Personal Autonomous Non-invasive Device for Assisting Kids (PANDA Kids)

Capstone Proposal

Group Members:
Mohib Azam
Spenser Cheung
Daravichea Lim
Spencer Pozder
Alyxandra Spikerman
Trung Tran

Advisor:
Bahram Shafai

Table of Contents

I. Abstract

The purpose of this project is to implement a machine-learning and AI driven approach to children monitoring and care. The system we intend to create will be robust enough to provide usefulness in a few different scenarios: aiding nurses, teachers, and guardians in monitoring the wellbeing of children. Our platform will be built on a strong foundation, employing advanced hardware that powers the main computational load of image processing and machine learning. It will also use other hardware elements to aid user interaction, rounding out the system and making it more user-friendly. The software that controls the system will be built using state-of-the-art machine learning paradigms and will have an interface building on robust UI technology capable of displaying useful and efficient information to the user. This proposal synthesizes our problem formulation with a detailed analysis of our solution and gives an extensive overview of what we plan to accomplish in the Fall semester.

II. Introduction

Hospitals and classrooms are very busy environments. Nurses and teachers don't have enough time to attend to each individual child at all moments of the day. Based off of prior personal experiences with interning in hospital and school environments, our team came to the conclusion that if we were able to attend to a child's needs without using an employee, then we could make both parties' lives much easier.

Child-care is a field that occurs around the world, regardless of culture and language barriers. Since our product aims to help individuals with child-care, our product will be usable in multiple scenarios and be usable in multiple locations. Based off of a 2014 study at the Vickie Milazzo Institute, nurses were rarely getting 7-8 hours of sleep and could be expected to work a 12 hour shift some days. Another article we found touched on the notion that nurses being put on these long shifts can lead to them making more mistakes with patients and to a decline in nurse's health. At the same time, as nurses are being affected, imagine the children they're tending to as well. In a recent study between Northeastern University and Boston Children's Hospital, researchers found that when children are in a positive mood, the likelihood of them recovering from an illness improves.

Figure 1: Hospital Staff Moving Quickly from Patient to Patient

Another sector that would benefit from our product would be daycares. A recent Center for American Progress study found that mothers who live in childcare deserts—areas with an undersupply of childcare—have lower rates of workforce participation than mothers in non-desert areas. This leads to stressful environments for both the parents and their children. It also causes economic strain on the family from the caregiver's perspective, taking care of multiple children in one setting can be a daunting task, let alone most states having a minimum requirement of 1:10 ratio class size. Some states don't even have a maximum limit for student teacher ratio; class sizes average about 20 students for most states. The ramifications from high turnover rate of 16% every year and teachers either changing to another school or leaving the education sector left students in disarray. The abrupt change of teachers any time in the year and issue with staffing are crucial to a child's performance. This is due to the lack of support in general from the school where they have to plan their own lessons and monitor each individual child coupled with a high student-teacher ratio. A study in OECD showed that the crucial stage is in early education, where younger students need more time and attention given to them to have a quality education. Datapoints by the U.S Bureau of the Census showed that a class size reduction from 23 to 20 students leads to a decrease of 4.9% in a teacher's turnover rate. However, with limited resources and finite teachers to allocate around, class size reduction alone might not be an optimal solution.

Our solution aims to help reduce a nurse's work-load by aiding a child when a nurse is not around—such as playing a video when a child is sad or detecting an emergency situation with the child—while also enabling the child to have a friendly companion. The method we plan to implement for determining the optimal response to a child's mood is purely camera-based. By doing so, we make our solution inexpensive, noninvasive—as no sensors or gadgets need to come in contact with the child—and usable with multiple children.

The Personal Autonomous Non-invasive Device for Assisting Kids (PANDA Kids) is a device that uses facial and voice recognition, hand in hand with artificial intelligence in order to interact with children between the age range of 4 to 12 in the childcare sector, such as hospitals, daycares, and bedrooms.

The intent is to use a camera that captures children's emotions to prompt the AI implemented through our developer's kit to play certain pre-loaded videos through a display module that will help improve the child's mood. Voice recognition will be used to identify simple prompts which will be used as a communication tool for a child in need of assistance.

As the product is intended to be used commercially, the different components will be enclosed in a friendly looking panda done through 3D-printing. Wirings and other background components will be kept inside the enclosure for safety and the only visible components will be the camera, microphone and screen.

III. Problem Formulation

As the internet of things (IoT) sector in consumer products is growing, our goal is to integrate facial recognition and voice processing with a consumer robot that children can interact with. We have three target markets: hospitals, classrooms, and daycare centers. Our product will have three main components: AI-driven facial recognition, AI-driven voice processing, and a robot with a display, microphone, camera, and speaker.

For the target markets, we chose each one for the issues that each industry faces. At hospitals, nurses are often running around their section, assisting each patient efficiently. Children are often searching for a companion to keep them company—a role which nurses are often not able to take. Children's hospitals can also have long-term patients who deserve a companion with them. PANDA Kids can take on this companion role to the child and assist them in displaying entertainment for the child based on detected emotions or voice commands. We also plan to build in emergency-detecting functionality to determine if the child is in need of emergency care. This could be with signals such as intense pain on the face or the child calling for the nurse, doctor, or their parents.

Figure 2: Child Patient on a Tablet

In classrooms and daycare centers, we plan to take a more educational approach to the robot. PANDA Kids will act as either a personal device or a device for the classroom, playing educational videos and interacting with the children in a fun and engaging way. PANDA Kids will have a vibrant graphical user interface (GUI) and its panda shape will appeal to children.

Figure 3: Children in a Classroom on Electronics

Finally, in children's rooms at home, PANDA Kids will act as a mixture of a baby monitor and an entertainment system. With the camera and microphone, PANDA Kids can detect an emergency situation with the child—as in the hospital environment—while also allowing the child to turn on or off entertainment. We can make the entertainment educationally focused and build in specific videos to display so the child doesn't wander into media that's too old for their age group. As a baby monitor for a target audience of 4-12 years old, the device will act more as a way for the parent to determine if the child is in need of immediate assistance.

Figure 4: Child on Tablet at Home

For our design, our robot will take the form of a panda enclosure. It will include a display for media, a microphone to listen to the child, a camera for facial recognition, and a speaker for both media and to provide feedback to the child. The robot’s software pipeline will include an AI system to detect facial emotions and do voice recognition as well as methods for producing appropriate outputs to the speakers and display.

IV. Analysis

IV.I I/O System Design

Our system has three different inputs: a Video Feed, an Audio Feed, and UI Interactions. These different inputs must then be processed, and various outputs are produced. The following diagram lays out the general "flow" of information through the system:

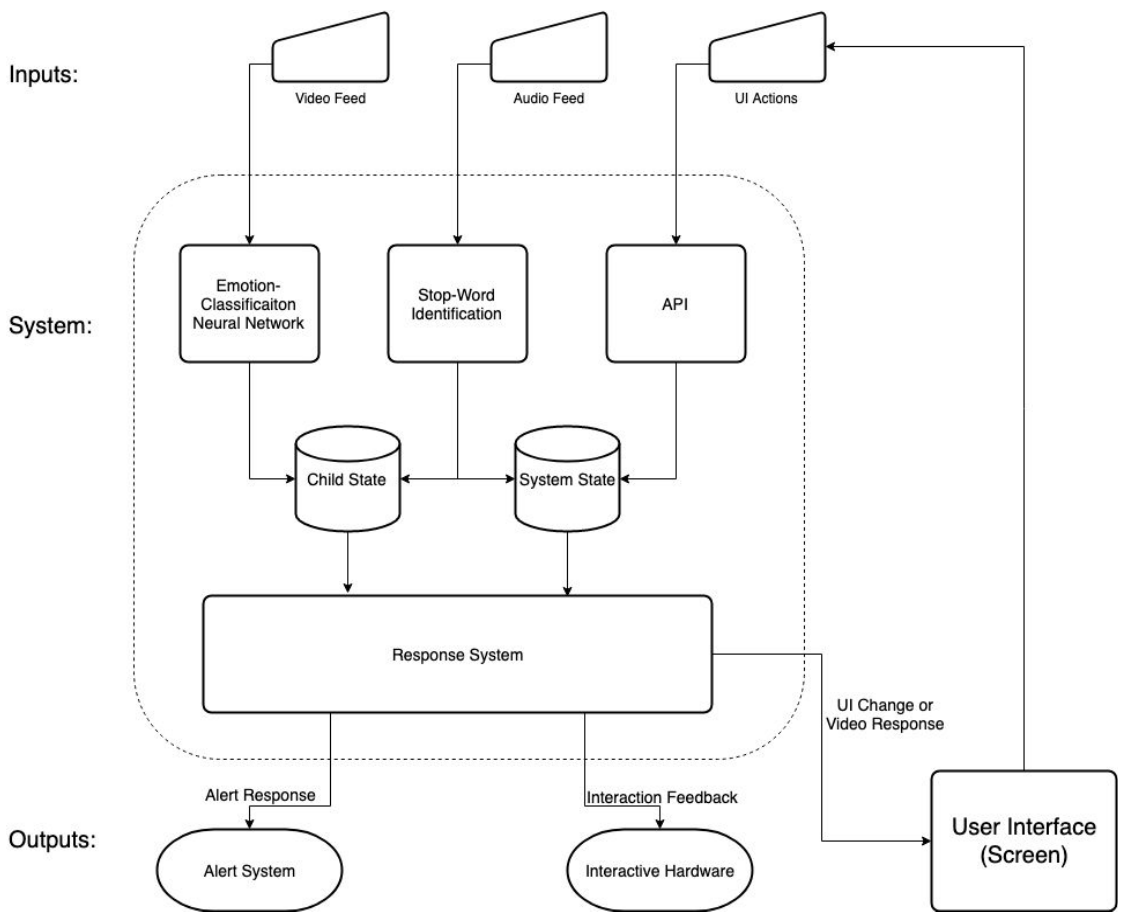

Figure 5: I/O System Design

IV.II Service Interaction

In order to manage the many inputs and outputs of our system, we have decided to design a modular, service-oriented system. This way, each input has a direct service to interact with, which then has a connection to an overall system "state" that impacts the different output-managing services. This is outlined in the following figure:

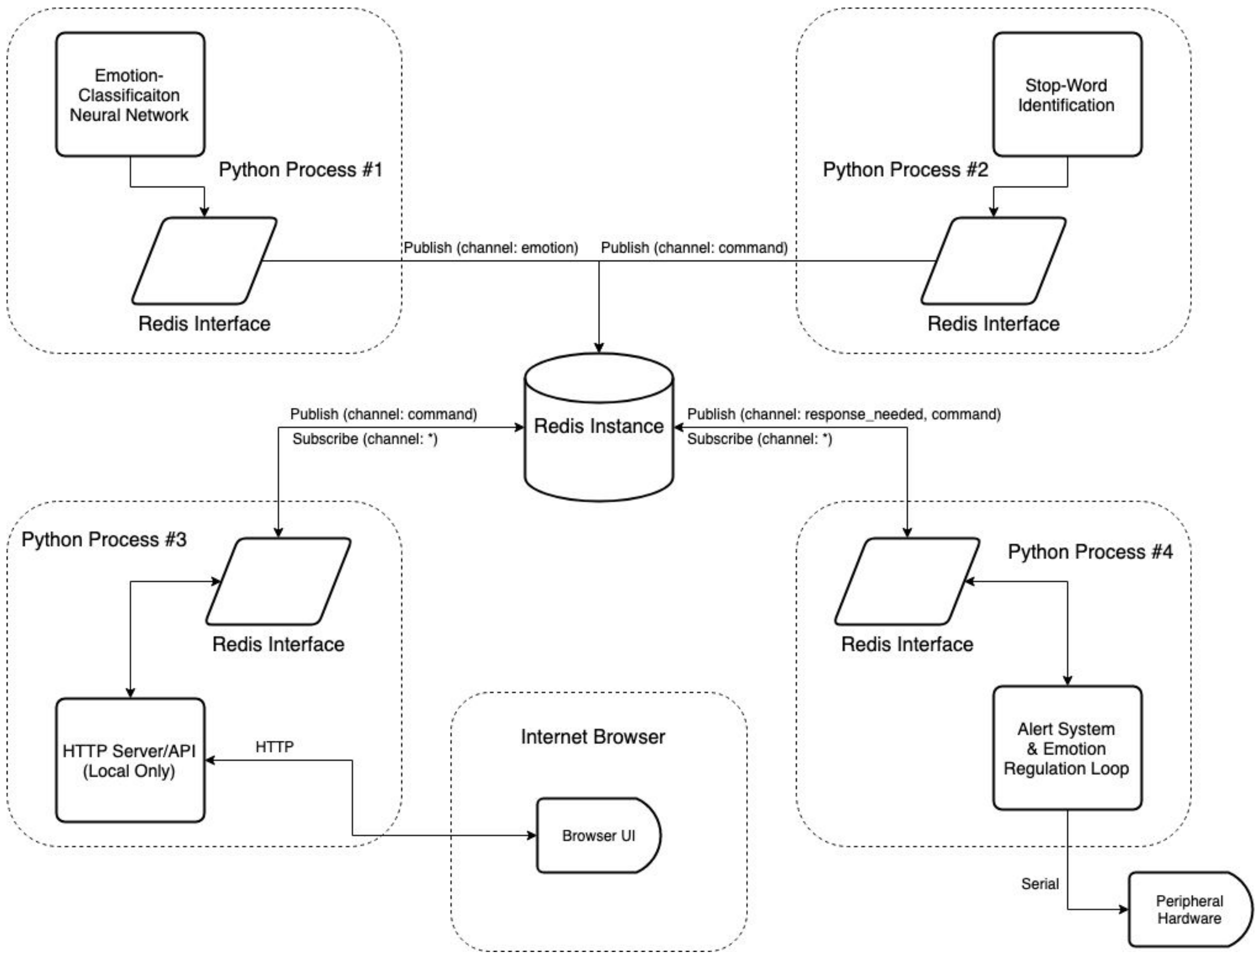

Figure 6: Response Handling, Service Interaction

In general, each Python process handles a separate service, only interacting with other services through a Publish/Subscribe model, handled by a local Redis instance. This way, each module can be separately developed, as long as the expected messages are published or received to/from the Redis server. This also allows us to take full advantage of our powerful computing hardware: we chose the TX2 for its quad-core Cortex-A57 processor, so each Python process can take advantage of a full core of computing power.

IV.III Response Flow

Within the “Alert System and Emotion Regulation Loop” defined in Figure 5, there is a specific response loop that must be maintained in order to correctly handle the different inputs and system states. The flow for this loop is shown in the figure below:

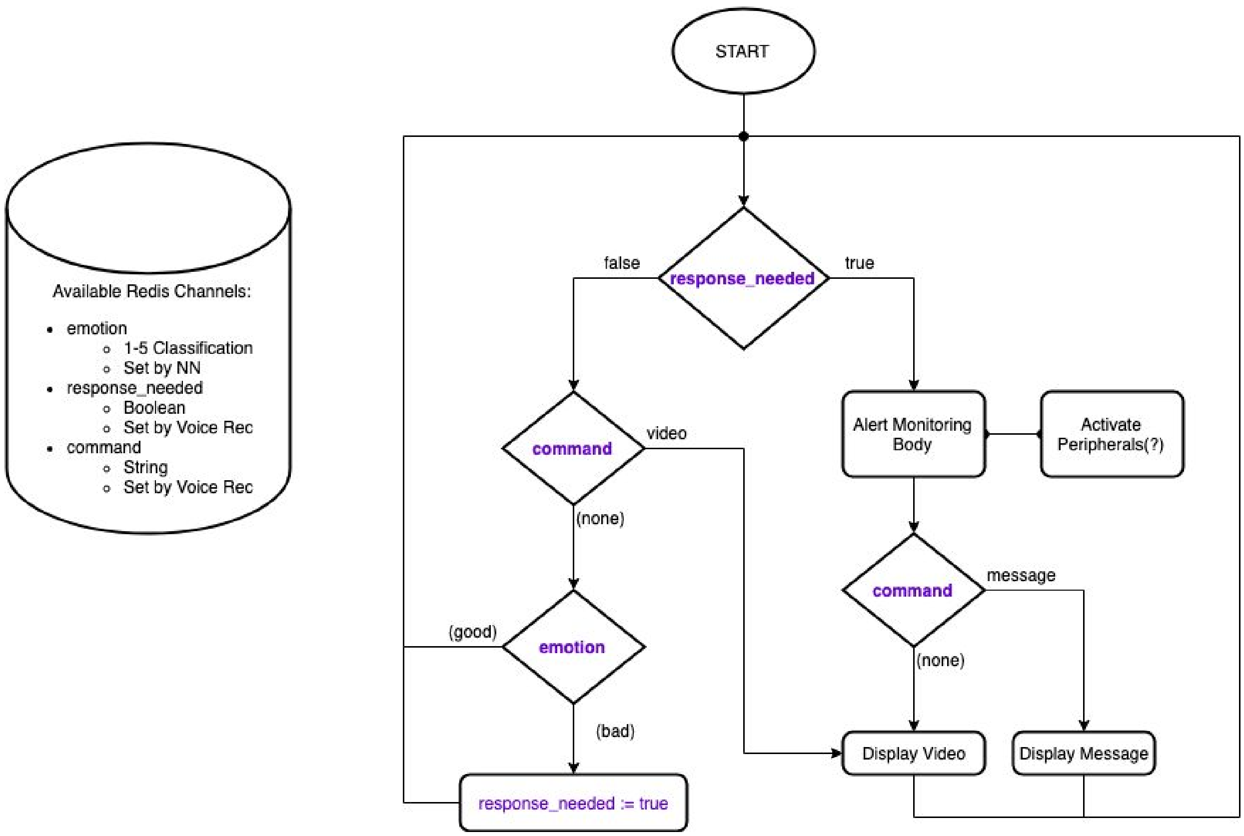

Figure 7: Alert System and Emotion Regulation Loop

Each available Redis channel is subscribed to by the regulator system Python process, allowing for these flow decisions to be made instantly. When another service changes the system state by publishing to one of these channels, the loop picks it up and is immediately able to act on it accordingly.

IV.IV Emotion Classification

The emotion classifier must be designed to be performant and accurate, so that it can drive device actions based on the mood of the child(ren) being monitored. Its architecture is defined in the figure below:

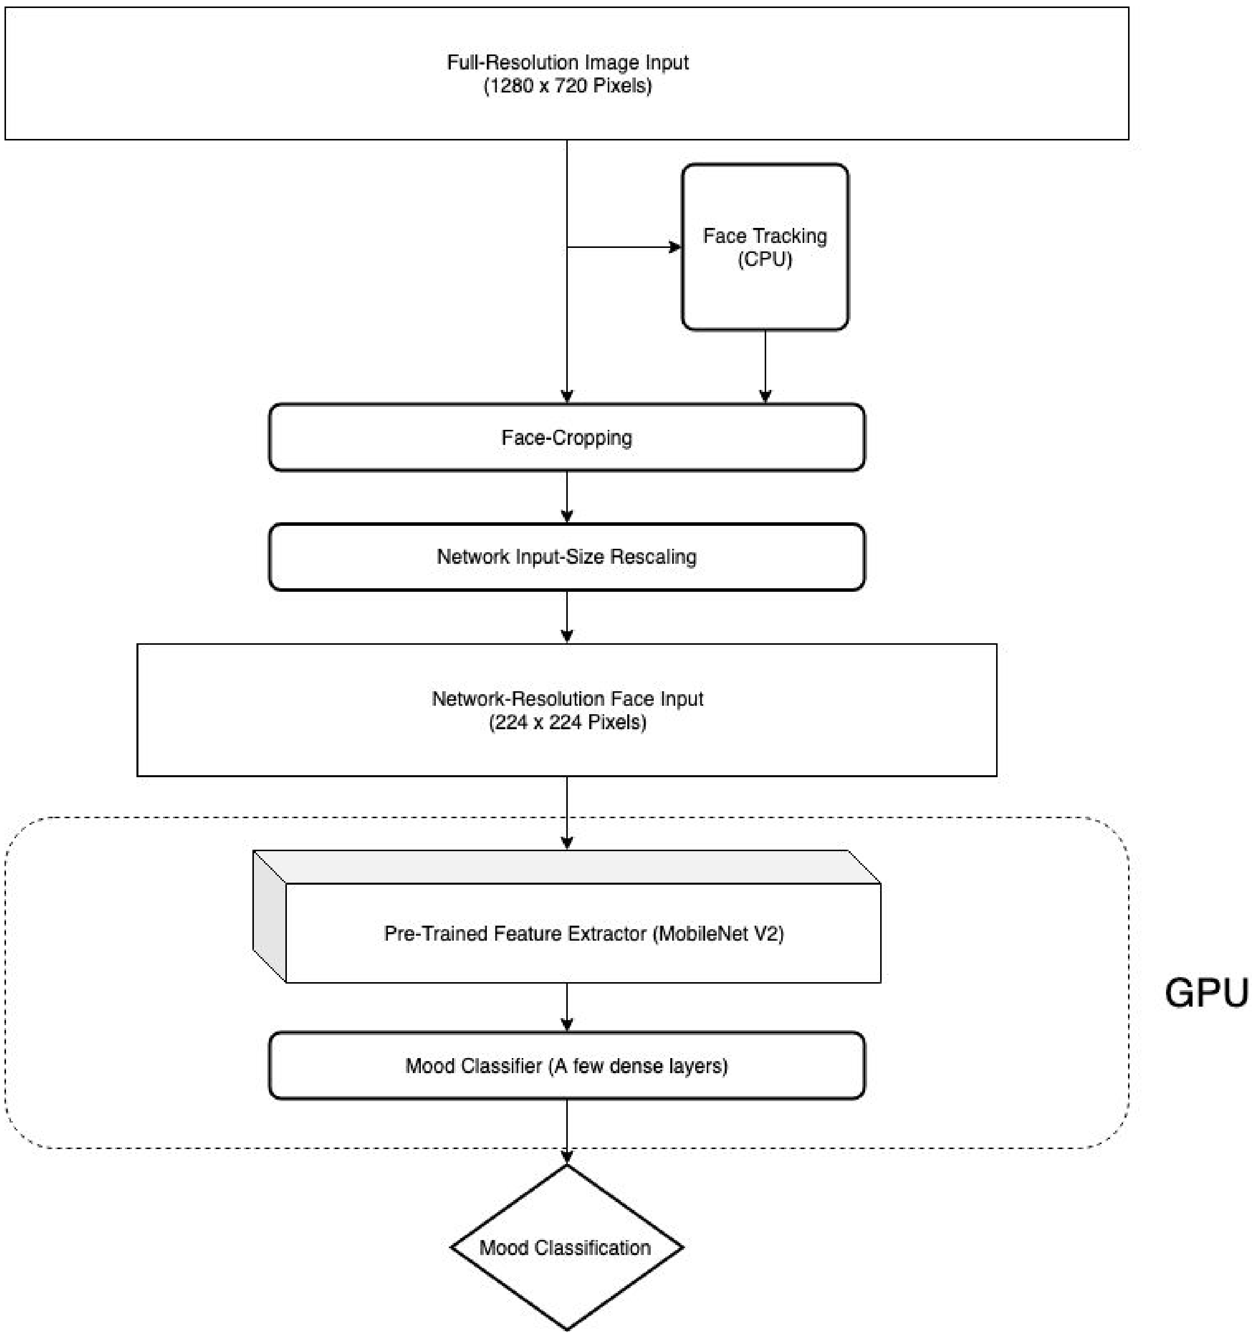

Figure 8: Emotion-Classification Neural Network Architecture

In order to keep the network performant, and to reduce the amount of training data required for its creation, we have decided to use Transfer Learning and incorporate the MobileNetV2 neural network model into our own network. This model was designed by Google to allow for fast and accurate feature extraction from images on embedded and mobile devices, which is perfect for our application. We will then follow this model with a small number of dense network layers we train ourselves, using data from public databases, to teach the network to classify face images into a set of emotional categories.

IV.V Hardware Architecture

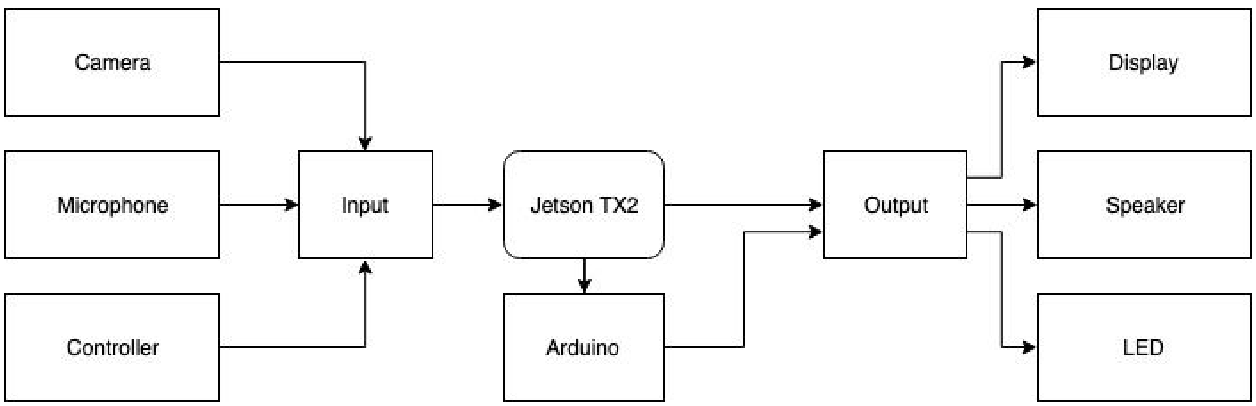

Figure 9: Hardware Architecture Workflow

Above is a top level scheme of the workflow of the hardware components and how they interact with each other. The input will be made up of the camera, microphone, and controller specifically for caregiver or parents to set up and modify the device. The information will then be funneled to the Jetson TX2 for further processes. After receiving input, the Jetson TX2 will then do image processing, voice recognition, and controls processes and send signals to the Arduino. Additional command outputs will then be sent to the Display and Speaker via the Jetson TX2 and LEDs via Arduino.

V. Design Strategies & Approach

V.I Software

Our software pipeline will be developed in Python. We've chosen it due to it being easy to use and having numerous resources for relevant processes we need to implement. For the AI, Computer Vision, and Machine Learning tasks we wish to implement, Keras, Tensorflow, and OpenCV are commonly used resources for object detection, facial recognition, and building and training neural networks. To tackle the voice recognition aspects of our project, we will build upon open source speech recognition code and design a method to pick up keywords (i.e. 'Mommy', 'Daddy', 'help', 'pause', 'play', etc.) If time is left over, we will tackle the issue of noise reduction and find a way to classify if a voice is coming from a child or not: one method we are considering is using a known frequency range of children's voices and creating a bandpass filter to eliminate any frequencies outside that range. To bring together our software components for a good UI/UX experience, we will be utilizing Flask to create a GUI in the form of a local web server. This allows us to test across multiple platforms and allows us to utilize our HTML and CSS knowledge to create responsive and colorful pages.

V.I.i. AI Facial Recognition

Our software will include an AI system for detecting children's faces and identifying their emotions accordingly. A package already exists for facial recognition in Python called face_recognition, which extracts relevant features on the face. Using these features, we will use a neural network to allow the AI to identify the emotions expressed by the faces in real-time. The network will be programmed using Keras and Tensorflow. A variety of datasets featuring both images of faces and their respective emotions exist for us to train the network.

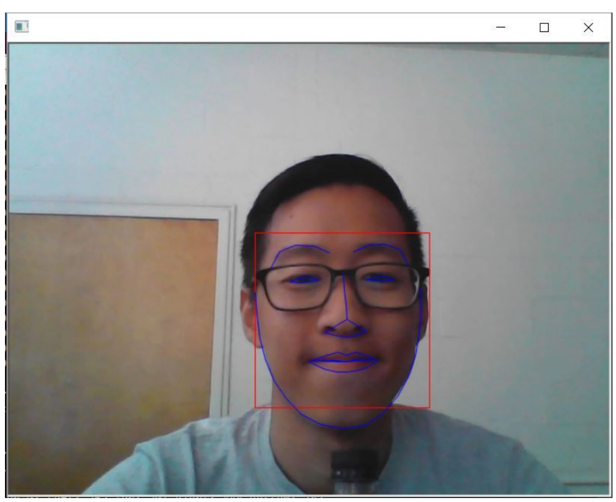

Figure 10: Detection output of dlib's facial_landmark_detection code

V.I.ii. Voice Recognition

Our software will include a way to handle voice recognition of a child. There is an existing speech recognition module for Python called SpeechRecognition that we'll be using to cover this aspect.

Our plan is to create responses in reaction to determined keywords for emergency situations or child vocal controls. The responses for the emergency situations will focus on notifying a nurse, doctor, teacher, and/or parent whereas the control responses will directly affect the GUI.

The voice recognition component is vital, as it will allow us to help detect emergency situations based on keywords that a child says. This also allows the child to control the device rather than interacting directly with the GUI (as we cannot assume that the child would either know how to or have the motor capabilities to use a mouse or other external hands-on device).

V.I.iii. GUI

As our device will be interacting with children, we will be building a bright and colorful GUI to display media. We will be using Flask to create a local web server. Flask is a lightweight web application framework that is very easy to learn and use to create a quick website. Flask will allow us to embed YouTube videos for display and create a variety of pages for different settings—such as a hospital versus a classroom versus a home. We will also integrate Bootstrap which is a CSS framework that allows for customization and makes it easy to build responsive and colorful web pages. Our server will constantly send information to the webpage to determine the display.

We also considered creating a GUI through Python with wxPython, but decided against it as wxPython is not as customizable as Flask. With Flask, HTML, and CSS, we can make any page come to life for the child and change the design based on where the device will be. Our team also has previous experience with Flask.

The GUI will be handling media display, as described above, and will handle parental control functionality in the form of a settings page. The parents or other caretakers will interact with the GUI via a mouse.

V.II Hardware Components

V.II.i Computer Component

The NVIDIA Jetson TX2 developer kit is the computer component of our device. It will be the cornerstone that connects the software and hardware aspects together. The Jetson TX2 has the computing power needed to process real time input in a reasonable time set and gives output adequately. The Jetson TX2 supports Nvidia Jetpack, which includes libraries for machine learning, computer vision, and GPU computing. The kit also includes I/O ports for other components required for the PANDA Kids system, such as an HDMI port, USB port, camera port and audio port. The kit has Wi-Fi and Bluetooth capability, allowing streaming of videos. The Jetson Nano and Raspberry Pi 4 were also considered, however given our budget and the scope of computation that the system must be able to do, these two products do not give us enough performance to achieve the goals of our application.

Figure 11: Jetson TX2 Developer Kit

V.II.ii Microphone

The microphone is used to detect voice inputs. The inputs will then be digitized to our system to capture the child's voice commands. The microphone will be outputted through a USB port available on the Jetson TX2 kit and be implemented into one of the eyes of the PANDA Kids. The requirements for the microphone includes the ability to recognize voice in the range of 300-1300Hz, which is the frequency children from 4-12 years old, and the ability to have a USB output in order for it to be implemented on to the kit. The Saramonic SR-ULM5 microphone meets these criteria. It is omnidirectional and has a frequency response of 35Hz to 18kHz while having a 2 meter long cable with a USB output.

Figure 12: Saramonic SR-ULM5 Microphone

V.II.iii Camera

The camera is used to capture the child's motions and facial expression. The camera is the most important aspect of the system in terms of input data, as it is crucial to our machine learning algorithm to be able to output the child's motions and current state. The most important requirement of the camera is that it must be compatible with the developer's kit in order for the facial recognition to work. Other criteria include being compact, having a wide angle of view and a decent frame rate that allows the facial recognition function to run smoothly. The IMX219-120 8-megapixel camera has been chosen as one of its main functions is that it is specifically made to be compatible with the chosen developer's kit. The camera has an angle of view of 120 degrees and a frame rate of 30fps. In addition, the camera is compact, having a dimension of only 25mm x 24mm, which will allow it to be implemented inside the panda enclosure. Considering the camera will be attached to one of the panda's eyes, this small camera is ideal. Another option that was available was the Eivotor 720P webcam which was already on hand. Unfortunately the camera is too large and will not fit into the enclosure.

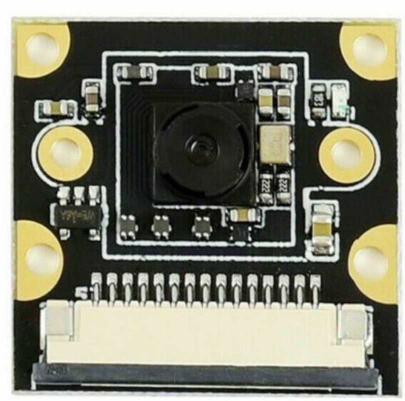

Figure 13: IMX219-120 Camera

Figure 14: Eivotor 720P Camera

V.II.iv Display Module

The display module is the communication instrument between the child and the system. The screen should be colored to compliment our immersive GUI. An ideal monitor would be small enough to be compact but also be large enough for easy-viewing for the child. In addition, a requirement is that the display module needs to have an HDMI port which will allow it to be connected to the developer's kit. After comparing two monitors, the Esoku 10.1" Monitor and the SunFounder 7" Monitor, we decided to pick the former. This option exceeds the criteria of an ideal monitor. Both monitors have the same display quality at HD 1024x600, an HDMI port and as a bonus, both have a built-in speaker. The difference is that the Esoku monitor is three inches larger while only being $5 more expensive than the SunFounder monitor.

Figure 15: Esoku 10.1" Monitor

Figure 16: SunFounder 7" Monitor

V.II.v Speaker

The speaker is used to output sound. It will be connected to the screen in which it will produce sound to match the display. Considering that it will be used indoors in small compact areas, the speaker will be small and only needs to be loud enough to be heard in a room. The built-in speaker for the Esoku monitor achieves this purpose and no additional external speaker needs to be purchased.

IV.II.vi Panda Enclosure

The enclosure casing is in the shape of a panda for commercial purposes to be visually appealing to the children. The enclosure will be the base that the hardware components attach to. The head of the panda will be used as a camera and a microphone while the main body will be used to hold up the screen in place along with the developer's kit along with the wiring. The top of the enclosure, the head of the panda, will be 3D printed (cost about $50) while the body needs to be strong and stable enough to withstand the weight of all the components—which will be approximately 3 lbs maximum.

Figure 17: Panda Enclosure Prototype

V.II.vii Arduino Board with LEDs

The main purpose of the arduino board will be to wire and power LED lights which will be used as an indication for interaction between the PANDA Kids and the user. For example, while the device is talking, the LED lights will light up when listening and responding to the child. The arduino board can also be used to implement any additional modules mentioned below.

V.III Future Enhancements

The scope of our project is defined in the previous sections, but there are additional future modifications that would benefit our product and its interactions with target audiences. A cooling mechanism to prevent overheating of the device when in use for long periods of time and the addition of a Servo that would rotate the enclosure's head to face the direction of input—either following the voice from the child in a direction off screen or the child's face as they move off screen

VI. Cost Analysis

Product	Quantity	Base Cost	Actual Cost
NVIDIA Jetson TX2 Developer Kit	1	$400	$299 (student discount)
Microphone	1	$29	$29
IMX219-120 Camera	1	$29	$29
Screen & Speaker	1	$61.89	$61.89
3D Printed Enclosure	1	$50	$50
Total Cost		$569.89	$468.89

VII. Division of Tasks

ID	Tasks	Mohib Azam (Software)	Spenser Cheung (Software)	Daravichea Lim (Hardware)	Spencer Pozder (Software)	Alyxandra Spikerman (Software, Team Leader)	Trung Tran (Hardware)
1	Facial Recognition Software	X	X		X		
2	Voice Processing Software	X	X		X		
3	Voice Responses (i.e. pre-determined responses)					X	
4	GUI				X	X	
5	CAD 3D Printing Head			X	X		X
6	CAD Body Enclosure Design			X	X		X
7	Arduino Board LED Implementation			X			X
8	Hardware Implementation & Assembly			X			X
9	Video Research	X	X	X		X	X
10	Final Testing	X	X	X	X	X	X
11	Code Reviews	X	X		X	X	
12	Project Write Up - Cleanup	X	X	X	X	X	X
13	Project Write Up - Abstract				X		
14	Project Write Up - Introduction			X		X	X
15	Project Write Up - Systems Analysis & Approach	X	X	X	X	X	X
16	Project Write Up - Conclusion					X	

VIII. Timeline

ID	Tasks	Week 1 (Sept 4-8)	Week 2 (Sept 9-15)	Week 3 (Sept 16-22)	Week 4 (Sept 23-29)	Week 5 (Sept 30-Oct 6)	Week 6 (Oct 7-13)	Week 7 (Oct 14-20)	Week 8 (Oct 21-27)	Week 9 (Oct 27-Nov 3)	Week 10 (Nov 4-10)	Week 11 (Nov 11-17)	Week 12 (Nov 18-24)	Week 13 (Nov 25-Dec 1)	Week 14-15 (Dec 2-13)
1	Facial Recognition Software														
2	Voice Processing Software														
3	Voice Responses (i.e. pre-determined responses)														
4	GUI														
5	CAD 3D Printing Head														
6	Arduino Board LED Implementation														
7	CAD Body Enclosure Design														
8	Hardware Implementation & Assembly														
9	Video Research														
10	Final Testing														
11	Code Reviews														
12	Project Write Up - Cleanup														
13	Project Write Up - Abstract														
14	Project Write Up - Introduction														
15	Project Write Up - Systems Analysis & Approach														
16	Project Write Up - Conclusion														

IX. Conclusion

We are proposing a device that uses facial and voice recognition, powered by artificial intelligence (AI) and machine learning (ML), as a tool to help children, specifically in the childcare sector and will use the NVIDIA Jetson TX2 as our base platform for development. This project aims to reduce stress and workload issues in the childcare sector. Ease of use and the consumer-friendly requirement of the product will be emphasized in our development and design. Our current solution to the problem is inadequate to address the systemic issue in the United States where resources of human capacity is scarce, but our product will address those voids and improve the childcare sector. It will be the first product of its kind to hit the market that uses AI and ML to fine-turn its response and read emotion tailored towards specific individuals. The final product will be an innovative solution to the current problems.

X. Appendices

X.i. Tech Specs for Jetson TX2 Module

The Tech Specs

JETSON TX2 MODULE

- NVIDIA Pascal™ Architecture GPU
- 2 Denver 64-bit CPUs + Quad-Core A57 Complex
- 8 GB L128 bit DDR4 Memory
- 32 GB eMMC 5.1 Flash Storage
- Connectivity to 802.11ac Wi-Fi and Bluetooth-Enabled Devices
- 10/100/1000BASE-T Ethernet

JETSON CAMERA MODULE

- 5 MP Fixed Focus MIPI CSI Camera

BUTTONS

- Power On/Off
- Reset
- Force Recovery
- User-Defined

I/O

- USB 3.0 Type A
- USB 2.0 Micro AB (supports recovery and host mode)
- HDMI
- M.2 Key E
- PCI-E x4
- Gigabit Ethernet
- Full-Size SD
- SATA Data and Power
- GPIOs, I2C, I2S, SPI, CAN*
- TTL UART with flow control
- Display Expansion Header*
- Camera Expansion Header*
    - *I/O expansion headers: refer to product documentation for header specification.

POWER OPTIONS

- External 19V AC Adapter

KIT CONTENTS

- NVIDIA Jetson TX2 Developer Board
- AC Adaptor
- Power Cord
- USB Micro-B to USB A Cable
- USB Micro-B to Female USB A Cable
- Rubber Feet (4)
- Quick Start Guide
- Safety Booklet
- Antennas to Connect to Wi-Fi-Enabled Devices (2)

XI. References

XI.I. Research

1. Schochet, Leila, and Rasheed Malik. “2 Million Parents Forced to Make Career Sacrifices Due to Problems with Child Care.” *Center for American Progress*, 13 Sept. 2017, www.americanprogress.org/issues/early-childhood/news/2017/09/13/438838/2-million-parents-forced-make-career-sacrifices-due-problems-child-care/.
2. Workman, Simon, and Steven Jessen-Howard. “Understanding the True Cost of Child Care for Infants and Toddlers.” *Center for American Progress*, www.americanprogress.org/issues/early-childhood/reports/2018/11/15/460970/understanding-true-cost-child-care-infants-toddlers/.
3. “Are 12-Hour Shifts Safe?” *American Nurse Today*, www.americannursetoday.com/blog/are-12-hour-shifts-safe/.
4. "Review education policies - Education GPS." OECD. 05 Aug. 2019 <http://gpseducation.oecd.org/revieweducationpolicies/#!node=41720&filter=all>.
5. Redding, Christopher. "Teacher turnover is a problem – here's how to fix it." The Conversation. 03 June 2019. 05 Aug. 2019 <http://theconversation.com/teacher-turnover-is-a-problem-heres-how-to-fix-it-101584>.
6. Isenberg, Emily Pas. "THE EFFECT OF CLASS SIZE ON TEACHER ATTRITION: EVIDENCE FROM CLASS SIZE REDUCTION POLICIES IN NEW YORK STATE." Feb. 2010. 05 Aug. 2019 <https://www2.census.gov/ces/wp/2010/CES-WP-10-05.pdf>.
7. Barnett, W. Steven, Karen Schulman, and Rima Shore. Dec. 2004. 05 Aug. 2019 <http://nieer.org/wp-content/uploads/2016/08/9.pdf>.
8. Milazzo, Vickey. “‘Are You Way Too Stressed Out?’ .” *Legal Nurse*, www.legalnurse.com/wp-content/uploads/2014/06/RN-Stress-Survey-Results-2014-VickieMilazzoInstitute.pdf.
9. “Are 12-Hour Shifts Safe?” *American Nurse Today*, www.americannursetoday.com/blog/are-12-hour-shifts-safe/.
10. Fliesler, Nancy. “'Huggable' Social Robot Eases Kids' Hospital Stay - Boston Children's Discoveries.” *Boston Childrens Discoveries*, 26 June 2019, discoveries.childrenshospital.org/huggable-social-robot-trial/.

XI.II. Images

1. “Public Poll Reveals Concern over NHS Staffing Figures: News: Royal College of Nursing.” *The Royal College of Nursing*, www.rcn.org.uk/news-and-events/news/nurses-rally-as-public-recognise-staffing-crisis.

2. Wagenen, Juliet Van, et al. “Phoenix Children's Hospital Boosts Patient Satisfaction with Centrally-Managed IPads.” *Technology Solutions That Drive Healthcare*, healthtechmagazine.net/article/2018/03/phoenix-childrens-hospital-boosts-patient-satisfaction-centrally-managed-ipads.
3. Steve. “How Is Tech in Schools Helping Your Child to Learn? - Trabasack - Lap Desk and Bag in One.” *Trabasack*, 5 June 2017, trabasack.com/how-is-tech-in-schools-helping-your-child-to-learn/.
4. CBS News. “Why Some High-Tech Parents Want Low-Tech Kids.” *CBS News*, CBS Interactive, 16 Sept. 2014, www.cbsnews.com/news/high-tech-parents-low-tech-kids-steve-jobs/.

XI.III. Software

1. “Face_recognition Python package,” *face_recognition*, https://github.com/ageitgey/face_recognition.
2. “Flask,” *The Pallets Projects,* https://palletsprojects.com/p/flask/.
3. “Getting Started,” *Bootstrap*, https://v4-alpha.getbootstrap.com/getting-started/download/#bootstrap-css-and-js.
4. “SpeechRecognition Python package,” *SpeechRecognition*, https://github.com/Uberi/speech_recognition#readme.

XI.IV. Hardware

XI.IV.i Camera

1. “HD Webcam, EIVOTOR PC Webcam 720P USB Mini Computer Camera Built-in Mic, Flexible Rotatable Clip, for Laptops and Desktop, Black.” *Amazon*, Amazon, www.amazon.com/EIVOTOR-Computer-Microphone-Laptops-Desktop/dp/B07HQDT6S9.
2. “IMX219-120 Camera for NVIDIA Jetson Nano Developer Kit 3280 × 2464 Resolution 8 Megapixels IMX219 Camera Module with 120 Degree FOV.” *Amazon*, Amazon, www.amazon.com/IMX219-120-Camera-Developer-Resolution-Megapixels/dp/B07SJYGTTP?th=1.

XI.IV.ii Microphone

3. “Saramonic SR-ULM5 USB Output Omnidirectional Clip-on Microphone for PC & Mac.” *Adorama*, www.adorama.com/sarsrulm5.html?utm_medium=cpc&utm_source=adl-gbase.

XI.IV.iii Display Module

4. “ESoku 10.1’ Inch Small CCTV Monitor - HD 1024x600 Portable Display LCD Color Monitors Screen with HDMI AV VGA Port Remote Control Built-in Speaker for DVR PC CCTV Security Camera.” *Amazon*, Amazon, www.amazon.com/ESoku-10-1-Small-Monitor-Built/dp/B07GCKJ15N/ref=sr_1_6?keywords=10%2Binch%2Bmonitor&qid=1564095563&s=gateway&sr=8-6.
5. “SunFounder 7 Inch Monitor HDMI - 1024x600 HD IPS LCD Screen Display AV VGA Input Built in Speaker for Raspberry Pi 3 Model B+ 3B 2B B+ A+ Zero CCTV Computer PC DVR Car.” *Amazon*, Amazon, www.amazon.com/SunFounder-Inch-Monitor-HDMI-Raspberry/dp/B073GYBS93/ref=asc_df_B073GYBS93/?tag=&linkCode=df0&hvadid=309751315916&hvpos=1o5&hvnetw=g&hvrand=9386997545077860890&hvpone=&hvptwo=&hvqmt=&hvdev=c&hvdvcmdl=&hvlocint=&hvlocphy=1018127&hvtargid=pla-623140180287&psc=1&ref=&adgrpid=67183599252.

XI.IV.iv Developer Kit

6. “Harness AI at the Edge with the Jetson TX2 Developer Kit.” *NVIDIA Developer*, 4 July 2019, developer.nvidia.com/embedded/jetson-tx2-developer-kit.