

# Omobot: a low-cost mobile robot for autonomous search and fall detection

Shihab Uddin Ahamad\*, Masoud Ataei\*, Vijay Devabhaktuni<sup>†</sup>, Vikas Dhiman\*

\*Electrical and Computer Engineering Department,  
University of Maine, Orono, ME, USA  
shihab.ahamad@maine.edu

<sup>†</sup>Electrical Engineering Department,  
Illinois State University, Normal, IL, USA

**Abstract**—Detecting falls among the elderly and alerting their community responders can save countless lives. We design and develop a low-cost mobile robot that periodically searches the house for the person being monitored and sends an email to a set of designated responders if a fall is detected. In this project, we make three novel design decisions and contributions. First, our custom-designed low-cost robot has advanced features like omnidirectional wheels, the ability to run deep learning models, and autonomous wireless charging. Second, we improve the accuracy of fall detection for the YOLOv8-Pose-nano object detection network by 6% and YOLOv8-Pose-large by 12%. We do so by transforming the images captured from the robot viewpoint (camera height 0.15m from the ground) to a typical human viewpoint (1.5m above the ground) using a principally computed Homography matrix. This improves network accuracy because the training dataset MS-COCO on which YOLOv8-Pose is trained is captured from a human-height viewpoint. Lastly, we improve the robot controller by learning a model that predicts the robot velocity from the input signal to the motor controller.

## I. INTRODUCTION

Fall prevention and management among the elderly has been a well-recognized problem since the 1990s [1], [2]. As the proportion of 65 and above population in the OECD countries has grown from 11.36% in 1990 to 17.96% in 2022 [3], the importance of fall prevention and management has grown accordingly. While preventing falls is the first line of defense, the second line of defense managing falls by reducing the response and rescue time [4]. Fall detection systems are essential for the elderly, as they play a critical role in detecting falls promptly and ensuring timely medical intervention. By providing continuous monitoring and immediate assistance, the safety and independence of elderly individuals can be enhanced, while their need for caregivers and healthcare resources can be reduced.

Fall detection technologies that alert caregivers [2] include user-activated alarm systems [5], passive fall detection systems [4]–[7] and mobile robots as fall detectors [8]–[10]. These technologies have complementary strengths and weaknesses, and a hybrid system can be customized based on the user’s needs and preferences. We compare these technologies in the related work section (Section II).

In this work, we focus using mobile robots as fall detectors. While many mobile robots have been proposed for fall detection [8]–[10], all of them formulate fall detection as a classification problem. This requires classifying images as

either ”fall” or ”no fall.” We instead formulate fall detection as an object detection problem. We note that there have been great improvements in the accuracy of object detection algorithms like YOLOv8 [11]. By formulating fall detection as a problem of object detection instead of one of image classification, we can identify multiple persons in the same image and classify each as fallen or not. This is especially useful when a dummy or statue is present in the same image as the elderly person that the robot is supposed to watch. Moreover, low-cost design has not been a priority for the above projects. To address this, we design and prototype a custom low-cost mobile robot designed expressly for the purpose of periodically surveying an area and checking for persons who have fallen.

The main contributions of this project are threefold. (1) We design and develop a low-cost, open-source mobile robot for indoor autonomous navigation and fall detection. Our robot is similar in cost to the popular robotics platform Turtlebot but differs in that it is equipped with Nvidia Jetson Nano, omnidirectional wheels, and autonomous wireless charging. (2) We improve the fall detection algorithm from the robot’s viewpoint. Our robot camera is installed at a height of 0.15m from the ground to maintain a small form factor, unlike fall detection datasets that are typically collected from a human viewpoint or higher. Because of this, machine-learning based fall detection and object detection algorithms have lower accuracy when image is taken from a robot’s perspective. We compensate for this by computing a homography transformation that adjusts the images from the robot’s perspective to match the typical human-height perspective in our dataset. Our experiments demonstrate improved fall detection performance by a factor of 6-12%. (3) Lastly, we improve the robot controller by system identification of the motors. We modified the controller by training a model that predicts the motor’s rotation velocity using the input Pulse Width Modulation (PWM) signal.

## II. RELATED WORK

### A. Fall detection using sensors

In this section, we compare the following fall detection technologies: user-activated alarm systems [5], passive fall detection systems [4]–[7] and mobile robots as fall detectors [8]–[10].

User-activated alarm systems [5] require the person to press a button to request help from community responders. One

limitation of this approach is that it requires the fallen person to be conscious and, furthermore, to decide to call for help. The hesitancy of the elderly to ask for help under such circumstances has pushed researchers to develop passive fall detection systems that do not require any action from the patients themselves [5]

Literature on passive fall detection systems has been reviewed multiple times in the last two decades [4]–[7]. These fall detectors can be classified based on the location of sensors as *wearable* or *ambient*. A *wearable* fall detector is worn by the person to be tracked. It typically uses IMU (Inertial Measurement Units) and health sensors to detect falls. An *ambient* fall detection technology is usually installed in the person’s home. These systems typically use pressure sensors, vibration sensors, or cameras to detect a fall and alert the caregivers. While wearable sensors have become less conspicuous, they can easily lead to false alarms. On the other hand, ambient sensors, while more accurate, are more costly. Admittedly, these sensors can be combined to complement each other. One limitation of all camera-based (vision-based) systems is that some consider them to invade the privacy of the individual being monitored.

These limitations, combined with advances in deep learning and autonomous robots, have led to a new type of ambient fall detectors: mobile robots as fall detectors [8]–[10], [12], [13]. These are ambient-mobile sensors, contrasting with the previous generation of ambient-static sensors. While most mobile robots do depend on camera-based fall detection, they may be preferred by some users over ambient camera-based systems. Because mobile robots perform only periodic monitoring (as opposed to ambient systems’ continuous monitoring), mobile robots are not all-seeing at all times. A mobile robot system can be configured to check on a person at regular intervals, such as every 30 minutes, and be triggered by wearable sensors or loud noises. The mobile robot can then search for the person in the house. If the person is found and identified as being in a fallen position, then an alert is generated. A trigger is generated if a room or a bedroom is found inaccessible. The advantage of this approach over ambient cameras is that a mobile robot checking on an elderly person is periodic, and the person being checked on is reminded of the robot’s presence. The requirement to keep the floor clear for the robot’s passage can also help prevent falls.

### B. Low cost mobile robots

Extensive research has been carried out in the field of intelligent autonomous mobile robots to develop advanced features such as obstacle avoidance, object detection, path planning, and map creation. In one such effort, Andruino-R2, a mobile robot, was developed and implemented by [14] with line-following navigation combining Arduino and Java-based ROS. [15] developed an integrated system named BeeButler that combines a flutter-based multi-platform mobile app and a robot to aid hotel guests in ordering amenities. It solves the previous problem of external devices by using onboard Jetson Nano and Arduino Mega 2560 to control the robot. Still, it

uses a line follower and RFID tag for localization, which is unsuitable for navigating an unknown environment without setting additional lines and RFID tags. A two-wheel home-assistive robot featuring a combination of omnidirectional wheels with differential driving was developed by author [16]. The author used an STM32 microcontroller and NVIDIA Jetson Nano to control the robot and demonstrated successful 2D and 3D SLAM experiments, autonomous navigation, and obstacle avoidance. The author [17] developed a ROS-based omnidirectional robot with mapping, localization, and navigation capabilities using a multi-sensor fusion on Raspberry Pi 4B, emphasizing accuracy and efficiency.

### III. ROBOT DESIGN

We briefly summarize the robot’s overall design, covering the mechanical, electronic, and software elements. For a detailed description of the hardware and software design, as well as videos, our dataset, and other supplementary material, please consult our GitHub repository<sup>1</sup>. We aim to design a robust, reliable, and low-cost robot. We were able to design the robot with parts costing less than \$700 USD. The cost breakdown of all the parts is displayed in Table I, which illustrates how the full robot is built for a budget of under 700 USD, ensuring affordability without compromising quality. The items were purchased and priced based on market price between November 2022 and December 2023. The overall hierarchy of system design is presented in Figure 1.

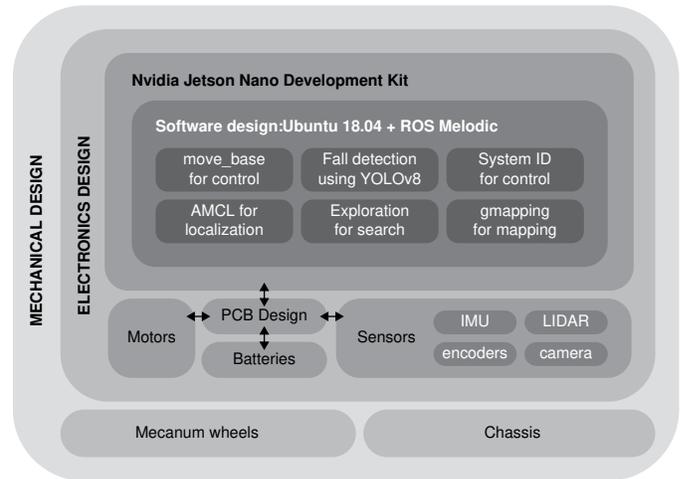


Figure 1. System overview.

#### A. Mechanical design

The mechanical design of the robot includes the base chassis, battery holder, mecanum wheels, wireless charger, and wireless power receiver. The base chassis provides structural support and holds together all the essential components of the robot. 3D-printed Polyethylene Terephthalate Glycol (PETG) filament was used to create the base chassis, mountings to hold sensors, input-output (I/O) devices, and joints. The

<sup>1</sup>[https://github.com/shihab28/omobot\\_js](https://github.com/shihab28/omobot_js)

components were assembled using M3 and M4 bolts of lengths 6mm, 10mm, 15mm, and 25mm. The robot’s 3D model was designed using Solidworks. The final assembly of the robot is shown in Figure 2. A 3D-printed battery holder was mounted at the bottom section of the base chassis to keep the center of mass low. The robot has a wireless recharging system with a 3D-printed charging dock and wireless transmission coil. We chose wireless charging because it avoids the problem of contact mismatch during docking. A 12.6V-1A wireless receiver is mounted under the battery holder for power transmission. The robot moves using *mecanum wheels* mounted using aluminum mounts. The mecanum wheels allow the robot to move sideways.

### B. Electronics design

The electronics design includes the choice of microcomputer, printed circuit board (PCB) design, battery management system (BMS), motors, sensors, and communication modules with complex interconnections as shown in Figure 1.

1) *Microcomputer*: We chose NVIDIA’s Jetson Nano 4GB Developer Edition A1 Kit as the robot’s brain because it balances computing efficiency, power consumption, and cost. It is responsible for receiving and processing sensor data, controlling motors, and sending emails. It is Robot Operating System (ROS)-compatible and has general-purpose input-output (GPIO) pins for sensors, making it ideal for mobile robots [18].

2) *PCB design*: The PCB is mounted on top of the base chassis, next to the Jetson Nano. The PCB includes the charging controller, two H-bridge motor drivers LM298N, and two Arduino Nano RP2040 microcontrollers are assembled into the PCB Figure 2.

We use two Arduino *microcontrollers*, one for high-frequency sensor input and another for high-frequency motor control. The first *Arduino-based sensor encoder* reads battery voltage, current flow, and rotary encoder pulses per rotation and sends the data to the Jetson Nano. We use eight pins on the Arduino for reading data from four rotary encoders and 20 interrupt pins to read and process the high-frequency pulses from the rotary encoders without delay. The second *Arduino-based motor controller* controls the motor by receiving the desired motor speed from the Jetson Nano and sending an 8-bit pulse width modulation (PWM) signal to the motors. The *Arduino-based sensor encoder* is connected using the UART-TX/RX pins, and the *Arduino-based motor controller* is connected using the USB port.

3) *Battery management system*: Our robot’s battery management system (BMS) consists of batteries and power splitters. We use a 3-cell, 12V-5200mAh Lithium Polymer (Li-Po) battery. We chose Li-Po batteries because they are small, lightweight, highly efficient, and have a high-energy density [19]. For power splitting, we use a DC-DC switching 12V-3A buck-boost converter to supply a constant 12V input to the H-bridge motor driver and a DC-DC 5V-4A buck converter to power the Jetson Nano and the Arduino.

Table I  
COST BREAKDOWN

Components	Cost (USD)
Jetson Nano (Dev-4GB) with accessories, and Camera	300
DC gear motor (12V) with encoder and mecanum wheels	90
Arduino (Nano RP2040)	42
LiDAR (LD19)	99
IMU (MPU-6050)	3
Lipo Battery (5200mAh)	29
Power supply and charging module (12V)	39
PCB Board (Dual Layer-Through Hole)	15
Wireless charging system (12V-2A DC)	36
3D printed parts	30
Misc (wire, bolts, nuts, connectors, etc.)	10
Total	693

4) *Motors*: We use four 12V DC geared motors with rotary encoders. These motors are easy to control, provide high torque, and are suitable for carrying payloads. The wheel rotary encoder measures the actual wheel rotation velocity, making it perfect for integrating feedback-control [20].

5) *Sensors and communication modules*: We equip the robot with several sensors and communication modules besides the wheel rotary encoders. The sensors include LiDAR-LD19, MPU-6050 Inertial Measurement Unit (IMU), and Sony IMX219-83 Stereo cameras. LiDAR-LD19 is connected to Jetson Nano via USB. The MPU-6050 IMU sensor is connected to the Jetson Nano using the SDA and SCL pins of the GPIO pins for I2C communication. The stereo cameras have a baseline of 80mm and are connected to Jetson Nano via two CSI interfaces.

For communication, we use the Waveshare AC8265 Wireless NIC Module, which supports 2.4GHz/5GHz Dual Band WiFi and Bluetooth-4.2, allowing the robot to connect to the internet.

### C. Software design

Software design is crucial for making our robot autonomous. We install all our software on NVIDIA’s Jetson Nano. We use the latest Linux kernel supported by Jetson Nano, NVIDIA L4T 32.5.2, a part of NVIDIA’s JetPack-4.5.1. This also includes ARM64-based Ubuntu-18.04 Linux distribution. We install Robotic Operating System (ROS) Melodic [21], the latest version of ROS supported by Ubuntu-18.04. Jetson Nano supports several communication protocol methods to communicate with peripheral devices and sensors, including universal asynchronous receiver transmitter (UART), universal serial bus (USB), I2C, Internet Protocol (IP), and Camera Serial Interface (CSI) [18]. We develop a complex software pipeline comprising several interconnected ROS software packages. The following section discusses the details of the software used and the custom software developed.

## IV. ROS-BASED SOFTWARE FRAMEWORK

The overall ROS-based software framework interacting with sensors, motors, joystick, and internet is shown in Figure 3. We discuss the novel software components, the Fall detection

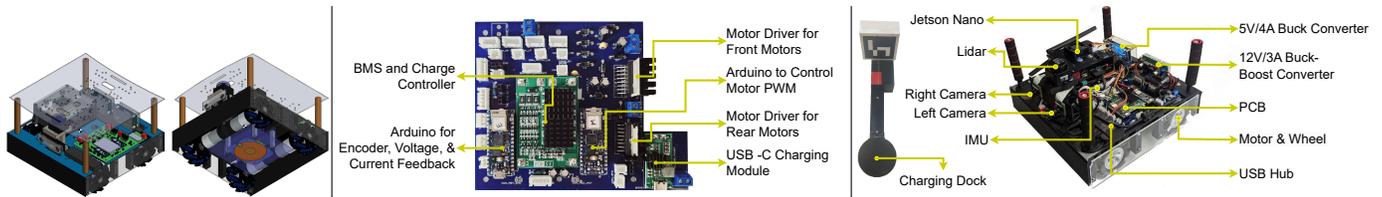


Figure 2. From left to right: 1) 3D model of the robot developed in Solidworks, 2) PCB Assembly with the major components, 3) the final assembly of the robot.

system in Section IV-C and the Motor system identification in Section V. Before that, we briefly discuss other software components.

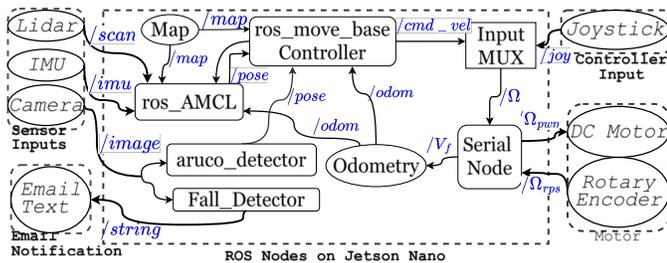


Figure 3. Communication among sensors and devices with ROS nodes running on Jetson Nano.

### A. Robot setup phase

Our robot system for autonomous search and fall detection requires a *setup phase* before its first use. The robot setup phase is the mapping phase, where the robot is introduced to a new environment. In this phase, an operator moves the robot around the house using a Joystick, and the `ros-slam_toolbox` builds a 2-dimensional occupancy grid-based map of the house. This map is then saved and manually marked with crucial landmark points on the map. For example, Figure 6 shows a map with seven locations marked from A to G. While this step can also be automated using robot exploration algorithms [22], the setup has to be done only once per environment, hence we leave automation of setup-phase for future work.

### B. Autonomous phase

After the setup phase, our robot is fully autonomous and requires no human intervention during regular operation. The autonomous charging module and `ros-slam_toolbox` module support the main autonomy.

The autonomous charging module uses an ARUCO marker (a QR-code-like marker) on the docking station. The left mono camera detects and locates the wireless charging dock using ARUCO-marker detection and a real-time video feed. Before use, the camera is calibrated using Python3 and OpenCV camera calibration to correct lens distortions and adjust intrinsic camera properties. The camera driver publishes the camera feed, and the `ros-aruco_detect` package in ROS detects ARUCO-markers and 3D poses from the camera feed,

providing positional and orientation data for localizing the wireless charging dock with respect to the robot. In practice, we found that the ARUCO localization accuracy is not enough for robust docking.

The `ros-slam_toolbox` module localizes the robot in the room using the map built during the setup phase. It depends only on the LiDAR-LD19 laser scans and wheel odometry data. A 2D map is created with the help of `ros-slam_toolbox` package to locate the robot and objects.

### C. Fall detection pipeline

The fall detection system contains three steps: preprocessing the frame, person detection and pose estimation, and fall detection module. The preprocessing phase includes resizing, rescaling, and possibly applying a homography transformation. Subsequently, the transformed image is fed to the person detection and pose estimation model (YOLOv8-Pose). The person detection step identifies a person in an image and provides a confidence score by estimating a bounding box around them. Simultaneously, the pose estimation locates the key points on the human body (nose, eyes, ears, shoulders, elbows, wrists, hips, knees, and ankles) on the preprocessed frame. We evaluate two methods of detecting falls from key points, a Rules-based method [23] and an MLP-based method.

1) *Pre-processing* : Our robot uses YOLOv8-Pose [11] for pose estimation. This model is trained on the MS-COCO dataset [24]. Like most photographs found online, images in the MS-COCO dataset are taken from the human point of view (POV). Given the extremely low POV of our robot (0.15m from ground), this reduces the accuracy of our robot’s ability to extract pose. Our experiment in Section VI-E2 confirms this intuition. One way to address this problem is by data augmentation during the model’s training [25]. In this work, we instead transform images from the robot point of view (POV) to an average human POV before pose estimation. This has the advantage of avoiding the additional step of retraining the model. We show an example of this transformation in Figure 7.

*Using homography for POV conversion:* We use a homography matrix to transform robot POV images into human POV. Homography refers to the transformation of images from one perspective projection into another. A homography transform preserves straight lines as straight lines but parallel lines may not be preserved as parallel.

We know that images of a 3D plane in two cameras are related by a homography transformation [26]. To perform this

transformation, consider a 3D location  $\mathbf{x}_1 \in \mathbb{R}^3$  in space. Project it to two cameras; name the pixel coordinates of that location  $\mathbf{u}_1 = [u_1, v_1, 1]^\top \in \mathbb{P}^2$  in the first image and  $\mathbf{u}_2 = [u_2, v_2, 1]^\top \in \mathbb{P}^2$  in the other. If the 3D points  $\mathbf{x}_1$  all lie on the same plane  $\mathcal{P}$ , a fixed Homography matrix,  ${}^2H_1 \in \mathbb{R}^{3 \times 3}$  maps all corresponding coordinates from the first image onto the second image [26] by,  $\alpha \mathbf{u}_2 = {}^2H_1 \mathbf{u}_1$ .

**Theorem 1.** *Let the plane  $\mathcal{P}$ , on which the points  $\mathbf{x}_1$  lie, be described by a unit normal vector  $\hat{\mathbf{n}} \in \mathbb{R}^3$  and the distance  $h \in \mathbb{R}$  from the origin so that the plane is defined as  $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^3 \mid \hat{\mathbf{n}}^\top \mathbf{x} = h\}$ . Let any point  $\mathbf{x}_1 \in \mathcal{P}$  on the plane projected to two cameras  $\lambda_1 \mathbf{u}_1 = K_1 \mathbf{x}_1$  and  $\lambda_2 \mathbf{u}_2 = K_2 ({}^2R_1 \mathbf{x}_1 + {}^2\mathbf{t}_1)$ , where  $K_1 \in \mathbb{R}^{3 \times 3}$  and  $K_2 \in \mathbb{R}^{3 \times 3}$ , are the intrinsic matrices of the cameras and  ${}^2R_1 \in SO(3)$  and  ${}^2\mathbf{t}_1 \in \mathbb{R}^3$  are relative rotation and translation from camera 1 to camera 2. Then, we can compute the homography matrix  ${}^2H_1(\hat{\mathbf{n}}, h)$  that maps a point  $\mathbf{u}_1$  in image 1 to image 2  $\alpha \mathbf{u}_2 = {}^2H_1 \mathbf{u}_1$  as,*

$${}^2H_1(\hat{\mathbf{n}}, h) = hK_2 {}^2R_1 K_1^{-1} + K_2 {}^2\mathbf{t}_1 \hat{\mathbf{n}}^\top K_1^{-1}. \quad (1)$$

The proof of above theorem is provided in the supplementary material due to space constraints. Since the homography depends on the location of the plane, therefore, we sample multiple homographies corresponding to different distances  $h$  between minimum and maximum of the LiDAR scan. Then we pick the homography that gives us the highest confidence detection by YOLOv8.

2) *Person detection and human pose estimation:* We use YOLOv8-Pose [27] for person detection and pose estimation. YOLOv8-Pose identifies a person in an image and provides a confidence score by estimating a bounding box around them. The bounding box includes the top corner of the bounding box and its width ( $w$ ) and height ( $h$ ). Simultaneously, the YOLOv8-Pose also estimates the human pose. The human pose is represented as the 2D location of the 17 keypoints  $(x_p, y_p)$  that correspond to different human body parts for example,  $p \in \{\text{shoulder, foot, hip, } \dots\}$ . YOLOv8-Pose models come in five sizes, from smallest to largest (n)ano, (s)mall, (m)edium, (l)arge, e(x)tra-large. Bigger models have more layers, thus more weights and thus more representation power, but at the cost of compute and memory resources.

3) *Fall detection from human pose:* We experiment with two approaches for fall detection modules, (a) Rules-based fall detection and (b) MLP-based fall detection.

a) *Rules-based fall detection:* Using the keypoints detected by YOLOv8-Pose, we can determine a fall  $F$  using (2) [23].

$$F = (y_{\text{shoulder}} > (y_{\text{foot}} - l)) \wedge (y_{\text{hip}} > (y_{\text{foot}} - l/2)) \wedge (y_{\text{shoulder}} > (y_{\text{hip}} - l/2)) \vee (h < w), \quad (2)$$

where  $l = \|(x, y)_{\text{shoulder}} - (x, y)_{\text{hip}}\|_2$ ,  $h$  and  $w$  are the height and width of the person bounding box.

b) *MLP-based fall detection:* We train a binary classifier (fallen or not-fallen) on each person detected by the YOLOv8-Pose. The 2D location  $(x_p, y_p)$  of all the human pose keypoints

(17 points) predicted from the YOLOv8-Pose [27] are concatenated into a 34 size vector and fed into a three-layer Multi-layer Perceptron (MLP), with 2 hidden layers, 20 hidden units each, ReLU activation function and the softmax output. The MLP is trained with the Binary Cross Entropy loss function on the FallDetectionDatabase [28] with YOLO parameters frozen.

## V. MOTOR SYSTEM IDENTIFICATION

We estimate the mapping from the motor input to the robot velocity in order to find the required Pulse Width Modulation (PWM) signal for driving the robot with a desired velocity. The process consists of two parts: (1) mapping the wheel velocities to the robot velocity and (2) mapping the PWM signal to the wheel velocities.

### A. Mapping wheel velocities to robot velocity

Let the desired robot velocity in 2D be  $\mathbf{v} = [v_x, v_y, \omega_z]^\top \in \mathbb{R}^3$ , where  $v_x$  and  $v_y$  are linear velocities along the X and Y axis, and  $\omega_z$  is the angular velocity along the Z-axis. Also, let the angular velocity of each of the four wheels denoted by  $\boldsymbol{\omega} = [\omega_{FL}, \omega_{FR}, \omega_{RL}, \omega_{RR}]$  corresponding to the front-left, front-right, rear-left, and rear-right wheels. For a desired control signal  $\mathbf{v}$ , the corresponding angular velocities ( $\boldsymbol{\omega}$ ) of the four wheel can be written as [29],

$$r\omega_{FL} = v_x + v_y - R\omega_z, \quad r\omega_{FR} = v_x - v_y + R\omega_z, \quad (3)$$

$$r\omega_{RL} = v_x - v_y - R\omega_z, \quad r\omega_{RR} = v_x + v_y + R\omega_z, \quad (4)$$

where  $r$  is the radius of the wheel, and  $R$  is the radius of the robot with respect to its center.

### B. Motor controller model

We use an 8-bit PWM signal to control the angular speed of the motor through the H-bridge motor driver. For a wheel  $i \in \{FL, FR, RL, RR\}$  and its angular velocity  $\omega_i$ , we assume the following inverse relationship with PWM signal  $u_{\text{pwm},i}$ ,

$$u_{\text{pwm},i} = \frac{b_i}{\omega_i} + c_i, \quad \text{for each } i \in \{FL, FR, RL, RR\}, \quad (5)$$

where  $b_i \in \mathbb{R}$  and  $c_i$  are unknown parameters. We collect the motor velocities corresponding to the PWM signal as a dataset. Then, we fit a least square model to estimate  $b_i$  and  $c_i$  for positive and negative motor direction separately. The estimated parameters  $b_i$  and  $c_i$  are shown in Table II. To drive

Table II  
HYPER PARAMETERS (B, C) FOR THE MOTOR SYSTEM MODEL

Motor Direction	Param	Wheel name $i$			
		FL	FR	RL	RR
Positive	$b_i$	130.63	132.61	137.10	129.74
Negative	$b_i$	-129.32	-133.68	-130.08	-132.54
Positive	$c_i$	-4096.39	-4446.28	-4431.14	-3954.66
Negative	$c_i$	-4089.71	-4585.02	-4080.27	-4167.72

the robot at a desired velocity  $\mathbf{v}$ , we first estimate the wheel velocities using (4) and then PWM signal using (5).

## VI. EXPERIMENTS AND RESULTS

### A. Effectiveness of motor system identification

To check the effectiveness of the motor system identification, described in Section V, we tasked the robot controller to follow a desired circular trajectory of radius 0.65 m with a constant angular velocity. We evaluated two controllers in their ability to follow the desired trajectory accurately, (a) controller with system identification and (b) controller without system identification. In the case of the *controller with system identification*, the robot controller used the estimated parameters  $b_i$  and  $c_i$  from Table II, and the robot's position per second was recorded. In the case of *controller without system identification*, the robot controller used a linear model from  $\omega_i$  to  $u_{pwm,i}$  with the slope estimated from the range of wheel velocities and corresponding range of the PWM signal. The trajectories followed by the robot are shown in Figure 4. The trajectory's root mean squared (RMS) deviation without system identification is 0.297 m, which is 5.6 times the deviation using our controller model with system identification (0.053 m). From the expected trajectory's radius of 0.65 m, the deviation is 8.15% for our controller model and 45.69% for the controller without using system identification. This shows that system identification allowed us to effectively follow desired trajectories with small errors that can be compensated by a high-level feedback loop.

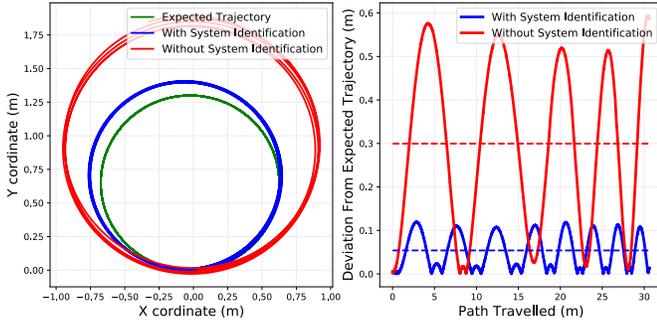


Figure 4. **Left:** Trajectories of the robot. **Right:** Deviation from the expected trajectory with and without using the parameters from motor system identification.

### B. Payload capacity and battery capacity

Two experiments were conducted to determine the robot's payload capacity and runtime under a single battery charge. To measure the payload capacity, weight was gradually added to test the robot's ability to move. If the robot was able to gain a non-zero velocity, then the payload was increased with intervals of 200g. We found that the maximum payload capacity is at least 5.4 kg, excluding the robot's mass.

To determine the robot's battery capacity, we programmed the robot to move to random goal points in the pre-built map of the room, and the battery voltage was recorded as a rosbag (ROS logging file format). We found the runtime of the robot to be 1 hour and 33 minutes under a single discharge of the 5200mAh battery (90% battery charge to 30% charge remaining).

### C. Odometry accuracy

We performed the mapping experiments in a  $12.45 \times 9.95$  sq m room at the University of Maine (Barrows Hall, Rm 201). The room was first mapped using the `ros-slam_toolbox` library. The `ros-slam_toolbox` uses LiDAR observations and the robot's position from wheel odometry to create a 2D map and update the map continuously. Loop closure is used to correct errors and ensure an accurate map of the environment. It corrects the odometry to align the observed LiDAR scans with the expected LiDAR scans from the map.

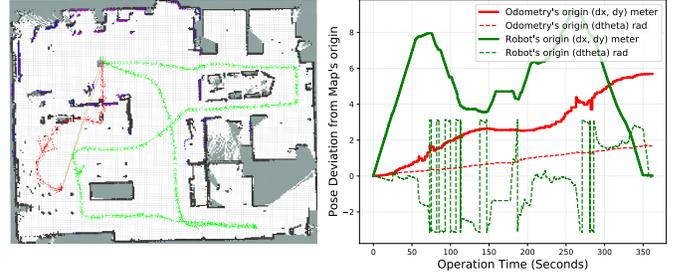


Figure 5. Mapping using LiDAR, Wheel Odometry, and `ros-slam_toolbox`. **Left:** robot's and odometry's trajectory. The green arrows show the trajectory of the robot's origin while creating the map, and the red arrows show the trajectory of the odometry's origin. **Right:** the deviation of the robot's and odometry's pose (translation in XY axis, orientation w.r.t Z axis) from the map's origin.

To compare wheel odometry with `ros-slam_toolbox` provided trajectory, we navigated the robot in a full loop around the room so that it comes back to the starting point. The map and the trajectories are shown in Figure 5. Black colors in the grid map represent space occupied by obstacles, white represents free space, gray represents space yet to be updated, and the purple dots represents the LiDAR scan. Note that the LiDAR scans align with the walls and obstacles very closely, which demonstrates the accuracy of `ros-slam_toolbox` trajectory. The wheel odometry trajectory is shown in red and corrected `ros-slam_toolbox` trajectory is shown in green. The odometry's estimated pose at the end of the trajectory deviated from the starting pose by 5.64 m in translation and 1.66 rad in orientation. From this experiment, we conclude that having an inexpensive LiDAR does not affect the performance. However, the odometry data generated using the inexpensive wheel encoder had notable errors, but `ros-slam_toolbox` is able to mitigate the errors without compromising the mapping accuracy in this room. It is well understood that LiDAR based mapping does not work well in locations with long featureless corridors. Under such situations, the wheel odometry will become more important, highlighting the trade-offs of using low-cost wheel encoders.

### D. Obstacle avoidance efficiency

For this experiment, a map of the room was built and manually marked with landmark locations as explained in the Setup phase (Section IV-A). We tasked the robot to navigate to random landmark locations in the map, and to detect and report any falls observed during the exploration. The

flowchart of map exploration, fall detection and reporting is shown in Figure 6 (Left). The landmark locations (A-G) and the trajectories of the robot during 30 minutes of operation are shown in green, shown in Figure 6 (Right). The robot encountered only *one collision* during 30 minutes of operation. The collisions may happen when the 2D map built by the LiDAR scan does not match the 3D reality of obstacles. For example, LiDAR might capture the stem of an office chair as the obstacle but the robot encounters the feet of the chairs that are wider than the stem.

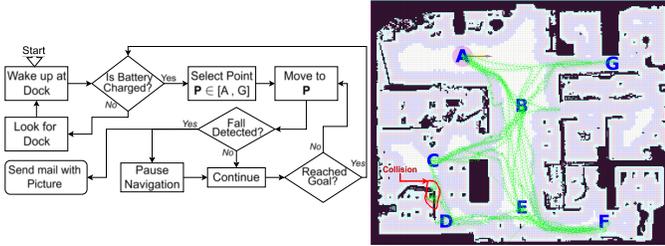


Figure 6. **Left:** Robots workflow. **Right:** Occupancy grid map of a room with landmarks A-G and robot trajectories. Black pixels indicate static obstacles, white pixels indicate free areas; purple pixels indicate inflated obstacles, A-G indicate manually marked landmark locations; and green lines indicate the robot trajectories during the 30-minute robustness of autonomous navigation experiment (see Section VI-D).

### E. Fall detection

We collected a human fall dataset consisting of 128 images captured from our robot’s POV, manually categorized into 62 images depicting human falls and 66 images depicting not-falls under similar lighting conditions. This dataset was used only for evaluating the methods, not training, and is available on the project GitHub page.

1) *Evaluating the fall detection pipeline:* The fall detection pipeline contains of 3 steps, (a) pre-processing (b) person detection and human pose estimation and (c) fall detection from human pose. We evaluate two options for the pre-processing step: (i) *robot POV*, where no homography is applied to the input image (ii) *human POV* where homographies are computed based on the minimum and maximum distance in the room, applied to the input image and highest confidence results are picked. We evaluate five options for the person detection step: n, s, m, l, x for each size of the YOLOv8-Pose model from the smallest to the largest. Lastly, we evaluate two options for the fall detection from human pose: (i) Rules-based fall detection and (ii) MLP-based fall detection. All model variations combined leads to 20 models. The accuracy of these models on the dataset is plotted in Figure 8 (Right). On average, across the five models, the homography improved the accuracy by 6-12 percent. We note that human POV (our contribution) improves accuracy over robot POV consistently. Also Rules-based model has a better accuracy than MLP-based model.

2) *Robustness to homography:* Our results from previous section show that changing robot POV to human POV using homography transform improves accuracy. This suggests that

YOLOv8 trained on MS-COCO is not robust to homographic transformations of the original image. We validate this hy-



Figure 7. Applying Homography on a sample image in the data set. **Left:** image taken from the robot’s POV. **Right:** transformed image to a human POV.

pothesis in this experiment. We evaluate the performance of the pre-trained YOLOv8n [27] model on homographic-variations of MS-COCO validation dataset. To compute these homographies, we vary the angle of view in the pitch direction (downwards) ( $\theta \in \{0^\circ, 5^\circ, \dots, 55^\circ\}$ ) from the original human viewpoint  $\theta = 0$  in the MS-COCO dataset. We assume a fixed distance of 3m from the person to the camera. Then we transform the MS-COCO validation images using the computed homographies. These images are fed into YOLOv8-Pose-nano for pose estimation. We compute F1-score (harmonic mean of precision and recall) under different confidence thresholds of and a fixed-distance threshold for pose estimation. The resulting plot is shown in Figure 8 (Left). The results show that homography-corruption of the MS-COCO dataset reduces the F1 score of pose estimation at all confidence levels. The mean F1 score of pose estimation drops from 0.789 to 0.745 and a drop in the mean average precision score (mAP50) for the predicted bounding box from 0.911 (original dataset) to 0.638 ( $\theta = 30^\circ$ ).

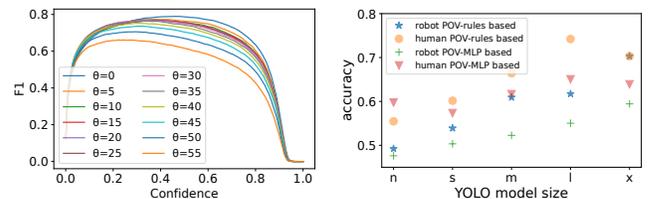


Figure 8. **Left:** F1 Score of Pose Estimation of MS-COCO validation dataset using YOLOv8-Pose-nano. **Right:** Comparison of fall detection accuracy from robot POV vs human POV. By transforming images from the robot POV to the human POV through our Homography matrix, we could improve fall detection performance by 6-12% across models. The x-axis indicates the 5 model sizes (n)nano, (s)small, (m)medium, (l)large, (x)tra-large of YOLOv8-Pose.

3) *Qualitative results:* Figure 9 shows the robot detecting multiple fallen people and sending an email alert with the image as an attachment. In this experiment, multiple individuals are within the view of the robot, and the fall detection module detects and reports the detected falls, accompanied by labeled images. We demonstrate the robot searching for and detecting a fallen person and then email a list of people to be alerted.

The overall demo of the entire system working together is shown as a video here<sup>2</sup>.

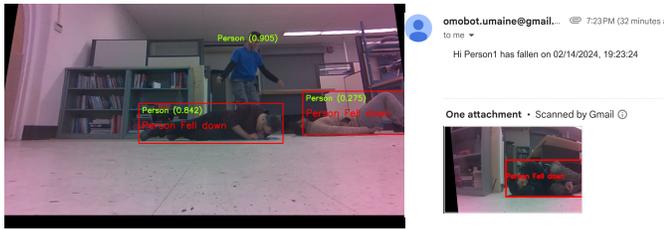


Figure 9. Left: Fall event detection in the presence of multiple people. Right: Fall detected by the robot, and sample email generated by the system.

## VII. CONCLUSION

We describe the design and development of a low-cost robot with a novel design and a customized fall detection pipeline that improves upon YOLOv8-Pose fall detection by a factor of 6-12%.

In the future, we will continue to improve the robustness and accuracy of obstacle avoidance, autonomous charging, and fall detection. We will also follow up on this work by testing this setup in an elderly person's house and getting user evaluations regarding their comfort level with this product.

## ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No 2218063.

## REFERENCES

- [1] M. Tinetti, "Prevention of falls and fall injuries in elderly persons: A research agenda," *Preventive Medicine*, vol. 23, no. 5, pp. 756–762, 1994. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0091743584711303>
- [2] K. Doughty, K. Cameron, and P. Garner, "Three generations of telecare of the elderly," *Journal of Telemedicine and Telecare*, vol. 2, no. 2, pp. 71–80, 1996.
- [3] "Elderly population (indicator)," OECD (2024) (Accessed on 05 Feb 2024). [Online]. Available: <http://dx.doi.org/10.1787/8d805ea1-en>
- [4] P. Rajendran, A. Corcoran, B. Kinoshian, and M. Alwan, *Falls, Fall Prevention, and Fall Detection Technologies*. Totowa, NJ: Humana Press, 2008, pp. 187–202. [Online]. Available: [https://doi.org/10.1007/978-1-59745-233-5\\_8](https://doi.org/10.1007/978-1-59745-233-5_8)
- [5] S. Chaudhuri, H. Thompson, and G. Demiris, "Fall detection devices and their use with older adults: a systematic review," *Journal of geriatric physical therapy (2001)*, vol. 37, no. 4, p. 178, 2014.
- [6] Z. Zhang, C. Conly, and V. Athitsos, "A survey on vision-based fall detection," in *Proceedings of the 8th ACM International Conference on Pervasive Technologies Related to Assistive Environments*, ser. PETRA '15. New York, NY, USA: Association for Computing Machinery, 2015. [Online]. Available: <https://doi-org.wv-o-ursus-proxy02.ursus.maine.edu/10.1145/2769493.2769540>
- [7] X. Wang, J. Ellul, and G. Azzopardi, "Elderly fall detection systems: A literature survey," *Frontiers in Robotics and AI*, vol. 7, 2020. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frobt.2020.00071>
- [8] C. Menacho and J. Ordoñez, "Fall detection based on cnn models implemented on a mobile robot," in *2020 17th International Conference on Ubiquitous Robots (UR)*, 2020, pp. 284–289.
- [9] W. H. Chin, N. Nuo Wi Tay, N. Kubota, and C. K. Loo, "A lightweight neural-net with assistive mobile robot for human fall detection system," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–6.
- [10] G. Chen and X. Duan, "Vision-based elderly fall detection algorithm for mobile robot," in *2021 IEEE 4th International Conference on Electronics Technology (ICET)*, 2021, pp. 1197–1202.
- [11] A. Wong, M. Famuori, M. J. Shafiee, F. Li, B. Chwyl, and J. Chung, "Yolo nano: A highly compact you only look once convolutional neural network for object detection," in *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*. IEEE, 2019, pp. 22–25.
- [12] Z. A. Mundher and J. Zhong, "A real-time fall detection system in elderly care using mobile robot and kinect sensor," *International Journal of Materials, Mechanics and Manufacturing*, vol. 2, no. 2, pp. 133–138, 2014.
- [13] L. Ciabattoni, F. Ferracuti, G. Foresi, A. Freddi, A. Monteriù, and D. P. Pagnotta, "Real-time fall detection system by using mobile robots in smart homes," in *2017 IEEE 7th International Conference on Consumer Electronics - Berlin (ICCE-Berlin)*, 2017, pp. 15–16.
- [14] F. M. Lopez-Rodriguez and F. Cuesta, "An android and arduino based low-cost educational robot with applied intelligent control and machine learning," *Applied Sciences*, vol. 11, no. 1, 2021. [Online]. Available: <https://www.mdpi.com/2076-3417/11/1/48>
- [15] A. A. S. Gunawan, B. Clemons, I. F. Halim, K. Anderson, and M. P. Adianti, "Development of e-butler: Introduction of robot system in hospitality with mobile application," *Procedia Computer Science*, vol. 216, pp. 67–76, 2023, 7th International Conference on Computer Science and Computational Intelligence 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1877050922021901>
- [16] P. He, H. Lv, H. Wu, and G. Yang, "Modeling and control of differential-drive chassis for a homecare assistive robot," in *2023 IEEE 32nd International Symposium on Industrial Electronics (ISIE)*, 2023, pp. 1–4.
- [17] X. Tan, S. Zhang, and Q. Wu, "Research on omnidirectional indoor mobile robot system based on multi-sensor fusion," in *2021 5th International Conference on Vision, Image and Signal Processing (ICVISP)*, 2021, pp. 111–117.
- [18] A. A. Süzen, B. Duman, and B. Şen, "Benchmark analysis of jetson tx2, jetson nano and raspberry pi using deep-cnn," in *2020 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, 2020, pp. 1–5.
- [19] T. de J. Mateo Sanguino and J. E. González Ramos, "Smart host microcontroller for optimal battery charging in a solar-powered robotic vehicle," *IEEE/ASME Trans. Mechatronics*, vol. 18, no. 3, pp. 1039–1049, 2013.
- [20] B. Joshi, R. Shrestha, and R. Chaudhary, "Modeling, simulation and implementation of brushed dc motor speed control using optical incremental encoder feedback," 10 2014.
- [21] Stanford Artificial Intelligence Laboratory et al., "Robotic operating system." [Online]. Available: <https://www.ros.org>
- [22] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Trans. Robot.*, vol. 39, no. 3, pp. 1686–1705, 2023.
- [23] E. Alam, A. Sufian, P. Dutta, and M. Leo, "Real-time human fall detection using a lightweight pose estimation technique," in *Computational Intelligence in Communications and Business Analytics*, K. Dasgupta, S. Mukhopadhyay, J. K. Mandal, and P. Dutta, Eds. Cham: Springer Nature Switzerland, 2024, pp. 30–40.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. Springer, 2014, pp. 740–755.
- [25] K. Wang, B. Fang, J. Qian, S. Yang, X. Zhou, and J. Zhou, "Perspective transformation data augmentation for object detection," *IEEE Access*, vol. 8, pp. 4935–4943, 2019.
- [26] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. New York, NY, USA: Cambridge University Press, 2003.
- [27] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics yolov8," 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [28] U. K. Kandagatla, "Fall detection dataset," Available at <https://www.kaggle.com/datasets/uttejkmkandagatla/fall-detection-dataset> (05/10/2024), 2022.
- [29] H. Taheri, B. Qiao, and N. Ghaeminezhad, "Kinematic model of a four mecanum wheeled mobile robot," *International Journal of Computer Applications*, vol. 113, pp. 6–9, 03 2015.

<sup>2</sup><https://youtu.be/wcP0rxz69o>