Manufacturer: Raspberry Pi Foundation
 Quantity: 1
 Cost: $75.00

32-inch LED Display
 Manufacturer: Dynex
 Quantity: 1
 Cost: $40.00

Acrylic Two-Way Mirror Panels
 Quantity: 6
 Cost: $45.00

Wiring, Connectors, Mounting Hardware
 Cost: $25.00

Total Estimated Parts Cost:
$36.34 + $75 + $40 + $45 + $25 = $221.34

4.1.3 Grand Total Cost

Total Labor Cost: $24,000
 Total Parts Cost: $221.34

Grand Total Project Value: $24,221.34

While the hardware prototype cost remains under the $300 course budget, the full engineering value of the project exceeds $24,000 when labor is included.

4.2 Schedule

The Smart Mirror project was completed over a 16-week semester. The development timeline is outlined below.

Weeks 1–3: System Architecture and Planning
 The overall system architecture was defined, including subsystem division between the STM32 microcontroller and Raspberry Pi. Display and mirror materials were selected, and responsibilities were assigned.

Weeks 4–6: Subsystem Design
 Power regulation circuitry and PCB schematics were designed. The computer vision and confidence-based facial detection pipeline was developed. Components were ordered.

Weeks 7–9: Hardware Assembly
 The PCB was assembled and validated. The Raspberry Pi, display, camera, and Time-of-Flight sensor were integrated. Enclosure construction began.

Weeks 10–12: System Integration
 UART communication between the STM32 and Raspberry Pi was implemented. Gesture mapping and display control were finalized. Subsystem-level testing transitioned to full-system integration testing.

Weeks 13–14: Testing and Refinement
 The system was tested under varying lighting conditions and distances. Confidence thresholds for facial matching were refined. Mechanical stability improvements were implemented.

Weeks 15–16: Final Validation and Demonstration Preparation
 Final system validation was performed, documentation was completed, and safety checks were conducted prior to demonstration.

Responsibilities were divided according to subsystem expertise during early design stages, with both team members collaborating closely during integration, debugging, and final validation.