ISSN No. 0976-5697

VOLUME 17 SPECIAL ISSUE 1, JULY 2026
International Journal of Advanced Research in Computer Science

RESEARCH PAPER

Available Online at www.ijarcs.info

Adapta: An Edge-Deployed Multimodal Smart Lighting System Using Computer Vision and Federated Machine Learning for Circadian-Aware Adaptive Illumination

Abhiraj Singh Panwar
Dept. of Computer Science & Engineering with specialization in Artificial Intelligence
Geetanjali Institute of Technical Studies
Udaipur, Rajasthan, India
abhirajpanwar110@gmail.com

Lavisha Gupta
Dept. of Computer Science & Engineering
Geetanjali Institute of Technical Studies
Udaipur, Rajasthan, India
lavishagupta378@gmail.com

Sarvagya Soni
Dept. of Computer Science & Engineering with specialization in Artificial Intelligence
Geetanjali Institute of Technical Studies
Udaipur, Rajasthan, India
12sarvagyasoni@gmail.com

Poshit Nagda
Dept. of Computer Science & Engineering
Geetanjali Institute of Technical Studies
Udaipur, Rajasthan, India
poshitn@gmail.com

Kanika Garg
Dept. of Computer Science & Engineering
Geetanjali Institute of Technical Studies
Udaipur, Rajasthan, India
kanika.garg@gits.ac.in

Abstract: Most smart lighting systems today still rely on passive infrared sensors that cannot distinguish between a person working at a desk and someone trying to rest. This paper introduces Adapta, a building-deployed smart lighting framework that uses an on-device RGB-D camera, a lightweight MobileNetV3-GRU hybrid network, and a federated machine learning pipeline to continuously recognise occupant activities and adjust both brightness and colour temperature of luminaires in real time. All inference runs locally on a Raspberry Pi 5 edge node; no raw video is ever transmitted to any external server. Illumination targets are computed from the Rea-Figueiro circadian stimulus (CS) model, so the system actively supports occupants biological day-night rhythms. Deployed across 120 volunteers in three building types over 90 days, Adapta achieved a 41.3 percent energy reduction over static baselines, 94.7 percent activity classification accuracy, and statistically significant improvements in daytime alertness and sleep quality. End-to-end actuation latency remained below 28 ms throughout the entire trial period.

Keywords: Smart Lighting; Computer Vision; Federated Learning; Circadian Rhythm; Human Activity Recognition; Edge AI; MobileNetV3-GRU; Neuromorphic Control; Energy Efficiency; Human-Centric Illumination

I. INTRODUCTION

Lighting is easy to take for granted. Flick a switch, the room brightens. That is more or less how most buildings still work, despite decades of research suggesting we could do substantially better. Globally, artificial lighting consumes somewhere in the neighbourhood of 15 percent of all electricity generated [1], and a disproportionate share of that is wasted on empty rooms or burning at full intensity for activities that demand far less.

Part of the problem is that conventional sensor-based automation only knows one thing about an occupant: whether or not they are present. A passive infrared (PIR) detector fires when someone walks through its field of view and goes quiet when they leave. That is genuinely useful, as occupancy-linked switching can trim energy bills by 20 to 30 percent with minimal disruption [4]. But it tells the controller nothing about whether the person is reading a dense report, exercising, holding a meeting, or trying to fall asleep. Each situation calls for a quite different lighting environment, and collapsing them into a single on/off decision is, at best, a rough approximation.

The biological dimension compounds the engineering one. Light does not simply let us see; it actively regulates physiology through intrinsically photosensitive retinal ganglion cells (ipRGCs) whose peak sensitivity sits around 490 nm, within the blue-white region of the visible spectrum [2]. These cells drive the daily oscillation of cortisol, melatonin, body temperature, and other systems governing alertness, mood, and sleep. Morning exposure to high-correlated-colour-temperature (CCT) light above 5000 K suppresses melatonin and raises alertness; the same light an hour before bed is counterproductive [3]. Commercial smart bulbs offer preset scenes, but they require occupants to activate them, which in practice rarely happens.

Adapta is designed to fill this gap autonomously. The specific contributions of this paper are as follows:

- A MobileNetV3-GRU hybrid architecture classifying nine distinct activity states from RGB-D streams in real time, achieving 94.7 percent accuracy with 21.4 ms inference time on a sub-10 W edge processor.

- A circadian stimulus mapping layer translating recognised activities into biologically grounded illumination profiles (illuminance, CCT, uniformity ratio, flicker index) using the Rea-Figueiro CS metric.
- A federated learning pipeline enabling the activity model to personalise to individual occupants without any raw video leaving the local device.
- A neuromorphic event-driven control interface that dispatches DALI-2 commands only when the model output distribution shifts meaningfully, reducing bus traffic by 73 percent versus fixed-rate polling.
- A 90-day real-world deployment across three building environments and 120 participants, with energy, accuracy, latency, and psychometric outcomes all measured and reported.

The remainder of the paper is organised as follows. Section II surveys related literature. Section III describes the system architecture. Section IV details the ML methodology and circadian model. Section V presents experimental results. Section VI addresses limitations and future work. Section VII concludes.

II. Related Work

A. *Occupancy-Driven and Daylight Harvesting Systems*

The dominant strand of lighting automation research has focused on occupancy detection paired with daylight harvesting. Singhvi et al. [4] demonstrated that rule-based threshold switching tied to PIR sensors could reduce lighting energy by roughly a quarter in typical office deployments. Zhang et al. [8] extended this with closed-loop photometric feedback that continuously adjusted output to maintain a setpoint illuminance, achieving a further 35 percent saving on top. Both approaches share a fundamental ceiling: they cannot optimise beyond the binary question of occupancy and know nothing about what the room is actually being used for.

B. *Computer Vision Approaches*

Camera-based systems offer a much richer window into room state. Erickson et al. [9] mounted fisheye cameras in ceiling tiles and applied Gaussian mixture model background subtraction for occupancy counting, reporting 29 percent HVAC energy reductions. Aghajan et al. [10] extracted skeleton keypoints to distinguish between activities and link classifications to lighting presets. The obvious concern is privacy: most prior systems stream raw video to a server for processing, which is a hard sell for occupants. Adapta addresses this by processing everything on-device and never transmitting a pixel beyond the local edge unit.

Deep learning has substantially outpaced classical computer vision for activity recognition [11]. MobileNetV2 [12] and EfficientNet [13] demonstrated viable accuracy-efficiency trade-offs for mobile deployment, though their designs did not target temporal classification tasks. The MobileNetV3-GRU hybrid proposed here is, to our knowledge, novel in the context of intelligent lighting control.

C. *Human-Centric Lighting*

The term Human-Centric Lighting (HCL) has become prominent in architectural research. Viola et al. [15] conducted a controlled trial in which office workers under blue-enriched white light reported better alertness, higher self-rated performance, and improved sleep quality compared to those under standard fluorescent light. Figueiro et al. [2] developed the Daysimeter CS metric as a biologically anchored measure of how much a given light exposure actually stimulates the circadian system, accounting for the melanopsin spectral sensitivity curve. Commercial products offer time-scheduled CCT variation, but these schedules are fixed and oblivious to actual room activity.

D. *Federated Learning in Building Systems*

Federated Learning was introduced by McMahan et al. [16] as FedAvg, a framework for training neural networks collaboratively without centralising raw data. Li et al. [19] proposed FedProx, adding a proximal penalty term to handle statistical heterogeneity across clients with different occupant populations. The building systems community has begun exploring FL for HVAC and thermal comfort applications [17, 18]. To the best of our knowledge, no prior work applies federated learning to personalised lighting control grounded in computer vision, which constitutes a primary novelty of this contribution.

III. System Architecture

Adapta is structured as four cooperating layers. From the physical hardware upward: a Perception Layer handles sensing and preprocessing; an Intelligence Layer runs the activity classification model; a Control Layer translates model outputs into luminaire commands; and a Federated Learning Layer enables personalisation while enforcing strict data locality.

A. *Hardware at Each Room Node*

Each room is equipped with a Raspberry Pi 5 (8 GB RAM) as the edge processor. An Intel RealSense D435i RGB-D camera provides colour and depth frames at 30 fps at 640 x 480 resolution. Ambient conditions are monitored by a Sensirion SHT40 temperature and humidity sensor and a TSL2591 high-dynamic-range lux meter. A Melexis MLX90640 far-infrared thermal array provides supplementary occupancy confirmation. Luminaire control runs over a DALI-2 bus supporting 16-bit dimming resolution and continuous CCT adjustment from 1800 K to 6500 K. Total per-node hardware cost was approximately INR 28,400 at time of assembly.

No raw video frame is ever transmitted beyond the local edge device. The only data crossing a network boundary are compressed, differentially private gradient updates, transmitted only during a nightly synchronisation window when occupants are typically absent from the space.

B. *Perception Layer*

The Perception Layer runs three concurrent processes. First, a DBSCAN clustering algorithm (epsilon = 0.4 m, minPts = 15) applied to the depth point cloud counts and localises occupant clusters in 3 to 4 ms and acts as a gate: if no clusters are detected, the system holds a low-power standby state and skips the heavier activity recognition pipeline. Second, when occupancy is confirmed, the RGB stream feeds the activity classifier. Third, the thermal array runs a parallel binary presence check for cases where an occupant is very still such as reading or sleeping.

C. *Intelligence Layer*

The Intelligence Layer hosts the trained MobileNetV3-GRU model, which takes a sliding window of $T = 16$ consecutive frames (approximately 0.53 seconds at 30 fps) and outputs a probability distribution over nine activity classes: Concentrated Work, Casual Reading, Video Conferencing, Physical Exercise, Social Interaction, Resting or Napping, Eating, Whiteboard Presentation, and Absent. The activity with highest posterior probability passes to the Control Layer, along with the full softmax vector used by the event-triggering logic described next.

D. Control Layer: Event-Driven Actuation

Rather than polling the model output at a fixed rate, Adapta uses an event-driven approach inspired by temporal coding in spiking neural networks. A DALI-2 command is dispatched only when the Kullback-Leibler divergence between the current softmax distribution and the one from the previous dispatch exceeds a calibrated threshold of delta-KL greater than 0.12 nats. Below that threshold, the model may detect momentary movement but the dominant activity has not changed meaningfully, so there is no reason to update the lights. This scheme issued 73 percent fewer DALI bus commands than a naive 1 Hz polling design with no perceptible difference in lighting responsiveness from the occupants perspective.

E. Federated Learning Layer

Overnight, each room node participates in a federated training round coordinated by a building-level server (Intel NUC 13 Pro). Occupants provide activity label corrections via a small touchscreen panel near the room entrance in a 90-second end-of-day interaction. The federation algorithm is FedProx [19] with proximal term mu = 0.01, chosen after comparing it with standard FedAvg on a simulated heterogeneity scenario. Each node runs 3 local gradient steps per round. Gradients are clipped to L2 norm C = 1.0 and perturbed with Gaussian noise calibrated to (epsilon = 1.2, delta = 10e-5) differential privacy before upload. Convergence was reached within 14 communication rounds across all buildings.

IV. Methodology

A. MobileNetV3-GRU Activity Classifier

The choice of base architecture involved a genuine trade-off analysis. We evaluated MobileNetV3-Small, MobileNetV3-Large, EfficientNet-B0, and ResNet-18 as spatial feature extractors, each paired with a bidirectional GRU of 128 hidden units as the temporal aggregator. Running on the Raspberry Pi 5 with INT8 quantisation, MobileNetV3-Small was the only configuration clearing the 28 ms end-to-end latency budget while maintaining accuracy above 93 percent. Training used ImageNet-pretrained weights fine-tuned for 60 epochs using AdamW with initial learning rate 2e-4, cosine annealing, and weight decay 1e-4.

After fine-tuning, the model was post-training quantised to INT8 using TensorFlow Lite, reducing the file from 4.9 MB to 1.3 MB and accelerating on-device inference 2.2 times via NEON SIMD instructions with only 0.4 percentage points accuracy loss. Knowledge distillation [22] using EfficientNet-B2 as teacher recovered 1.8 percentage points of accuracy gap without changing the model size or speed.

B. The Adapta-Activity Dataset

Existing public datasets such as NTU RGB+D, HMDB-51, and Kinetics cover broad activity ranges under controlled studio conditions unsuitable for real building lighting contexts. We therefore built and will publicly release the Adapta-Activity dataset, collected from June through September across seven rooms in three building types. Forty-two volunteers participated under ethics approval (IIT Jodhpur IRB 2024-0047). Sessions covered morning, afternoon, and evening slots across all seasons. Only skeleton keypoints extracted by MediaPipe BlazePose were archived, so the dataset contains no personally identifiable visual information.

Table I. **Adapta-Activity Dataset: Class Distribution**

Activity Class	Clips	Avg Dur (s)	Participants	Environments
Concentrated Work	4,820	142.3	38	Office, Residential
Casual Reading	3,612	118.7	35	All three
Video Conferencing	2,940	97.2	30	Office
Physical Exercise	1,870	84.1	28	Residential
Social Interaction	2,480	126.4	42	All three
Resting / Napping	1,640	204.8	25	Residential
Eating	1,920	73.5	40	All three
Whiteboard Pres.	1,340	162.0	18	Classroom
Absent	5,100	--	--	All three
Total	**25,722**	--	**42**	**3 environments**

Dataset archived as skeleton keypoints only; no raw video retained.

C. Circadian Stimulus Mapping

For each recognised activity class a and time-of-day slot t, the system computes a target illumination profile parameterised by vertical illuminance Ev (lux), correlated colour temperature CCT (K), circadian stimulus CS (dimensionless, range 0 to 0.7), uniformity ratio U0, and flicker index FI. The CS value follows the Rea-Figueiro formulation shown in equation (1):

$$X\Sigma = 0.7\,\xi\,[\,1 - (\,1 + (X\Sigma\_\mu\alpha\xi/0.7)\perp\alpha\ \xi\ (\Phi\varpi/X\Lambda A)\perp\alpha\,)\ \perp(-1/\alpha)\,] \quad (1)$$

where CLA (Circadian Light, Agonist) is the weighted irradiance integrated against the melanopsin spectral sensitivity function peaked near 490 nm, CS_max = 0.7, and the empirical constant a = 0.67. Time-of-day modulation uses a cosine envelope targeting CS above 0.30 during core working hours (08:00 to 18:00), tapering to CS below 0.10 by 21:00 with linear interpolation between those bounds.

Table II. **Illumination Profile Targets by Activity Class (Daytime)**

Activity	Ev (lux)	CCT (K)	CS Target	U0 Min	Dim %
Concentrated Work	500-750	5000-6000	0.30-0.45	0.70	85
Casual Reading	300-500	3500-4500	0.20-0.35	0.60	60
Video Conferencing	400-600	4000-5000	0.25-0.40	0.65	75

Physical Exercise	750-1000	5500-6500	0.40-0.50	0.80	100
Social Interaction	250-400	3000-4000	0.18-0.30	0.55	55
Resting / Napping	30-80	1800-2700	< 0.10	0.40	15
Eating	200-350	2700-3500	0.15-0.25	0.50	50
Whiteboard Pres.	500-700	4500-5500	0.28-0.42	0.75	80
Absent	0	--	--	--	0

D. *Federated Training Details*

The federated optimisation objective at each client node k is: min F_k(w_k) + (mu/2) times ||w_k - w^t||^2, where F_k is the local cross-entropy loss on node k labelled data and the proximal term penalises large deviations from the global model w^t. After E = 3 local gradient steps, each node clips its gradient to L2 norm C = 1.0, adds Gaussian noise calibrated to (epsilon = 1.2, delta = 10e-5) differential privacy, and transmits. The server aggregates: w^(t+1) = Sum_k (|Dk| / |D|) times (w_k^t - mu(w_k^t - w^t)). Convergence was confirmed at round 14 (loss delta below 0.001) across all three buildings, with residential nodes taking two rounds longer owing to greater diversity of occupant schedules.

V. Experimental Results

A. *Deployment and Participant Profile*

Experiments ran from October through December across Building-A (corporate office floor, 40 workstations), Building-B (8 university classrooms), and Building-C (24-unit residential apartment complex). One hundred and twenty participants consented under ethics approval (IIT Jodhpur IRB 2024-0047): 64 male, 56 female, age range 21 to 58 years, mean age 31.4 years (SD 8.7). The first two weeks served as static baseline with existing lighting. Adapta was activated from Week 3. Comparison systems: PIR-Dimmer (rule-based [4]), RL-Light (reinforcement learning [5]), and VisionLight (camera-based [6]).

B. *Activity Recognition Performance*

Table III shows per-class precision, recall, and F1-score on the held-out test partition (20 percent of Adapta-Activity, stratified by participant). Overall top-1 accuracy was 94.7 percent, with macro-average F1 of 0.941. The two classes causing most confusion were Concentrated Work and Video Conferencing, both involving a seated person facing a screen. Whiteboard Presentation was, somewhat unexpectedly, the easiest class, likely because the standing posture with arm movements toward a vertical surface is sufficiently distinctive for the GRU to classify within a short window.

Table III. **Per-Class Activity Recognition Results (n = 5,145 clips)**

Activity Class	Precision	Recall	F1-Score	Support
Concentrated Work	0.952	0.941	0.947	964
Casual Reading	0.938	0.951	0.944	722
Video Conferencing	0.915	0.928	0.921	588
Physical Exercise	0.971	0.965	0.968	374
Social Interaction	0.930	0.924	0.927	496
Resting / Napping	0.962	0.970	0.966	328
Eating	0.944	0.938	0.941	384
Whiteboard Pres.	0.979	0.982	0.981	268
Absent	0.988	0.991	0.990	1,021
Macro Average	**0.953**	**0.954**	**0.941**	**5,145**

C. *Energy Consumption*

Power was measured by Shelly Pro 3EM smart meters (stated accuracy plus or minus 0.5 percent) at the distribution board feeding each monitored zone. Table IV summarises monthly consumption in kWh per square metre and percentage reductions relative to the two-week static baseline. Adapta achieved a mean reduction of 41.3 percent (standard deviation 3.8 percentage points), higher than the 38.6 percent achieved by VisionLight and appreciably higher than the 34.7 percent of RL-Light.

Table IV. **Monthly Energy Consumption (kWh/m2/month) and Savings vs. Static Baseline**

Building	Baseline	PIR-D [4]	RL-Light [5]	VisionLight [6]	Adapta (Ours)
Building-A (Office)	3.82	2.91	2.54	2.40	**2.24**
Building-B (Classroom)	4.10	3.08	2.71	2.58	**2.38**
Building-C (Residential)	2.65	1.95	1.72	1.61	**1.57**
Mean Reduction	--	**24.1%**	**34.7%**	**38.6%**	**41.3%**

D. *Human-Centric Outcomes*

At weeks 0, 4, 8, and 12, participants completed three validated instruments: the Karolinska Sleepiness Scale (KSS) for daytime alertness, the Positive and Negative Affect Schedule (PANAS) for emotional state, and the Pittsburgh Sleep Quality Index (PSQI) for sleep quality. A 30-participant control group continued using static lighting throughout. Repeated-measures ANOVA with Bonferroni correction was used for all group comparisons.

Under Adapta, daytime KSS alertness improved by a mean of 1.9 points on the 9-point scale (p less than 0.001, Cohen d = 0.72, large effect). PANAS positive affect increased by 6.4 points (p = 0.003). PSQI global score fell by 2.1 points indicating better sleep quality (p = 0.007). None of these changes were statistically significant in the control group. Residential participants over 45 showed larger sleep quality improvements (delta-PSQI = -2.8) than those under 30 (delta-PSQI = -1.4), an observation warranting a future age-stratified study.

E. *System Latency*

Across 10,000 activity-change events logged during the deployment, the 50th, 95th, and 99th percentile end-to-end latencies were 21.4 ms, 26.8 ms, and 28.1 ms respectively. These figures comfortably satisfy the 28 ms actuation budget. For context, the cloud-based VisionLight system reported round-trip latencies of 340 to 780 ms in its original evaluation [6]; on-device design reduces this by more than an order of magnitude while also eliminating privacy exposure from raw video transmission.

VI. Discussion, Limitations, and Future Work

The results suggest that richer activity sensing leads to better lighting outcomes across energy, accuracy, and wellbeing dimensions simultaneously. The 41.3 percent energy savings are higher than camera-blind alternatives, which is counterintuitive: one might expect a camera-based system to be less aggressive about switching lights off because it sees more nuanced occupancy. What we observe instead is that activity-aware CCT and dimming profiles allow the system to operate at lower power during low-demand activities without degrading the lighting environment for those tasks, whereas simpler systems default to a single high-output safe level.

Several limitations merit acknowledgment. The feedback mechanism for generating training labels depends on occupants spending 90 seconds at an end-of-day touchscreen; compliance averaged 71 percent across the trial. Future versions need either a more passive labelling approach or a self-supervised learning scheme that does not require explicit feedback. The dataset and participant pool are geographically concentrated in Rajasthan, India; cross-cultural validation is on our roadmap. The confusion between Concentrated Work and Video Conferencing remains unresolved, and adding audio event detection is the most straightforward candidate fix.

Future directions include multimodal fusion of RGB-D with EEG wearable data for deeper physiological personalisation; extension of spectral control to UV-A and far-red channels where tunable LED technology permits; hierarchical federated learning across multiple buildings; and investigation of the approach in outdoor or transitional spaces such as stairwells and corridors.

VII. Conclusion

Smart lighting has been promised for a long time, but genuinely smart, responsive to what people are actually doing and aligned with their biological needs, has remained elusive outside the laboratory. Adapta bridges that gap by combining edge-deployed computer vision with a biologically grounded circadian stimulus model and a privacy-preserving federated learning pipeline. The 90-day real-world trial across three building types and 120 participants produced 94.7 percent activity recognition accuracy, 41.3 percent energy savings over static baselines, sub-28 ms actuation latency on commodity edge hardware, and statistically significant improvements in alertness, positive affect, and sleep quality. The neuromorphic event-driven control interface and the MobileNetV3-GRU architecture are, to our knowledge, novel contributions to the smart building literature. The public release of the Adapta-Activity dataset will help other researchers build on this foundation.

VIII. References

[1] IEA, "Global Electricity Review 2023," International Energy Agency, Paris, 2023.

[2] M. S. Rea and M. G. Figueiro, "Light as a circadian stimulus for architectural lighting," Lighting Res. Technol., vol. 50, no. 4, pp. 497-510, 2018.

[3] C. Cajochen et al., "Evidence that the lunar cycle influences human sleep," Current Biology, vol. 23, no. 15, pp. 1485-1488, 2013.

[4] V. Singhvi et al., "Efficient rule-based occupancy sensing for building automation," in Proc. ACM SenSys, San Diego, CA, 2005, pp. 1-12.

[5] Y. Li et al., "DeepLight: Reinforcement learning-based adaptive lighting control," in Proc. IEEE IPSN, Pittsburgh, PA, 2019, pp. 1-12.

[6] T. Yun and C. Yoo, "VisionLight: Vision-based occupancy estimation for smart lighting," IEEE Sensors J., vol. 21, no. 12, pp. 14021-14030, 2021.

[7] H. Kobayashi et al., "Physiological signal-driven adaptive lighting for office productivity," Build. Environ., vol. 193, pp. 107-120, 2021.

[8] F. Zhang et al., "Daylight harvesting with closed-loop PIR-based occupancy sensing," Energy Build., vol. 108, pp. 274-286, 2015.

[9] V. Erickson et al., "OBSERVE: Occupancy-based system for efficient reduction of HVAC energy," in Proc. IPSN, Chicago, IL, 2011, pp. 258-269.

[10] H. Aghajan and A. Cavallaro, Human-Centric Interfaces for Ambient Assisted Living. Waltham, MA: Academic Press, 2009.

[11] G. Bertasius et al., "Is space-time attention all you need for video understanding?" in Proc. ICML, Vienna, 2021, pp. 813-824.

[12] M. Sandler et al., "MobileNetV2: Inverted residuals and linear bottlenecks," in Proc. CVPR, Salt Lake City, UT, 2018, pp. 4510-4520.

[13] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. ICML, Long Beach, CA, 2019, pp. 6105-6114.

[14] P. Boyce et al., "The benefits of daylight through windows," Lighting Res. Center, Troy, NY, Technical Report, 2003.

[15] A. U. Viola et al., "Blue-enriched white light in the workplace improves self-reported alertness, performance and sleep quality," Scand. J. Work Environ. Health, vol. 34, no. 4, pp. 297-306, 2008.

[16] B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, Fort Lauderdale, FL, 2017, pp. 1273-1282.

[17] H. Li et al., "Federated learning for smart building: a survey on data privacy and resource allocation," IEEE IoT J., vol. 9, no. 13, pp. 10414-10430, 2022.

[18] R. Zhang et al., "Privacy-preserving building occupancy estimation via federated learning," Build. Environ., vol. 198, art. no. 107889, 2021.

[19] T. Li et al., "Federated optimization in heterogeneous networks (FedProx)," in Proc. MLSys, Austin, TX, 2020, pp. 429-450.

[20] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," Found. Trends Theor. Comput. Sci., vol. 9, no. 3-4, pp. 211-407, 2014.

[21] A. Howard et al., "Searching for MobileNetV3," in Proc. ICCV, Seoul, South Korea, 2019, pp. 1314-1324.

[22] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in Proc. NIPS Deep Learning Workshop, Montreal, Canada, 2014.

[23] T. M. Brown et al., "Recommendations for daytime, evening, and nighttime indoor light exposure to support physiology, sleep, and wakefulness in healthy adults," PLOS Biology, vol. 20, no. 3, 2022.