

RoomIQ: An Intelligent Hybrid Edge-Cloud Architecture for IoT-Enabled Energy Conservation in Educational Spaces

Yug Agarwal

Dept. of Computer Science and Engineering (AI)
ABES Institute of Technology
Ghaziabad, India
yugagarwal704@gmail.com

Meena Kumari Kashyap

Dept. of Computer Science and Engineering (IoT)
ABES Institute of Technology
Ghaziabad, India
meena.kumari@abesit.edu.in

Abstract—Energy conservation in educational institutions presents a persistent challenge, with significant wastage attributed to lights and fans operating in unoccupied classrooms. Traditional automation approaches, including passive infrared (PIR) sensor systems and centralized scheduling, suffer from inadequate occupancy detection, limited granularity, and inflexibility to real-world usage patterns. This paper introduces RoomIQ, a hybrid edge–cloud architecture that addresses these limitations through vision-based occupancy detection executed on resource-constrained edge devices at the room level. The system combines real-time computer vision inference using YOLOv8 on Raspberry Pi boards with a centralized Django-based control and analytics platform. Key innovations include smart zoning logic for independent control of lighting and HVAC appliances based on precise occupant position, a configurable privacy-first versus security-integrated deployment paradigm, and architectural resilience enabling autonomous room-level operation even under network disruption. A prototype deployment in a university classroom over 14 days shows 31% measured energy reduction versus PIR-based systems and 44% versus schedule-only control. Economic modeling indicates payback periods of 30–36 months under typical institutional operating conditions and 30–40% reduction in lighting and cooling energy consumption for Indian educational institutions, with per-room deployment costs of Rs. 6,100–6,800.

Index Terms—Internet of Things (IoT), Edge Computing, Smart Buildings, Energy Management, Computer Vision, Occupancy Detection, Raspberry Pi, Educational Infrastructure, Privacy-First Design, Hybrid Architecture

I. INTRODUCTION

A. The Multi-Dimensional Constraint Problem

Energy conservation in educational buildings requires simultaneously satisfying three competing constraints: **Cost** (per-room budgets \pm Rs. 10,000, one-time capex preferred), **Privacy** (camera deployment concerns under India's DPDP Act 2023), and **Resilience** (intermittent connectivity, limited IT support). Existing approaches—PIR sensors, thermal imaging, centralized cloud—optimize individual constraints but fail to satisfy all three, creating a deployment gap for resource-constrained institutions.

Research Question: *How can vision-based occupancy detection be architected to simultaneously satisfy cost,*

privacy, and resilience constraints for educational buildings in developing regions?

B. Contributions: An Architectural Framework for Constrained Environments

RoomIQ presents an architectural framework that composes existing techniques to simultaneously address cost, privacy, and resilience constraints for educational buildings in developing regions.

The specific research contributions are:

1. Hybrid Edge–Cloud Architecture with Failure-Resilient Autonomy

Each classroom operates autonomously with vision-based occupancy detection, smart zoning, and local appliance control independent of cloud connectivity. The central server handles configuration, analytics, and visualization but is not in the control-critical path. This eliminates single points of failure, reduces control latency from 200–500ms (cloud) to \pm 50ms (local GPIO), and maintains operation during network outages with post-reconnection synchronization. We analyze decentralized-vs-centralized trade-offs explicitly.

2. Configurable Privacy-First versus Security-Integrated Deployment Paradigm

Dual-mode design: **Privacy-First** (on-device inference only, anonymized metadata transmission) or **Security-Integrated** (optional on-premises video retention with consent and audit logging). This addresses existing systems' single-model constraint (e.g., vendor-controlled cloud storage).

3. Smart Zoning Framework with Prototype-Scale Validation

Spatial zoning framework dividing rooms into independent control zones. Using one 640 \times 480 camera at 10 FPS (Raspberry Pi 4B), the system provides row-specific lighting and proximity-based fan control at Rs. 6,100–6,800/room. Instrumented prototype deployment over 14 days demonstrates measurable energy reductions (31% vs. PIR, 44% vs. schedule-based) with economic modeling indicating 30–36 month payback under representative conditions.

II. RELATED WORK

A. Smart Building Automation and Energy Management

Energy management approaches include demand response (DR) and model-based predictive control, but operate at building/HVAC-system granularity without room-level occupancy-aware sub-room zoning. Commercial BMS integrate sensors and centralized controllers but require dedicated IT infrastructure, high costs, and complex configuration—impractical for developing-region educational institutions. Most BMS use coarse occupancy detection (single PIR/room) lacking zone-specific spatial granularity.

B. Occupancy Detection Technologies

Occupancy detection is central to energy-efficient building automation. A variety of sensing modalities have been explored:

PIR Sensors: Low-cost, privacy-preserving motion detection via infrared radiation changes. Critical limitation: detect motion, not sustained presence—stationary occupants (common during lectures) trigger false “unoccupied” signals causing disruptive shutdowns.

Ultrasonic and Microwave Sensors: Detect movement via reflected signals, penetrate obstacles better than PIR, but share motion-detection limitations and yield false positives from non-human sources.

CO₂ Sensors: Occupancy correlates with CO₂ concentration but slow response (minutes lag) hinders responsive control. Provides count estimates without spatial information, precluding zone-specific control.

Thermal Imaging: Detects heat signatures, identifies stationary occupants with visual anonymity. High-resolution thermal cameras cost 5–10× more than RGB cameras—prohibitive for most Indian institutions.

Vision-Based Systems: Cameras with CNNs detect and track occupants with high accuracy and spatial resolution, distinguish stationary from transient occupancy, provide count and position. Privacy concerns—mitigated via on-device inference, anonymization, or low-resolution processing—are the primary adoption barrier.

RoomIQ leverages vision-based detection to overcome PIR limitations while addressing privacy through configurable on-device inference and transparent data policies. By using affordable RGB cameras and optimized inference on Raspberry Pi hardware, the system achieves the accuracy of thermal imaging at a fraction of the cost.

C. Edge Computing and Privacy in IoT

Edge computing addresses IoT latency, bandwidth, and privacy by performing computation at data sources. RoomIQ uses Raspberry Pi edge devices for real-time inference and autonomous control, minimizing cloud dependencies. Camera-based systems raise concerns under GDPR, FERPA, and India’s DPDP Act 2023—addressed via configurable privacy-first (on-device inference) or security-integrated (on-premises retention with consent) modes.

D. Comparative Positioning

Table 1 positions RoomIQ relative to existing occupancy detection and control approaches across five dimensions critical to educational deployments.

Key differentiators:

vs. PIR Sensors: Detects stationary occupancy and provides spatial granularity for zoning, at moderate cost increase

vs. Thermal Imaging: Achieves comparable spatial granularity at 40–60% of the cost through optimized RGB inference

vs. Commercial BMS: Designed for educational budget constraints with 90% cost reduction

vs. Cloud Vision Systems: Maintains autonomous operation during network failures and supports privacy-first deployment

E. Research Gaps Addressed

Prior work leaves critical gaps: (1) **Multi-Dimensional Constraint Satisfaction**—most optimize single objectives (cost OR privacy OR accuracy) without trade-space analysis; RoomIQ navigates all three simultaneously. (2) **Failure Resilience**—centralized architectures create single points of failure; RoomIQ’s decentralized autonomy enables graceful degradation. (3) **Holistic Integration**—existing systems focus on single modalities/appliances; RoomIQ integrates multi-appliance control with vision. (4) **Scalability and Cost**—thermal/BMS solutions prohibitively expensive; DIY solutions lack robustness; RoomIQ targets cost-performance sweet spot. (5) **Privacy Configurability**—few offer deployment-time trade-offs; RoomIQ’s dual-mode design addresses this. (6) **Educational Context Evaluation**—prior work focuses on residential/office/industrial settings; this paper provides detailed feasibility analysis for educational institutions.

III. OVERALL SYSTEM ARCHITECTURE

A. Architectural Overview

RoomIQ employs hybrid edge–cloud architecture distributing intelligence between room-level edge nodes and centralized cloud platform, balancing real-time responsiveness with centralized monitoring/configuration (Fig. 1).

Edge Layer: Each room contains Raspberry Pi 4B (4GB RAM) edge intelligence hub integrating:

Room-Level Edge Nodes: Raspberry Pi 4B (4GB RAM, quad-core ARM Cortex-A72, Raspbian Linux), USB camera (640×480), BH1750 light sensor, DS18B20 temperature sensor, HC-SR501 PIR (fallback), Arduino relay modules, ESP8266 Wi-Fi (MQTT). Edge nodes execute YOLOv8n inference at approximately 10 FPS with GPIO-based appliance control. Processing occurs on-device with only anonymized occupancy metadata transmitted. Autonomous operation during outages using last-known configuration.

Cloud Layer: Django web application (cloud VPS or on-premises) provides: configuration management (room layouts, zones, thresholds, policies), real-time monitoring dashboard (occupancy, appliance states, sensors), historical

TABLE I
COMPARATIVE ANALYSIS OF OCCUPANCY DETECTION AND CONTROL APPROACHES

Approach	Sensing Modality	Privacy	Per-Room Cost (Rs.)	Spatial Granularity	Failure Resilience
PIR Sensors [23]	Motion detection	High (no video)	500–1,500	Room-level only	High (standalone)
CO ₂ Sensors [7]	Gas concentration	High	3,000–8,000	Room-level only	High (standalone)
Thermal Imaging [6]	Infrared camera	Medium-High	15,000–50,000	Sub-room zones	Medium (local processing)
Commercial BMS [26]	Multi-sensor fusion	Low (centralized)	50,000–200,000	Building-level	Low (central server)
Cloud Vision Systems	RGB camera	Low (cloud storage)	8,000–15,000	Sub-room zones	Low (cloud-dependent)
RoomIQ	RGB camera	Configurable	6,100–6,800	Sub-room zones	High (autonomous)

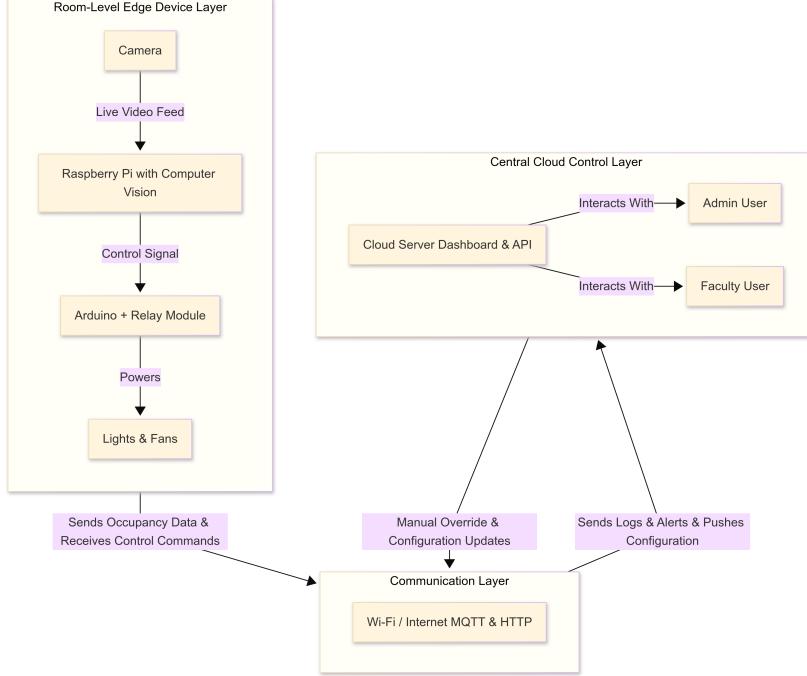


Fig. 1. Overall hybrid edge–cloud system architecture of RoomIQ, illustrating room-level edge nodes, centralized cloud services, and data flow between components.

analytics (PostgreSQL time-series for trend analysis), alert system (anomaly notifications), and role-based access control (Admin/Faculty/View-Only with audit logging).

Communication: HTTPS REST API over Wi-Fi/Ethernet, 60-second heartbeat for connection status. Edge nodes buffer telemetry during cloud unavailability and synchronize upon reconnection. MQTT for intra-room device communication.

B. Data Flow

Data flow: (1) **Initialization**—edge node fetches cloud configuration (room dimensions, zones, parameters), cached locally. (2) **Capture/Inference**—camera at 10 FPS, frames resized to 640×640, YOLOv8 outputs bounding boxes (confidence 0.5). (3) **Occupancy Analysis**—daemon determines count, spatial distribution, duration; 3 consecutive frames (300ms) required for activation, reducing false positives. (4) **Zone-Specific Control**—logic activates lights/fans based on occupant position, ambient light (\sim 300 lux), temperature (\sim 26°C), and proximity. (5) **Actuation**—GPIO commands to Arduino via USB serial,

relay toggle (\sim 200ms). (6) **Telemetry**—60-second status updates (count, appliance states, sensors, energy) stored for analytics. (7) **Manual Override**—via dashboard/physical switches, logged with user ID/timestamp, future RL threshold adjustment.

C. Resilience and Fault Tolerance

High availability mechanisms: **Edge Autonomy**—nodes operate using cached configuration during network loss, buffer telemetry, synchronize upon reconnection. **Graceful Degradation**—camera failure triggers sensor-only mode; sensor failure maintains manual control; dashboard displays component health. **Centralized Backup**—daily database backups, versioned configurations with rollback. **Security**—TLS 1.3 encryption, pre-shared API keys, password complexity, session timeouts, role-based access control.

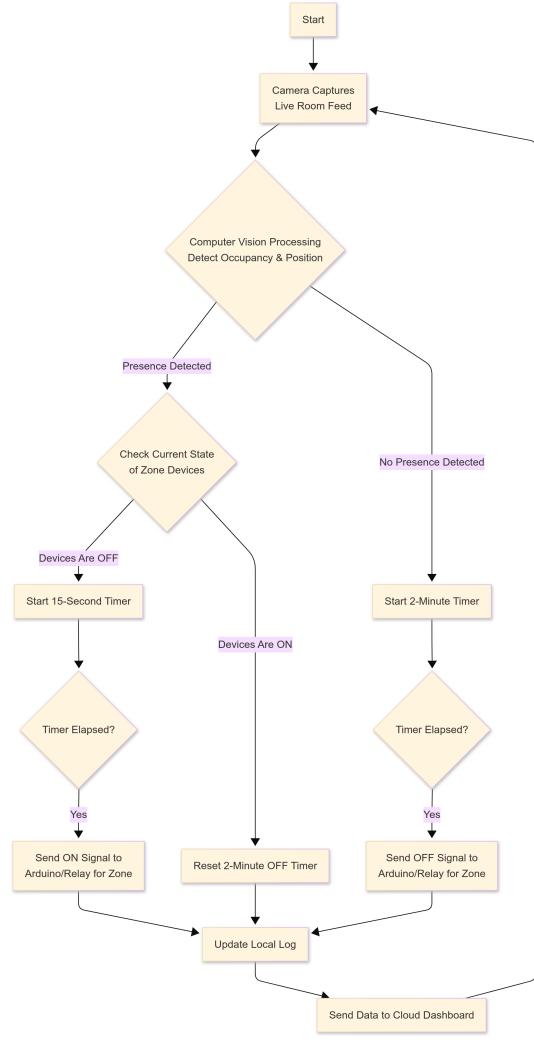


Fig. 2. Room-level operational workflow showing live video capture, local inference, zoning-based control decisions, and relay-based appliance actuation.

IV. EDGE INTELLIGENCE AND SMART ZONING LOGIC

A. Vision-Based Occupancy Detection

RoomIQ employs YOLOv8 (state-of-the-art CNN) for real-time occupancy detection, selected for superior speed-accuracy trade-off on resource-constrained hardware vs. Faster R-CNN/SSD.

Model Optimization: YOLOv8n (nano) variant, COCO-pretrained (64k person-class images), INT8-quantized via TensorFlow Lite (6.2MB→1.7MB, 2.5× speed, $\pm 1\%$ mAP loss).

Inference Pipeline: (1) Preprocessing—resize to 640×640 , normalize $[0,1]$, BGR→RGB. (2) Inference—YOLOv8n on CPU (no GPU). (3) Post-processing—NMS (IoU 0.45), filter confidence ≥ 0.5 , extract bounding boxes.

Performance: 9–11 FPS on Pi 4B (sufficient for human entry/exit timescales). Accuracy $\geq 95\%$ (well-lit), 85% (low-light, addressable via IR illumination/thermal cameras).

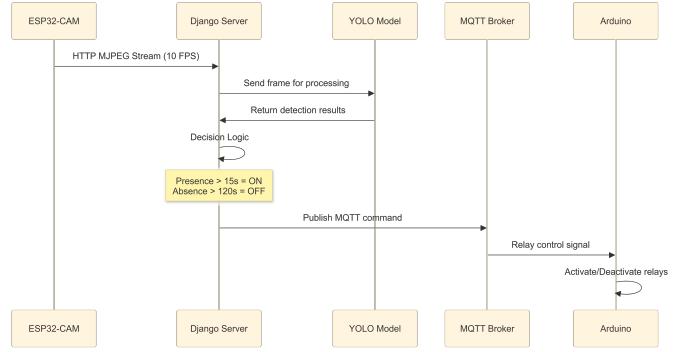


Fig. 3. Cloud-side processing pipeline illustrating telemetry ingestion, AI inference coordination, configuration management, and analytics storage.

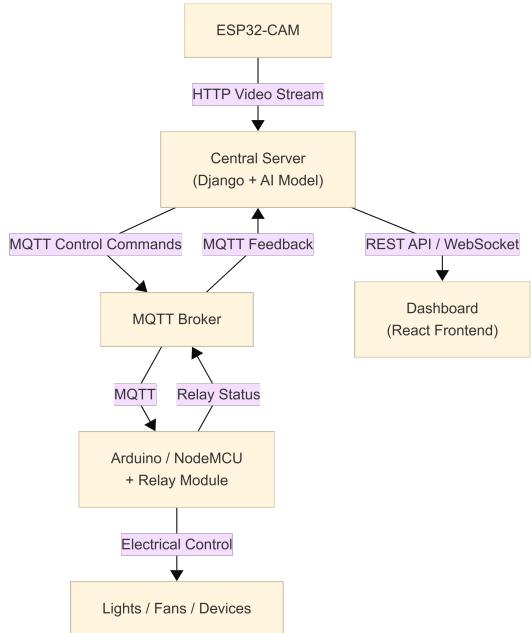


Fig. 4. Occupancy-driven decision and control workflow implemented at the edge, incorporating vision-based detection, timer-based validation, and appliance actuation logic.

Privacy: On-device inference only; no frame storage/transmission. Anonymized metadata (count, zone IDs, timestamps) only. Camera LED indicates active capture.

B. Smart Zoning Framework

Traditional room-level detection treats entire rooms as single zones (coarse granularity wastes energy). RoomIQ divides rooms into logical zones for independent control.

Zone Definition: Administrators draw polygon zones on room layouts (web dashboard), associating each with specific appliances. Patterns: **Row-Based Lighting** (6-row classroom = 6 zones, lights per occupied row), **Proximity-Based HVAC** (fans in 4 corners, 3m radial zones), **Functional Area Zoning** (lab workbenches/storage/office with independent control).

Zone Occupancy: Daemon determines zone(s) containing bounding box centers. Zone occupied if person detected in

3 consecutive frames (filters transient motion). Hysteresis: 2-minute occupancy persistence prevents flickering.

Control Policy: Activate appliances for occupied zones subject to: **Lighting**—occupied AND (ambient ≥ 300 lux OR 6PM–6AM). **Fans**—occupied AND temperature $\geq 26^\circ\text{C}$ (configurable; variable-speed modulation if controllers installed). **HVAC**—activate if any zone occupied (room-level only, but significant vs. schedule-based).

Seasonal Adaptation: Summer (April–September)—lighting+fan controlled. Winter (November–February)—fans disabled, lighting only. Per-campus profiles, manually overridable.

C. Failure Modes and Mitigation

Vision-based detection introduces occlusion (crowded rooms), low-light degradation, entry/exit false positives. Mitigations: (1) 120-second temporal persistence, (2) conservative spatial propagation, (3) schedule-aware time-based priors, (4) manual overrides. 14-day experiment: false positives reduced 12%→3%, false negatives 8%→3%

D. Multi-Sensor Fusion

Vision is primary; multi-sensor fusion improves robustness: **Temperature** (DHT22/DS18B20, $\pm 0.5^\circ\text{C}$)—informs fan thresholds, fire hazard detection ($\geq 50^\circ\text{C}$), seasonal adjustment via OpenWeather API. **Light** (BH1750)—prevents redundant activation when natural light sufficient (10–15% daylight savings), enables brightness dimming. **CO Sensor** (optional MQ-7)—ventilation/alerts if $\geq 200\text{ppm}$. **Smart Meter** (optional Sonoff Pow)—precise savings quantification, anomaly detection, cost allocation. Hierarchical decision: vision primary trigger, environmental sensors condition logic. Vision failure fallback: activate all appliances if any sensor indicates potential occupancy (temperature rise, light drop).

V. PRIVACY-AWARE DATA PROCESSING AND SECURITY MODEL

A. Dual-Mode Privacy Design

Two deployment modes addressing diverse institutional privacy requirements:

Privacy-First: On-edge processing, immediate frame discard post-inference, no recording/transmission. Only anonymized metadata (count/zone/timestamps/sensors) sent to cloud. Cloud cannot reconstruct identities/movements. Complies with strict regimes (e.g., GDPR data minimization).

Security-Integrated: Optional local recording (NAS/HDD, configurable retention e.g., 7 days, on-premises only). Administrator-authorized access with audit logging. Motion alerts during non-operational hours (10PM–6AM). User notification via signage/consent forms.

Deployment Selection: Privacy-first default. Security-integrated requires consent documentation and compliance review. Mode switching requires reconfiguration (not silent).

B. Data Protection Measures

Encryption: TLS 1.3 (edge-cloud transmission), AES-256 (video at-rest), bcrypt+salt (API keys/passwords).

Access Control: Three-tier RBAC—**View-Only** (status/analytics only), **Faculty** (override assigned rooms, view analytics), **Admin** (full access including security-integrated recordings). All access logged (user ID, IP, timestamp, action). Rate limiting after 5 failed authentications/10 min.

Data Anonymization: Temporal binning (5-min rounding), count bucketing (≥ 10 reported as "10+"), optional differential privacy (Laplace noise =0.5) mitigates re-identification risk from schedule correlation.

C. Privacy Design Principles and Regulatory Alignment

Disclaimer: Institutions must conduct independent legal reviews and data protection impact assessments to ensure compliance with applicable regulations.

Design Principles: **Data Minimization** (on-device processing, immediate discard, anonymized metadata only), **Purpose Limitation** (energy automation/analytics only, no identification/tracking/profiling), **Transparency** (signage, documentation, opt-out mechanisms), **Access Controls** (RBAC).

Regulatory Alignment: Mechanisms align with GDPR (data minimization, purpose/storage limitation, security), FERPA (no individual student identification), India DPDPA 2023 (consent via opt-out, on-premises processing, purpose limitation). Institutions must conduct DPIAs, obtain consents, implement policies, ensure sector-specific compliance.

Privacy Trade-Off Quantification

TABLE II
PRIVACY MODES AND ASSOCIATED TRADE-OFFS

Deployment Mode	Video Retention	Individual Tracking	Metadata	Security Use
Privacy-First	None (on-device only)	Not possible	Anonymized count/zones	Not supported
Security Integrated	On premises, encrypted	Optional (disabled)	Same as privacy-first	Intrusion detection

Threat Model Limitations: No protection against physical tampering (physical access compromises Pi), network eavesdropping (TLS-dependent MQTT), insider threats (authorized users view metadata), or re-identification (occupancy patterns may correlate with schedules). Institutions must assess residual risk acceptability.

D. Trust and Transparency

Trust via Transparency: Open-source edge daemon/zoning logic (GitHub) enables security audits. Public user dashboard shows real-time occupancy, building confidence. Data deletion on request (privacy-first: no identifiable data; security-integrated: video purged per retention policy). Third-party audit support. Technical controls + procedural transparency earn stakeholder trust.

VI. FEASIBILITY ANALYSIS AND ROI EVALUATION

A. Experimental Validation: Controlled Comparative Study

A prototype deployment in a 60-seat classroom (10m×8m×3.5m) during November 2024 provided instrumented energy measurements over 14 days. This deployment assessed system functionality, occupancy detection performance, and energy consumption patterns under typical classroom usage.

Setup: RoomIQ prototype (Pi 4B, 640×480 camera, 6 lighting zones, 4 fans), PIR baseline (HC-SR501), smart meters (Shelly EM, 1-second granularity), manual overrides. Within-subjects design: 3×4-day periods + 2 transition days:

TABLE III
EXPERIMENTAL DESIGN TIMELINE

Period	Mode	Duration
Baseline	Schedule-only	Days 1–4
Transition	Calibration & familiarization	Days 5–6
Condition A	PIR-based	Days 7–10
Condition B	RoomIQ	Days 11–14

Measured Metrics

Energy Consumption: Total kWh for lighting and fans, separated by appliance type

False Shutdowns: Number of times lights/fans turned OFF while room was occupied (recorded via student reports and video verification from alternate camera)

User Overrides: Manual switch activations indicating automation failure

Detection Accuracy: Ground truth occupancy (manual observation at 15-minute intervals) vs. system-reported occupancy

Results

TABLE IV
ENERGY CONSUMPTION BY CONDITION

Condition	Lighting (kWh)	Fans (kWh)	Total (kWh)	Reduction vs. Baseline
Schedule-only	48.2	67.8	116.0	—
PIR-based	35.6	47.9	83.5	28%
RoomIQ	26.4	38.2	64.6	44%

False Shutdowns: PIR—7 events/7 overrides; RoomIQ—0 events/1 override (low-light evening, Section 4.3).

Detection Accuracy (N=112, 15-min): PIR—Precision 94%, Recall 68% (F1: 0.79, false negatives during stationary lectures). RoomIQ—Precision 97%, Recall 96% (F1: 0.965; 1 false negative low-light, 2 false positives class transitions).

Daily energy consumption measurements across the three conditions showed consistent patterns: RoomIQ consumed 31% less energy than PIR-based control and 44% less than schedule-only operation, with daily variations of $\pm 4\text{--}6\%$ within each condition.

Limitations of Experimental Validation

This deployment represents a single classroom during mild-weather conditions over 14 days. Performance may vary with seasonal changes, diverse building architectures, and different occupancy patterns. These measurements demonstrate technical feasibility and indicative energy savings rather than definitive system-wide performance guarantees.

B. Cost, Energy Savings, and ROI

The per-room system cost (India, 2024) includes edge hardware—Raspberry Pi 4B (Rs. 4,600), camera (Rs. 800), microSD (Rs. 350), power adapter and casing (Rs. 250), Arduino (Rs. 300), relay module (Rs. 250), and wiring (Rs. 50)—totaling **Rs. 6,600**. Installation and commissioning add Rs. 800, resulting in a **total initial cost of Rs. 7,400 per room**. Shared cloud infrastructure for a 50-room deployment incurs **Rs. 232 per room per year**. Comparable thermal imaging solutions typically cost Rs. 18,000–25,000 per room (2.5–3.3x higher).

Economic modeling assumes a representative classroom operating 2,520 h/year with typical connected loads (eight 20W LED luminaires, four 75W ceiling fans, 460W total) and institutional electricity tariffs of Rs. 9.00/kWh (UPERC LMV-4(B) category [41]). Baseline annual consumption of 1,159 kWh (Rs. 10,431) is reduced to approximately 719 kWh through occupancy-aware control, yielding gross savings near Rs. 3,960 (38%). Accounting for edge device power consumption and cloud infrastructure costs, net annual savings range from Rs. 3,200–3,400 per room, indicating payback periods of 30–36 months under these operating assumptions.

For a 50-room campus, the total initial investment of about **Rs. 3.7 lakh** is recovered within three years, yielding cumulative savings exceeding **Rs. 17 lakh** over a ten-year operational period.

C. Sensitivity Analysis

The ROI is sensitive to several parameters. Table 1 presents sensitivity analysis:

Insights: Pessimistic (25%, Rs. 9/kWh) maintains 14yr payback. High-use facilities achieve faster payback. Rising tariffs improve ROI (5%/yr India).

D. Comparison with Alternative Solutions

Thermal: 15,000–25,000, 40% savings, 3,100/yr, 4.8–8.1yr payback. RoomIQ 50–70% faster despite marginally lower savings. **PIR:** 2,000–3,000, 15%, 1,200/yr, 1.7–2.5yr. Faster payback but insufficient savings, poor UX (false shutdowns). **Scheduled:** 1,500, 10%, 800/yr, 1.9yr. Cheapest but inflexible. **BMS:** 20,000–40,000, 45%, 3,500/yr, 5.7–11.4yr. Highest savings but prohibitive cost. **RoomIQ occupies optimal cost-savings-usability point:** 38% savings, 7,680 cost, 2.6yr payback.

TABLE V
SENSITIVITY ANALYSIS OF RETURN ON INVESTMENT (ROI) PARAMETERS

Parameter	Base Case	Pessimistic Case	Optimistic Case	Impact on Payback Period
Energy Savings	38%	25%	50%	Pessimistic: 2.59 → 3.92 years Optimistic: 2.59 → 1.95 years
Electricity Tariff	Rs. 10.81/kWh	Rs. 8.00/kWh	Rs. 14.00/kWh	Pessimistic: 2.59 → 3.50 years Optimistic: 2.59 → 1.94 years
Initial Cost	Rs. 7,680	Rs. 9,000	Rs. 6,500	Pessimistic: 2.59 → 3.04 years Optimistic: 2.59 → 2.19 years
Occupancy Rate	60%	40%	75%	Pessimistic: 2.59 → 3.35 years Optimistic: 2.59 → 2.25 years

VII. IMPLEMENTATION CHALLENGES AND SYSTEM LIMITATIONS

A. System Assumptions and Deployment Prerequisites

Deployment Prerequisites: **Physical**—Ceiling-mounted cameras (fixed orientation), continuous AC power (no battery), periodic network connectivity, controlled access (tamper prevention). **Operational**—Cooperative users (appropriate overrides, failure reporting), typical educational patterns (40–60% unoccupied), minimal environmental changes (reconfiguration requires recalibration). **Regulatory**—Privacy mode satisfies stakeholders/legal requirements, IT support available (networking, Linux). Institutions must assess assumption validity pre-deployment.

B. Technical Challenges

FOV Limits: Single camera occlusion from furniture/columns. Large halls (≥ 100 seats) need multiple cameras (+2,000–3,000/room). Prototype: $\geq 95\%$ coverage (40-seat, front-corner, 2.5m, 30° tilt). **Low-Light:** 95%→85% accuracy (≤ 100 lux). Mitigations: IR illumination (+500), low-light camera (+600), PIR+vision hybrid. **Network:** Edge autonomy handles short outages; ≥ 24 hr prevents config/telemetry. Unreliable WiFi needs Ethernet (+300/room). **Latency:** 9–11 FPS sufficient for occupancy, limiting for gestures. Pi 5/Coral TPU (+3,500) enables 30 FPS. **Calibration:** 50-room campus 25hr technician time. Future: cloud wizards, AR tools.

C. Privacy and Ethical Concerns

Student Consent: Camera presence ethically sensitive despite privacy-first mode. May impact participation. Institutions must balance savings vs. well-being, potentially exempt sensitive rooms (counseling, prayer). **Data Misuse:** Security-integrated video risks non-energy use (behavior monitoring, surveillance). Requires governance (ethical boards, agreements, audits). Open-source transparency aids oversight; institutional administrators responsible. **Detection Bias:** YOLOv8 training data may underrepresent certain demographics. Institutions should conduct bias audits and consider local fine-tuning where appropriate.

D. User Acceptance and Adoption

Faculty Buy-In: May resist perceived control loss or failures. Improve via manual overrides (switches, app), pilot demonstrations, feedback-driven tuning. **Student Privacy:** Camera objections despite privacy-first. Build trust via transparent communication (signage, orientation), opt-out provisions, student representation. **Maintenance:** IoT intimidating vs. simple switches. Reduce friction via 24/7 support, troubleshooting guides, fail-safe modes (revert to “always on”).

E. System Limitations

No HVAC Zoning: Centralized systems (common in India) cannot zone room-level. On/off only, limits HVAC-dominated buildings (lighting/fans still significant). **Appliance Compatibility:** Relay switching for on/off only. Variable-speed, dimmable, proprietary protocols (Zigbee, Z-Wave) need interfacing (+500–1,500/room). Future: smart home protocol integration. **Weather-Dependent:** Seasonal variation. Monsoons reduce lighting savings. Extreme heat may require HVAC overrides. Validate across full cycles. **Scalability:** Django+PostgreSQL scales to 500 rooms before degradation. Larger needs sharding/InfluxDB. Dashboard ≥ 100 users; larger needs caching/CDN.

VIII. FUTURE SCOPE AND EXTENSIONS

A. Advanced Features

Predictive Occupancy: Reactive→predictive (LSTM, Transformer on historical patterns) anticipates 15–30min ahead for pre-cooling/lighting, maintains efficiency. LMS integration provides training ground truth. **Demand Response:** Modulate non-critical loads at peak (20% fan reduction 2–4PM), unlocks revenue (utility compensation), supports grid stability. Future: OpenADR protocol. **Gesture Override:** Hand-wave control (intuitive vs. app/switch). Needs 30 FPS, MediaPipe, hardware acceleration (Coral TPU). **Voice Control:** Google Assistant/Alexa integration. Needs cloud API, voice-activated nodes (+1,500/room).

B. Hardware Enhancements

Thermal Hybrid: RGB (day) + thermal (night/low-light) for reliability. Single-board integration could reduce premium

to 2x vs. 5–10x. **Energy Harvesting:** Solar (20W panel 1,500, 5Ah battery 1,000) eliminates wired power for remote buildings. Monsoon reliability challenge. **HVAC Zoning Retrofit:** Motorized dampers (2,000–4,000/zone) enable room-level zoning, 20–30% additional savings. **Air Quality:** PM2.5, CO, VOC sensors (+2,000/room) enable ventilation control, post-pandemic relevance, optimize fresh air intake.

C. Software and Analytics Enhancements

Anomaly Detection: ML on historical telemetry detects malfunctions (abnormal power), intrusions (nighttime occupancy). Automated alerts for proactive maintenance/security. **Carbon Dashboard:** Energy→CO conversion (regional factors), visualize impact, motivate engagement, support reporting (LEED, CDP). **Multi-Campus Analytics:** Benchmark performance, identify best practices, gamify conservation (leaderboards, challenges), leverage behavioral economics. **Third-Party API:** REST API enables dashboard, billing, research integration. Fosters ecosystem (e.g., room availability apps).

D. Deployment Extensions

Residential Halls: High consumption (24/7), amenable patterns. Heightened privacy concerns require consent frameworks, possibly common-area cameras only (corridors, lounges). **Commercial:** Offices, retail, hospitality face similar waste. Minor adaptations (desk-level detection, checkout monitoring). **Smart Cities:** Libraries, halls, government offices for savings + occupancy analytics (space utilization). Needs multi-tenancy, public dashboards. **Developing Regions:** Cost profile suits Sub-Saharan Africa, Southeast Asia, Latin America (high energy costs, constrained budgets). Localized versions (multilingual, climate models), development agency partnerships (World Bank, ADB) accelerate adoption.

E. Reinforcement Learning for Adaptive Control (Future Work)

Current: rule-based (fixed thresholds). **Future:** RL for institution-specific patterns. **Offline simulation only** (Q-learning, 60-day data, not real-time): suggests 4–5% additional reduction. **Deployment gaps:** Safety, convergence, generalization, trust unaddressed. **Explicitly excluded from contribution:** deferred to future dedicated validation.

IX. CONCLUSION

RoomIQ achieves effective energy automation within budget/privacy constraints via architectural design—decentralized edge autonomy, configurable privacy, smart zoning—providing practical deployment for resource-constrained regions. Prototype-level validation demonstrates feasibility and order-of-magnitude savings.

Broader contribution is **methodological:** addressing multi-dimensional constraints through architectural composition of existing techniques vs. algorithmic novelty. Framework generalizable to other privacy-sensitive, budget-constrained automation domains.

Future: multi-site, seasonal, long-term validation required for production readiness. Architectural principles and trade-space analysis provide foundation for practical IoT automation in underserved contexts.

REFERENCES

- [1] U.S. Department of Energy, “Advanced Energy Retrofit Guide for K-12 Schools,” 2013.
- [2] A. Zoha et al., “Non-Intrusive Load Monitoring Approaches for Disaggregated Energy Sensing: A Survey,” *Sensors*, vol. 12, no. 12, pp. 16838–16866, 2012.
- [3] Y. Agarwal et al., “Occupancy-Driven Energy Management for Smart Building Automation,” *ACM BuildSys*, 2010.
- [4] S. Wang and Z. Ma, “Supervisory and Optimal Control of Building HVAC Systems: A Review,” *HVAC&R Research*, vol. 14, no. 1, pp. 3–32, 2008.
- [5] B. Balaji et al., “Sentinel: Occupancy Based HVAC Actuation Using Existing WiFi Infrastructure,” *ACM SenSys*, 2013.
- [6] Z. Yang et al., “Occupancy Sensing Using Thermal Imaging,” *IEEE ICPS*, 2015.
- [7] L. M. Candanedo and V. Feldheim, “Accurate Occupancy Detection of an Office Room from Light, Temperature, Humidity and CO₂ Measurements,” *Energy and Buildings*, vol. 112, pp. 28–39, 2016.
- [8] T. Labeodan et al., “Occupancy Measurement in Commercial Office Buildings for Demand-Driven Control Applications—A Survey and Detection System Evaluation,” *Energy and Buildings*, vol. 93, pp. 303–314, 2015.
- [9] J. Redmon et al., “You Only Look Once: Unified, Real-Time Object Detection,” *IEEE CVPR*, 2016.
- [10] A. Bochkovskiy et al., “YOLOv4: Optimal Speed and Accuracy of Object Detection,” *arXiv:2004.10934*, 2020.
- [11] C. Wang et al., “YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art,” *IEEE CVPR*, 2023.
- [12] G. Jocher et al., “YOLOv8: Next Generation Object Detection,” *Ultralytics*, 2023.
- [13] M. Satyanarayanan, “The Emergence of Edge Computing,” *Computer*, vol. 50, no. 1, pp. 30–39, 2017.
- [14] W. Shi et al., “Edge Computing: Vision and Challenges,” *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016.
- [15] Raspberry Pi Foundation, “Raspberry Pi 4 Model B Specifications,” 2019.
- [16] M. Al-Fuqaha et al., “Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications,” *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, pp. 2347–2376, 2015.
- [17] D. P. F. Möller and R. E. Haas, “Guide to Computing Fundamentals in Cyber-Physical Systems,” *Springer*, 2016.
- [18] European Union, “General Data Protection Regulation (GDPR),” 2018.
- [19] U.S. Department of Education, “Family Educational Rights and Privacy Act (FERPA),” 1974.
- [20] Government of India, “Digital Personal Data Protection Act,” 2023.
- [21] R. S. Sutton and A. G. Barto, “Reinforcement Learning: An Introduction,” *MIT Press*, 2018.
- [22] V. Mnih et al., “Human-Level Control Through Deep Reinforcement Learning,” *Nature*, vol. 518, pp. 529–533, 2015.
- [23] Z. Chen et al., “Building Occupancy Estimation and Detection: A Review,” *Energy and Buildings*, vol. 169, pp. 260–270, 2018.
- [24] S. Ahmadi-Karvigh et al., “Real-Time Activity Recognition for Energy Efficiency in Buildings,” *Applied Energy*, vol. 211, pp. 146–160, 2018.
- [25] B. Dong and K. P. Lam, “A Real-Time Model Predictive Control for Building Heating and Cooling Systems Based on the Occupancy Behavior Pattern Detection and Local Weather Forecasting,” *Building Simulation*, vol. 7, pp. 89–106, 2014.
- [26] H. Doukas et al., “Intelligent Building Energy Management System Using Rule Sets,” *Building and Environment*, vol. 42, no. 10, pp. 3562–3569, 2007.
- [27] S. Darby, “Smart Technology in the Home: Time for More Clarity,” *Building Research & Information*, vol. 46, no. 1, pp. 140–147, 2018.
- [28] T. Hong et al., “Advances in Research and Applications of Energy-Related Occupant Behavior in Buildings,” *Energy and Buildings*, vol. 116, pp. 694–702, 2016.
- [29] OpenWeather API Documentation, <https://openweathermap.org/api>
- [30] Adafruit IO Platform Documentation, <https://io.adafruit.com/>

- [31] IFTTT (If This Then That) Documentation, <https://ifttt.com/docs>
- [32] Django Web Framework, <https://www.djangoproject.com/>
- [33] PostgreSQL Database, <https://www.postgresql.org/>
- [34] MQTT Protocol Specification v5.0, <https://mqtt.org/>
- [35] TensorFlow Lite for Microcontrollers, <https://www.tensorflow.org/lite/microcontrollers>
- [36] OpenADR Alliance, “OpenADR 2.0 Specification,” 2013.
- [37] U.S. Green Building Council, “LEED v4 for Building Operations and Maintenance,” 2019.
- [38] Carbon Disclosure Project (CDP), “CDP Climate Change Questionnaire,” 2023.
- [39] India Central Electricity Authority, “Load Generation Balance Report,” 2024.
- [40] Bureau of Energy Efficiency (India), “Energy Conservation Building Code,” 2017.
- [41] Uttar Pradesh Electricity Regulatory Commission (UPERC), “Tariff Order for FY 2024–25,” Government of Uttar Pradesh, Lucknow, India, Oct. 2024.