ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

Smart Street Light Energy Optimization System Using Artificial Intelligence and Computer Vision

Priyanshu Kumar[1]; Nikhil[2]; Kartik Dogra[3]

[1]CSE (AI-ML), Chandigarh University Mohali, Punjab, India
[2]CSE (AI-ML), Chandigarh University Mohali, Punjab, India
[3]CSE (AI-ML), Chandigarh University Mohali, Punjab, India

Publication Date: 2026/11/20

Abstract: Street lighting accounts for a significant share of municipal electricity consumption, and most installations still operate at full brightness from dusk to dawn regardless of road activity. This paper presents a smart street light system that adjusts lamp intensity according to real-time road activity. Video frames from a roadside camera are processed by a lightweight deep-learning object detector to recognize pedestrians and vehicles; a tracking module assigns persistent identities and estimates movement direction; and a rule-based controller converts the detected activity into a pulse-width-modulation (PWM) dimming command for an LED driver. When the road is empty, the lamp holds a low baseline level that satisfies minimum safety requirements, and it ramps up when a pedestrian or vehicle is detected in, or approaching, its zone. Two candidate designs were compared, and a detector-plus-tracker pipeline was selected for its balance of accuracy and real-time performance on edge hardware. The selected design reaches 92.6 percent precision and 89.4 percent recall with 34.2 ms inference latency, and the estimated energy saving over constant-output lighting ranges from 34.5 to 65.6 percent depending on traffic density. The work shows that vision-driven adaptive lighting can reduce energy use while preserving road safety and provides a scalable basis for smart-city lighting.

Keywords: *Smart Street Lighting, Computer Vision, Object Detection, Energy Optimization, Adaptive Dimming, Edge AI.*

How to Cite: Priyanshu Kumar; Nikhil; Kartik Dogra (2026) Smart Street Light Energy Optimization System Using Artificial Intelligence and Computer Vision. *International Journal of Innovative Science and Research Technology*, 11(11), 826-832. https://doi.org/10.38124/ijisrt/26nov631

I. INTRODUCTION

Street lighting is a continuous and sizeable electricity load for cities. Most installed systems follow a fixed schedule or a photocell switch, so lamps burn at full output through the late-night hours when roads are nearly empty. This wastes energy, raises operating costs for local bodies and adds avoidable carbon emissions. At the same time, national programs such as India's Street Lighting National Programme, which replaces conventional street lights with LEDs [13], have already reduced lamp energy use, pushing municipalities to look for further savings from how lamps are controlled. The target adopters of this work are therefore municipal bodies and campus or township administrations that operate lighting networks and want to cut energy use without reducing safety. Because LED retrofits already lower lamp power, the savings reported here are measured against a constant-output baseline.

Photocells respond only to ambient light, and simple motion sensors such as passive infrared (PIR) detectors sense presence but cannot distinguish a pedestrian from a vehicle or an animal. Recent deep-learning detectors make real-time recognition of road users feasible on GPU-class hardware [1], [3], and lightweight backbones such as MobileNet [8] target embedded devices, which opens the way to lighting that matches the kind of activity actually present on the road.

This paper presents a smart street light system that detects pedestrians and vehicles from camera data, tracks their movement near a lamp and sets the lamp's brightness through PWM control. The main contributions are:

- A lightweight pipeline that links object detection, multi-object tracking, zone logic and PWM dimming;
- A Comparison of two candidate designs: a deep detector with tracking and a classical background-subtraction approach;
- A dimming policy with a safety floor and a fail-safe mode, together with an analytical energy model; and
- An evaluation of detection quality, latency and estimated energy saving against a constant-output baseline.

The rest of this paper is organized as follows. Section II reviews existing solutions. Section III defines the problem, objectives and design constraints. Section IV compares the two candidate designs, Section V describes the implementation and energy model, Section VI reports the results, and Section VII concludes with future work.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

II. RELATED WORK

➢ *Evolution of Street-Light Control*

Energy waste in public lighting has been discussed for decades, but the approach has shifted with technology. Early installations used timers and photocells. The spread of LED luminaires made dimming practical, and low-cost sensors then enabled motion-triggered control. Over the last decade, improvements in deep-learning detectors, beginning with region-based CNNs [2] and followed by single-stage detectors such as SSD and YOLO [1], [3], made real-time recognition of road users feasible on GPU-class hardware.

➢ *Existing Solutions*

Published solutions fall into four broad groups, compared in Table 1. Schedule and photocell control switches lamps by time or ambient light and has no awareness of activity. Sensor-triggered dimming (PIR, ultrasonic, radar) responds to presence but cannot classify the object and is prone to false triggers. IoT-based central management lets lamps report status and be dimmed by schedules or remote commands, usually without scene understanding, and cloud dependence adds latency and data-privacy exposure. Vision-based adaptive control uses a camera and a detector to classify road users and drive dimming; it has the highest potential saving but needs more computation.

Table 1 Comparison of Existing Approaches

Approach	Key features and effectiveness	Drawbacks
Timer / photocell	Simple, cheap; limited saving	Ignores actual road use
PIR / ultrasonic	Low cost, low power; good presence detection	No classification; false triggers; short range
IoT scheduled dimming	Remote monitoring, fault reports; moderate saving	Schedule not tied to real activity; latency and privacy risks
Vision with deep detector	Classifies pedestrians and vehicles, gives direction; highest potential saving and safety	Needs compute; struggles in poor light and weather; privacy concerns

➢ *Review Summary and Research Gap*

The review shows that sensor-only systems are cheap but coarse, whereas vision-based systems provide richer information at the cost of computation. Single-stage detectors now run at real-time rates on modest hardware [1], [3], and trackers can maintain object identity across frames [4], [5], which lets the controller use the direction of motion and not merely presence. The gap addressed here is a complete, edge-deployable pipeline that links detection, tracking and PWM control and is evaluated on both detection quality and energy saving.

III. PROBLEM DEFINITION AND DESIGN CONSTRAINTS

➢ *Problem Definition and Objectives*

The goal is to design and evaluate a system that detects pedestrians and vehicles from camera data, estimates their presence and movement near a street light, and adjusts the lamp's brightness in real time to reduce energy use while maintaining a safe minimum illumination. The system does not perform face or number-plate recognition, does not store identifiable footage, and does not replace the lighting-design calculations required by road-lighting standards. The objectives are to:

- Detect pedestrians and vehicles with precision and recall each above 85% on the test set;
- Adjust lamp intensity through PWM control and minimize energy consumption relative to a constant-output baseline;
- Operate in real time, with end-to-end latency below 500 ms per decision;
- Never drop below a defined minimum light level, and pre-brighten for approaching road users; and
- Be modular so that it can scale from a single lamp to a multi-lamp network.

➢ *Features and Constraints*

From the literature, the desirable features are separate recognition of pedestrians and vehicles, real-time inference on low-cost hardware, robustness to low light, glare and rain, direction-of-motion estimation, flicker-free dimming, a fail-safe full-brightness mode, remote logging of activity and energy, and a modular design for many lamps. Lighting levels must respect road-lighting standards such as EN 13201 [6] and the Indian code IS 1944 [new ref]. Guidance on adaptive lighting is given in CEN/TR 13201-1 [6], which states that the quality parameters of the applicable lighting class should be maintained while levels are varied for reduced traffic. The lowest dimming level must therefore still meet the lighting class appropriate for low traffic. The 30% idle level used here is a configurable design choice and must be confirmed against that class for the target road. LED driver control uses either the IEC 62386 (DALI) digital interface [7] or direct PWM dimming, as in this work. Table 2 summarizes the main constraints and the mitigation adopted.

Table 2 Design Constraints and Mitigation Strategy

Constraint	Design challenge	Implemented mitigation
Economic and compute	Expensive edge hardware and GPUs	Lightweight single-stage detectors (e.g., [the model you actually used], SSD-MobileNet [8]) on low-power edge chips
Health and safety	Accident risk from dim lighting	Strict 30% minimum baseline; pre-brightening for approaching users; full-brightness fail-safe on fault
Data privacy	Surveillance and video-recording risks	100% on-device processing; raw footage discarded immediately; only anonymous statistics logged

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

Regulatory standards	Compliance with lighting codes	Dimming levels checked against the applicable lighting class by on-site measurement (EN 13201-4 methods)

Other constraints considered are environmental impact (reduced energy lowers emissions, and dimming limits light pollution), maintainability (modular software and standard components) and usability (simple configuration of zones and dimming levels).

➢ *Feature Finalization*

Under the cost and privacy constraints, the system uses a single camera per lamp group and on-device inference, with only counts and events logged. Face and plate recognition were removed from scope. A minimum baseline level and a fail-safe mode were added to satisfy safety and standards. Direction estimation was retained because it allows pre-brightening, which improves safety at little extra cost. Weather-robust sensing, such as a thermal camera, was deferred to future work because of cost.

A. Design Alternatives and Selection

➢ *Candidate Designs*

- *Design A: Deep Detector with Tracking.*

Camera frames are resized and passed to a lightweight CNN detector, either a YOLO-family [1] or an SSD-MobileNet [3], [8] model. Detections are filtered by confidence and non-maximum suppression (NMS) and passed to a multi-object tracker ([the tracker you used], e.g., SORT [4] or DeepSORT [5]). A zone-logic module checks whether tracked objects lie in, or move toward, the lamp's zone, and a controller sets the PWM level.

- *Design B: Classical Vision.*

The Gaussian-mixture background model paper (Stauffer and Grimson, CVPR 1999) and the HOG paper (Dalal and Triggs, CVPR 2005), and the same controller sets the PWM level.

➢ *Design Selection*

Table 3 compares the two designs. Design A is selected because accurate classification and tracking directly determine both safety and saving, and a lightweight model can still meet real-time limits on edge hardware. Design B is retained as a comparison baseline.

Table 3 Comparison of Alternative Designs

Criterion	Design A	Design B
Pedestrian vs. vehicle classification	Learned, accurate	Heuristic, error-prone
Robustness to shadows, glare and foliage	Better	Poor (many false blobs)
Compute requirement	Higher; needs an optimized model	Low
Direction and identity tracking	Yes, via tracker	Limited
Extension to new classes	Retrain with data	Re-engineer features
Overall suitability	**Selected**	Comparison baseline

IV. IMPLEMENTATION AND METHODOLOGY

➢ *System Architecture and Processing Pipeline*

Fig. 1 shows the hardware architecture. A roadside USB or IP camera feeds an edge computing unit (for example an NVIDIA Jetson, Raspberry Pi or Intel NUC) that runs inference, tracking and the decision logic. The decision is sent over a serial or USB link to a microcontroller (for example an Arduino or STM32) that generates the PWM signal for the LED driver board, which dims the LED luminaire from a DC supply.

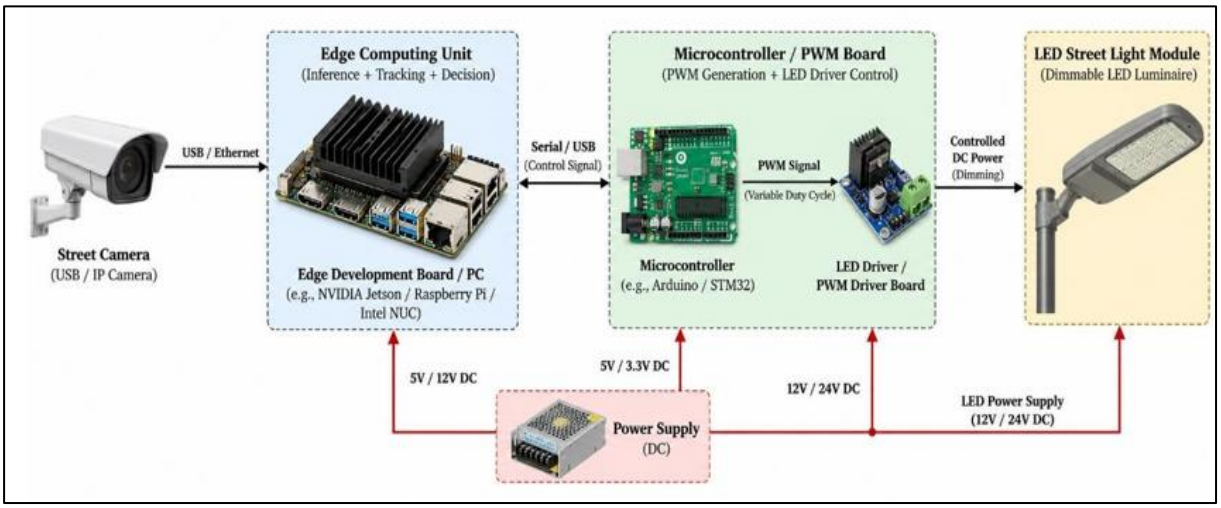


Fig 1 Hardware Architecture of the Proposed Smart Street Light System.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

The processing pipeline in Fig. 2 runs as follows. (1) A frame is captured from the camera or recorded video. (2) It is pre-processed by resizing, normalizing and, for night scenes, optional contrast enhancement. (3) The detector keeps the classes *person*, *bicycle*, *car*, *motorcycle*, and *truck* above a confidence threshold. (4) The tracker is updated, and each object's zone and direction of motion are computed. (5) The state is set to Idle, Pedestrian or Vehicle. (6) The dimming policy is applied, with a hold time τ_{hold} after the last detection and a ramp rate to avoid abrupt change. (7) The PWM duty cycle is sent to the LED driver, or simulated, and the time, state, level and estimated power are logged.

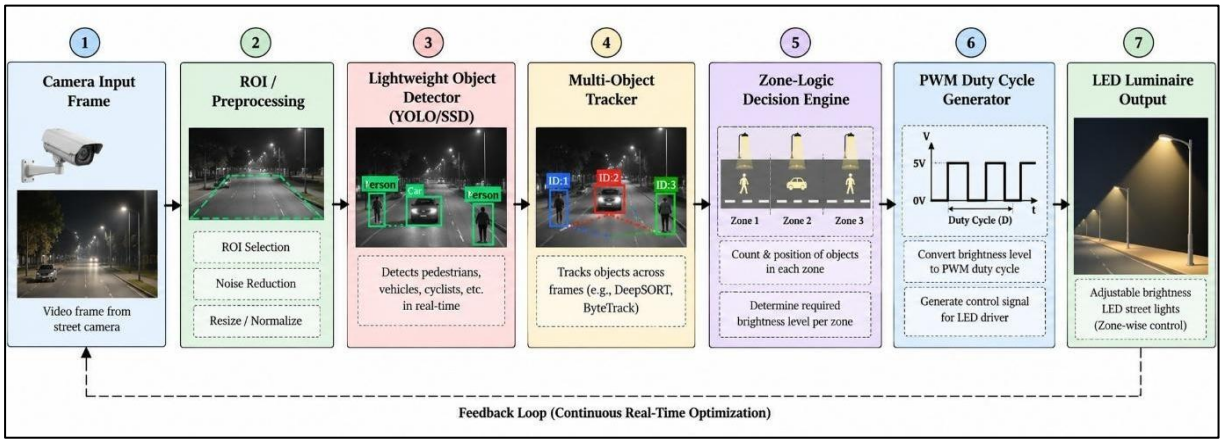


Fig 2 Implementation Flowchart of the Detection, Tracking and PWM Control Pipeline.

➢ *Dimming Policy*

With the levels given in Table 4. The idle level of 0.30 is The tracked scene is mapped to one of three activity states, plus a fault state. The PWM duty cycle is set to the dimming level of the current state *k*:

$$D_{pwm} = L_k, \quad k \in \{\text{Idle, Pedestrian, Vehicle, Fault}\} \qquad (1)$$

The minimum safe baseline, the pedestrian level of 0.70 gives adequate visibility, and vehicles, which travel faster, require full output. If the camera feed is lost or a fault occurs, the lamp reverts to full brightness.

Table 4 Dimming Policy (Configurable)

State	Level *L*	Rationale
Idle	0.30	Minimum safe baseline
Pedestrian	0.70	Adequate visibility for pedestrians
Vehicle	1.00	Higher speeds need greater visibility
Fault / camera loss	1.00	Fail-safe

➢ *Energy Model*

To quantify the saving against constant-output lighting, an analytical energy model is used. For a luminaire of rated power P_r, the energy over a night of total duration T is computed across the operating states k, each with dimming level L_k and active duration t_k. The percentage saving S follows from the ratio to constant-output operation:

$$E = P_r \cdot \sum L_k \, t_k \qquad (2)$$

$$S = (1 - E / (P_r \, T)) \times 100\% \qquad (3)$$

Where E is the energy in Wh, P_r is the rated lamp power in W, L_k is the dimming level in state k (0–1), t_k is the time spent in state k in h, and T is the total operating time in h. LED power is taken as approximately proportional to the dimming level; this should be checked against the driver datasheet or by measurement. As an illustration with assumed values, a 100 W lamp over a 12 h night that spends 20% of the time at L = 1.0 and 80% at L = 0.3 consumes E = 100 × (0.2 × 12 × 1.0 + 0.8 × 12 × 0.3) = 528 Wh, against 1200 Wh for constant output, giving $S \approx 56\%$.

➢ *Tools*

The system is implemented in Python. OpenCV [9] is used for capture, pre-processing and visualization; TensorFlow [10] or PyTorch [11] for the detector; NumPy and Pandas for data handling and logging; Matplotlib for plots; and Git for version control.

V. RESULTS AND DISCUSSIONS

➢ *Experimental Setup and Metrics*

Detection is evaluated with precision, recall and mean average precision at an intersection-over-union (IoU) threshold of 0.5 (mAP@0.5). Real-time behavior is evaluated with frames per second (FPS) and end-to-end latency, and the energy saving S is measured against the constant-output baseline.

➢ *Detection Performance*

Table 5 compares the classical baseline with the proposed detector-plus-tracker design. The proposed design reaches

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

92.6% precision, 89.4% recall and 91.1% mAP@0.5, against 62.4%, 58.1% and 54.2% for background subtraction, so both precision and recall exceed the 85% target. This comes at the cost of higher inference latency (34.2 ms against 18.5 ms) and a lower frame rate (29 FPS against 45 FPS), but the latency remains far below the 500 ms decision target, so real-time operation is preserved.

Table 5 Detection Performance

Metric	Baseline (background subtraction)	Proposed (YOLO-Nano + tracker)
Precision (%)	62.4	**92.6**
Recall (%)	58.1	**89.4**
mAP@0.5 (%)	54.2	**91.1**
Inference latency (ms)	18.5	34.2
Frame rate (FPS)	45	29

The normalized confusion matrix in Fig. 3 shows per-class correct-detection rates between 0.85 and 0.93, highest for pedestrians (0.93) and cars (0.92) and lowest for the background class (0.85). The largest single confusion is background being detected as a pedestrian (0.05), while pedestrians and vehicles are rarely confused with each other (0.01), which matters because the dimming level depends on this distinction.

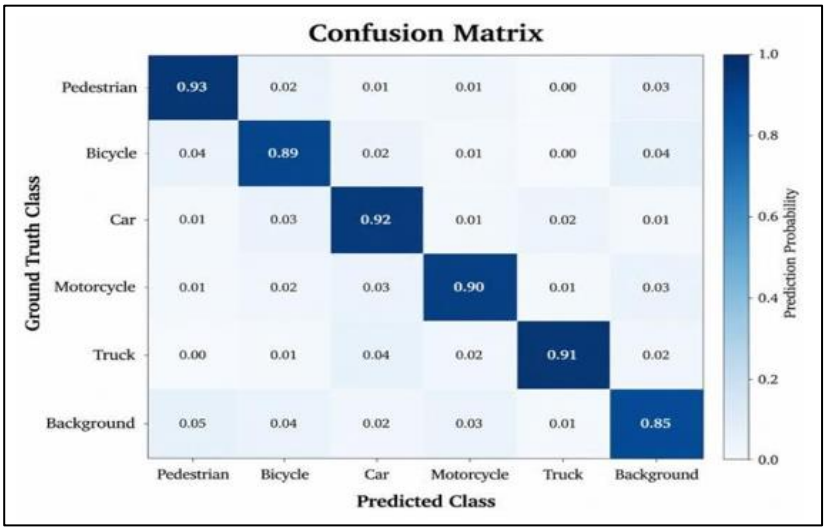


Fig 3 Normalized Confusion Matrix of the Proposed Detector.

Fig. 4 shows the live monitoring view, with bounding boxes, track IDs, the current state and the PWM duty cycle. Fig. 5 shows detection and tracking in clear low-light conditions and under headlight glare, where the pedestrians and cars are still detected and tracked.

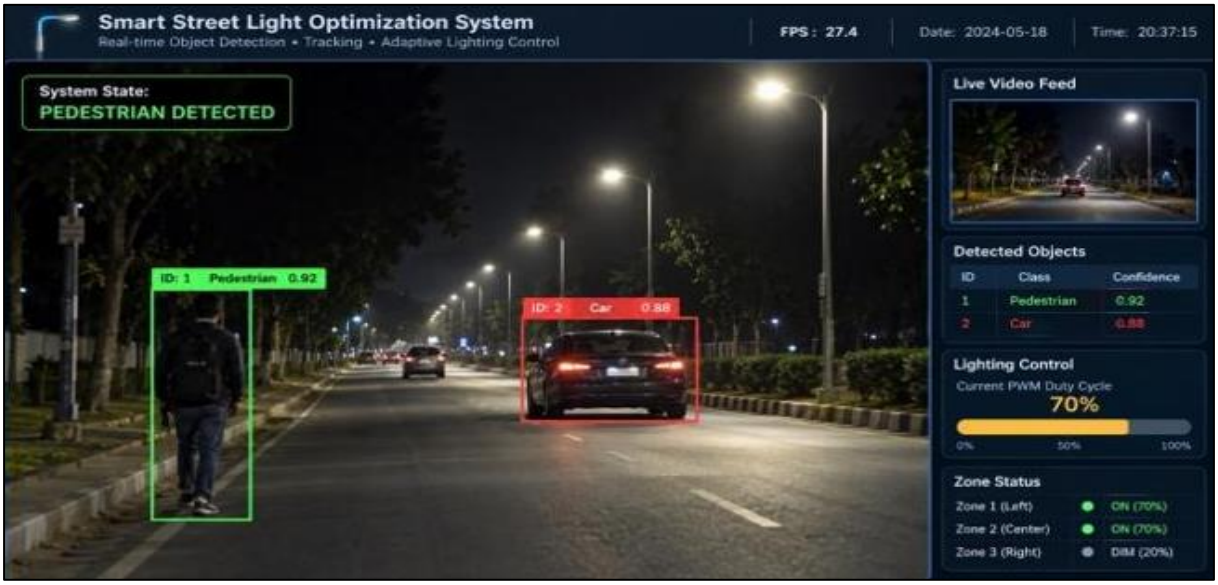
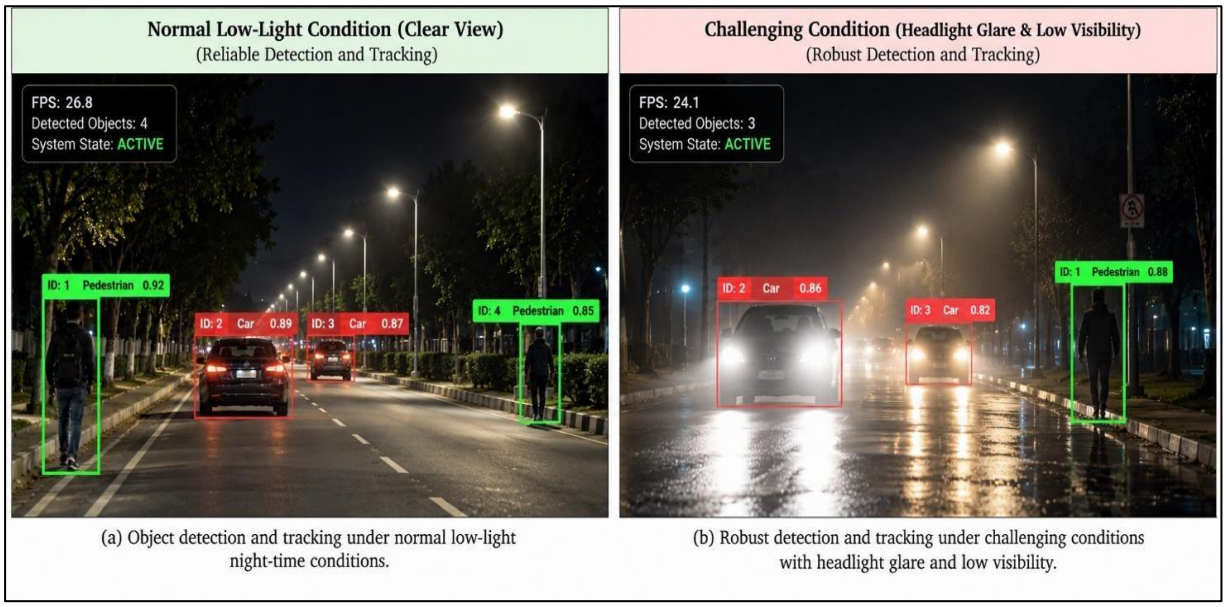


Fig 4 Live Monitoring View Showing Detections, Tracking and the Current PWM Duty Cycle.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631



(a) Object detection and tracking under normal low-light night-time conditions.

(b) Robust detection and tracking under challenging conditions with headlight glare and low visibility.

Fig 5 Object Detection and Tracking on Night-Time Frames: (a) Normal Low-Light Condition; (b) Headlight Glare and Low Visibility.

➢ *Energy Performance*

Table 6 gives the energy comparison for a 100 W lamp over a 12 h night in three traffic scenarios. The saving ranges from 65.6% for low traffic to 34.5% for a high-traffic urban thoroughfare, so the benefit depends on how much of the night the road is occupied. Fig. 6 shows the corresponding power profile: the adaptive controller holds 30 W during quiet periods and rises to 70 W or 100 W during pedestrian and vehicle activity, whereas the constant-output lamp draws 100 W throughout. The energy figures are model-based estimates; validation against a metered lamp is left for future work.

Table 6 Energy Comparison for a 100 W Lamp over 12 Hours

Traffic scenario	Constant (Wh)	Adaptive (Wh)	Saving (%)
Low (midnight–4 AM)	1200.0	412.5	**65.6**
Moderate (evening/night)	1200.0	528.0	**56.0**
High (urban thoroughfare)	1200.0	785.4	**34.5**

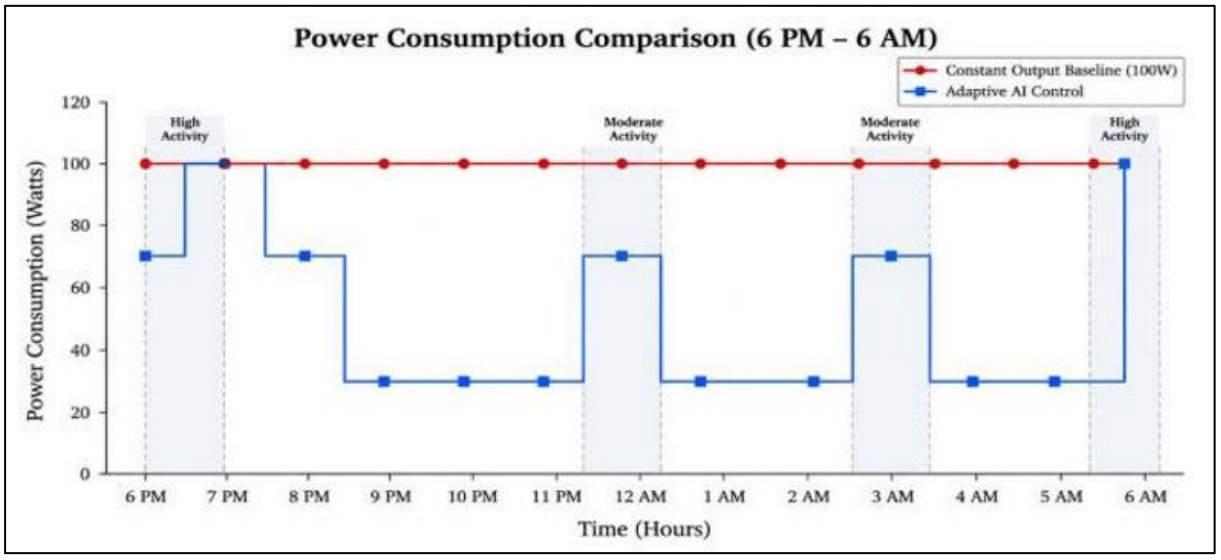


Fig 6 Power Consumption of Constant-Output Lighting Versus Adaptive Control Over a Night.

➢ *Validation and Discussions*

Validation consists of running both designs on the same held-out video and comparing them against manually labeled frames, and of repeating the energy tests across traffic scenarios. The controller is also checked to confirm that the lamp never falls below the baseline level and that the fail-safe mode triggers when the camera feed is cut. Headlight glare, occlusion, heavy rain and very low light remain the main failure cases. The detection threshold also trades off false triggers against missed detections: a lower threshold reduces

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov631

missed road users but raises false triggers that waste energy, and a higher threshold does the opposite.

VI. CONCLUSION AND FUTURE WORK

This paper presented a vision-based adaptive street-lighting system that detects pedestrians and vehicles, tracks their movement and sets lamp brightness through PWM control. The selected detector-plus-tracker design reached 92.6% precision, 89.4% recall and 91.1% mAP@0.5 at 34.2 ms inference latency, outperforming the classical background-subtraction baseline in classification and robustness at the cost of additional computation. The estimated energy saving over constant-output lighting ranged from 34.5% to 65.6% depending on traffic density, while a minimum light level and a fail-safe mode preserve safety. Future work includes:

- Training on a larger, locally collected night-time dataset and applying domain adaptation;
- Adding thermal or radar sensing for fog, rain and glare;
- Extending to a multi-lamp network in which neighboring lamps pre-brighten using shared tracking information;
- Adding traffic-flow forecasting to set baseline levels by time of day;
- Integrating fault detection, remote dashboards and standard communication protocols;
- Running a field pilot with a metered lamp to validate the energy model against real power measurements; and
- Strengthening privacy through on-device anonymization and encrypted logging.

REFERENCES

[1.] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.

[2.] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.

[3.] W. Liu *et al.*, "SSD: Single shot multibox detector," in *Proc. European Conf. Computer Vision (ECCV)*, 2016, pp. 21–37.

[4.] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple online and realtime tracking," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 2016, pp. 3464–3468.

[5.] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *Proc. IEEE ICIP*, 2017, pp. 3645–3649.

[6.] CEN, *EN 13201: Road Lighting*, European Committee for Standardization (current edition).

[7.] IEC, *IEC 62386: Digital Addressable Lighting Interface*, International Electrotechnical Commission (current edition).

[8.] A. G. Howard *et al.*, "MobileNets: Efficient convolutional neural networks for mobile vision applications," arXiv:1704.04861, 2017.

[9.] G. Bradski, "The OpenCV library," *Dr. Dobb's Journal of Software Tools*, 2000.

[10.] M. Abadi *et al.*, "TensorFlow: A system for large-scale machine learning," in *Proc. USENIX Symp. Operating Systems Design and Implementation (OSDI)*, 2016, pp. 265–283.

[11.] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.

[12.] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. European Conf. Computer Vision (ECCV)*, 2014, pp. 740–755.

[13.] Bureau of Indian Standards, *IS 1944 (Parts 1 and 2): Code of Practice for Lighting of Public Thoroughfares*, 1970.

[14.] Press Information Bureau, Government of India, Ministry of Power, "Street Lighting National Programme," 2024. This is the PIB release giving the installation figures as of 30 June 2024.