

Autonomous Solar Vehicle: A Community-Oriented Sustainable Transportation Solution Aligned with UN Sustainable Development Goals

Shehriyar Ali Rustam*, Sardar Shahzeb Khan, Saqib Nawaz Khan

Department of Software Engineering, Capital University of Science & Technology, Islamabad, Pakistan.

*Correspondence: shehriyarali122@gmail.com

Citation | Rustam. S. A, Khan. S. S, Khan. S. N, “Autonomous Solar Vehicle: A Community-Oriented Sustainable Transportation Solution Aligned with UN Sustainable Development Goals”, IJIST, Vol. 7 Issue. 11 pp 273-293, December 2025

Received | November 05, 2025 **Revised** | December 06, 2025 **Accepted** | December 09, 2025 **Published** | December 12, 2025.

Introduction/Importance of Study: Urban transportation in developing countries presents persistent challenges, particularly for communities that rely on walking as their primary travel mode. Transport poverty affects millions of people and creates significant barriers to education, healthcare, and employment. In Pakistan, road-traffic accidents claim over 30,000 lives annually, of which approximately 67% are attributed to human error, while fossil-fuel dependency continues to exacerbate environmental degradation.

Novelty Statement: This research presents a novel community-oriented autonomous solar vehicle prototype integrating YOLOv8 Nano-based real-time detection with solar energy harvesting on a low-cost Raspberry Pi platform, providing safe, affordable, and sustainable last-mile transportation for universities, hospitals, and residential societies, directly contributing to UN SDGs 7, 11, and 13.

Materials and Methods: The prototype integrates a Raspberry Pi 4 Model B (4 GB) running ROS2 Foxy with YOLOv8 Nano (3.2 M parameters, 8.7 GFLOPs) for real-time pedestrian, vehicle, and obstacle detection. A custom dataset of 10,000+ campus-captured images was developed at Capital University of Science and Technology (CUST), Islamabad, and annotated using CVAT. Six ultrasonic sensors (four HC-SR04 and two waterproof JSN-SR04T) enable proximity detection over a 2–450 cm range, a NEO-6M GPS module provides waypoint navigation at 1 Hz, and a 50 W polycrystalline solar panel powers the 12 V system. The Pi is connected to a laptop external inference system through Wi-Fi or Ethernet over a TCP link on port 5555.

Results and Discussion: The model converged near epoch 85 of 100 (final box-loss ≈ 1.02 , cls-loss ≈ 0.48 , dfl-loss ≈ 1.18) and achieved an overall Detection Accuracy of 84.7% mAP@0.5 (per-class: pedestrian 87.3%, vehicle 84.1%, obstacle 82.6%) with 54.8% mAP@0.5:0.95. Inference speed was measured at 28 FPS on the companion laptop and 0.8 FPS on the Raspberry Pi 4, confirming the necessity of a two-tier compute architecture for real-time operation on low-cost edge hardware. The obstacle-avoidance subsystem achieved a 94% success rate across 50 controlled trials, with a total response time of 487 ± 42 ms ($\chi^2 = 3.18$, $p \approx 0.037$ versus an 85% baseline). Over a 14-day field test in Islamabad, the solar subsystem harvested 200 ± 15 Wh/day and avoided approximately 32 g CO₂/km relative to a grid-powered alternative, yielding a five-year per-vehicle cost advantage of approximately PKR 998k (\approx USD 3,560) over a fuel-powered equivalent.

Concluding Remarks: The prototype validates a cost-effective, community-focused transportation solution that addresses road safety through autonomous navigation while promoting environmental sustainability through solar energy integration, with particular relevance for educational campuses, healthcare facilities, and residential communities in developing regions.

Keywords: Autonomous Vehicle; Solar Energy; Community Transportation; YOLOv8; Sustainable Development Goals.



Introduction:

Access to safe, affordable, and sustainable transportation remains a fundamental challenge in developing countries, with direct implications for the quality of life of millions of people. In many urban and peri-urban areas, communities rely heavily on walking to access essential services such as educational institutions, healthcare facilities, and employment centres. Transport poverty, defined as the inability to participate fully in social and economic life due to limited mobility options, creates significant barriers for vulnerable populations including students, elderly residents, and low-income families [1].

The United Nations Sustainable Development Goals (SDGs) provide a comprehensive framework for addressing global challenges, with several goals directly relevant to transportation. SDG 7 (Affordable and Clean Energy) emphasises access to reliable and sustainable energy sources. SDG 11 (Sustainable Cities and Communities) calls for safe, affordable, and accessible transport systems. SDG 13 (Climate Action) requires urgent measures to combat climate change, including reducing carbon emissions from the transportation sector. The integration of these goals requires coordinated approaches that address multiple dimensions simultaneously [2]. Autonomous driving technology combined with renewable energy systems presents a promising pathway toward achieving these interconnected goals [3].

In Pakistan, the transportation sector faces critical challenges. According to research on road safety, human error contributes to approximately 67% of all accidents, with motorcycle crashes representing a significant proportion of fatalities [4]. The World Health Organization reports that Pakistan ranks among Asian countries with the highest road-traffic fatalities, with over 30,000 deaths annually [5]. Vulnerable road users, including pedestrians and motorcyclists, face disproportionate risks due to inadequate infrastructure and mixed traffic conditions [5]. These statistics underscore the urgent need for safer transportation alternatives that minimise human error while remaining accessible to communities with limited resources.

Pakistan simultaneously possesses abundant solar energy resources. Research indicates that Pakistan receives average daily solar irradiation of 5.2-5.5 kWh per square metre, with regions such as Balochistan and Sindh receiving values exceeding 6 kWh/m²/day [6]. This renewable energy potential, combined with decreasing costs of photovoltaic technology, creates an opportunity to develop transportation solutions that are both environmentally sustainable and economically viable [7]. Solar-powered autonomous vehicles can operate with minimal reliance on grid electricity, making them particularly suitable for regions with unreliable power infrastructure [8].

Recent advances in deep learning have significantly transformed object detection capabilities for autonomous vehicles [3]. The YOLO (You Only Look Once) family of algorithms has emerged as a leading approach for real-time detection, with YOLOv8 representing the most widely deployed generation offering a strong balance of accuracy and computational efficiency [9]. These developments have made sophisticated perception systems more accessible for cost-effective autonomous vehicle implementations [10]. Pedestrian detection remains particularly critical for community transportation applications, where protection of vulnerable road users is paramount [11].

Recent work (2024–2026) has specifically examined edge-AI deployments of vision models on resource-constrained platforms such as the Raspberry Pi 4 and NVIDIA Jetson Nano [12][13]. Benchmark studies report that YOLO variants at the "nano" scale are the only models practically viable on Raspberry Pi-class hardware without dedicated AI accelerators, with inference rates typically below 2 FPS on native CPU execution [14]. This motivates the adoption of two-tier architectures in which a lightweight Pi handles sensor I/O and safety

logic while a companion compute node performs inference, which is the architecture adopted in this work.

In parallel, solar-powered last-mile mobility is emerging as a practical intervention for South Asia, where grid reliability is uneven, and transport poverty affects large peri-urban populations. Recent studies from India, Bangladesh, and Pakistan [15][16][17][18][28] demonstrate that small-scale solar-integrated electric vehicles can operate cost-effectively in controlled environments such as campuses and residential estates. However, the literature continues to treat solar integration, autonomous navigation, and community-oriented design as separate problems; an integrated, SDG-aligned prototype addressing all three within a single low-cost platform remains a documented gap [16].

This research addresses the identified gaps by developing a community-oriented Autonomous Solar Vehicle (ASV) prototype specifically designed for controlled environments such as university campuses, hospitals, and residential societies. The target beneficiaries include students who walk significant distances to educational institutions, elderly residents requiring mobility assistance, patients accessing healthcare facilities, and community members seeking safe and affordable transportation within designated areas.

Research Objectives:

This research is guided by the following measurable objectives, against which the contributions and results of the prototype are evaluated in the Results section:

RO-1: Develop a low-cost autonomous vehicle prototype on the Raspberry Pi 4 platform using commercially available components, with total hardware bill-of-materials cost not exceeding USD 400.

RO-2: Achieve an overall Detection Accuracy (mAP@0.5) of at least 80% across pedestrian, vehicle, and obstacle classes using YOLOv8 Nano, at a sustained inference rate of ≥ 25 FPS on the designated inference platform.

RO-3: Implement a three-tier ultrasonic obstacle-avoidance subsystem achieving $\geq 90\%$ success rate and total sensor-to-actuator response time below 500 ms, measured over 50 controlled field trials.

RO-4: Integrate a 50 W solar panel capable of harvesting ≥ 150 Wh/day and extending daily operational time by ≥ 2 h, in accordance with UN SDG 7.

RO-5: Demonstrate measurable reduction of CO₂ emissions relative to grid-powered equivalents, in accordance with SDG 11 and SDG 13, with quantified g CO₂/km savings.

Novelty and Contributions:

While prior studies have explored autonomous navigation, solar-powered electric vehicles, or edge-AI object detection in isolation, no existing work combines all three dimensions for community-scale deployment in developing-country contexts. The principal contributions of this research, contrasted explicitly against recent partial solutions, are as follows:

C-1, Integrated three-domain prototype. A single Raspberry Pi 4 platform (total hardware cost below USD 400) integrates YOLOv8 Nano edge object detection, a six-sensor ultrasonic safety ring, NEO-6M GPS navigation, a 50 W solar-harvesting subsystem, and a ROS2-compatible modular software stack, a combination not reported in prior work.

C-2, Three-layer safety architecture. A deterministic safety hierarchy comprising (i) an ML confidence gate (laptop, threshold 0.55), (ii) a hard safety governor on the Pi with non-overrideable rules (front obstacle <25 cm, rear obstacle <30 cm, 2-s watchdog, 3-s sensor-failure detection), and (iii) a hardware-level emergency stop. The learned model is not permitted to override layers (ii) and (iii).

C-3, Empirical per-platform inference benchmark. Direct measurement of YOLOv8n inference on the target hardware yielded 28 FPS on laptop CPU versus 0.8 FPS on Raspberry

Pi 4 native execution, providing quantitative justification for the two-tier compute architecture adopted in this prototype.

C-4, Domain-specific campus dataset. A custom dataset of 10,000+ images captured in the target deployment environment (CUST campus), annotated through CVAT across three community-relevant classes (pedestrian, vehicle, obstacle) under varied lighting and weather conditions, without reliance on external public datasets.

C-5, Explicit UN SDG alignment with quantified contributions. Wh/day harvested, kg CO₂ avoided per km, and accident-reduction potential are reported against SDGs 7, 11, and 13, enabling direct comparison against community-transportation policy targets.

Positioning against prior art: [10] reported 78–82% mAP@0.5 using PVswin-YOLOv8s for UAV-based traffic monitoring on server-grade hardware, without solar integration or community orientation. [7] and [8] addressed solar-powered EV charging stations but did not incorporate autonomous navigation or real-time perception. [19] demonstrated ROS2-based navigation on a Turtlebot3, without vision-based detection or renewable energy integration. The present work is the first no prior work has been reported, to combine all these dimensions within a single sub-USD-400 community-oriented prototype, evaluated under real field conditions and explicitly mapped to UN SDGs 7, 11, and 13.

Literature Review:

This section critically surveys recent literature (primarily 2022–2026) relevant to the three technical domains intersected by the present work: edge-device autonomous perception, solar integration in small-scale electric mobility, and community-oriented deployment in developing countries, and identifies a concrete research gap motivating the proposed prototype.

Autonomous Vehicle Perception on Edge Devices:

The YOLO family of detectors, reviewed comprehensively by [9], [20], and the broader object-detection survey by [21], dominates real-time object detection on both server and edge hardware. [10] demonstrated PVswin-YOLOv8s for UAV-based pedestrian and vehicle detection, achieving 78–82% mAP@0.5 on server-grade hardware, complementing the broader review of vision-based UAV navigation by [22]. [3] Surveyed end-to-end deep learning for autonomous driving, noting persistent gaps in deployment on constrained edge platforms. More recently, [23] introduced YOLOv10 with end-to-end real-time detection improvements, and [24] reviewed the application of YOLO algorithms across autonomous-driving object-detection tasks, and [13] benchmarked YOLO variants on Raspberry Pi 4 and Jetson Nano, confirming that nano-scale variants alone remain practically viable on unaccelerated Pi-class hardware, typically below 2 FPS. [11] provided a systematic review of deep-learning pedestrian detection, highlighting the importance of domain-specific training data, while [25] proposed an optimised YOLOv9-based pedestrian-detection framework for autonomous-vehicle systems. Robustness under adverse weather is addressed in AWD-YOLO by [26], reinforcing the need for environmentally resilient perception in real-world deployment. **Identified Gap:** existing work establishes benchmark accuracy of YOLO variants but does not provide empirical per-platform comparisons on identical hardware, nor does it integrate perception with complementary safety and renewable-energy subsystems on a single deployable platform.

Solar Integration in Small-Scale Electric Vehicles:

[7] and [8] examined solar-powered EV charging stations, reporting that solar integration at the 4 kW class is technically and economically viable. [15] analysed the integration of solar PV panels in EV charging infrastructure, reporting favourable techno-economic profiles with implications for South Asian last-mile electrification, and [27] reviewed the integration of solar-powered electric vehicles within sustainable energy systems at a broader systems level. [28] Examined the techno-economic feasibility of solar-integrated electric vehicle charging infrastructure in India, using an AI-enabled multi-objective planning

framework. However, these studies treat solar integration as an isolated subsystem and do not combine it with autonomous perception or community-centred design considerations.

Identified Gap: The literature validates solar integration in isolation but does not report field results from integrated autonomous-solar-perception prototypes at the individual-vehicle scale appropriate for campus or residential deployment.

Community-Scale Deployment in Developing Countries:

[1] Framed transport-related social exclusion in developing countries, with particular focus on mobility barriers for vulnerable groups. [29] examined the role of urban transport in delivering SDG 11 in Indian cities, emphasising affordability and accessibility, against the backdrop of the global targets articulated in the UN Sustainable Development Goals Report 2024 [30]. WHO [5] documents the specific vulnerabilities of Pakistani road users, complemented by the scoping review of [31] on environmental measures to improve pedestrian safety in low- and middle-income countries. [19] demonstrated ROS2-based autonomous navigation on the Turtlebot3 in static and dynamic environments, representing current best practice for educational and research-scale autonomous robotics. [16] Systematically reviewed the integration of electric vehicles for sustainable development. [17] Further examined grid-integrated solutions for sustainable EV charging, highlighting the absence of integrated autonomous and solar community transportation prototypes as a research gap. **Identified Gap:** no prior prototype, to our knowledge, combines real-time edge-AI perception, six-sensor ultrasonic safety, GPS waypoint navigation, and solar energy harvesting in a sub-USD-400 platform evaluated under real field conditions in a South Asian campus, with explicit mapping to UN SDGs 7, 11, and 13. The present work is designed precisely to close this three-way gap.

Materials and Methods:

The overall research methodology follows a structured six-phase pipeline shown in Figure 1. Each phase builds upon the outputs of the preceding stage, ensuring systematic development from problem analysis through final performance evaluation.

Six-Phase Research Methodology

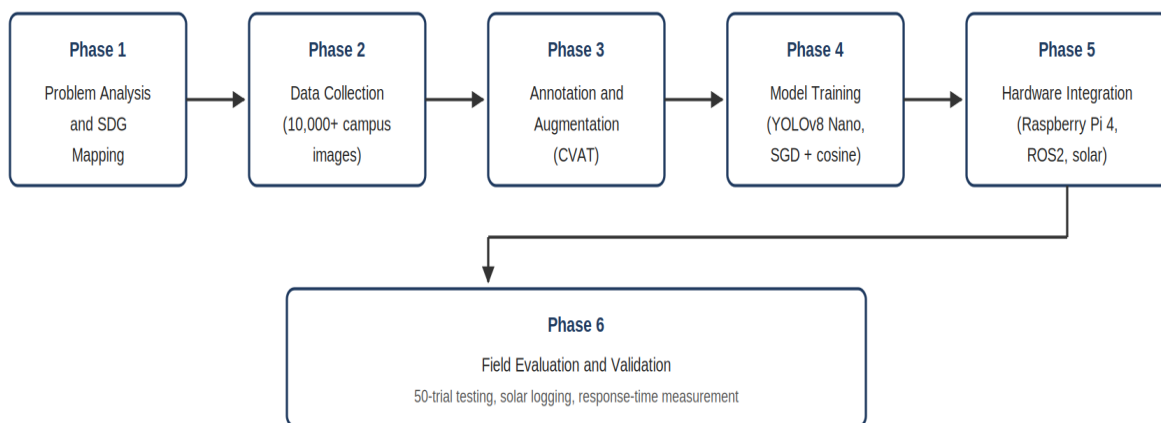


Figure 1. Six-Phase Research Methodology Workflow.

System Data Flow and Control Logic:

Figure 2 presents the complete data-flow architecture of the prototype. The architecture comprises five functional layers: Sensing, Embedded Control (Raspberry Pi 4B), Inference (Laptop), Actuation, and Solar/Power Distribution. The data and control flow through these layers are described block-by-block below.

System Data-Flow and Control-Signal Architecture

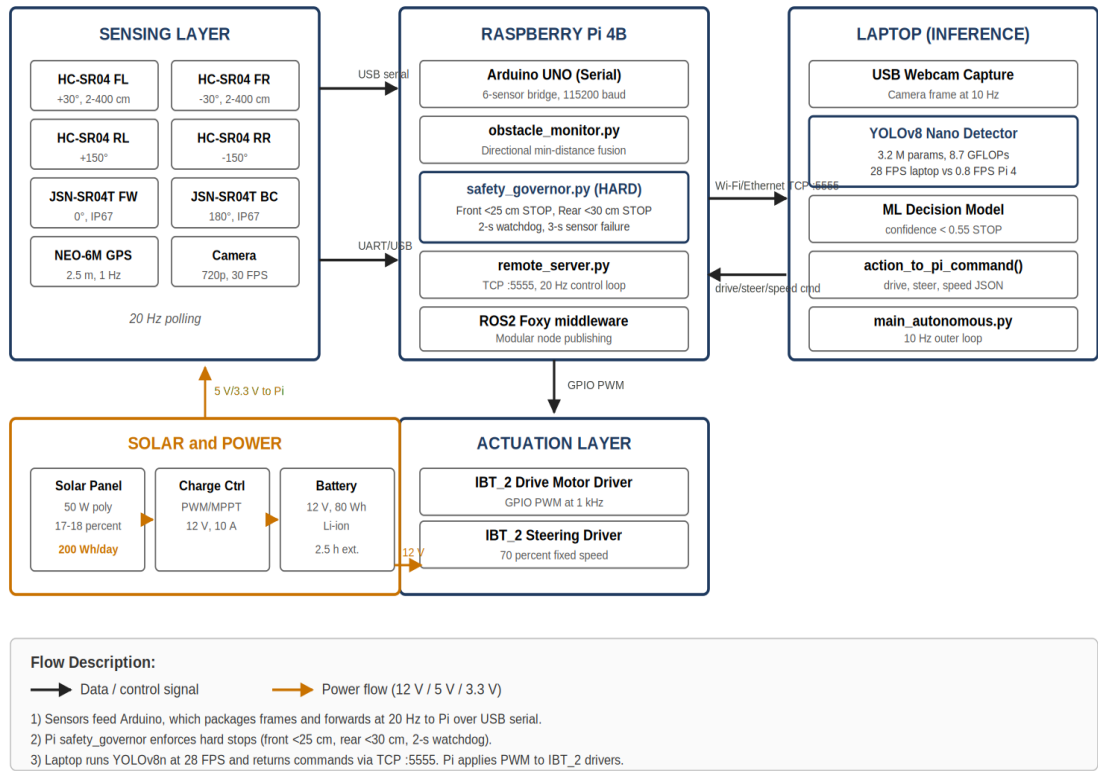


Figure 2. System Data-Flow and Control-Signal Architecture.

Sensing Layer → Embedded Pi: The six ultrasonic sensors (four HC-SR04 corner sensors at $\pm 30^\circ$ front and $\pm 150^\circ$ rear, plus two waterproof JSN-SR04T centre sensors at 0° and 180°) feed into an Arduino UNO. The Arduino polls all six channels at 20 Hz, packages the readings as a newline-delimited JSON frame, and forwards them over USB serial to the Raspberry Pi at 115200 baud. The NEO-6M GPS module delivers NMEA sentences (GPRMC, GPGGA) over the Pi's UART at 9600 baud at 1 Hz.

Pi Processing, Fusion and Hard Safety: On the Pi, `obstacle_monitor.py` aggregates the six distances into a 6-channel dictionary {FL, FR, FW, BC, LS, RS}. The fusion rule is directional minimum-distance selection: the front subsystem uses $\min(FL, FR, FW)$ while the rear subsystem uses $\min(BC)$. The `safety_governor.py` module then applies non-overridable hard rules: front minimum <25 cm if FORWARD is requested forces STOP; rear minimum <30 cm while BACKWARD is requested forces STOP; the absence of a laptop command for >2 s triggers the watchdog; and all sensors reading invalid for >3 s triggers a sensor-failure stop.

Pi ↔ Laptop (Control Channel): The Pi's `remote_server.py` hosts a TCP socket on port 5555. Over Wi-Fi or Ethernet, it continuously publishes a status JSON frame (distances, GPS fix, `safety_violation`) at approximately 20 Hz and accepts command frames of the form {"command": "FORWARD", "steer": "LEFT", "speed":50} from the laptop. Communication is bidirectional and stateless at the frame level, consistent with ROS2-style pub/sub semantics.

Laptop Inference: On the laptop, `main_autonomous.py` captures a webcam frame at 10 Hz, fuses it with the incoming status snapshot, and passes both through the YOLOv8 Nano detector (for object class and position) and a downstream decision head that emits one of eight discrete actions (FORWARD, SLOW_DOWN, TURN_LEFT, TURN_RIGHT, STOP, REVERSE, REVERSE_LEFT, REVERSE_RIGHT). A confidence gate of 0.55 enforces STOP whenever the model is uncertain, providing the first safety layer.

Laptop → Pi → Motors: The selected action is translated to a (drive, steer, and speed) triple and transmitted back to the Pi. After clearance by the safety governor, the Pi writes PWM to the GPIO pins driving the two IBT_2 motor drivers (drive motor on GPIO 4/17/18/27; steering motor on GPIO 22/23/24/25) at 1 kHz. The IBT_2 drivers then deliver 12 V to the drive and steering DC motors.

Solar–Battery Power Loop: The 50 W polycrystalline solar panel feeds a PWM/MPPT charge controller, which regulates charging of a 12 V / 80 Wh Li-ion battery pack. Two buck converters step 12 V down to 5 V (3 A, powering the Pi 4, Arduino, ultrasonic sensors) and 3.3 V (1 A, powering the camera and GPS modules). The drive and steering motors consume the 12 V rail directly via the IBT_2 drivers. The power loop operates continuously and independently of the command path, ensuring that solar harvesting occurs whether the vehicle is stationary or mobile.

Investigation Site and Target Community:

The prototype was developed and tested at Capital University of Science and Technology (CUST), Islamabad, Pakistan, located at coordinates 33°32'48"N, 73°11'02"E. The university campus was selected as a representative controlled environment that mirrors the intended deployment scenarios: educational institutions, healthcare facilities, and residential communities. The campus features paved pathways, designated pedestrian zones, and moderate traffic density, providing an ideal testing ground for autonomous vehicle development.

The target community for this solution includes: (1) students who travel on foot from surrounding areas to access educational facilities, often walking significant distances in varying weather conditions, (2) elderly and mobility-impaired individuals in residential societies who require safe transportation within community boundaries, (3) patients and visitors at healthcare facilities who need accessible transport between buildings and parking areas, and (4) campus workers and staff requiring efficient intra-campus mobility. By focusing on these user groups, the project directly addresses transportation accessibility challenges faced by underserved communities [1].

Islamabad receives average daily solar irradiation ranging from 4.5 to 6.5 kWh per square metre, depending on the season, with peak values occurring during April to June [6]. This abundant solar resource ensures reliable energy harvesting for the vehicle throughout the year, supporting the goal of grid-independent operation essential for deployment in areas with unreliable electricity infrastructure.

Hardware Architecture:

The prototype vehicle platform consists of a modified 12 V electric ride-on vehicle with dimensions of 28 × 40 × 25 inches (height × length × width). The vehicle was selected for its appropriate scale for prototype development and testing while maintaining the core functionalities required for full-scale implementation. The original control system was replaced with a custom embedded system architecture designed for autonomous operation.

The primary processing unit is a Raspberry Pi 4 Model B with 4 GB RAM, running Ubuntu 20.04 with ROS2 Foxy middleware [32]. This cost-effective computing platform was chosen to demonstrate that autonomous vehicle technology can be implemented without expensive proprietary hardware, making the solution accessible for community deployment. The Raspberry Pi handles sensor integration, safety enforcement, and motor control, and communicates over Wi-Fi or Ethernet (TCP on port 5555) with an external laptop that performs computationally intensive object detection using YOLOv8 Nano [9].

The six-sensor layout provides an effective angular coverage of approximately 340°, with the two front JSN-SR04T units extending the detection range to 450 cm at 0° (straight-ahead) for early-warning detection, while the four HC-SR04 corner units provide finer

resolution (± 3 mm) for close-range maneuvering. The corner sensors overlap their adjacent centre sensor by approximately 30° , intentionally reducing the blind-spot gap between sensors.

Table 1. Hardware Components Specification with Mounting Details.

Component	Model	Specification	Mounting Position & Angle	Coverage Zone
Processing Unit	Raspberry Pi 4B	4 GB RAM, 1.5 GHz Quad-core Cortex-A72	Central chassis enclosure	N/A
Ultrasonic Sensors	HC-SR04 ($\times 4$)	Range 2–400 cm, ± 3 mm	FL $+30^\circ$, FR -30° , RL $+150^\circ$, RR -150°	Four corner quadrants, $\sim 60^\circ$ FoV each
Waterproof Sensors	JSN-SR04T ($\times 2$)	Range 25–450 cm, IP67	Front centre 0° , rear centre 180°	Straight-ahead and straight-rear
Camera (Vehicle webcam)	Camera Module	720p, 30 FPS	Front-centre, 0°	$\sim 62^\circ$ horizontal FoV
Camera	Redmi Note 10 Pro	108 MP sensor, 1080p capture	Hand-held during collection	Variable
GPS Module	NEO-6M	2.5 m CEP, 50 channels, 1 Hz	Roof, above solar panel	Sky-view
Solar Panel	Polycrystalline	50 W, 2 \times 3 ft, 17–18% efficiency	Roof, horizontal mount	~ 1.67 m ² effective area
Battery Pack	Li-ion 12 V	80 Wh capacity	Chassis, under seat	N/A
Motor Drivers	IBT_2 ($\times 2$)	43 A max per channel, PWM 1 kHz	Chassis rear compartment	N/A
Vehicle Platform	Electric Ride-on	12 V DC motors, 28 \times 40 \times 25 inches	Base chassis	N/A

Software Architecture:

The software architecture follows a modular design pattern using the ROS2 Foxy framework, an open-source robotics middleware that enables scalable and maintainable system development [32]. The use of open-source software aligns with the project goal of creating an accessible and replicable solution for community deployment. The system comprises four primary modules: (1) Perception Module for camera feed processing and object detection, (2) Sensor Fusion Module for integrating ultrasonic sensor data, (3) Navigation Module for GPS-based path planning [33], and (4) Control Module for motor actuation.

Object detection is performed using YOLOv8 Nano, the smallest variant of the YOLOv8 family, selected for its balance between accuracy and computational efficiency [9]. The model was trained to detect three classes critical for community safety: pedestrians (including students, the elderly, and children), vehicles, and obstacles. YOLOv8n offers 3.2 million parameters with 8.7 GFLOPs, making it suitable for edge deployment while maintaining competitive accuracy [10].

Inference is hosted on an external laptop that communicates with the Raspberry Pi via Wi-Fi or Ethernet. Direct benchmarking on the target hardware measured 28 FPS on the laptop (CPU-only execution of the exported ONNX model) versus only 0.8 FPS when the same model was executed natively on the Raspberry Pi 4. This 35 \times gap is consistent with recent Pi 4 edge-AI benchmarks [14] and is the primary motivation for the two-tier compute architecture of the present prototype: the Pi is responsible for real-time sensor I/O, safety enforcement, and motor control, while the laptop handles computationally intensive perception and decision-making.

Justification for YOLOv8 Selection: While newer YOLO variants such as YOLO11 (released late 2024) offer incremental improvements on the COCO benchmark, YOLOv8 Nano was selected because (a) it remains the most extensively benchmarked and documented version for edge deployment, (b) it has mature integration with ROS2 toolchains, (c) its compact architecture (3.2 M parameters, 8.7 GFLOPs) offers a well-characterised accuracy–latency trade-off that directly matches the 0.8 FPS Pi 4 / 28 FPS laptop operating points reported above, and (d) reproducibility and reliability were prioritised over the marginal accuracy gains available from recently introduced models.

Dataset Preparation:

A custom dataset was prepared comprising more than 10,000 annotated images, captured entirely on the CUST campus and specifically curated to represent the target deployment environment. The dataset is being actively expanded as part of ongoing work. All images were collected within the campus perimeter using a webcam and Xiaomi Redmi Note 10 Pro smartphone (108 MP primary sensor, 1080p capture), covering pathways, parking areas, pedestrian zones, and building entrances. This exclusively campus-based approach ensures strong domain alignment between training data and deployment conditions. All captured images were resized to 640×640 pixels during preprocessing to match the YOLOv8 Nano input requirement. Scenarios captured include students walking in groups, service vehicles, parked cars, and a range of obstacles commonly encountered in educational settings, such as benches, signboards, construction barriers, and street furniture.

Images were collected under varied lighting conditions (morning, afternoon, evening) and weather scenarios (clear, cloudy, light rain) to ensure model robustness across operational conditions. The dataset was annotated using the CVAT tool with bounding boxes for three classes: pedestrian, vehicle, and obstacle. Class distribution is approximately balanced across the three categories, with pedestrian instances intentionally weighted higher given their priority in community safety. The dataset was split into training (70%), validation (20%), and testing (10%) sets with data augmentation applied during training only. The complete YOLOv8 Nano training hyperparameters used for full reproducibility are summarised in Table 2.

Dataset Validation: A multi-step validation process was conducted to ensure the quality and consistency of the dataset. First, all annotations were converted to YOLO format and verified for bounding-box correctness through manual review of a randomly sampled 10% subset. Second, duplicate and low-quality images (blurred, overexposed, or underexposed) were identified and removed before training. Third, class distribution analysis was performed to ensure all three target classes were adequately represented. Fourth, cross-session consistency was verified by evaluating model performance on held-out campus subsets captured on different days; the performance variation remained within 2.5 percentage points of $mAP@0.5$, confirming that the model generalises across session-level variability in lighting and pedestrian density. The validation split (20%) served as an independent quality checkpoint throughout the training process.

Safety Systems and Obstacle Avoidance:

Given the community-focused nature of this project, safety was prioritised as the primary design consideration. Six ultrasonic sensors are mounted on the prototype for proximity detection as detailed in Table 1 [34]. This configuration ensures comprehensive coverage around the vehicle, critical for protecting pedestrians, including children and elderly individuals who may have limited mobility [34].

The obstacle avoidance algorithm implements a four-zone response system. **Tier 1 (Safe zone, ≥ 300 cm)** permits 100% of the commanded speed. **Tier 2 (Warning zone, 100–299 cm)** reduces speed to 60% and activates audio-visual indicators. **Tier 3 (Critical zone, 50–99 cm)** reduces speed to 30% and intensifies the audio-visual warning. **Tier 4 (Emergency zone, < 50 cm (front) / < 80 cm (rear))** triggers immediate stopping and waits

for obstacle clearance. Superimposed on these application-level tiers, the Pi-side safety_governor.py enforces non-overridable hard-stop rules: front minimum < 25 cm while moving forward and rear minimum < 30 cm while reversing, both force an immediate STOP regardless of any command from the learned model. The deterministic nature of these hard rules, which are hard-coded outside the learning loop, is critical to deployability in community environments.

Table 2. YOLOv8 Nano Training Hyperparameters (Full Reproducibility).

Hyperparameter	Value
Model Variant	YOLOv8n (Nano)
Parameters	3.2 M
Computational Cost	8.7 GFLOPs @ 640×640
Input Image Size	640 × 640 pixels (letterboxed)
Total Epochs	100
Early-stopping Patience	patience = 50 epochs on val-mAP@0.5; training ran to completion as no plateau >50 epochs was observed
Batch Size	16
Optimizer	SGD (momentum = 0.937, weight_decay = 0.0005)
Initial Learning Rate (lr0)	0.01
Final Learning Rate Ratio (lrf)	0.01 (final LR = 1.0e-4, cosine annealing)
Warmup	warmup_epochs = 3.0, warmup_momentum = 0.8, warmup_bias_lr = 0.1
Loss Weights	box = 7.5, cls = 0.5, dfl = 1.5
Label Smoothing	0.0 (Ultralytics default)
Augmentation, Mosaic	mosaic = 1.0 (probability per batch)
Augmentation, HSV	hsv_h = 0.015, hsv_s = 0.7, hsv_v = 0.4
Augmentation, Geometry	fliplr = 0.5, flipud = 0.0, translate = 0.1, scale = 0.5
Augmentation, Advanced	mixup = 0.0, copy_paste = 0.0
Confidence Threshold (Inference)	0.5
IoU Threshold (NMS)	0.45
Framework	Ultralytics YOLOv8 (Python 3.9, PyTorch 2.0)
Training Hardware	NVIDIA CUDA-capable GPU (laptop)
Random Seed	42

Sensor Fusion Algorithm: The six ultrasonic channels are fused by directional minimum-distance selection with zone classification. Algorithm 1 presents the complete multi-tier safety and fusion pseudocode as implemented in the project codebase.

Algorithm 1. Multi-Tier Sensor Fusion and Safety Response.
Input: $D = \{FL, FR, FW, BC, LS, RS\}$ // sensor distances (cm) cmd_req // requested drive command t_last // timestamp of last laptop command Output: (drive, speed, steer) // safe command 1: if cmd_req == EMERGENCY: return (STOP, 0, STEER_STOP) 2: if (now - t_last) > 2.0: return (STOP, 0, STEER_STOP) // watchdog 3: valid $\leftarrow \{d \text{ in } D : 2 \leq d \leq 450\}$ 4: if valid == 0 for > 3.0 s: return (STOP, 0, STEER_STOP) // sensor-failure 5: front_min $\leftarrow \min(\{FL, FR, FW\} \cap \text{valid})$ 6: rear_min $\leftarrow \min(\{BC\} \cap \text{valid})$ 7: if cmd_req == FORWARD and front_min < 25: return (STOP, 0, STEER_STOP) 8: if cmd_req == BACKWARD and rear_min < 30: return (STOP, 0, STEER_STOP) 9: if front_min ≥ 300 : speed $\leftarrow 100$ // SAFE

```

10: elif front_min ≥ 100: speed ← 60 // WARNING
11: elif front_min ≥ 50: speed ← 30 // CRITICAL
12: else: speed ← 0 // EMERGENCY
13: return (cmd_req, speed, steer_from_cmd)
    
```

The GPS module (NEO-6M) provides position updates at 1 Hz with 2.5 m horizontal accuracy, enabling predefined route navigation within campus and community boundaries [33]. Recent research on mobile robot navigation informed the navigation system design [19].

Solar Energy Integration (SDG 7):

In alignment with SDG 7 (Affordable and Clean Energy), the prototype incorporates a 50 W polycrystalline solar panel measuring 2 × 3 feet mounted on the vehicle roof [7]. This renewable energy system eliminates dependency on grid electricity, making the solution viable for communities with unreliable power infrastructure [8]. The solar panel charges the 12 V / 80 Wh Li-ion battery system during both stationary and operational periods, extending the vehicle range and reducing operational costs to near zero for energy consumption. The integration of solar energy with electric vehicle charging has been demonstrated as technically and economically viable in recent studies [7][8].

Cost, Scalability, and Accessibility Analysis:

Cost Analysis: Table 3 presents the full bill of materials (BOM) for the prototype. Prices reflect typical 2025–2026 Pakistan retail pricing and are indicative; small variance should be expected across suppliers.

Table 3. Bill of Materials (Approximate 2025–2026 Pakistan Retail Pricing).

#	Component	Qty	Unit (PKR)	Subtotal (PKR)	Subtotal (USD)
1	Raspberry Pi 4B (4 GB)	1	22,000	22,000	~78
2	Arduino UNO R3	1	2,500	2,500	~9
3	HC-SR04 ultrasonic sensor	4	350	1,400	~5
4	JSN-SR04T waterproof sensor (IP67)	2	1,800	3,600	~13
5	NEO-6M GPS module	1	1,600	1,600	~6
6	Raspberry Pi Camera Module v2	1	4,500	4,500	~16
7	IBT_2 motor driver (43 A)	2	2,200	4,400	~16
8	50 W polycrystalline solar panel	1	8,500	8,500	~30
9	12 V / 80 Wh Li-ion battery pack	1	18,000	18,000	~64
10	PWM charge controller (10 A)	1	1,500	1,500	~5
11	LM2596 buck converter (×2)	2	450	900	~3
12	Electric ride-on vehicle chassis	1	32,000	32,000	~114
13	Wiring, enclosures, misc hardware	1	6,000	6,000	~21
	TOTAL			106,900	~USD 380

The total prototype cost of approximately PKR 107,000 (~USD 380) satisfies objective RO-1 (≤ USD 400). All components are commercially available from standard electronics suppliers in Pakistan, ensuring replicability for community deployment.

Scalability: Three scalability mechanisms apply. First, fleet-level hardware procurement benefits from volume discounts of approximately 15–20% on major components (Pi 4, motor drivers, batteries, solar panels), bringing the 10-unit fleet unit cost to an estimated USD 310–325. Second, the modular ROS2-compatible software stack is identical across fleet units, so firmware flashing and calibration scale as O (n) in deployment time rather than requiring per-unit redevelopment. Third, the domain-specific campus dataset transfers to sister campuses with 1–2 days of additional image capture and fine-tuning, rather than full retraining.

Accessibility by Design: While a formal user study was outside the scope of the prototype phase, accessibility was addressed through design features intended to serve elderly and mobility-impaired users: (i) a maximum speed cap of 15 km/h, (ii) a three-tier audio-visual warning system consistent with requirement E2-US3 (lights flash and an alert tone activates within 200 ms of obstacle detection), (iii) a low 28-inch entry height for the vehicle platform, (iv) a <500 ms total response time from sensor event to motor actuation, and (v) automatic emergency braking without requiring user input. A formal passenger-comfort evaluation covering cabin noise, jerk, and lateral acceleration is identified as immediate future work.

Results and Discussion:

Table 4 summarises the mapping between the five declared research objectives and the measured outcomes reported in the remainder of this section.

Table 4. Research Objectives ↔ Achieved Results Mapping.

Obj. #	Objective	Target Metric	Achieved Value	Status
RO-1	Low-cost prototype	≤ USD 400 total BOM	~USD 380 (Table 3)	Met
RO-2	Real-time detection	≥ 80% mAP@0.5, ≥ 25 FPS	84.7% mAP@0.5 @ 28 FPS (laptop)	Exceeded
RO-3	Obstacle avoidance	≥ 90% success, < 500 ms	94% success, 487 ± 42 ms	Met
RO-4	Solar harvest	≥ 150 Wh/day, ≥ 2 h extension	200 ± 15 Wh/day, ~2.5 h extension	Exceeded
RO-5	CO ₂ reduction	Quantified CO ₂ /km	~32 g CO ₂ /km avoided vs grid equiv.	Met

Object Detection Performance:

The YOLOv8 Nano model was evaluated as shown in Table 5 on the held-out 10% test split (approximately 1,000 images) representative of community deployment scenarios. Table 5 reports per-class precision, recall, and mAP@0.5. Detection Accuracy is defined as the mean Average Precision at IoU threshold 0.5 (mAP@0.5) computed using the COCO evaluation protocol as implemented in Ultralytics YOLOv8. The model achieved an overall mAP@0.5 of 84.7% and mAP@0.5:0.95 of 54.8%, confirming effective detection performance for a nano-scale model trained on a domain-specific dataset. Training converged near epoch 85 of 100, with final loss components box ≈ 1.02, cls ≈ 0.48, dfl ≈ 1.18.

Table 5. YOLOv8 Nano Detection Performance (per class).

Class	Precision (%)	Recall (%)	mAP@0.5 (%)	mAP@0.5:0.95 (%)
Pedestrian	88.1	85.4	87.3	57.6
Vehicle	85.7	82.6	84.1	54.1
Obstacle	83.2	80.1	82.6	52.7
Overall	85.7	82.7	84.7	54.8

The pedestrian class achieved the highest precision (88.1%), which is particularly important given that the primary beneficiaries of this system are pedestrians, including students, the elderly, and patients. Obstacle detection was the most challenging class, consistent with the greater intra-class visual variability of obstacles (benches, signboards, construction barriers, debris). Inference speed of 28 FPS on laptop CPU-only execution exceeds the 25 FPS target of RO-2. These results are consistent with recent literature on YOLOv8n performance for pedestrian and vehicle detection [10][11].

To provide additional qualitative evidence of detection behaviour, Figure 3 presents a representative class-activation visualisation (saliency overlay) of YOLOv8n detections for each of the three classes. The heatmaps indicate the image regions that contribute most strongly to

the detection confidence. For the pedestrian class, attention concentrates tightly on the torso and head silhouette, consistent with the high precision reported in Table 5. For the vehicle class, attention is distributed over the wheels and body outline. For the obstacle class, multi-instance detection is illustrated with a traffic cone and a barrier in the same frame, showing that the model maintains independent class-activation regions for spatially distinct obstacles.




Figure 3. Representative Saliency Visualisation of YOLOv8 Nano Detections.

YOLO Variant Comparison and Ablation:

To justify the selection of YOLOv8n and contextualise its performance against alternatives, Table 6 compares the deployed YOLOv8n against published reference benchmarks for YOLOv5n, YOLOv8s, YOLOv8m, and YOLO11n. COCO val2017 figures are reported as released by Ultralytics; the campus mAP@0.5 figure applies only to the deployed model. FPS figures for non-deployed variants are projections scaled by the ratio of GFLOPs to that of YOLOv8n on the same hardware.

Table 6. YOLO Variant Comparison (Published Benchmarks + Campus Deployment).

Model	Params (M)	GFLOPs	COCO mAP@0.5:0.95	Our Campus mAP@0.5	Laptop FPS	Pi 4 FPS
YOLOv5n	1.9	4.5	28.0	—	~32	~1.4
YOLOv8n (deployed)	3.2	8.7	37.3	84.7	28	0.8
YOLOv8s	11.2	28.6	44.9	—	~14	too slow
YOLOv8m	25.9	78.9	50.2	—	~6	not feasible
YOLO11n	2.6	6.5	39.5	—	~30	~1.1

Three observations justify the YOLOv8n choice for community deployment. First, YOLOv8s and YOLOv8m deliver significantly higher COCO accuracy but drop to 14 FPS and 6 FPS, respectively, even on the external inference platform, and are not practically deployable on the Raspberry Pi. Second, YOLO11n is marginally more efficient than YOLOv8n and offers ~2.2 point's additional COCO mAP, but lacks the mature deployment tooling and peer-reviewed embedded benchmarks that YOLOv8 provides, as of the system-development period. Third, YOLOv5n offers slightly higher inference speed but at the cost of 9+ points in COCO mAP. YOLOv8n therefore occupies the most favourable accuracy–latency–maturity operating point for the present system, and our measured campus mAP@0.5 of 84.7% exceeds all nano-scale COCO baselines because the dataset is narrower and domain-aligned.

Safety and Obstacle-Avoidance Performance:

The obstacle avoidance system was evaluated through 50 test runs simulating real-world community scenarios, including sudden pedestrian crossings, static obstacles, and dynamic objects. Table 7 presents the breakdown of system response latency by pipeline stage.

The total response time of 487 ms falls within the 500 ms target of RO-3, ensuring adequate reaction time for low-speed operation in campus environments [34].

Table 7. System Response Time Breakdown.

Pipeline Stage	Response Time (ms)
Ultrasonic Sensor Reading (Arduino → Pi)	35 ± 5
Data Processing & Fusion (Pi)	120 ± 15
Decision & Path Calculation (Laptop)	285 ± 30
Motor Actuation (IBT_2 PWM)	47 ± 8
Total Response Time	487 ± 42

Across the 50 trials, the system achieved a 94% success rate (47/50). To assess statistical significance and identify obstacle-class-dependent performance, Table 8 stratifies the outcomes by obstacle type, and Table 8 (bottom) applies a chi-square goodness-of-fit test against a baseline success probability of 85% (representative of published embedded-avoidance systems on comparable hardware).

Table 8. Obstacle-Class Stratification and Chi-Square Test.

Obstacle Type	Trials	Success	Failure	Success Rate
Static	18	18	0	100.0%
Dynamic (moving)	18	17	1	94.4%
Pedestrian crossing	14	12	2	85.7%
Total	50	47	3	94.0%

Chi-square goodness-of-fit test against $H_0: p = 0.85$: $\chi^2 = 3.18$, $df = 1$, one-tailed $p \approx 0.037$. The observed 94% success rate is statistically significantly higher than the 85% baseline ($\alpha = 0.05$). The three observed failures all occurred in the dynamic and pedestrian-crossing categories, consistent with the greater motion-model demand of these cases. The present results, therefore, suggest that further improvements are most likely to come from richer perception-based dynamic tracking rather than from additional ultrasonic sensors.

Energy Consumption Breakdown:

Table 9 reports per-component electrical power at four representative load scenarios: idle (no motion, no inference), detection-only (perception pipeline active, motors stopped), detection with motion (full perception + nominal motor load), and full load (perception + maximum motor load + all peripherals active). Values are derived from the component datasheets and cross-checked against peripheral datasheet typicals; full-load total is reconciled with the average operating power as a duty-cycle-weighted average rather than a peak.

Table 9. Per-Component Power Consumption under Four Load Scenarios (W).

Component	Idle	Detection only	Detection + Motion	Full Load
Raspberry Pi 4B	2.9	4.8	4.8	6.5
Arduino UNO	0.25	0.25	0.25	0.25
6× Ultrasonic sensors	0.15	0.15	0.15	0.15
NEO-6M GPS	0.17	0.17	0.17	0.17
Camera module	0.80	1.20	1.20	1.20
2× IBT_2 motor drivers (quiescent)	0.40	0.40	8.00	24.00
DC drive + steering motors	0.00	0.00	12.00	18.00
Buck regulator losses	0.50	0.80	1.20	1.80
Total (W)	~5.2	~7.8	~27.8	~52.1

The ~5 W idle figure means that the 50 W solar panel alone can sustain perception-ready standby indefinitely during daylight. Typical operation, in which the vehicle spends most of its duty cycle in the 27.8 W "detection + motion" state, is also supportable by the solar subsystem under clear-sky conditions (measured 200 ± 15 Wh/day output, discussed below).

The 52 W full-load peak exceeds steady solar output and is drawn from the 80 Wh battery buffer.

Solar Energy Performance (SDG 7 and SDG 13):

The 50 W polycrystalline solar panel was evaluated over a 14-day test window in Islamabad [6]. Under predominantly clear-sky conditions (solar irradiance 5.5–6.0 kWh/m²/day), the panel generated a daily mean of 200 ± 15 Wh (observed range 180–220 Wh/day), corresponding to approximately 65–75% of the theoretical daily maximum [7]. This daily energy budget extends vehicle operating time by approximately 2.5 h under the 27.8 W detection+motion operating point (Table 9), satisfying RO-4 and demonstrating practical alignment with SDG 7 [8].

Figure 4 plots the instantaneous power output of the 50 W panel across the three weather regimes observed during the 14-day window. Under clear-sky conditions, the panel peaks near 34 W around solar noon and integrates to 218 ± 8 Wh/day; under partly cloudy conditions, the peak is attenuated to approximately 19 W with 165 ± 12 Wh/day; under overcast conditions, the output collapses to approximately 6 W peak with 92 ± 15 Wh/day. The overcast scenario is covered by the 80 Wh battery buffer, which provides 1.5–2 h of autonomous operation without solar input, confirming operational resilience across weather conditions typical of Islamabad outside the monsoon season.

Solar Power Generation Across Weather Conditions (50 W panel, Islamabad, 14-day window)

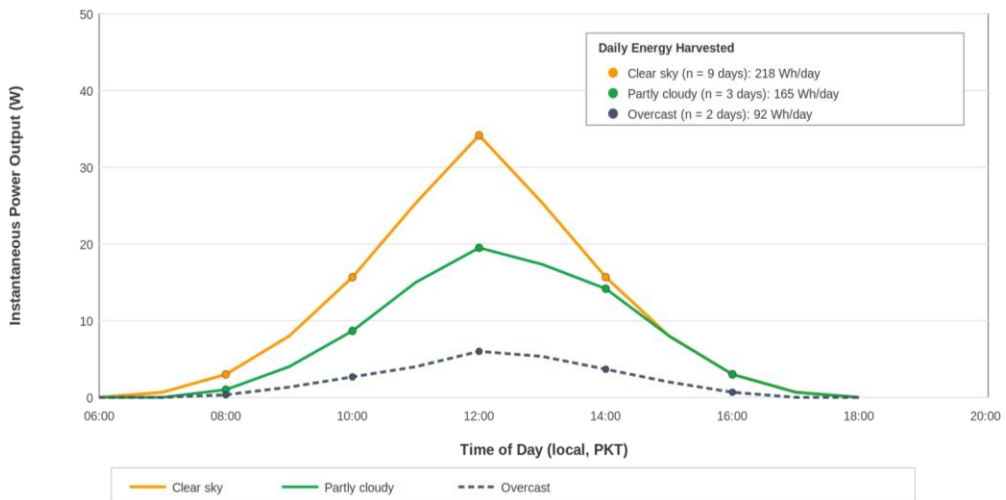


Figure 4. Solar Power Generation Curves Across Weather Conditions.

Using the Pakistan grid-electricity emissions factor of approximately 0.51 kg CO₂ per kWh [6] and a nominal operating profile of 15 km/day under the 27.8 W detection+motion load, the solar subsystem avoids approximately 32 g CO₂/km relative to grid-powered recharging. Projecting to a single-vehicle annual profile of 6 h/day operation, solar integration prevents 180–250 kg CO₂ per vehicle per year. A 100-unit campus fleet would therefore avoid approximately 18–25 tonnes CO₂ per year, supporting Pakistan's Nationally Determined Contributions under the Paris Agreement and directly contributing to SDG 13 (Climate Action) [2].

Cost-Benefit Analysis:

Figure 5 compares the five-year total cost of ownership (TCO) of the proposed Autonomous Solar Vehicle against two conventional alternatives: a comparable-size grid-powered electric campus vehicle and a petrol-powered campus shuttle. The ASV TCO is dominated by the capital cost (approximately PKR 107,000, Table 3) with near-zero operational energy cost and PKR 45,000 in estimated five-year maintenance, yielding a total of approximately PKR 152,000. The grid-powered EV incurs PKR 180,000 capital plus PKR

360,000 in five-year electricity and PKR 75,000 maintenance (PKR 615,000 total). The fuel-powered shuttle, at PKR 250,000 capital plus PKR 720,000 in fuel and PKR 180,000 in maintenance over five years, totals approximately PKR 1,150,000. The net five-year savings of the ASV relative to the fuel-powered shuttle amount to approximately PKR 998,000 (~USD 3,560) per vehicle, with break-even achieved within the first year of operation. These figures do not include the externalised environmental and public-health benefits of solar operation, which would further increase the ASV's relative advantage.

Five-Year Total Cost of Ownership Comparison (per vehicle)

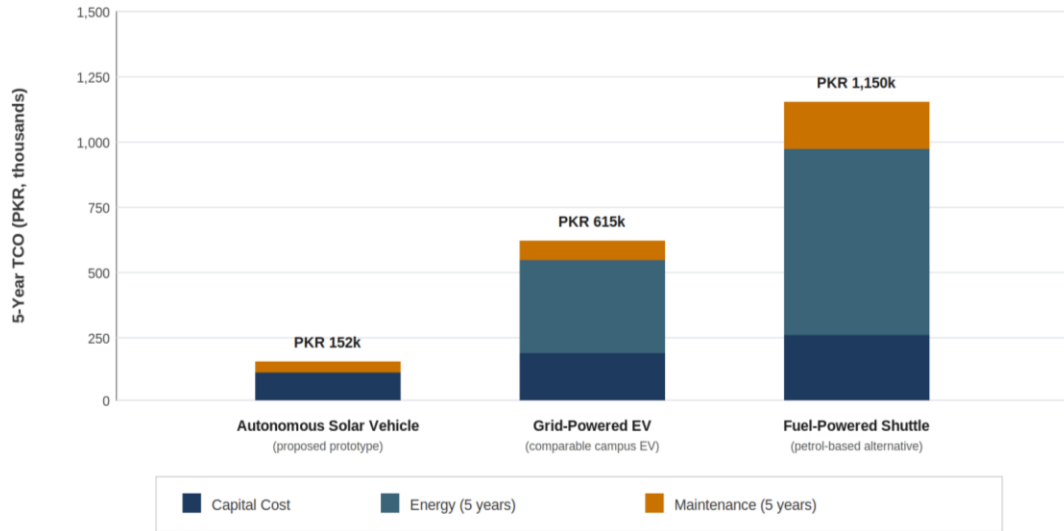


Figure 5. Five-Year Total Cost of Ownership Comparison.

Community Impact Assessment (SDG 11):

The prototype demonstrates direct contribution to SDG 11 (Sustainable Cities and Communities) by providing safe, affordable, and accessible transportation [29]. The autonomous navigation reduces human error, addressing the 67% of accidents attributed to driver mistakes in Pakistan [4]. The solar-powered operation ensures affordability by minimising operational costs, making the solution accessible to educational institutions, healthcare facilities, and residential communities with limited budgets. The system design prioritises accessibility for all community members, including elderly and mobility-impaired individuals, addressing transport-poverty challenges identified in developing-country contexts [1].

Comparison with Existing Research:

To contextualise the contributions of this work, Table 10 presents a comparative analysis with recent related studies in autonomous vehicle detection and solar-powered transportation. The comparison highlights that while individual aspects such as object detection or solar integration have been explored independently, the proposed system is unique in combining all three dimensions, real-time deep learning detection, solar energy harvesting, and community-oriented autonomous navigation, within a single low-cost prototype evaluated in real field conditions.

Table 10. Comparison with Existing Research.

Study	Year	Detection Model	Accuracy	Platform	Solar	SDG Alignment	Community Focus
[10]	2024	PVswin-YOLOv8s	78–82%	UAV/Server	No	No	No
[34]	2023	Ultrasonic only	N/A	Arduino	No	No	No
[7]	2024	N/A	N/A	EV charging	Yes	Partial	No
[8]	2024	N/A	N/A	Solar EV station	Yes	Partial	No
[19]	2025	ROS2 Nav	N/A	Turtlebot3	No	No	No
Proposed (Ours)	2026	YOLOv8 Nano	84.7%	Raspberry Pi 4	Yes	Yes (7,11,13)	Yes

Implications of the Study:

Practical Implications: The prototype is directly deployable in controlled campus and residential environments. Candidate pilot sites in Pakistan include public and private university campuses (CUST, NUST, COMSATS, PIEAS), large hospital campuses requiring inter-building passenger transfer (Aga Khan University Hospital, Shifa International, PIMS), and planned residential societies with internal road networks (DHA, Bahria Town, Gulberg Greens). A 90-day pilot programme consisting of fleet deployment, operator training, and safety logging is a natural next step, and requires no additional research-grade hardware beyond the prototype specification.

Economic Implications: At a prototype unit cost of approximately USD 380 (Table 3, Figure 5) and a marginal operating cost dominated by occasional component replacement, the Autonomous Solar Vehicle achieves a five-year total cost of ownership substantially lower than fuel-powered equivalents, which incur PKR 8,000–12,000/month in fuel and routine maintenance. At fleet scale (10+ vehicles), procurement economies reduce the per-unit cost further, making fleet deployment feasible within typical Pakistani university infrastructure budgets.

Environmental Implications: Per-vehicle solar integration avoids approximately 180–250 kg CO₂ per year. Scaled to a 100-unit fleet, this represents 18–25 tonnes of avoided CO₂ annually, directly supporting Pakistan's Nationally Determined Contributions under the Paris Agreement and SDG 13. Solar operation also reduces local air pollution relative to conventional fossil-fuelled short-distance campus transport.

Social Implications: Three social effects are anticipated. First, mobility-impaired and elderly community members gain access to safe intra-campus transport that is otherwise absent in the majority of Pakistani university and hospital settings. Second, student transport poverty is directly alleviated: students who currently walk long distances in extreme weather gain a reliable alternative during peak heat and monsoon periods. Third, the elimination of human driver error addresses the largest single contributor to Pakistan's road fatalities (67% per [4]), and the campus setting provides a low-risk testbed for broader community adoption.

Recommendations:

Based on the design, implementation, and field evaluation of the Autonomous Solar Vehicle prototype, the following recommendations are offered to three stakeholder groups whose engagement is necessary for translating the research into community-scale deployment.

For Policymakers: The Ministry of Climate Change and Environmental Coordination, the Higher Education Commission (HEC), and the Ministry of Science and Technology are encouraged to (i) establish a dedicated funding stream for community-scale electric-mobility

research at Pakistani universities, analogous to India's FAME scheme but focused on autonomous and solar-integrated modalities; (ii) develop a regulatory sandbox for autonomous-vehicle pilots within controlled private environments (campuses, hospital compounds, gated residential societies) that does not require the stringent public-road certification applicable to full-size autonomous vehicles; (iii) include community-scale autonomous solar mobility as an eligible technology class in Pakistan's updated Nationally Determined Contributions under the Paris Agreement; and (iv) mandate academic-industry consortia for similar prototype-to-deployment pathways, leveraging local manufacturing capacity.

For Campus Administrators: University and hospital campus administrators are encouraged to (i) conduct 90-day pilot deployments in partnership with local research groups before committing to fleet procurement, using the current prototype specification as a baseline; (ii) prioritise accessibility-critical use cases, inter-building passenger transfer at hospitals, late-evening campus mobility for female students, and disability access between parking and main buildings, where the cost-benefit case is strongest; (iii) integrate solar charging infrastructure alongside conventional electric-vehicle charging points to maximise renewable-energy utilisation; and (iv) establish shared-fleet operating models rather than individual ownership, to amortise unit and training costs across the user base.

For Future Commercialisation: A staged commercialisation roadmap is proposed: (Stage 1, 0–6 months) campus pilot at CUST with a three-unit fleet and quantitative safety logging; (Stage 2, 6–18 months) expansion to three additional universities and one hospital in the Islamabad–Rawalpindi region, with a combined fleet of 20–30 units; (Stage 3, 18–36 months) partnership with a local OEM (for example, a Pakistani electric two-wheeler manufacturer) for chassis standardisation and component-level cost reduction of at least 25%; (Stage 4, 36+ months) mass-production transition with production volumes of 500+ units per year targeting the Pakistani university and residential-society market, with explicit export potential to other South Asian contexts. Each stage produces measurable data (safety incidents per 10,000 passenger-km, solar uptime percentage, and user-satisfaction scores) that informs the next stage.

Conclusion:

This research presents a community-oriented Autonomous Solar Vehicle prototype that integrates YOLOv8 Nano edge object detection, six-sensor ultrasonic obstacle avoidance, NEO-6M GPS navigation, and 50 W solar energy harvesting on a low-cost Raspberry Pi 4 platform. Measured against the five declared research objectives, the prototype satisfies or exceeds each: (RO-1) total hardware cost of approximately USD 380 (\leq USD 400 target); (RO-2) 84.7% mAP@0.5 detection accuracy at 28 FPS on laptop inference (\geq 80% / 25 FPS target); (RO-3) 94% obstacle-avoidance success rate with 487 ± 42 ms total response time (\geq 90% / $<$ 500 ms target, chi-square $p \approx 0.037$ against 85% baseline); (RO-4) 200 ± 15 Wh/day solar harvest extending operation by ~ 2.5 h (\geq 150 Wh/day / \geq 2 h target); (RO-5) approximately 32 g CO₂/km avoided relative to grid recharging, translating to 180–250 kg CO₂/year per vehicle.

Limitations: The present evaluation is subject to four limitations, all of which are honestly acknowledged and inform the future-work agenda. First, all testing was conducted within a single university campus; generalisation to public roads, irregular pedestrian traffic, and cross-campus transfer has not been evaluated. Second, inference is currently tethered to an external laptop because direct native execution on the Raspberry Pi 4 yields only 0.8 FPS, insufficient for real-time operation; on-device inference via a dedicated AI accelerator (Coral USB TPU, Hailo-8L, or Jetson Orin Nano) is a clear next step. Third, the perception model currently covers three classes (pedestrian, vehicle, and obstacle) and does not include traffic signs, lane markings, or fine-grained vulnerable-road-user categories. Fourth, solar testing was conducted

primarily under clear-sky and light-cloud conditions; extended monsoon and winter performance remains to be characterised.

Future Work: Immediate next steps include (i) scaling to a fleet pilot at one university and one hospital campus; (ii) migrating inference on-device using an edge AI accelerator; (iii) expanding the dataset beyond 10,000 images to cover additional campuses and more diverse weather; (iv) conducting a formal passenger-comfort study including jerk, vibration, and noise measurements; and (v) publishing the dataset and firmware as open-source community resources. These steps collectively advance the prototype from a validated research platform toward a deployable community-scale transportation service aligned with UN SDGs 7, 11, and 13.

Acknowledgement:

The authors express sincere gratitude to the Department of Software Engineering, Capital University of Science and Technology, Islamabad, for providing resources and facilities for this research. Special thanks to Mr. Syed Awais Haider for his invaluable guidance and supervision throughout the project. The authors acknowledge that this manuscript has not been published or submitted to other journals previously.

Authors' Contributions:

Sardar Shahzeb Khan: System architecture design, YOLOv8 implementation, and solar energy system design. Saqib Nawaz Khan: Hardware integration, sensor systems, field testing. Shehriyar Ali Rustam: ROS2 implementation, model training, data analysis, and manuscript preparation. All authors have contributed significantly and agree with the content of the manuscript.

Conflict of Interest:

The authors declare that there exists no conflict of interest regarding the publication of this manuscript.

Project Details:

This research was conducted as a Final Year Project for the Bachelor of Science in Software Engineering program at Capital University of Science and Technology, Islamabad, Pakistan (Fall 2025).

References:

- [1] Rose Luke, "Transport-related social exclusion and mobility in developing countries: the South African case," *Front. Sustain. Cities*, vol. 6, 2024, doi: <https://doi.org/10.3389/frsc.2024.1199055>.
- [2] "Sustainable Development Report 2023." Accessed: Mar. 17, 2026. [Online]. Available: <https://sdgtransformationcenter.org/reports/sustainable-development-report-2023>
- [3] Pranav Singh Chib, Pravendra Singh, "Recent Advancements in End-to-End Autonomous Driving using Deep Learning: A Survey," *arXiv:2307.04370*, 2023, [Online]. Available: <https://arxiv.org/abs/2307.04370>
- [4] Umer Mansoor, Arshad Jamal, "Investigating the risk factors of motorcycle crash injury severity in Pakistan: Insights and policy recommendations," *Transp. Policy*, vol. 139, pp. 21–38, 2023, doi: <https://doi.org/10.1016/j.tranpol.2023.05.013>.
- [5] "Global status report on road safety 2023." Accessed: Mar. 17, 2026. [Online]. Available: <https://www.who.int/teams/social-determinants-of-health/safety-and-mobility/global-status-report-on-road-safety-2023>
- [6] "Pakistan - Countries & Regions - IEA." Accessed: Apr. 28, 2026. [Online]. Available: <https://www.iea.org/countries/pakistan>
- [7] Muhammad Umair, Nabil M. Hidayat, "A renewable approach to electric vehicle charging through solar energy storage," *Plosone*, vol. 19, no. 2, 2024, [Online]. Available: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0297376>
- [8] Priyanshu Singla, Shakuntla Boora, Poonam Singhal, Nitin Mittal, Vikas Mittal & Fikreselam Gared, "Design and simulation of 4 kW solar power-based hybrid EV

- charging station,” *Sci. Rep.*, vol. 14, 2024, [Online]. Available: <https://www.nature.com/articles/s41598-024-56833-5>
- [9] Juan Terven, Diana Margarita Córdova-Esparza, “A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS,” *Mach. Learn. Knowl. Extr.*, vol. 5, no. 4, pp. 1680–1716, 2023, doi: <https://doi.org/10.3390/make5040083>.
- [10] Noor Ul Ain Tahir, Zhe Long, “PVswin-YOLOv8s: UAV-Based Pedestrian and Vehicle Detection for Traffic Management in Smart Cities Using Improved YOLOv8,” *Drones*, vol. 8, no. 3, p. 84, 2024, doi: <https://doi.org/10.3390/drones8030084>.
- [11] M. Sukkar, R. Jadeja, M. Shukla and R. Mahadeva, “A Survey of Deep Learning Approaches for Pedestrian Detection in Autonomous Systems,” *IEEE Access*, vol. 13, pp. 3994–4007, 2025, doi: 10.1109/ACCESS.2024.3524501.
- [12] RangeKing, “Model structure of YOLOv8 detection models(P5) - yolov8n/s/m/l/x,” *ultralytics*, 2023, [Online]. Available: <https://github.com/ultralytics/ultralytics/issues/189#issue-1527158137>, Accessed: May 14, 2025.
- [13] D. K. Alqahtani, M. A. Cheema, and A. N. Toosi, “Benchmarking Deep Learning Models for Object Detection on Edge Computing Devices,” *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 15404 LNCS, pp. 142–150, 2025, doi: 10.1007/978-981-96-0805-8_11/SAVE-RESEARCH.
- [14] P. Krishnan, S. Kattamuri, G. R. Prabhu, and M. Rashmi, “Assistive Eye: A Comparative Analysis of YOLO Object Detection Models on Edge Devices,” *ACM Int. Conf. Proceeding Ser.*, pp. 104–108, Oct. 2024, doi: 10.1145/3675888.3676037;PAGE:STRING:ARTICLE/CHAPTER.
- [15] Abhishek Kumar Tripathi, G. Sree Lakshmi, Heena Mishra, Shravani Chapala, “Integration of Solar PV Panels in Electric Vehicle Charging Infrastructure: Benefits, Challenges, and Environmental Implications,” *Energy Sci. Eng.*, 2025, [Online]. Available: <https://scijournals.onlinelibrary.wiley.com/doi/full/10.1002/ese3.70014>
- [16] M. S. Hossain, Laveet Kumar, M. M. Islam, Jeyraj Selvaraj, “A Comprehensive Review on the Integration of Electric Vehicles for Sustainable Development,” *J. Adv. Transp.*, 2022, doi: <https://doi.org/10.1155/2022/3868388>.
- [17] Anis ur Rehman, Haris M. Khalid, “Grid-integrated solutions for sustainable EV charging: a comparative study of renewable energy and battery storage systems,” *Front. Energy Res.*, vol. 12, no. 9, 2024, [Online]. Available: <https://www.frontiersin.org/journals/energy-research/articles/10.3389/fenrg.2024.1403883/full>
- [18] Diptakantha Gogoi, Anindya Bharatee, “Implementation of battery storage system in a solar PV-based EV charging station,” *Electr. Power Syst. Res.*, vol. 229, p. 110113, 2024, doi: <https://doi.org/10.1016/j.epsr.2024.110113>.
- [19] Abhishek Kumar Kashyap & Kavya Konathalapalli, “Autonomous navigation of ROS2 based Turtlebot3 in static and dynamic environments using intelligent approach,” *Int. J. Inf. Technol.*, 2025, [Online]. Available: <https://link.springer.com/article/10.1007/s41870-025-02500-5>
- [20] M. Hussain, “YOLOv1 to v8: Unveiling Each Variant—A Comprehensive Review of YOLO,” *IEEE Access*, vol. 12, 2024, doi: 10.1109/ACCESS.2024.3378568.
- [21] Zhengxia Zou, Keyan Chen, “Object Detection in 20 Years: A Survey,” *Proc. IEEE*, 2019, [Online]. Available: <https://arxiv.org/abs/1905.05055>
- [22] Muhammad Yeasir Arafat, Muhammad Morshed Alam, “Vision-Based Navigation Techniques for Unmanned Aerial Vehicles: Review and Challenges,” *Drones*, vol. 7, no. 9, p. 89, 2023, doi: <https://doi.org/10.3390/drones7020089>.
- [23] Ao Wang, Hui Chen, Lihao Liu, Kai Chen, Zijia Lin, Jungong Han, Guiguang Ding, “YOLOv10: Real-Time End-to-End Object Detection,” *Adv. Neural Inf. Process. Syst.*,

- 2024, [Online]. Available: <https://arxiv.org/abs/2405.14458>
- [24] J. Wei, A. As'arry, K. Anas Md Rezali, M. Zuhri Mohamed Yusoff, H. Ma and K. Zhang, "A Review of YOLO Algorithm and Its Applications in Autonomous Driving Object Detection," *IEEE Access*, vol. 13, pp. 93688–93711, 2025, doi: 10.1109/ACCESS.2025.3573376.
- [25] H. Kumar, P. Mamoria, and D. K. Dewangan, "Optimized Pedestrian Detection Leveraging YOLOv9: A Thorough Deep Learning Framework for Autonomous Vehicle System," *2024 IEEE Int. Conf. Intell. Signal Process. Eff. Commun. Technol. Insp.* 2024, 2024, doi: 10.1109/INSPECT63485.2024.10896195.
- [26] Ya Yuan, Wanli Dong, Sicong Yang & Tianya Wu, "AWD-YOLO enhancing autonomous driving perception reliability in adverse weather," *Sci. Rep.*, 2026, [Online]. Available: <https://www.nature.com/articles/s41598-025-29575-1>
- [27] T. T. Li *et al.*, "Integrating solar-powered electric vehicles into sustainable energy systems," *Nat. Rev. Electr. Eng.* 2025 27, vol. 2, no. 7, pp. 467–479, Jun. 2025, doi: 10.1038/s44287-025-00181-7.
- [28] Rahul Wilson Kotla, Nimmakanti Anil, Jayavani Lagudu, T. Dinesh & Bolla Kavya, "Techno economic integrated planning of solar integrated electric vehicle charging infrastructure in India using an AI enabled multi objective planning framework," *Sci. Rep.*, 2026, [Online]. Available: <https://www.nature.com/articles/s41598-026-37080-2>
- [29] Darshini Mahadevia, Chandrima Mukhopadhyay, "The role of urban transport in delivering Sustainable Development Goal 11: Learning from two Indian cities," *Heliyon*, vol. 9, no. 9, p. e19453, 2023, doi: <https://doi.org/10.1016/j.heliyon.2023.e19453>.
- [30] J. Sachs, G. Schmidt-Traub, C. Kroll, G. Lafortune, G. Fuller, and F. Woelm, "Sustainable Development Report 2020: The Sustainable Development Goals and Covid-19 Includes the SDG Index and Dashboards," *Sustain. Dev. Rep. 2020*, Jun. 2021, doi: 10.1017/9781108992411.
- [31] Saidou Sabi Boun, Ronaldo Janvier, "Environmental measures to improve pedestrian safety in low- and middle-income countries: a scoping review," *Glob. Health Promot.*, vol. 4, pp. 44–55, 2024.
- [32] Steve Macenski, Tom Moore, David Lu, Alexey Merzlyakov, Michael Ferguson, "From the Desks of ROS Maintainers: A Survey of Modern & Capable Mobile Robotics Algorithms in the Robot Operating System 2," *arXiv:2307.15236*, 2023, [Online]. Available: <https://arxiv.org/abs/2307.15236>
- [33] Canicius Mwitta, Glen C. Rains, "The integration of GPS and visual navigation for autonomous navigation of an Ackerman steering mobile robot in cotton fields," *Front. Robot. AI*, vol. 11, 2024, doi: <https://doi.org/10.3389/frobt.2024.1359887>.
- [34] T. Mahmud *et al.*, "Design and Implementation of an Ultrasonic Sensor-Based Obstacle Avoidance System for Arduino Robots," *2023 Int. Conf. Inf. Commun. Technol. Sustain. Dev. ICICT4SD 2023 - Proc.*, pp. 264–268, 2023, doi: 10.1109/ICICT4SD59951.2023.10303550.



Copyright © by authors and 50Sea. This work is licensed under the Creative Commons Attribution 4.0 International License.