# GreenCam: Solar-Powered Smart Camera for Traffic Condition Monitoring

Chengbin Lei
Nanyang Technological University
Singapore

Zimo Ma
Nanyang Technological University
Singapore

Rui Tan
Nanyang Technological University
Singapore

## Abstract

This paper presents GreenCam, a solar-powered smart traffic camera design. It applies deep reinforcement learning to learn a policy that controls battery charging and switches *in situ* inference among various deep neural network models with different accuracy and power profiles in the presence of power intermittency. Its cost function accounts for the amortized embodied carbon emissions carried by the added photovoltaic panel and battery. Evaluation driven by real data traces shows that, compared with the conventional on-grid design, GreenCam achieves similar quality of service, reduces carbon emission by 72%, and thus curtails three to four trees to be planted for neutralizing the carbon emission of each traffic camera installation. GreenCam also outperforms other solar-powered designs that adopt straightforward, heuristic, or ablated policies.

## CCS Concepts

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Hardware** → *Energy generation and storage*; • **Social and professional topics** → Sustainability.

**ACM Reference Format:**
Chengbin Lei, Zimo Ma, and Rui Tan. 2026. GreenCam: Solar-Powered Smart Camera for Traffic Condition Monitoring. In *The 24th Annual International Conference on Mobile Systems, Applications and Services (MobiSys Workshop '26), June 21–25, 2026, Cambridge, United Kingdom.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3812836.3815165

## 1 Introduction

Intelligent transportation systems (ITS) are essential to smart cities. As a fundamental component of ITS, camera-based traffic monitoring enables traffic signal control, infrastructure support for pedestrians and autonomous vehicles, congestion prevention and resolution, road network planning and retrofitting. Singapore is now operating around 1,100 traffic cameras [12]. Recently, convolutional neural network-based artificial intelligence (AI) has been integrated into cameras to achieve automated traffic analysis [7, 9].

However, the mission criticality and the power overhead of the *in situ* AI computing require a stable power supply and engender a negative environmental implication due to the related carbon emissions. First, limited grid coverage, especially in suburban areas, restricts the deployment of smart traffic cameras. The power grid access requirement will also limit temporary, swift deployments for event traffic control. Second, regarding the carbon overhead, considering a low power of 13.7 W per device (e.g., the HoloSens SDC [9]) and the electric grid carbon intensity of 480 $gCO_2$/kWh in Singapore [4], each device emits 57.6 $kgCO_2$ annually. To neutralize this emission, several trees would be needed given each tree's 25 $kgCO_2$/year average absorption rate [1]. However, it takes years to grow trees and achieve the required absorption capability.

The above issues generate an inquiry regarding the feasibility and benefits of solar-powered, battery-assisted smart traffic camera. Although renewable energy-based systems provide the general "greener" promise, for this smart traffic camera application, the potentially undermined serviceability due to the inherent intermittency of solar power and the possible *embodied carbon* pitfall due to say the overprovisioning of photovoltaic (PV) panel size and battery capacity for guaranteed serviceability, are two new factors requiring a quantitative investigation. Note that both the added PV panel and battery carry embodied carbon emissions, which are produced during their manufacturing, transportation, and disposal. For the investigation, we propose *GreenCam*, a solar-powered smart traffic monitoring system designed to optimize a metric jointly capturing carbon emissions and serviceability that is modeled by the product of monitoring accuracy and availability. With GreenCam, we can evaluate the carbon-versus-serviceability Pareto frontier under various realistic configurations of PV panel size and battery capacity, in the context of a specified solar irradiance profile.

GreenCam considers both *workload adaptation* and *battery management*. The former switches among multiple AI models of different sizes based on the battery energy level. Following from the general trend that larger models offer higher accuracy but require more energy, this AI model switching enables GreenCam to balance traffic monitoring accuracy and availability under finite energy budgets in the pursuit of serviceability. The latter, i.e., battery management through the control of the battery charging switch, aims at reducing the *amortized embodied carbon emission* (AECE) by extending the battery's operational lifespan. This follows from the facts that the charging and discharging processes generate heat due to the battery's internal resistance and the battery's operating temperatures affect its capacity degradation and thus lifespan [10]. Specifically, extreme battery temperatures would increase AECE. Therefore, GreenCam aims at managing the battery temperature through charging process control for greening the system.

Workload adaptation and battery management jointly affect GreenCam's carbon emission and serviceability. Specifically, the charging strategy determines AECE and the battery energy level, which constrains model selection for monitoring accuracy and modulates service availability under intermittent solar power. Moreover, owing to battery storage and the stochasticity of the solar power and traffic, the problem is a decision process for a dynamic system. GreenCam employs deep reinforcement learning (DRL) to

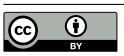

This work is licensed under a Creative Commons Attribution 4.0 International License.
*MobiSys Workshop '26, Cambridge, United Kingdom*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2712-2/2026/06
https://doi.org/10.1145/3812836.3815165

learn a joint policy for regulating AI model switching and battery charging control, in order to optimize the joint metric capturing carbon emission and serviceability. We conduct simulations based on realistic settings and real-world traces to evaluate the proposed GreenCam and four baselines in the context of Singapore. The evaluation results show that an NVIDIA Jetson Orin-based GreenCam reduces carbon emission by 72% compared with the conventional on-grid design. This translates to planting three to four fewer trees in neutralizing the carbon emission of each traffic camera installation. GreenCam also outperforms three off-grid designs due to its optimal policy. Sensitivity analysis shows that 1.5 $m^2$ PV panel and 0.5 kWh battery capacity are good configurations for Jetson Orin-based GreenCam in Singapore's context.

*Paper organization:* §2 and §3 present the design and evaluation results of GreenCam, respectively. §4 discusses related work. §5 concludes this paper.

## 2 Design of GreenCam

GreenCam is an off-grid, AI-driven traffic monitoring system designed to operate solely on solar energy. It integrates solar energy harvesting, battery storage, and on-device AI inference. It employs a DRL controller to manage energy usage and inference workload online, in order to achieve continuous and low-carbon operations. In this section, §2.1 overviews the system design; §2.2 analytically models the system; §2.3 presents the DRL controller design.

### 2.1 System Overview

*2.1.1 Hardware components.* GreenCam integrates three hardware components: 1) a smart camera with an RGB sensor and a computing unit for image stream processing, 2) a PV panel for solar energy harvesting, and 3) a lithium-ion battery for energy storage. We consider a traffic monitoring task of real-time object detection and tracking. For profiling and evaluation, we employ the YOLO family as the object detectors. The PV panel is properly oriented and connected to the battery via a controllable charging switch. The battery serves as a power intermediary and an energy buffer between the PV panel and the smart camera for system's continuous operations, given the fluctuating solar irradiance conditions.

*2.1.2 DRL controller.* GreenCam employs a software controller to jointly decide two operational actions: *battery charging* and *AI model switching*. The controller is executed at each *control period*'s beginning. Based on the measured ambient temperature, battery temperature, battery energy level, and the estimated current output of the PV panel, the controller (i) decides whether to charge the battery and (ii) selects an inference model from a predefined candidates set with varying accuracy metrics and energy usages. The controller dynamically regulates the battery charging action to mitigate the battery capacity degradation. Meanwhile, it adapts the inference workload to manage energy usage and improve service continuity by switching between models of different sizes. However, due to the complex impact of the operational actions on the system's AECE and serviceability, it is non-trivial to derive a policy for the controller to follow. Thus, we adopt DRL to learn the policy with a lightweight neural network. This learning is performed in offline simulations driven by solar irradiance and labeled traffic traces, as well as analytical system models presented in §2.2.

### 2.2 System Modeling

*2.2.1 Scope of carbon emission analysis.* The carbon emission of a conventional on-grid smart camera comprises the embodied carbon of the smart camera and the carbon emission arising from the electricity usage during operations (referred to as *grid carbon*). In contrast, although GreenCam is free of the power grid electricity usage, it incurs extra embodied carbon of the PV panel and the battery. Therefore, the carbon footprint comparison between the conventional on-grid design and GreenCam nails down to the comparison between the former's grid carbon and the latter's embodied carbon carried by the PV panel and the battery. Due to PV panel's long operational lifespan, we consider a constant AECE for PV panel. In contrast, the battery is subject to dynamic degradation driven by charging and discharging processes. Therefore, we model the battery degradation and the dependent AECE. In reality, an under-performing battery that cannot provide sufficient output voltage due to the degradation needs to be replaced with a new battery. As such, the total AECE related to PV panel and battery is simply the arithmetic sum of their individual AECEs.

*2.2.2 Battery thermal model.* The battery temperature $T_{\text{bat}}$ evolves due to internal heat generation and external heat exchange. Specifically, the charging and discharging currents, denoted by $I_c$ and $I_d$, respectively, induce heat via internal resistance (Joule heating), while the ambient temperature governs the convective heat dissipation rate. We characterize this thermal evolution using a first-order discrete-time recursive equation [3]: $T_{\text{bat}}[k+1] - T_{\text{bat}}[k] = \frac{Q_c[k]+Q_d[k]-Q_b[k]}{m_{\text{bat}} \cdot c_{\text{bat}}}$, where $T_{\text{bat}}[k]$ denotes the battery temperature in the $k^{\text{th}}$ control period. The $Q_c[k] = \tau R_c I_c^2[k]$ accounts for the Joule heating produced during charging, where $\tau$ is the control period, $R_c$ is the internal resistance. Correspondingly, $Q_d[k]$ captures the thermal contribution from discharging. The battery's thermal inertia is governed by its physical properties characterized by the battery's mass $m_{\text{bat}}$ and specific heat capacity $c_{\text{bat}}$. The $Q_b[k]$ represents the heat dissipated from the battery to the ambient in the $k^{\text{th}}$ control period. It is modeled as: $Q_b[k] = \frac{\tau \cdot (T_{\text{bat}}[k] - T_{\text{amb}}[k])}{R_{\text{thm}}}$, where $T_{\text{amb}}[k]$ represents the ambient temperature in the $k^{\text{th}}$ control period, $R_{\text{thm}}$ represents the effective thermal resistance regarding the heat transfer between the ambient air and the battery. Therefore, the battery temperature $T_{\text{bat}}$ is directly influenced by three variables: ambient temperature $T_{\text{amb}}$, charging current $I_c$, and discharging current $I_d$. The $I_c$ and $I_d$ are modeled as: $I_c[k] = \frac{P_s[k]}{V_c}$, $I_d[k] = \frac{P_d[k]}{V_d}$, where $P_s[k]$ denotes the solar charging power, $P_d[k]$ denotes the total power of discharging, and $V_c$ and $V_d$ denote the charging voltage of the battery and the operating voltage of the smart camera, respectively. The $P_s[k]$ is modeled as $P_s[k] = \frac{E_s[k]}{\tau}$, where $E_s[k]$ represents the energy harvested by the PV panel during the $k^{\text{th}}$ period. The $P_d[k]$ is modeled as $P_d[k] = P_m[k] + P_{\text{cam}}$, where $P_m[k]$ represents the inference power of selected AI model at $k^{\text{th}}$ control period, $P_{\text{cam}}$ denotes the camera sensor power.

*2.2.3 Carbon emission models.* For the conventional grid-powered smart camera, the grid carbon can be modeled by $C_g[k] = \text{CI} \cdot E_g[k]$, where CI denotes the carbon intensity of the power grid and $E_g[k]$ represents the total grid energy used. GreenCam's AECE can be modeld by $C_e[k] = C_{e,\text{bat}}[k] + C_{e,pv}$, where $C_{e,\text{bat}}[k]$ and

$C_{e,pv}$ denote the AECE of the battery and PV panel, respectively. The $C_{e,\text{bat}}[k]$ is modeled by $C_{e,\text{bat}}[k] = \Gamma[k] \cdot C^{\text{bat}}_{\text{emb}}$, where $C^{\text{bat}}_{\text{emb}}$ is the total embodied carbon of the battery and $\Gamma[k]$ represents the degradation factor. The $\Gamma[k]$ captures the combined effects of charging and discharging, i.e., $\Gamma[k] = \gamma_c[k] \cdot E_c[k] + \gamma_d[k] \cdot E_d[k]$, where $\gamma_c[k]$ and $\gamma_d[k]$ denote the degradation rates associated with charging and discharging, respectively, and $E_c[k]$ and $E_d[k]$ denote the amount of energy charged into and discharged from the battery during the $k^{\text{th}}$ control period. The $E_c[k]$ is modeled by $E_c[k] = E_s[k] - Q_c[k]$. The $E_d[k]$ is modeled by $E_d[k] = P_d[k] \cdot \tau$. In addition, the harvested solar energy is modeled by $E_s[k] = G_s[k] \cdot A \cdot \mu \cdot \eta_s$, where $G_s[k]$ denotes the solar irradiance, $A$ denotes the area of the PV panel, $\mu$ denotes the conversion efficiency of the PV panel, and $\eta_s$ denotes the system loss coefficient.

The degradation rates depend on the battery operating conditions, including battery temperature, charging and discharging currents. Specifically, $\gamma_c[k] = q(T_{\text{bat}}[k], I_c[k], E_{\max}[k])$ and $\gamma_d[k] = q(T_{\text{bat}}[k], I_d[k], E_{\max}[k])$, where the $E_{\max}[k]$ denotes the maximum battery capacity in the $k^{\text{th}}$ control period and the detailed form of $q(\cdot)$ can be the empirical degradation model obtained in [20]. The $E_{\max}[k]$ progressively decreases, as degradation accumulates over time, which is modeled by $E_{\max}[k] = E_{\max}[0]\left(1 - \sum_{i=0}^{k} \Gamma[i]\right)$, where $E_{\max}[0]$ is the initial maximum capacity of the battery.

The AECE of the PV panel is modeled by $C_{e,pv} = \frac{C^{\text{emb}}_{pv}}{N_{pv}}$, where $C^{\text{emb}}_{pv}$ denotes the total embodied carbon and $N_{pv} \in \mathbb{Z}^+$ is the number of control periods over its lifespan. In our evaluation (§3), we set $N_{pv}$ based on a 20-year lifespan following industry benchmarks [11].

## 2.3 DRL Controller Design

*2.3.1 Markov decision process (MDP) formulation. State:* The system state in the $k^{\text{th}}$ control period is $s_k = (T_{\text{bat}}[k], T_{\text{amb}}[k], E[k])$, where $T_{\text{bat}}[k]$ and $T_{\text{amb}}[k]$ denote the battery and the ambient temperatures, $E[k]$ denotes the battery energy. The $T_{\text{amb}}[k]$ governs the heat exchange rate, directly affecting $T_{\text{bat}}[k]$, while $T_{\text{bat}}[k]$ determines battery degradation process $\Gamma[k]$. The $E[k]$ characterizes the system's energy reserve. Together, these three variables capture the thermal condition of the battery and the system's energy state for control decision-making.

*Action:* The action in the $k^{\text{th}}$ control period is defined as $a[k] = (a_c[k], a_m[k])$, where $a_c[k] \in \{0, 1\}$ denotes the charging decision (charging or no charging), and $a_m[k]$ denotes the selected inference model from a predefined set.

*Reward:* We consider three factors, i.e., carbon emissions, inference accuracy, and system downtime over the system lifespan. Moreover, we incorporate downtime into accuracy, yielding *effective accuracy* $\alpha[k]$ as the serviceability metric. When the system is unavailable due to battery depletion, $\alpha[k] = 0$; otherwise, $\alpha[k]$ is the accuracy of the selected inference model. The carbon emission $C[k]$ is $C_g[k]$ for the conventional on-grid smart camera, or $C_e[k]$ for GreenCam. The reward is defined as $r[k] = \lambda \cdot \alpha[k] - C[k]$, where $\lambda$ is a fixed weighting coefficient that balances effective accuracy and carbon emission. A larger $\lambda$ emphasizes accuracy and availability, while a smaller $\lambda$ prioritizes carbon reduction. We empirically select $\lambda$ to achieve a balanced trade-off, which preserves effective accuracy while significantly reducing carbon emissions.

**Table 1: Default settings of system parameters**

| Parameter | Setting | Parameter | Setting |
|---|---|---|---|
| $\tau$ | 1 h | $E_{\max}[0]$ | 0.48 kWh |
| $\mu$ | 0.225 [13] | $A$ | 1.5 $\text{m}^2$ |
| $\eta_c$ | 0.75 | $C_{e,pv}$ | 1.05 g $CO_2$ [6] |
| $V_c$ | 14.4 V [18] | $C^{\text{bat}}_{\text{emb}}$ | 72 kg $CO_2$ [23] |
| $V_d$ | 12 V [18] | $\eta_{\text{bat}}$ | 0.7 |
| $P_{\text{cam}}$ | 6 W [8] | $\lambda$ | 30 |
| CI | 402 g $CO_2$/kWh [5] | $R_{\text{thm}}$ | 4 [14] |
| $m_{\text{bat}}$ | 4 kg | $c_{\text{bat}}$ | 900 (J/kg · K) [14] |

*Objective:* We aim to learn a policy $\pi \in \Pi$ that maximizes the expected cumulative reward over the system lifetime: $\pi^* = \arg\max_{\pi \in \Pi} \mathbb{E}\left[\sum_{k=0}^{K_{\text{term}}} r[k]\right]$, where $K_{\text{term}}$ denotes the number of control periods before episode termination, and the expectation is taken over stochastic exogenous inputs.

*Termination:* The episode terminates when the battery condition degrades to a predefined threshold $\eta_{\text{bat}}$, formally stated as $\frac{E_{\max}[k]}{E_{\max}[0]} \leq \eta_{\text{bat}}$. This termination design aligns with battery replacement criterion and ensures that the optimization focuses on the effective operational lifetime of the system.

*2.3.2 DRL-based policy learning.* We adopt proximal policy optimization (PPO) [19] to solve the MDP. PPO is chosen for its stable training and efficient policy updates. The policy is learned via offline training. We construct a digital replica of the system using real-world traces of exogenous inputs and system parameters, and train the PPO agent through extensive simulations driven by these traces. After training, the learned policy is deployed for online control, where the model parameters remain fixed. When environmental conditions change (e.g., varied solar irradiance profile), the agent can be retrained offline with updated data.

# 3 Evaluation

We conduct trace-driven simulations with real-world environmental and solar data traces to compare GreenCam and several baselines. We also perform sensitivity analysis on key system parameters.

## 3.1 Experiment Setup and Methodology

*Ambient temperature, solar irradiance, and other settings:* We obtain Singapore's ambient temperature and raw solar irradiance data in 2025 from the NASA POWER database [16]. The traces capture full daytime-nighttime cycles. Environmental conditions, such as time-of-day and weather, are implicitly captured through these traces. They directly affect both energy availability and thermal dynamics, thus enabling the learned policy to adapt its decisions under varying conditions, even under identical battery states. Table 1 summarizes the default parameter settings. To improve robustness of the trained DRL controller, we augment the data by adding zero-mean Gaussian noise with a standard deviation of 5% of the original value.

*Object detection models:* We adopt five YOLOv5 variants (i.e., n, s, m, l, x) as the candidates for the model switching, which have increasing model size and computational complexity. These models are trained on the UA-DETRAC dataset [21], which is a vehicle detection benchmark. Table 2 summarizes their model size,

accuracy (in mAP@50 and recall), and power overhead measured on an NVIDIA Jetson Orin 32GB (MAXN mode). The mAP@50 is the mean average precision at an intersection over union (IoU) threshold of 0.5, while recall is the fraction of ground-truth objects correctly detected.

**Table 2: Performance of YOLOv5 variants**

| Model | Params (M) | mAP@50 | Recall | Power ($W$) |
|---|---|---|---|---|
| YOLOv5n | 1.9 | 0.566 | 0.543 | 11.9 |
| YOLOv5s | 7.2 | 0.643 | 0.608 | 13.3 |
| YOLOv5m | 21.2 | 0.695 | 0.658 | 16.7 |
| YOLOv5l | 46.5 | 0.727 | 0.702 | 20.2 |
| YOLOv5x | 86.7 | 0.731 | 0.710 | 24.9 |

*GreenCam and baselines:* GreenCam controller's policy network is a fully connected neural network consisting of an input layer, two hidden layers, and an output layer. Each hidden layer has 64 neurons with rectified linear unit (ReLU) activation. The network is lightweight, containing only 9,608 parameters. We adopt the following four baseline designs:

- **On-grid** uses power grid as the sole energy source and always runs the largest model, i.e., YOLOv5x.
- **Model-X** has the same hardware setup as GreenCam but always charges the battery and runs YOLOv5x.
- **Heuristic** replaces GreenCam's DRL controller with a rule-based controller. It performs charging when $T_{\text{bat}} \in [10°\text{C}, 30°\text{C}]$, which is the sweet spot for minimum battery degradation [20]. For workload adaptation, it follows a greedy strategy by adopting the largest YOLOv5 variant that can be supported by the available energy in the upcoming control period.
- **GreenCam0** is an ablation of GreenCam that performs workload adaptation but keeps charging the battery. Comparison with GreenCam0 will show the contribution of the battery management to the system's carbon emission reduction.

*System initialization:* To eliminate the unnecessary battery depletion during the initial operation phase of off-grid designs (GreenCam, GreenCam0, Model-X, and Heuristic), we assume the battery is fully charged before deployment. The pre-charged energy accounts for only 0.06% of the total harvested energy over the battery lifetime. Its impact on the comparison with On-grid is negligible.

## 3.2 Comparison with Baselines

Fig. 1 compares GreenCam with four baselines on average carbon emission, mAP@50, recall, and downtime rate. The downtime rate is the percentage of time during which the battery depletes and the service is unavailable.

**On-grid vs. GreenCam.** GreenCam reduces carbon emissions by 72% compared with On-grid, with only marginal accuracy loss (< 3%). The reduction in carbon emissions is roughly equivalent to the absorption of 3.5 trees. This confirms the benefit of solar-powered, battery-assisted smart traffic camera by jointly considering both accuracy and carbon emission.

**Off-grid vs. GreenCam.** Among all off-grid designs, GreenCam achieves the best overall trade-off between carbon emission and serviceability (i.e., accuracy and availability). It reduces carbon emission by up to 21% compared with Model-X, while achieving higher mAP@50, recall, and lower downtime rate. This is because Model-X always runs the largest detection model, which leads to frequent outages and high carbon emissions. Although Heuristic achieves slightly lower carbon emission due to its careful charging, it incurs high downtime rate and therefore significantly lower mAP@50 and recall (up to 10%). GreenCam0, by contrast, maintains high accuracy and low downtime rate (both within 3%) through continuous charging. However, it leads to a 24% increase in carbon emission, compared with GreenCam. These results demonstrate that GreenCam balances energy usage and workload adaptation to minimize carbon emission without sacrificing serviceability, compared to naive strategies. Stronger baselines, such as model predictive control (MPC) with environmental forecasts or oracle policies with future knowledge, are left as future work, as our focus is to demonstrate the effectiveness of the proposed DRL-based system.

## 3.3 Sensitivity Analysis

To evaluate the impact of PV panel size and battery capacity, we train 25 DRL agents for all combinations of 5 PV panel sizes (from 0.5 $\text{m}^2$ to 2.5 $\text{m}^2$) and 5 battery capacities (from 0.24 kWh to 1.2 kWh), as shown in Fig. 2. From Fig. 2(a), carbon emission increases with the PV panel size due to higher amortized embodied carbon $C_{e,pv}$. In contrast, larger battery capacities reduce degradation rates and thus suppress carbon emission growth with PV panel size. Figs. 2(b), 2(c), and 2(d) present resulting accuracy and availability. When the battery capacity is below 0.48 kWh or the PV panel size is smaller than 1.5 $\text{m}^2$, insufficient energy supply leads to frequent outages, which results in high downtime rate and low accuracy. Once the PV panel size exceeds 1.5 $\text{m}^2$ and the battery capacity exceeds 0.48 kWh, the system achieves saturated performance, i.e., both mAP@50 and recall approaching their upper bounds and downtime rate approaching zero. Therefore, 1.5 $\text{m}^2$ and 0.5 kWh are good configurations for Jetson Orin-based GreenCam for traffic object detection in Singapore's context. Note that the 0.5 kWh battery capacity is practical. For instance, a 0.512 kWh battery product [17] has a volume of $196 \times 135 \times 172\ \text{mm}^3$, comparable to a small shoebox.

# 4 Related Work

*Energy-efficient object detection systems.* Existing works have explored the energy-efficient deployment from a model-level perspective. Lightweight architectures, such as MobileViT [15], reduce model redundancy while maintaining competitive mobile performance. Neuromorphic approaches, such as spiking neural networks [2], improve efficiency through event-driven computation and sparse activations. However, these prior efforts mainly focus on intra-model optimization and overlook system-level orchestration.

*Degradation-aware battery-powered system optimization.* Prior work incorporates degradation-aware optimization in PV-battery systems to determine optimal battery capacity and minimize lifecycle cost [22]. It has also been applied to grid-scale energy storage systems for battery technology selection, where degradation-aware models evaluate the suitability of different battery chemistries for specific services [24]. However, these studies often overlook the dynamic coupling between energy availability and workloads.

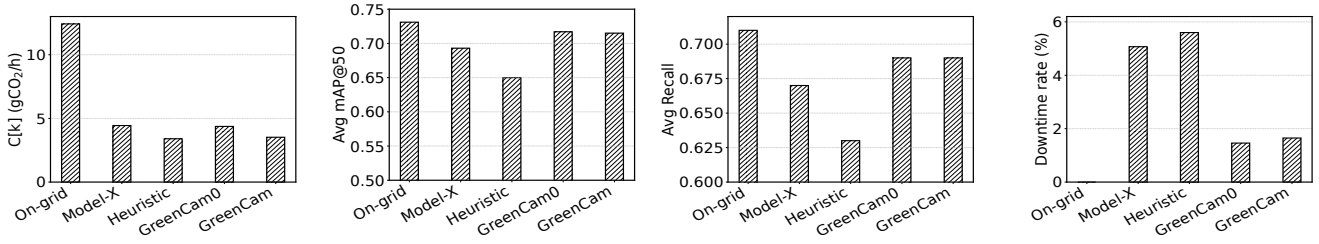


**(a) Average carbon emission** **(b) Average mAP@50** **(c) Average recall** **(d) Downtime rate**

**Figure 1: Performance of GreenCam and various baseline designs.**

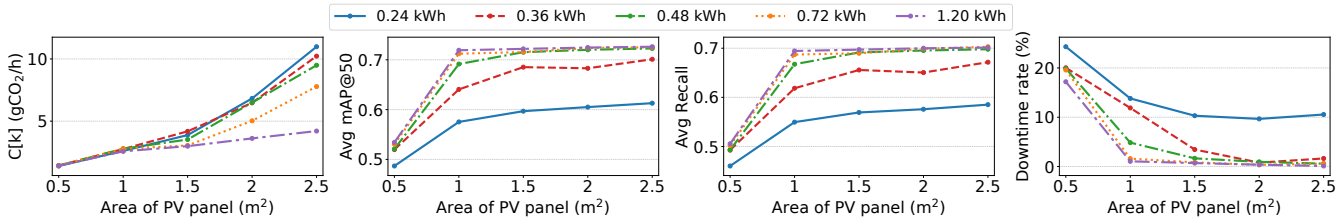


**(a) Average carbon emissions** **(b) Average mAP@50** **(c) Average Recall** **(d) Downtime rate**

**Figure 2: Sensitivity analysis for PV panel size (in x-axis) and battery capacity (in legends).**

## 5 Conclusion

We propose GreenCam, an off-grid AI-enabled smart traffic monitoring system that achieves low carbon emissions while maintaining continuous, high-quality object detection. By integrating solar power and battery storage, GreenCam eliminates operational carbon emissions. It further reduces embodied carbon through workload adaptation and temperature-aware battery management. Experimental results demonstrate its effectiveness in maintaining accuracy and reducing carbon emissions.

## Acknowledgments

This work is supported in part by Singapore Ministry of Education under AcRF Tier-1 RT14/22 and in part by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG4-GC-2023-006-1B).

## References

[1] 2025. *Huawei HoloSens iClient Product Documentation.* https://bit.ly/4vpliKf
[2] Lennard Bodden, Duc Bach Ha, Franziska Schwaiger, Lars Kreuzberg, and Sven Behnke. 2024. Spiking CenterNet: A Distillation-boosted Spiking Neural Network for Object Detection. In *2024 IJCNN*. 1–9. doi:10.1109/IJCNN60899.2024.10650418
[3] Yunus A. Cengel and Michael A. Boles. 2015. *Thermodynamics: An Engineering Approach* (8 ed.).
[4] Electricity Maps. 2026. *Real-time Carbon Intensity Data: Singapore (Daily Average).* Tomorrow. https://bit.ly/4cotAcV
[5] Energy Market Authority. 2024. Singapore Energy Statistics 2024: Electricity Generation and Consumption. https://bit.ly/4mqmWHJ.
[6] Yu Gan, Amgad Elgowainy, Zifeng Lu, Jarod C Kelly, Michael Wang, Richard D Boardman, and Jason Marcinkoski. 2023. Greenhouse gas emissions embodied in the U.S. solar photovoltaic supply chain. *Environ. Res. Lett.* 18, 10 (2023). doi:10.1088/1748-9326/acf50d
[7] Hikvision. 2024. DS-2XA8045E-A(B) AI Open Platform Bullet Camera. https://bit.ly/4dGVvH8.
[8] Hikvision. 2026. 2 MP ColorVu Network Camera DS-2CD3T26(D)WDV3-L. https://bit.ly/4ef6IyQ.
[9] Huawei Technologies Co., Ltd. [n. d.]. *Huawei HoloSens iClient Product Documentation.* Huawei Enterprise.
[10] Isidor Buchmann. 2021. BU-808: How to Prolong Lithium-based Batteries. Battery University.
[11] D. C. Jordan and S. R. Kurtz. 2013. Photovoltaic Degradation Rates—an Analytical Review. *Prog. Photovolt: Res. Appl.* 21, 1 (2013). doi:10.1002/pip.1182
[12] Land Transport Authority (LTA). 2026. Intelligent Transport Systems. https://bit.ly/48l8Knx.
[13] Ltd. LONGi Solar Technology Co. 2026. Hi-MO X6 Max Explorer (LR7-72HTH 605–615W) solar module. https://bit.ly/3Q1rJ6e.
[14] Hossein Maleki, Said Al Hallaj, J. Robert Selman, Ralph B. Dinwiddie, and H. Wang. 1999. Thermal Properties of Lithium-Ion Battery and Components. *Journal of The Electrochemical Society* 146, 3 (1999). doi:10.1149/1.1391704
[15] Sachin Mehta and Mohammad Rastegari. 2021. MobileViT: Light-weight, General-purpose, and Mobile-friendly Vision Transformer. *ArXiv* (2021).
[16] NASA POWER Project. 2025. Prediction of Worldwide Energy Resources (POWER) Data Access Viewer. https://power.larc.nasa.gov/.
[17] PowerTech Systems. 2026. 12V Lithium Battery 40Ah – PowerBrick+. https://bit.ly/41rcjuy.
[18] Renogy. 2025. Smart Lithium Iron Phosphate Battery RBT100LFP12S-US User Manual. https://bit.ly/4cjXZc6.
[19] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms.
[20] Andreas Thingvad, Lisa Calearo, Peter Bach Andersen, and Mattia Marinelli. 2021. Empirical Capacity Measurements of Electric Vehicles Subject to Battery Degradation From V2G Services. *IEEE Trans. Veh. Technol.* (2021). doi:10.1109/TVT.2021.3093161
[21] Longyin Wen, Dawei Du, Zhaowei Cai, Zhen Lei, Ming-Ching Chang, Honggang Qi, Jongwoo Lim, Ming-Hsuan Yang, and Siwei Lyu. 2020. UA-DETRAC: A new benchmark and protocol for multi-object detection and tracking. *Comput. Vis. Image Underst.* 193 (2020). doi:10.1016/j.cviu.2020.102907
[22] Yaling Wu, Zhongbing Liu, Jiangyang Liu, Hui Xiao, Ruimiao Liu, and Ling Zhang. 2022. Optimal battery capacity of grid-connected PV-battery systems considering battery degradation. *Renewable Energy* 181 (2022).
[23] Öivind Andersson and Pål Börjesson. 2021. The greenhouse gas emissions of an electrified vehicle combined with renewable fuels: Life cycle assessment and policy implications. *Applied Energy* 289 (2021). doi:10.1016/j.apenergy.2021.116621
[24] Maciej Świerczyński, Daniel Ioan Stroe, Ana-Irina Stan, Remus Teodorescu, and Dirk Uwe Sauer. 2014. Selection and Performance-Degradation Modeling of $LiMO_2/Li_4Ti_5O_{12}$ and $LiFePO_4$/C Battery Cells as Suitable Energy Storage Systems for Grid Integration With Wind Power Plants: An Example for the Primary Frequency Regulation Service. *IEEE Trans. on Sust. Energy* (2014). doi:10.1109/TSTE.2013.2273989