Depth Perception Haptic System

Group 40

Christa Lawrence CpE
Cristopher Matos CpE, Psych
Kathryn Pagano EE
Chad Pauley EE

Outline

- Goals and Requirements
- Hardware Discussion and Component Considerations
- Software Discussion
- Successes and Difficulties
- Administrative Details

Goals

- Wearable device that allows for hands-free navigation of indoor and outdoor spaces for people with impared vision
- In particular:
  - Avoid large obstacles (e.g., walls, poles, trash cans)
  - Properly respond to sidewalk curbs

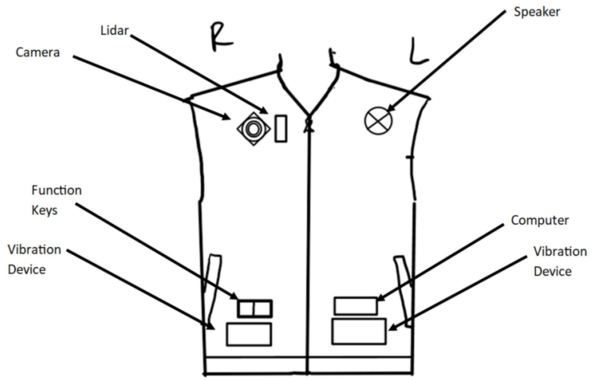

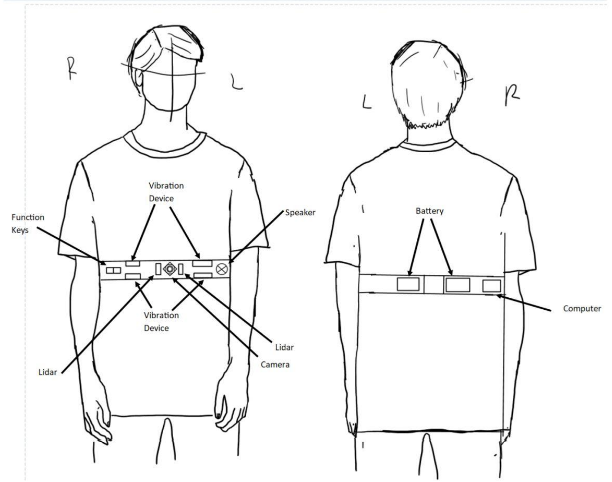

Requirement Specifications

Statement	Description	Value	Unit
Lightweight	Total weight of electronics, housings, and cloth.	< 10	lb
Compact	Thickness along all axes normal to the body.	< 2	in
Full Range of Motion (ROM)	ROM of shoulder joints before and after putting on device.	< 10	% diff
Large Buttons	Surface area large enough to contain braille.	> 0.7	in^2
Long Lasting, Rechargeable Battery	Time from full charge to depletion at 80% feedback level and all sensors active.	> 4	hrs
Haptic feedback that encodes distance	Provides a relative indication of distance (via feedback intensity) of arbitrarily shaped objects visible to the camera over a discrete range of intensities (including zero).	≥ 3	ct
Alert to proximal objects	Objects near the device trigger an unambiguous alert.	< 3	ft
Quick Response Time	Time it takes for the device to send a haptic response from a spontaneously appearing obstacle.	< 900	ms

Goals

Category	Goals
Physical Attributes	Lightweight, compact, durable
Usage & Maintenance	Responsive, rechargeable, long battery life, full range of motion (ROM)
User Interface	Simple, accessible, quick to learn
General Navigation	Encode distance and location of obstacles, alert to emergencies
Outdoor Navigation	Detect and describe curb presence, orientation, and magnitude

Solution Approach - Obstacle Avoidance

- Combine computer vision with LiDAR to map objects in the environment to the user’s body via haptic (vibration) feedback
  - Minimizes # of sensors, maximizes input data

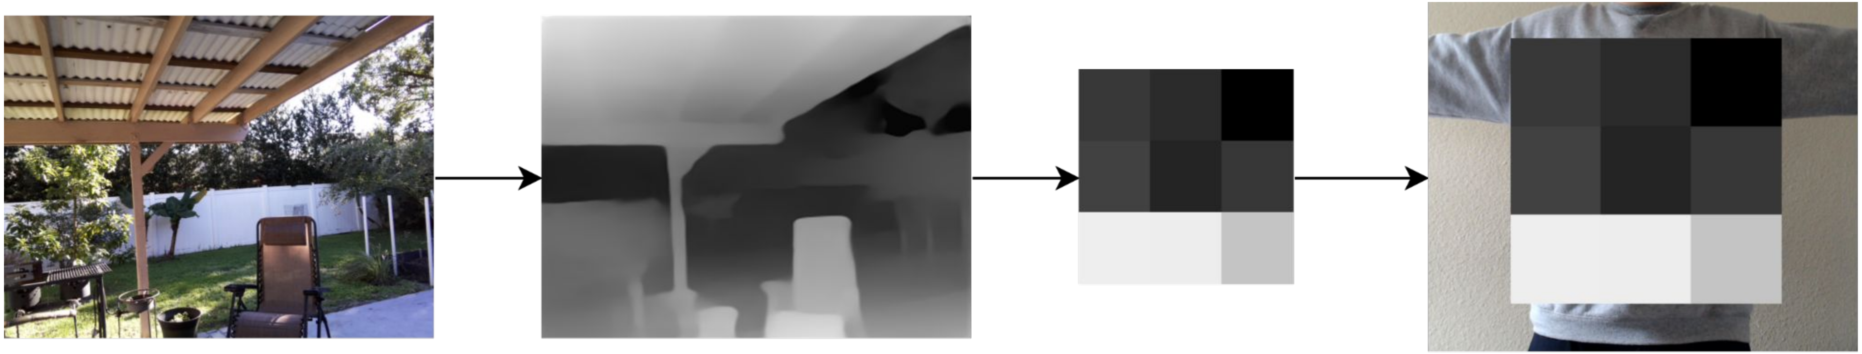

UCF

Overall Block Diagram

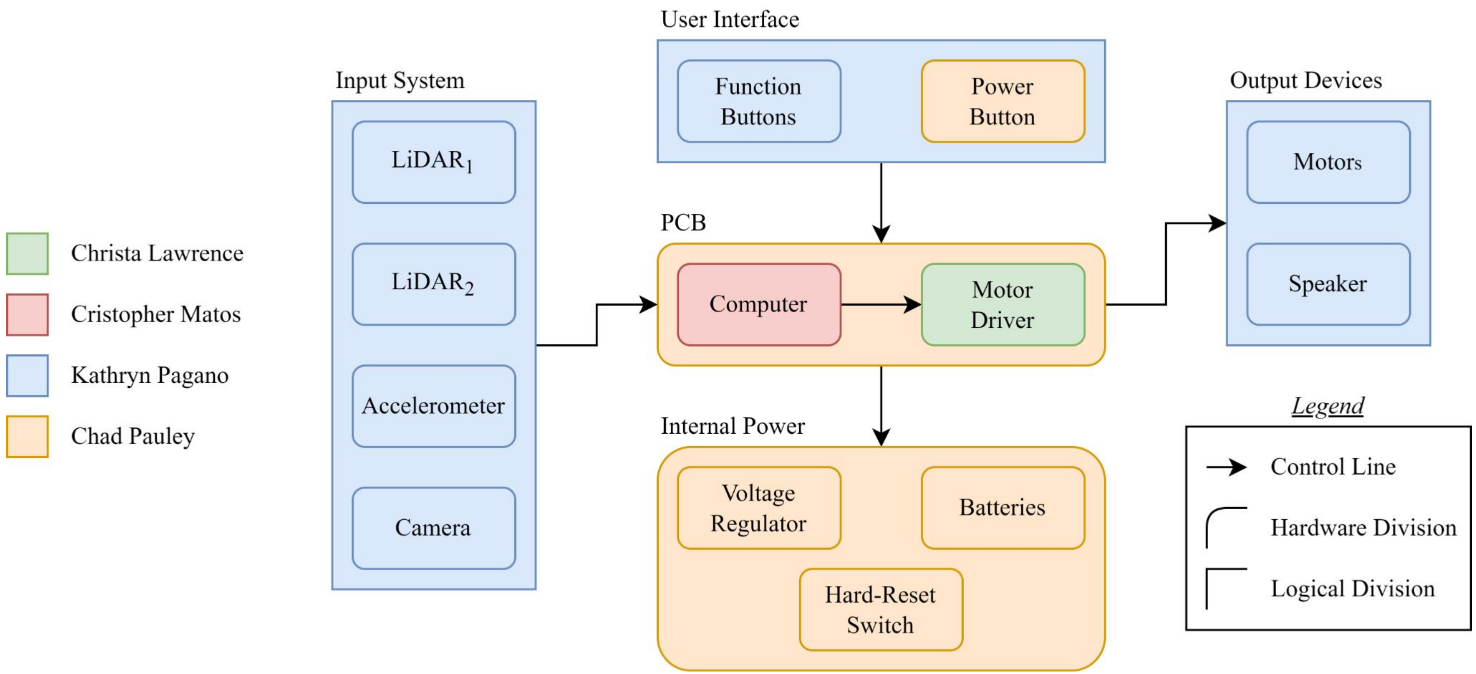

Hardware Design Approach

Always plan to make a prototype, because you'll end up building one anyway

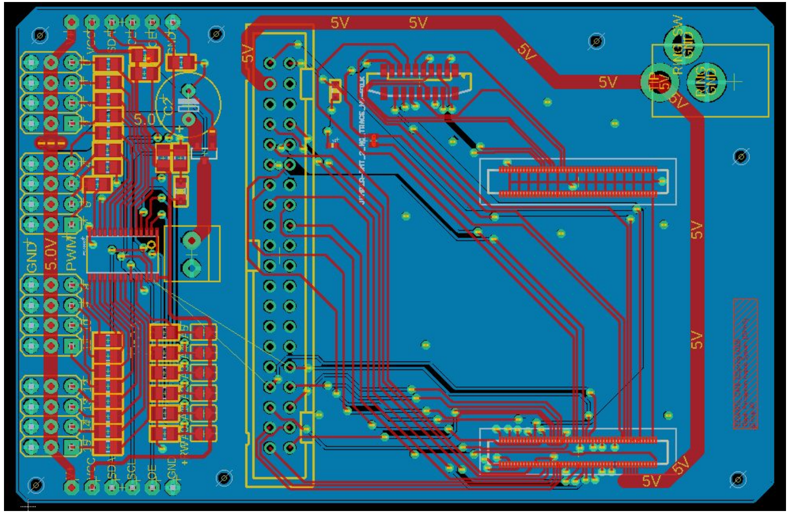

Design Considerations

- Minimize power consumption
- Low resolution output
- Minimize number of sensor units while maximizing input data (CV + LiDAR),
- Ensuring all clock and input data is coherent
- Enough outputs to ensure future development and testing potential

Major Component Considerations

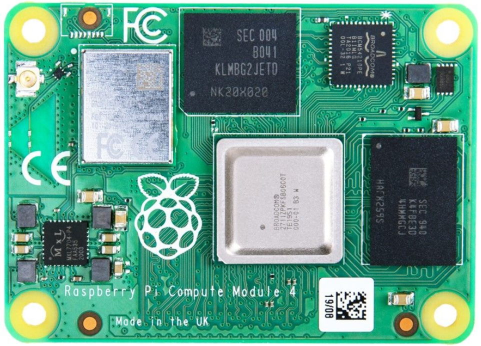

- Raspberry Pi Compute Module 4 (CM4)
  - Price
  - Hardware independence
    - Entire project run on ARM CPU (and integrated GPU)
  - Required RAM and embedded memory
  - Availability of schematics
  - Ability to modify carryout board
  - GPIO availability (up to 5x UART, 5x I2C, 1x SDIO, 1 PCM, 2x PWM, 3x GPCLK, two MIPI CSI channels and multiple voltage output lanes including 3.3V and 1.8V)

Raspberry Pi Camera

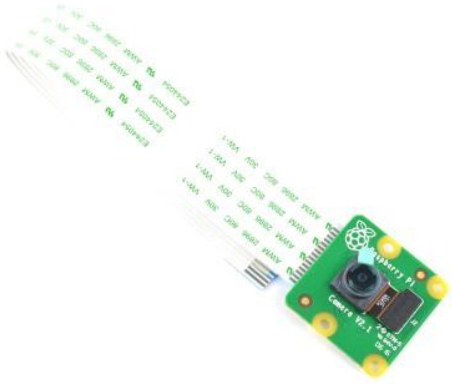

Model:Raspberry Pi Camera Module 2
Resolution
- 8-megapixel
- 3280 x 2464 still frame

Video mode
- 1080p30, 720p60, 480p90

Price
- $25.00

Board size
- 23.86 x 25 x 9mm

Onboard processor
- No

The camera module is incredibly small and can easily be concealed.
This camera is compatible with the Raspberry Pi board and can be connected using either a 15 pin connector or 22 pin adapter ribbon to work with the board.
It uses the Camera Serial Interface Type 2 (CSI-2).
Software Controlled directly by board, no special drivers needed.

UCF

PCB Housing

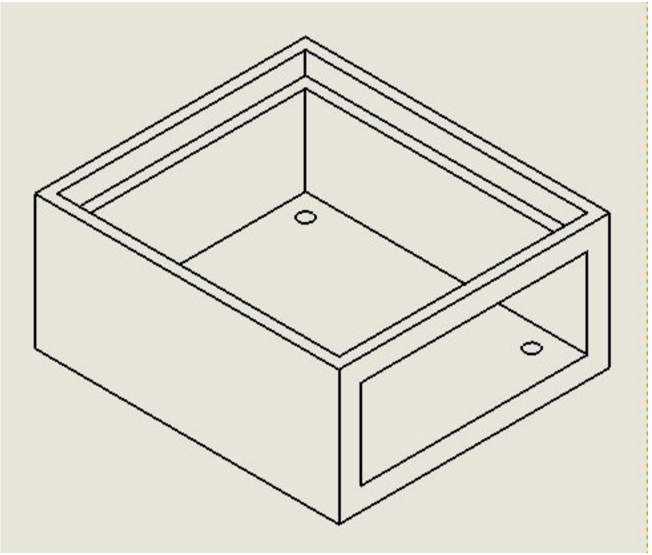

Raspberry Pi Case Angled View

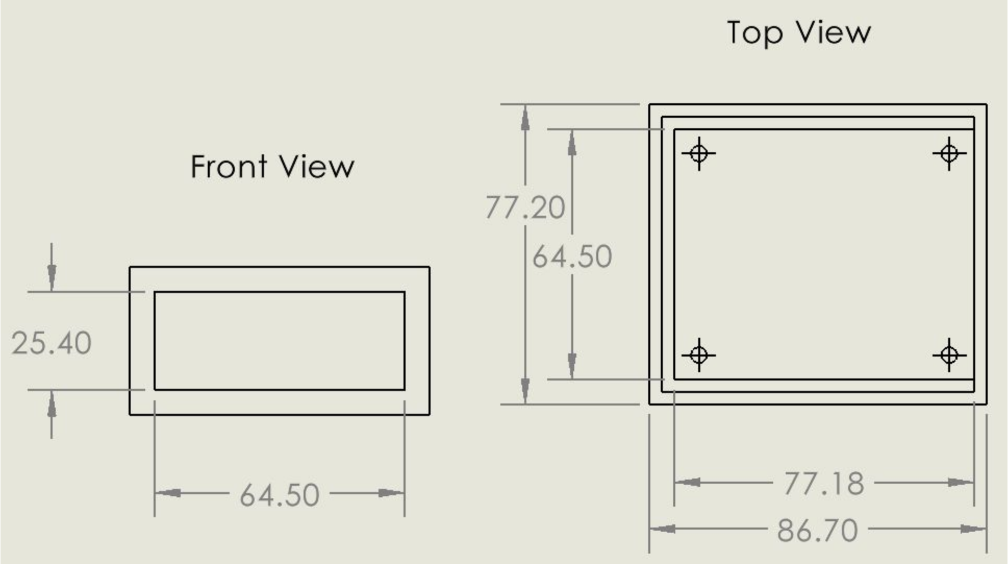

Raspberry Pi Case Measurements

Raspberry Pi Camera Case

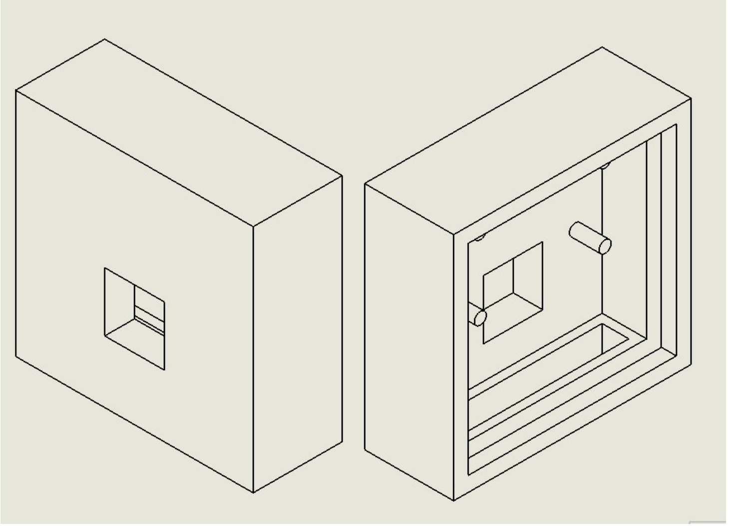

Raspberry Pi Camera Case Angled View

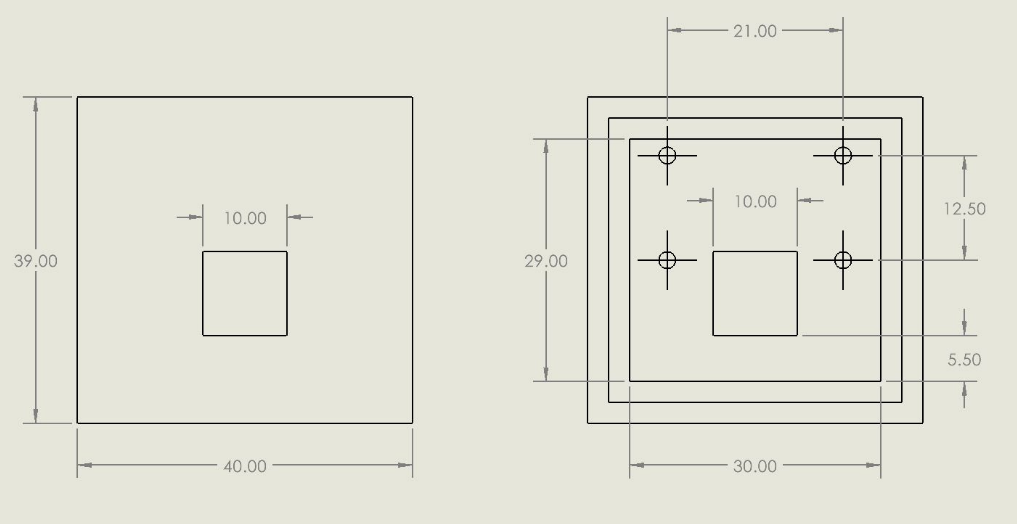

Raspberry Pi Camera Case Measurements

Keeps moisture and dirt out of the camera’s sensitive components and ensures durability.

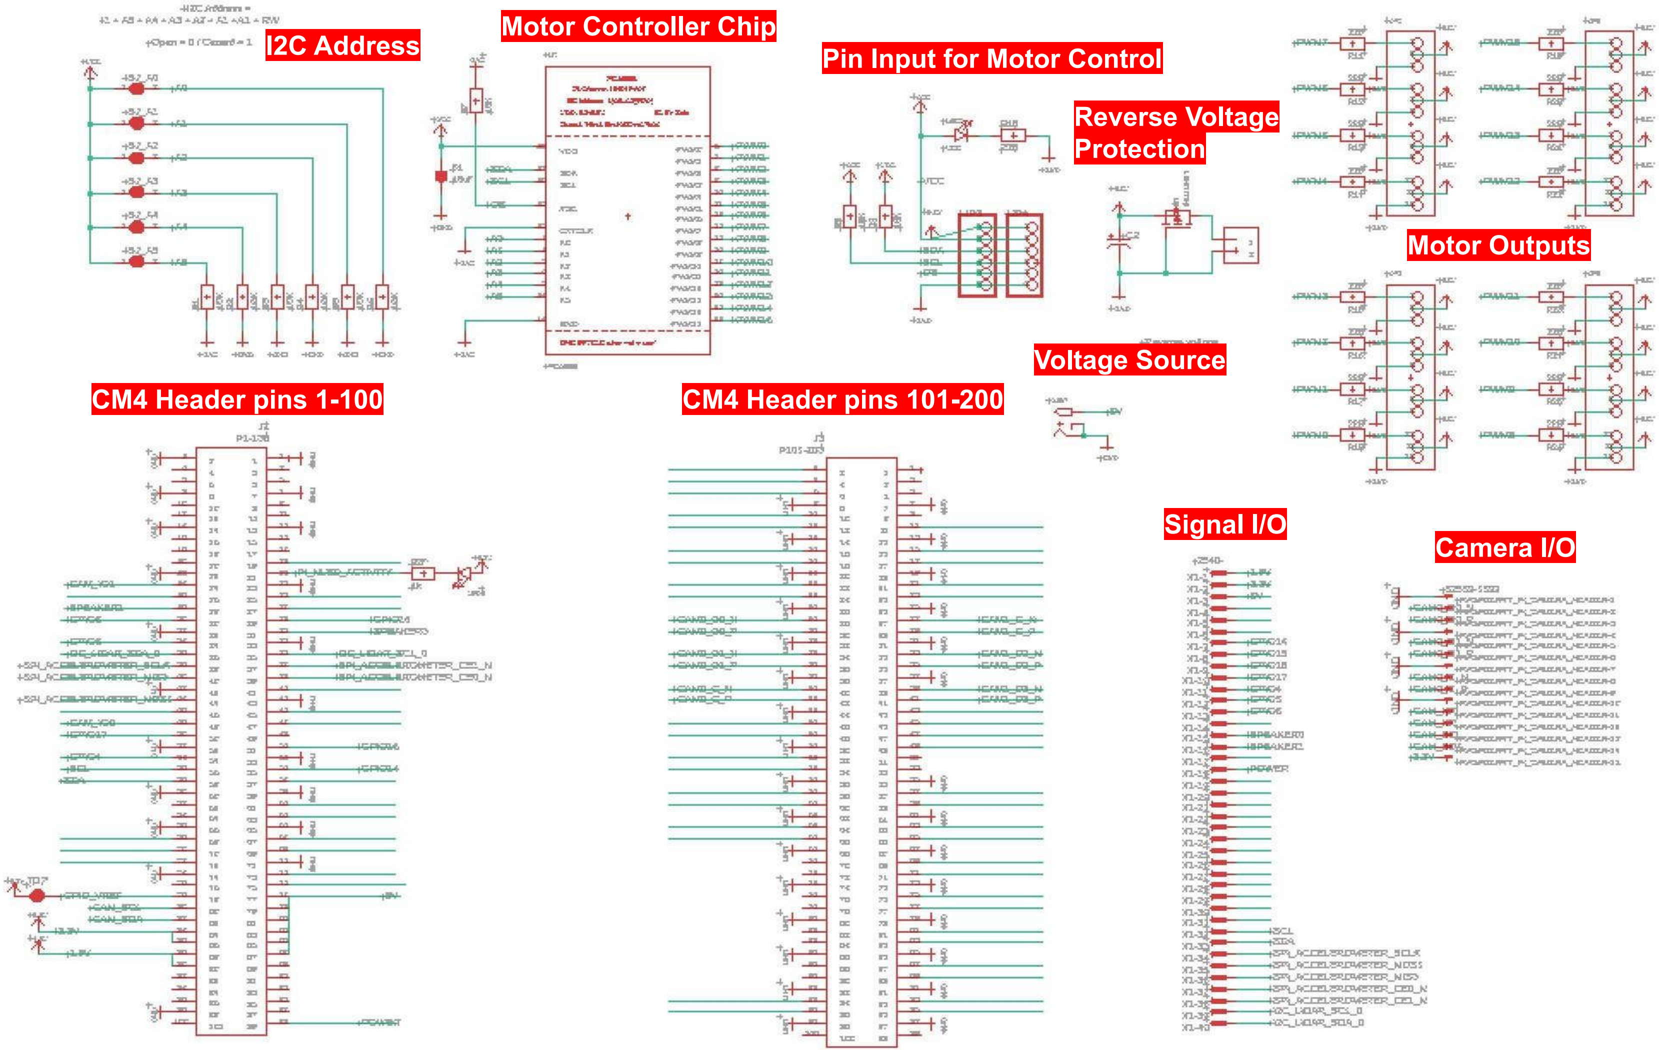
I2C Address
Motor Controller Chip
Pin Input for Motor Control
Reverse Voltage Protection
Motor Outputs
Voltage Source
CM4 Header pins 1-100
CM4 Header pins 101-200
Signal I/O
Camera I/O

Motor Output Schematics

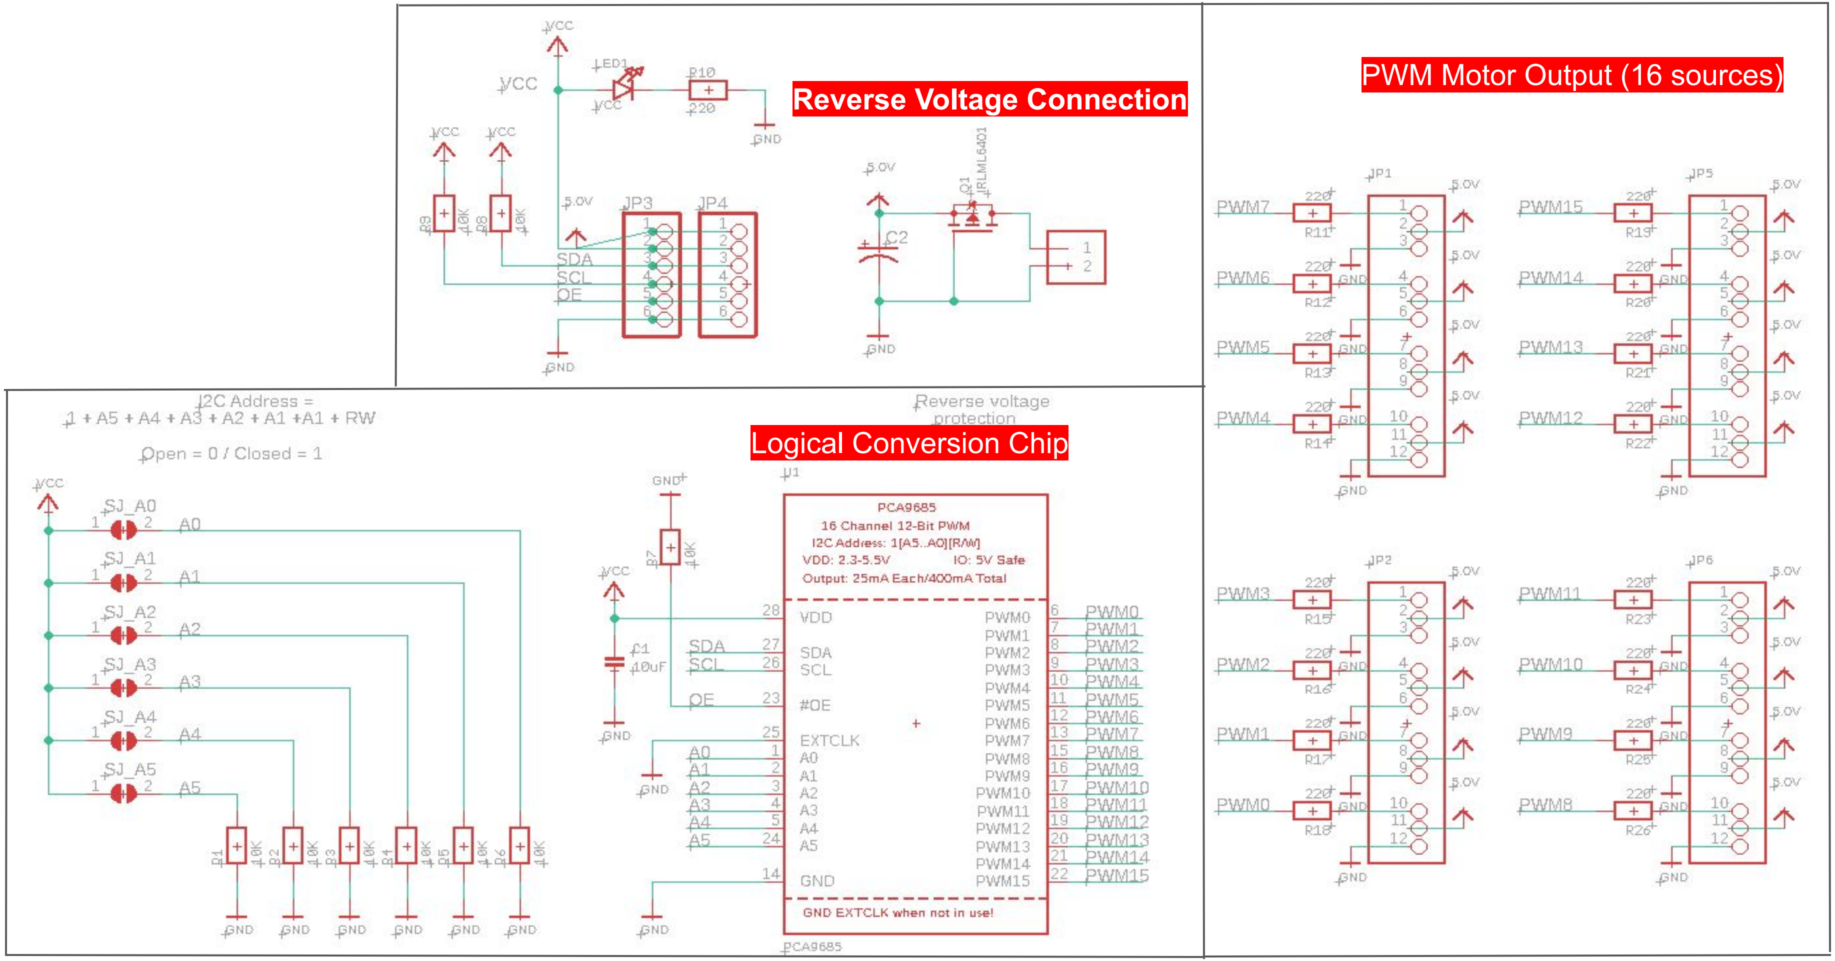

General I/O

CM4 Header pins 1-100

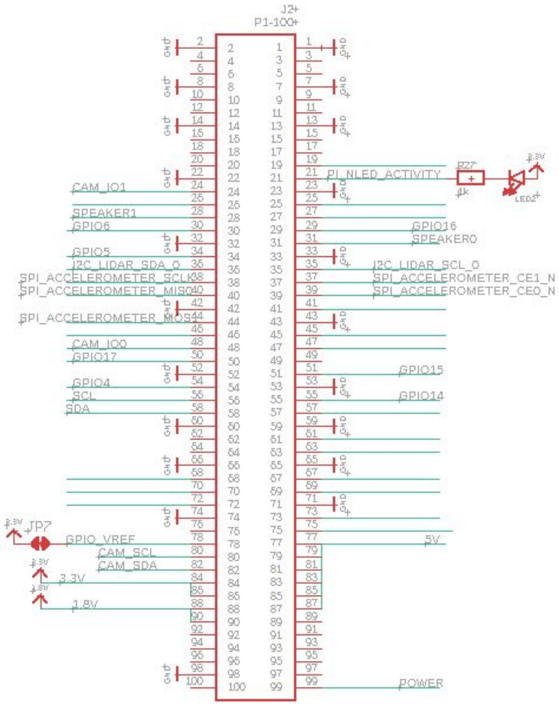

CM4 Header pins 1-100

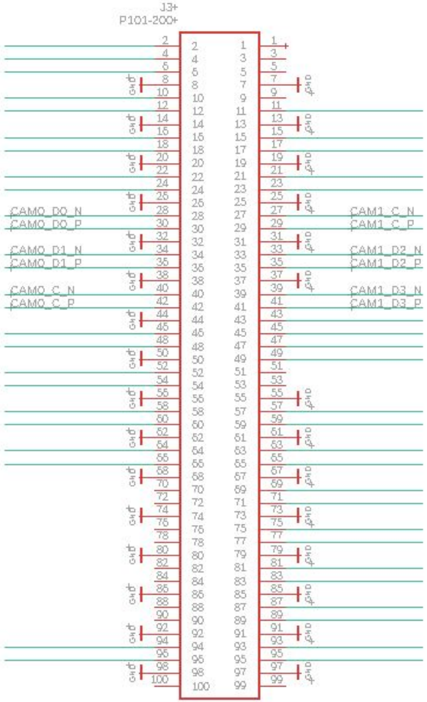

- 2 LiDARs (I2C)
- Accelerometer (SPI)
- Speaker (PWM)
- Soft power off (button tied to GND)
- Raspberry Pi Camera (15-pinout)
- 5 programmable GPIO pins for buttons
- Motor Controller Chip (I2C -> PWM)

Board layout

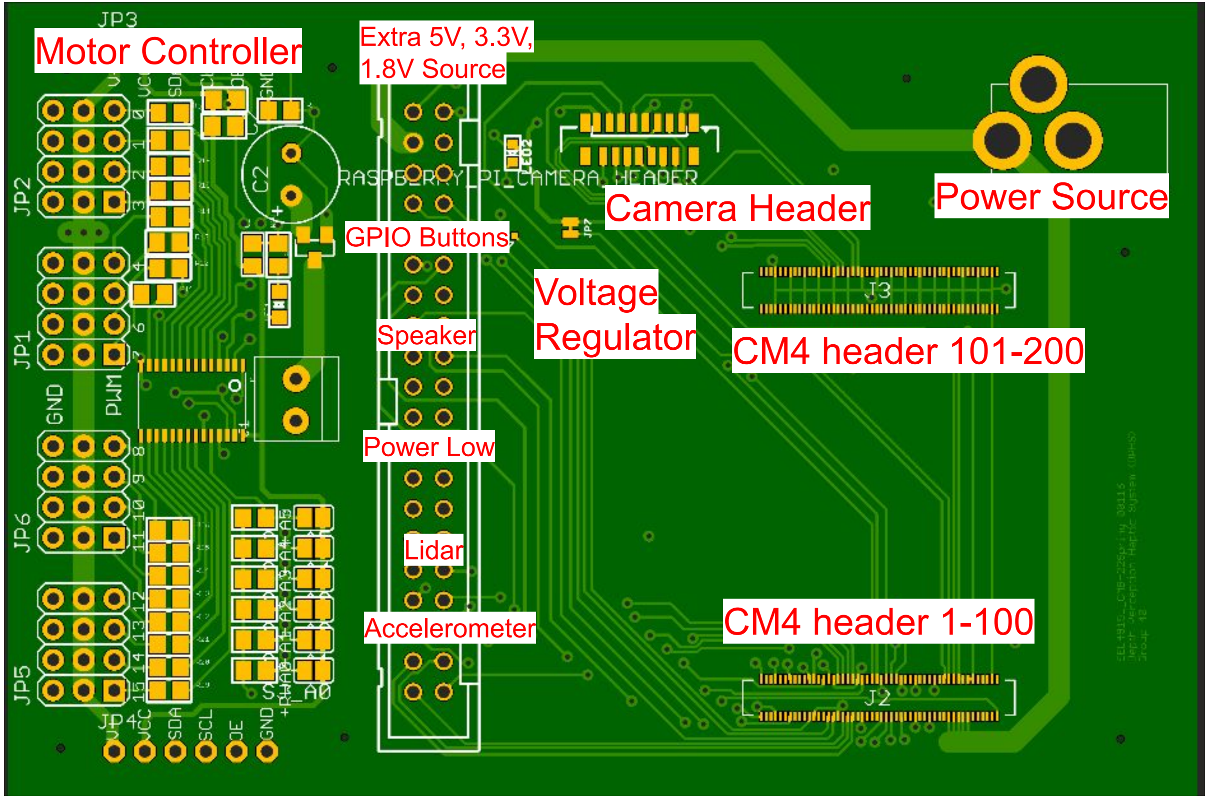

5 Programmable Buttons

1) Resting mode
2) Reset device due to error
3) Pathway Finder
4) Lower level of vibration
5) Increase level of vibration

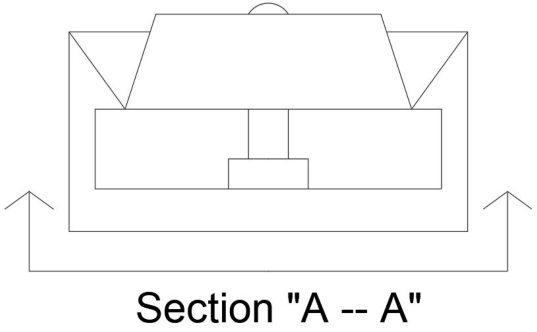

Button Cross-section

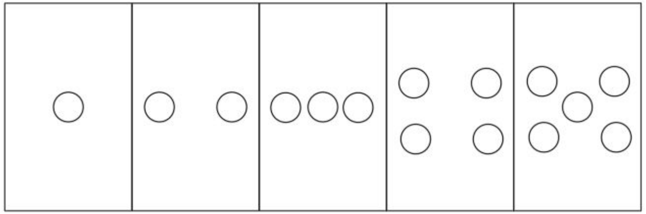

Button Domino Pattern

Three Levels of Power Control

- Idle Power Off (Software Controlled)
  - Controlled by accelerometer input value and determination by CM4; the device will go into a resting mode reducing power usage to a minimal amount (around 8mA).
- Button Power Off (Software Controlled)
  - Button connected to GPIO GLOBAL_EN set driven low sets device to minimum power usage (around 15μA).
- Switch Power Off (Hardware Controlled)
  - Physical switch between the batteries and voltage regulator to allow the batteries and device to maintain longevity as well as function as hard reset for any potential errors.

Voltage Regulator

Batteries will be wired in parallel to provide the longest battery longevity possible and allowing the device to maintain a steady current in the device for any needs.

Mounted separately to batteries to have 5V source instead of 9V+. This prevents the user from getting shocked or burned by the device when the user is sweating and the overall concern of the wear and tear of the device.

A TAMAYA plug or Barrel Power Jack will be used to charge the device (as shown below). User preference will be noted in a future study.

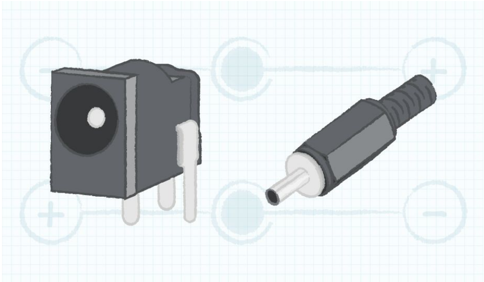

PCB of voltage regulator will be mounted to the batteries

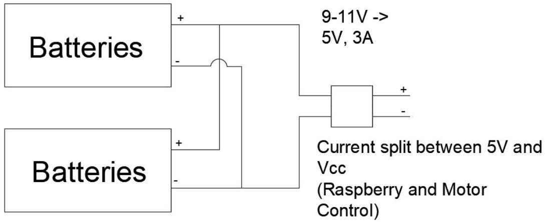

Power Considerations

- Battery must be able to supply 3A with a voltage source between 5V and 30V.
- Due to design constraints in creating a power bus; the power for the CM4 and PWM motor control are differentiated as a primary 5V pin and VCC pin respectively on the board to evenly split current and effectively half the needed width for the bus channels on the PCB.

Product Class	Product Class Description	Rated Battery Energy (Ebatt)	Special characteristics or battery voltage	Maximum UEC (kWh/yr)
1	Low-Energy	≤ 5 Wh	Inductive Connection	3.04
2	Low-Energy, Low-Voltage	<100 Wh	<4V	$.1440*E_{bat}+2.95$
3	Low-Energy, Medium-Voltage		4-10 V	For E_{batt} <10 Wh, 1.42 kWh/y $E_{batt} \geq 10$ Wh, $0.0255 * E_{batt} + 1.16$
4	Low-Energy, High-Voltage		>10V	$0.11 * E_{batt} + 3.18$

Electronic Code of Federal Regulations

- The device must be less than or equal to the maximum unit of energy consumption (UEC) for the appropriate product class and battery rating

UCF

Hardware Design Changes Made Over Time

- Removal of HDMI interface as later noted as unnecessary and to reduce size of PCB
- Removal of SD memory storage due to the need for more GPIO pins
- Removal of voltage regulator chip on primary PCB board due to concern of electric shock
- Change from using PWM with MOSFET connection which would control voltage of haptic motors with duty cycle to a direct connection with a simple diode and capacitor controlled by the programmable values set by analog

Software Block Diagram

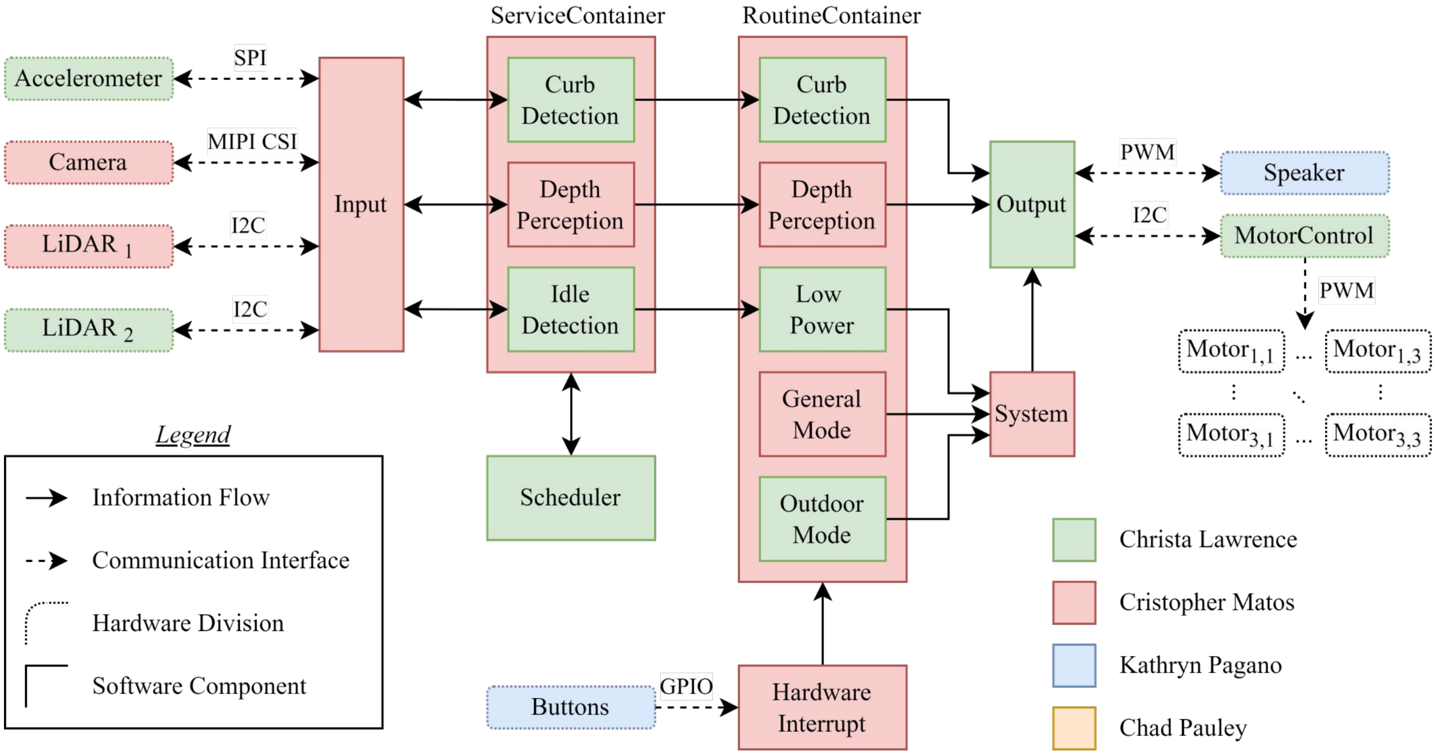

UCF

Depth Sensing Approach

- Problem: Need distances of objects located in front of the user
  - Existing products use:
    - Multiple distance sensors placed along body
    - One distance sensor, actively manipulated by user
  - Conflicts with our goals:
    - Few sensors (For full ROM and comfort)
    - Hands-free operation
- Solution: Combine computer vision (CV) with LiDAR to estimate depth

Depth Sensing Approach Motivation

- Explosion of research and development in monocular depth estimation for mobile devices (Ignatov et al., 2021)
  - Some models (including MIT Licensed) already exist online
- Prospect of combining depth inferencing with LiDAR

Inference using Intel’s Midas/v2_1_small on TensorFlow Lite

Depth Sensing Procedure

1. Use CV to create a depth map of the space in front of the user
   a. LiDAR used to approximate true values
2. Divide map into regions
3. Find one representative value per region
4. Set motor intensity based on value

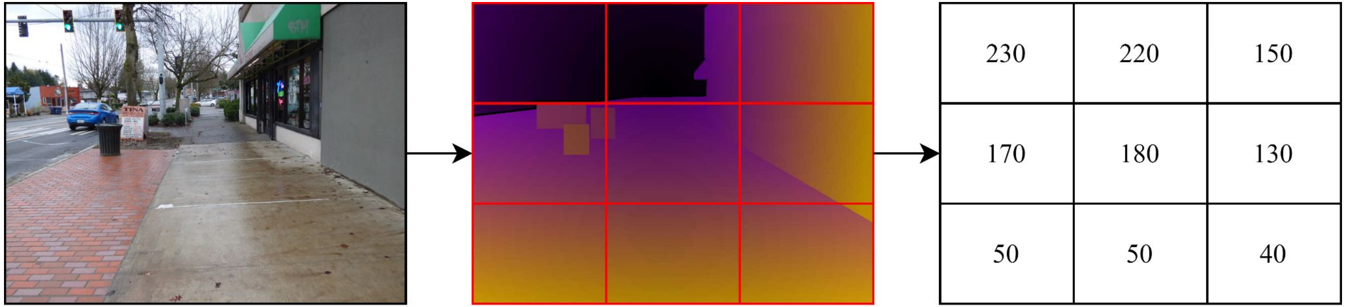

Raw Image | Image after Depth Inference | Reduced Resolution Depth Map

Raw Image sourced from (SDOT, 2015)

Depth Sensing Problems

- Inference time (exposed chip, no fan, 21℃, all on CPU)
  - Arithmetic Mean: 912 ms, Sample Std Dev: 11 ms, Maximum: 967 ms
- Temperature
  - Throttle temperature reached within 700 iterations (~ 12 mins)

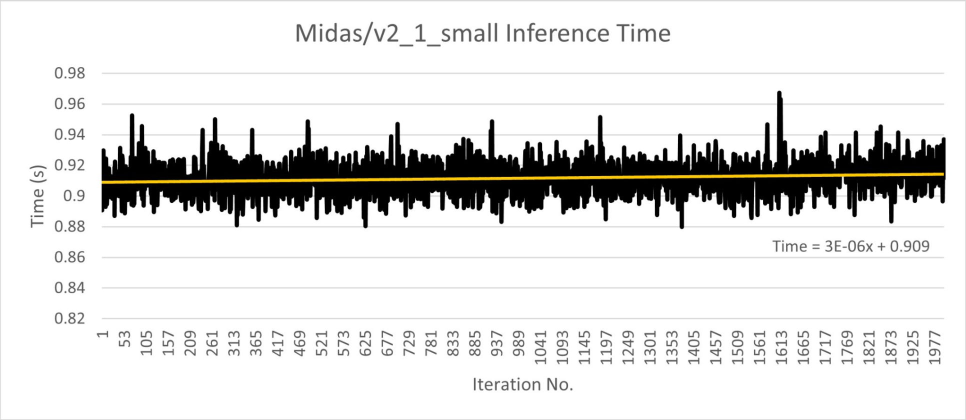

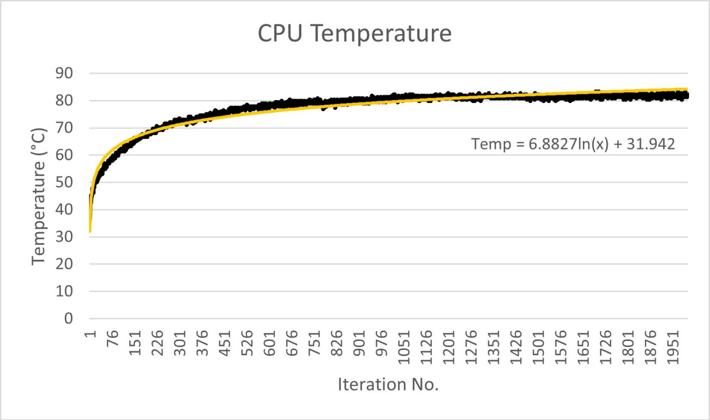

Depth Sensing Solutions Considered

- NVIDIA Jetson Nano?
  - 2x Price
  - Research shows ARM CPU is all that is needed for fast depth estimation
- Overclocking CPU?
  - Energy consumption and device lifetime
  - Intended to operate over several hours
- Start from scratch?
  - Lack time and resources to train robust model
  - Models tend to underperform on datasets not used in training/validation (Ranftl et al., 2020)
- Another pretrained model?
  - MIT Licensed models tested significantly slower (seconds)

Depth Sensing Solutions

- Quantization and Open-Source Optimization tools
- Knowledge distillation
  - Use representative dataset to train new model with Midas as teacher
  - Some approaches show inference time as low as 97 ms (Zhang et al., 2021)
- On-board GPU acceleration
  - Some libraries available (beatmup, ncnn)
- Introduce a heat sink
- Worst-case scenario: Forfeit robustness for speed
  - Train new model on a subset of datasets

UCF

Other Software

- Curb Detection
  - Use a LiDAR unit angled downward to determine magnitude and direction of curbs
- Idle Detection
  - Accelerometer used to estimate velocity of user
- Speaker announces mode changes, power-on, power-off, and emergencies
- Buttons control mode of system

Successes and Difficulties

Part Availability (Voltage Regulators)

Shipping Dates

PCB arrangement on 2-layer design

CV Runtime Reduction

Budget

Item Type	Quantity	Actual Cost	Description
LiDAR	2	$59.99	Garmin LiDAR
Accelerometer	1	$18.95	ADXL345
Camera	1	$25	Raspberry Pi Camera
Speaker	1	$10.69	MakerHawk Speaker
Haptic feedback	Set acquirement (20)	$16.04	Tatoko flat coin motors
Compute Module	1	$65	Raspberry Pi Compute 4 Module 4GB
Shipping		$37.93	Shipping on parts
Current Total (As of 11/19/2021)		=293.60	

UCF

Progress by Category

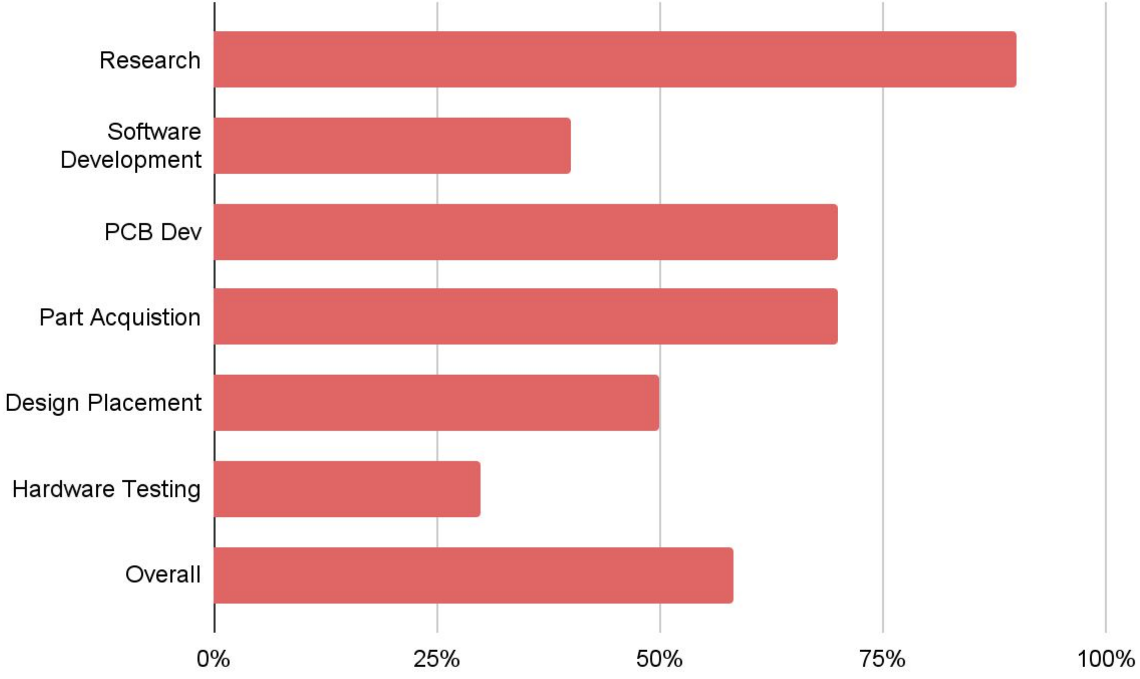

Steps Toward Completion - Software

Step	Deadline
Faster (< 900 ms) depth estimation on RPi4 Faster (< 500 ms) depth estimation on RPi4	February 27 March 13
UI, general mode, outdoor mode code complete	March 6
Idle detection, lower power mode, speaker code complete	March 20
Initial software testing and refactoring complete	April 3
Final software testing and optimization complete	April 10

Steps Toward Completion - Hardware

Step	Deadline
User interface constructed, tested All devices integrated in development environment	March 6
PCB final version submitted	March 20
All devices integrated in fully functional prototype	April 3
Assembly of final project complete	April 10

Milestones

Milestone	Due Date
Basic end-to-end functionality (CV + LiDAR integration, motor output)	February 25
Full initial prototype (buttons, all functionalities, motor output, improved CV performance, battery connection)	March 13
Final prototype complete (fully functional vest)	April 3
Final testing and optimization complete	April 10

Progress Photos

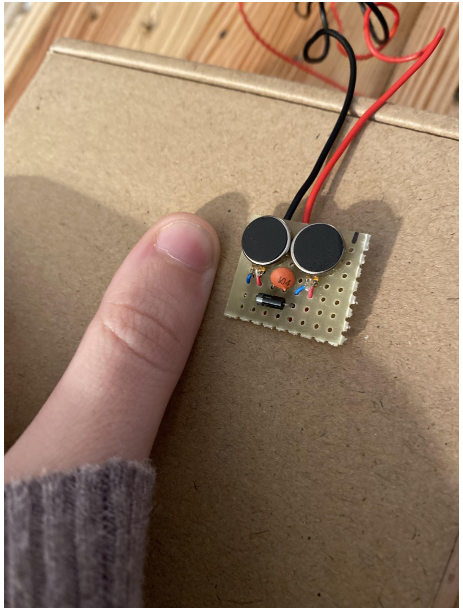

References

Allan, Alasdair. “Benchmarking TensorFlow Lite on the New Raspberry Pi 4, Model B.” Hackster.io, 2019, https://www.hackster.io/news/benchmarking-tensorflow-lite-on-the-new-raspberry-pi-4- model-b-3fd859d05b98.

Ignatov, Andrey, et al. “Fast and Accurate Single-Image Depth Estimation on Mobile Devices, Mobile AI 2021 Challenge: Report.” ArXiv, 17 May 2021, https://arxiv.org/pdf/2105.08630.pdf.

Ranftl, Rene, et al. “Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-shot Cross-dataset Transfer.” IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020, https://arxiv.org/abs/1907.01341.

SDOT Photos. “Columbia City Sidewalk Repair Project.” flickr, 16 Dec. 2015, https://www.flickr.com/photos/sdot_photos/23168490614. Image. Used in its resized form, and colored over with gradients to depict depth. Licensed under CC BY-NC 2.0.

Zhang, Ziyu, et al. “A Simple Baseline for Fast and Accurate Depth Estimation on Mobile Devices.” 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2021, https://openaccess.thecvf.com/content/CVPR2021W/MAI/papers/Zhang_A_Simple_Baseline_for_Fast_and_Accurate_Depth_Estimation_on_CVPRW_2021_paper.pdf.

Questions?