Structural Durability & Health Monitoring

Tech Science Press

https://doi.org/10.32604/sdhm.2026.077137

ARTICLE

AI-Driven Object Detection Framework for Live Load Monitoring and Structural Optimization

Luis Sánchez Calderón*, David Valverde Burneo and Walter Hurtares Orrala

Faculty of Engineering in Earth Sciences, ESPOL Polytechnic University, ESPOL, Campus Gustavo Galindo, Km. 30.5 Vía Perimetral, Guayaquil, Ecuador

*Corresponding Author: Luis Sánchez Calderón. Email: luisanch@espol.edu.ec

Received: 03 December 2025; Accepted: 09 February 2026; Published: 18 May 2026

ABSTRACT: Accurate characterization of live load histories remains critical for structural safety and efficient design; however, traditional codes often overestimate in-service loads. This study introduced an AI-driven framework integrating YOLOv8 object detection and DeepFace gender classification with continuous video surveillance to monitor live loads in academic buildings. Gender classification used local anthropometric data (77 kg males, 61 kg females) for precise load estimation, with privacy ensured via local processing and anonymized metadata only. Observed peaks were substantially below Eurocode and IBC provisions, confirming code conservatism. Uncertainty propagation from detector errors (recall 0.57, ±0.02 Kn/m^2) minimally impacted projections. These findings demonstrate the potential of computer vision for data-driven structural optimization and sustainable design.

KEYWORDS: Live loads; computer vision; Convolutional Neural Networks (CNN); video surveillance; structural engineering; YOLOv8; DeepFace; occupancy monitoring

1 Introduction

In building design, live load forces generated by occupants, furniture, and movable objects remain among the most uncertain structural actions due to their spatial and temporal variability. While design codes such as the IBC and Eurocode prescribe conservative load values to account for this unpredictability, empirical studies have repeatedly shown that actual loads in various building types often fall significantly below these standards, raising concerns over material inefficiency and overdesign.

Conventional data sources, such as short-term field measurements and occupancy surveys, are often limited in duration, spatial coverage, and real-time responsiveness. These constraints hinder the development of more precise, context-sensitive live load models. Recent research has explored statistical and probabilistic approaches to better reflect observed occupancy patterns; however, these methods still rely heavily on limited datasets and simplified assumptions about occupant behavior.

One-stage YOLO-family detectors (e.g., YOLOv8) offer a robust method for identifying and quantifying occupants in video frames. In this study, we propose a hybrid framework that combines YOLOv8 object detection and DeepFace gender attribution with continuous video surveillance data to generate detailed live load time histories and histograms in an academic building setting. This approach allows for continuous monitoring of occupant presence and behavior in predefined zones, translating visual data into estimated live loads.

By integrating computer vision with empirical measurement of static furniture weights, we provide a non-stochastic, data-driven method to characterize live loads in specific architectural spaces. This not only contributes to more informed design decisions but also opens new possibilities for smart infrastructure applications in real-time structural health monitoring and load forecasting [1].

Structural design categorizes loads based on nature as specified by codes such as the International Building Code (IBC) and Eurocode [2,3]. Live loads—dynamic forces due to occupants and movable objects—present unique challenges because of their spatial and temporal variability. In residential buildings, for instance, live loads vary in magnitude and distribution over time due to fluctuating occupancy patterns, furniture arrangements, and periodic reconfigurations of interior spaces [4].

It could be argued that monitoring spaces with historically low occupancy offers limited insight. However, confirming that actual service loads are consistently a fraction of the design codes (e.g., 10% of the prescribed value) is structurally significant. It provides the empirical evidence required to transition from prescriptive, conservative design to performance-based optimization. In an era demanding sustainable construction, identifying and quantifying this 'hidden safety factor' enables engineers too reduce material consumption and carbon footprint without compromising safety. Therefore, this study does not merely aim to measure loads, but to quantify the magnitude of potential overdesign in educational infrastructure.

1.1 State of the Art in Live Load Research

Accurate characterization of live loads is critical for structural safety and cost-effective design. Traditional methods such as short-term on-site inspections often miss true load conditions, risking over- or under-designed structures [5]. Numerous empirical investigations have consistently shown that actual live loads often deviate significantly from standardized design values, indicating potential material inefficiency and overdesign [6–9]. For instance, Kumar [8] identified peak occupancies of 3.5 kN/m^2 in offices, surpassing the standard design load of 2.5 kN/m^2, and surveys in library stack areas found local loads up to 7 kN/m^2 [9]. Early empirical surveys of office buildings also reported systematic discrepancies between measured live loads and code-prescribed values, reinforcing the need for context-specific load characterization.

Numerical and probabilistic modeling have refined load characterizations, improving precision through structural simulations and statistical analyses [6,8,10]. Despite advancements, inconsistencies remain, particularly systemic overestimations noted in Eurocode models [11]. Recent methods utilizing VR panoramas and real estate data have significantly expanded datasets [5]; however, their limited duration still fails to capture essential multi-year load variations, particularly crucial in residential and governmental buildings.

To better understand these deviations, theoretical frameworks such as those established by the CIB [12] decompose live loads into two distinct stochastic processes: a 'sustained' component (e.g., furniture, heavy equipment) that changes infrequently, and an 'intermittent' or 'extraordinary' component (e.g., human crowding) characterized by short-duration, high-magnitude spikes. While modern codes like ASCE 7 and Eurocode 1 are calibrated based on these probabilistic models, recent research suggests that the assumed magnitude of these intermittent spikes may be overly conservative for controlled environments like academic buildings, directly impacting the embodied carbon of the structure.

Technological advancements have recently enabled the transition from static surveys to continuous monitoring. Early occupancy monitoring approaches typically relied on discrete sensors such as Passive Infrared (PIR) systems, which were limited to binary presence detection and lacked the spatial resolution required for accurate structural load estimation.

More recent developments in computer vision have enabled vision-based occupancy detection using standard RGB cameras, allowing for continuous monitoring with spatial and temporal resolution sufficient

to capture transient occupancy patterns. Zhang et al. [13] demonstrated the feasibility of deep-learning-based vision systems for real-time occupancy detection in high-density indoor environments, highlighting their relevance for operational and building-performance analysis.

Despite these advances, the direct application of vision-based occupancy sensing to structural live-load characterization and safety assessment remains comparatively limited, particularly in terms of long-term, zone-specific empirical datasets.

1.2 Convolutional Neural Networks for Occupancy Detection

Deep convolutional neural networks (CNNs) have become the *de facto* standard for computer-vision tasks such as people counting and occupancy analytics [14]. Early detectors—including R-CNN and its successors Fast R-CNN and Faster R-CNN—delivered high accuracy but relied on two-stage region-proposal pipelines that incur substantial latency [15,16]. For real-time surveillance, latency is as critical as accuracy; consequently, recent research has shifted toward single-stage, fully convolutional detectors.

1.3 Rise of YOLO-Family One-Stage Detectors

Single-stage object detectors of the "You Only Look Once" (YOLO) family perform object localization and classification in a single forward pass, enabling substantially higher processing speeds than two-stage detection frameworks. This characteristic is particularly relevant for continuous monitoring applications, where latency constraints directly limit the temporal resolution and duration of usable datasets.

In the context of live-load monitoring based on video surveillance, sustained near–real-time inference is more critical than marginal gains in detection accuracy. Continuous analysis of weeks or months of CCTV footage requires detectors capable of processing large data volumes without prohibitive computational cost.

For this reason, YOLOv8 was selected as the primary person-detection backbone in this study. Compared with two-stage architectures such as Mask R-CNN, YOLOv8 offers a substantially higher throughput under typical desktop hardware configurations while maintaining competitive detection accuracy for human targets. Two-stage detectors, although accurate, incur significantly higher per-frame processing times, making them impractical for long-term, multi-camera monitoring scenarios.

Accordingly, the proposed framework adopts YOLOv8 for person detection and couples it with a conservative DeepFace-based gender attribution module, enabling the transformation of raw occupancy observations into continuous live-load time histories, as detailed in the Methodology section.

2 Objectives

2.1 Main Objectives

To propose and validate a scalable, continuous monitoring framework combining computer vision (YOLOv8 and DeepFace) and structural engineering principles to estimate live loads in academic buildings.

2.2 Specific Objectives

- Empirical Data Generation: To collect a continuous, large-scale empirical dataset of live load histories and occupancy in four distinct zones (hallways, stairwell, laboratory) over approximately one month.
- Code Comparison: To compare the observed peak live loads and the projected 50- and 100-year return level loads against the conservative values prescribed by the International Building Code (IBC) and Eurocode.

- Methodological Transparency: To detail the full methodology, including the detector configuration and threshold tuning, to ensure the approach is transparent and reproducible for future load calibration studies.

These novel contributions distinguish this work: (1) a scalable YOLOv8-DeepFace pipeline for continuous live load estimation in low-occupancy academic zones, prioritizing safety-biased thresholds; (2) empirical evidence from one-month monitoring showing 80%–90% code overestimation, with robust extreme-value projections; and (3) a transparent methodology bridging computer vision and structural engineering for performance-based design optimization.

3 Methodology

The methodology integrates video surveillance, frame extraction, object detection using a convolutional neural network (YOLOv8), and empirical load computation to generate live load time histories and histograms.

This dataset provides the necessary empirical basis to develop a non-stochastic model to calculate live loads for specific types of monitored spaces. This model is crucial for selecting appropriate loads in structural design.

The proposed framework is summarized in the workflow depicted in Fig. 1:

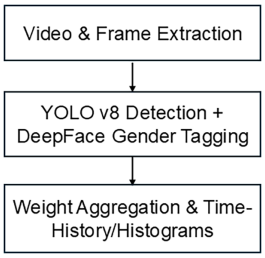

Figure 1: A workflow diagram illustrates the process from video surveillance to load history and histogram generation.

3.1 Data and Storage

Video acquisition was carried out through the ESPOL 911 surveillance system, which manages campus-wide security monitoring. Under current ESPOL 911 policy, the cameras operate on a basic storage plan that retains recordings for up to 15 days before automatic overwriting. Consequently, all footage had to be processed within this retention window to comply with institutional policies. Expanded cloud storage or large-scale microSD rotation were not options available within ESPOL 911's operational framework, which set a practical constraint on the monitoring campaign.

3.2 Cameras

The cameras used in the present study were manufactured by EZVID and have a resolution of 2.0 MP, recording color video at 30 frames per second. They operate continuously, recording 24 h per day. The output files are in MP4 format and are divided into 6-h recordings; each file is approximately 1 GB in size.

3.3 Monitored Spaces

All monitored spaces are located in Building 13H of the Faculty of Geosciences Engineering (FICT) on the Gustavo Galindo campus of the Escuela Superior Politécnica del Litoral (ESPOL), in Guayaquil, Ecuador.

Building 13H is a two-story reinforced concrete structure with a total area of 1100 m^2, constructed in 2008. Its location within the campus is shown in Fig. 2.

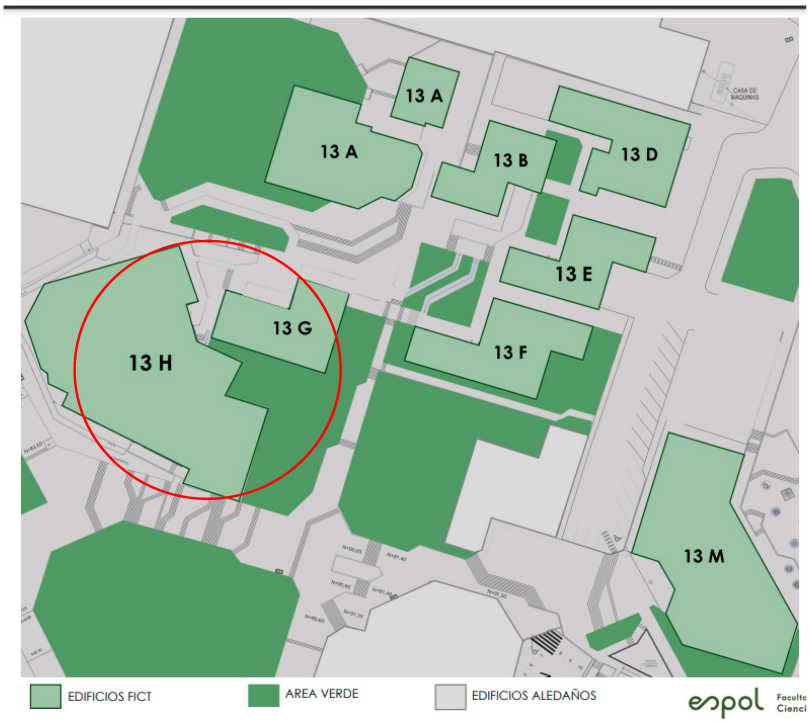

Figure 2: ESPOL campus map indicating the location of monitored Building 13H.

The building comprises 10 classrooms, each measuring around 50 m^2, along with 2 computer laboratories and a multipurpose hall known as "Aula Satelital". It serves to host classes across four undergraduate programs offered by the faculty: Civil Engineering, Mining Engineering, Geology, and Petroleum Engineering. During the 2023 academic year, these programs collectively enrolled 663 students. Additionally, there are 4 master's programs available, which account for 67 students in the same academic year.

Fig. 3 illustrates the floor plan of the second level of Building 13H, along with a color-coded segmentation of the monitored spaces.

Academic activities within the studied building are structured around undergraduate classes held Monday to Friday. Notably, undergraduate programs do not hold classes on Fridays. As of the 2024–2025 academic year, classes are scheduled from 7:00 AM to 5:00 PM.

On weekends, master's program students occasionally attend in-person classes during the mornings. A front elevation of Building 13H is shown in Fig. 4.

Note. Each monitored zone is identified by an official camera code from ESPOL's surveillance system "ESPOL 911" (e.g., *13H-PA-PASILLO 1*). For clarity, English names are used throughout the text, while the original Spanish codes are retained in parentheses to ensure traceability of footage and datasets.

The four monitored spaces are:

- Hallway 1 (camera code: 13H-PA-PASILLO 1, RED), connecting the stairwell to the first floor.
- Stairwell connector (camera code: 13H-PA-ESCALERA, BLUE), covering circulation between the two levels.

- Computer Laboratory 2 (camera code: 13H-PA-LABORATORIO 2, GREEN), used for software-based classes.
- Hallway 2 (camera code: 13H-PA-PASILLO 2, PURPLE), connecting deeper classrooms on the second floor.

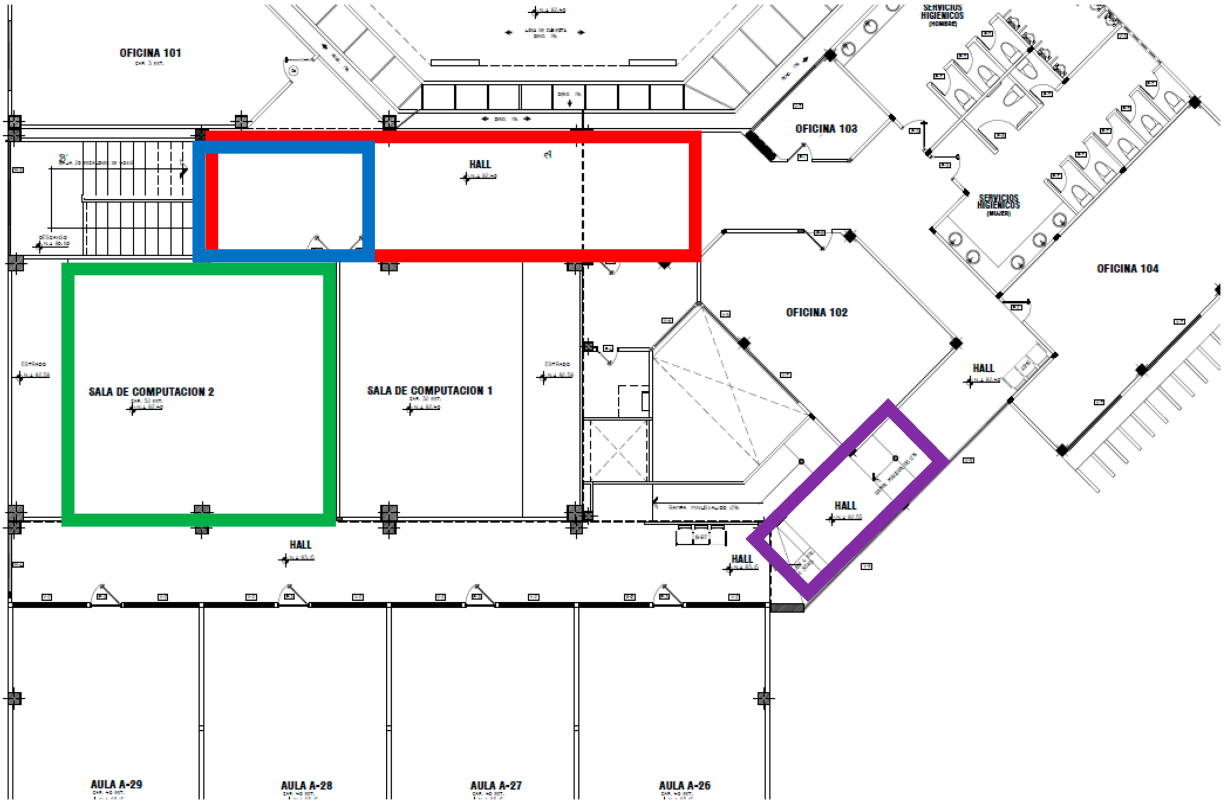

Figure 3: Floor plan for the second level of Building 13H, highlighting the four monitored areas.

Figure 4: Exterior front view of building 13H at the Gustavo Galindo campus.

For clarity, throughout this paper the term "monitored zone" refers to each camera-covered portion of the building treated as an independent dataset for analysis, whereas "area" denotes the physical floor space (in m^2) used to compute load densities.

The four cameras view the four monitored spaces from the perspectives shown in Figs. 5–8.

Figure 5: Camera perspective covering Hallway 1 (camera code: 13H-PA-PASILLO 1).

Figure 6: Camera perspective covering the stairwell connector (camera code: 13H-PA-ESCALERA).

Figure 7: Camera perspective covering Computer Laboratory 2 (camera code: 13H-PA-LAB 102).

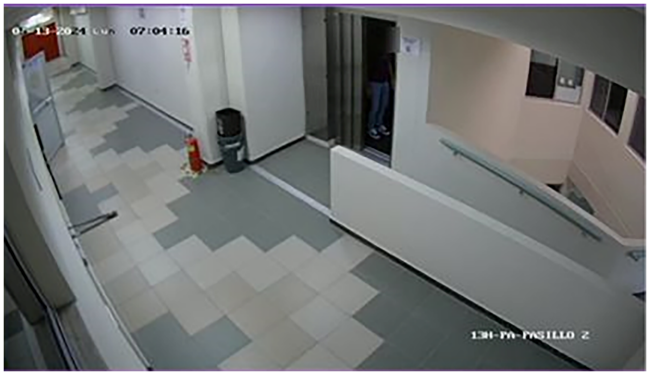

Figure 8: Camera perspective covering Hallway 2 (camera code: 13H-PA-PASILLO 2).

3.4 Video Processing

Video processing was divided into three steps: Image/Frame Extraction, Object Detection, and Load Estimation, as detailed in the following subsections. The one-month monitoring campaign across four cameras generated approximately 2.88 TB of video data for analysis.

3.5 Image Frame Extraction

A Python script (MoviePy) extracted frames with an adaptive interval based on expected activity. During active periods, one frame was sampled every 30 s; during passive periods, one frame every 6 h. Across a week (168 h), 72 h were considered active. This strategy captures potential peaks while reducing computation and storage during low-occupancy periods. These criteria are summarized and presented in Table 1.

Table 1: Frame-extraction rates applied to all monitored spaces in this study.

Period	Time Frame	Rate
Weekdays (Monday to Friday)	Active (7:00–17:00)	2 frames/1 min
	Passive (17:00–7:00)	1 frame/6 h
Weekends (Saturday and Sunday)	Active (8:30–13:00)	2 frames/1 min
	Passive (13:00–8:30)	1 frame/6 h

3.6 Detection and Live Load Computation

Each extracted frame was analyzed with YOLOv8; detections classified as person were counted and, when a face was visible, passed to DeepFace for gender attribution. Loads were computed by assigning 77 kg to male detections and 61 kg to female detections, based on the local anthropometric study [17] for 21–22-year-old university students in coastal Ecuador.

Based on the integration of static and dynamic components proposed in this framework, we defined the instantaneous live load density $L(t)$ in kN/m^2 for a given monitored zone at time t as follows:

$$L(t) = \frac{W_{static} + \sum_{i=1}^{N_m(t)} w_{male} + \sum_{i=1}^{N_f(t)} w_{female}}{A_{zone}} \quad (1)$$

where:

- W_{static} is the total weight of fixed furniture and equipment (kN).
- $N_m\,(t)$ and $N_f\,(t)$ are the number of detected males and females at time t.
- w_{male} and w_{female} are the average weights (0.755 and 0.598 kN, respectively).
- A_{zone} is the floor area of the monitored space m^2.

The summed occupant mass was added to the static furniture load, which was obtained by physically weighing all movable items in each monitored space prior to the campaign. For each frame, the timestamp and total live load were recorded. Regarding spatial distribution, loads were computed as zone-average densities; within-zone clustering (e.g., crowding near doorways) was not modeled in this pilot study.

3.7 Detector Configuration and Performance Evaluation

A two-stage computer vision pipeline was implemented, comprising: (1) YOLOv8-m person detection and (2) DeepFace gender attribution. Both components were threshold-tuned with a safety-oriented priority: in live-load estimation, false negatives (missed people) are more critical than false positives. Occlusion and partial visibility remain dominant sources of error in vision-based occupancy detectors, as extensively reported in controlled and real-world building environments [18].

3.7.1 YOLOv8 Confidence-Threshold Tuning

Precision, recall, and F1-scores were evaluated across thresholds ranging from 0.01 to 0.90 on the WiderPerson validation set (Fig. 9; Table 2). Recall reaches its peak between 0.01 and 0.10, gradually declining thereafter, while precision increases steadily above 0.30. The F1-score peaks in the 0.05–0.10 range. This behavior confirms that thresholds in the 0.03–0.10 band offer the optimal compromise for safety-critical applications by maximizing recall while retaining acceptable precision.

Selected threshold was set as ThresholdYOLOv8 = 0.07, the midpoint of this plateau, where recall remains >0.55 and precision ≈0.71—prioritizing recall to reduce the risk of underestimating occupancy.

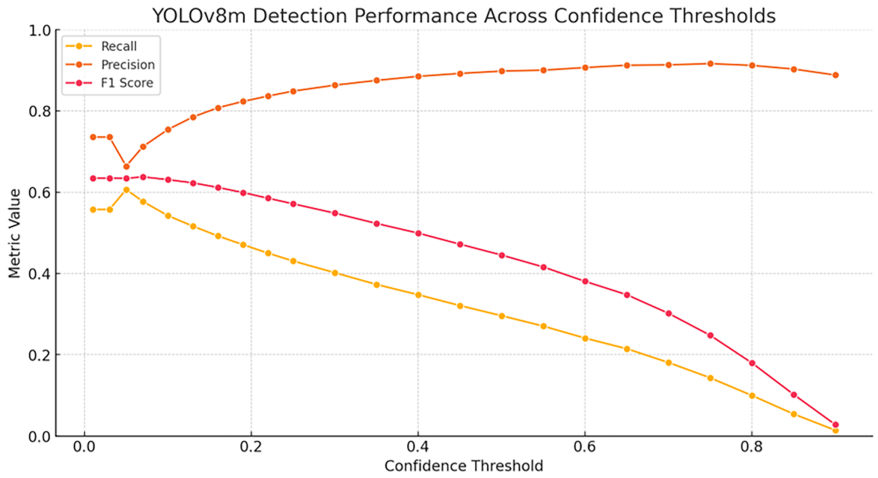

Figure 9: Precision, recall and F1 score vs. confidence threshold on the WiderPerson validation subset.

Table 2: YOLOv8-m detection metrics across confidence thresholds.

Threshold YOLOv8	Precision	Recall	F1 Score
0.010	0.736	0.558	0.634
0.030	0.736	0.558	0.634
0.050	0.664	0.607	0.634
0.070	0.712	0.577	0.638
0.100	0.754	0.542	0.631

3.7.2 DeepFace Gender-Threshold Tuning

On the FairFace validation set, ThresholdDeepFace = 0.30 yielded near-perfect male recall (~0.99) while keeping overall accuracy ≈ 0.72 (Table 3; Fig. 10). As shown in Fig. 10, female recall increases gradually with higher thresholds, while male recall drops. Since misclassifying a male as female would underestimate load, the decision boundary was intentionally set at 0.30 to bias ambiguous cases toward male labeling, thereby ensuring a conservative, safety-oriented estimate.

Table 3: DeepFace gender metrics across decision thresholds.

Threshold DeepFace	Female Recall	Male Recall	Accuracy
0.300	0.28	0.99	0.72
0.400	0.33	0.98	0.73
0.500	0.34	0.96	0.69
0.600	0.39	0.95	0.70
0.700	0.41	0.92	0.70

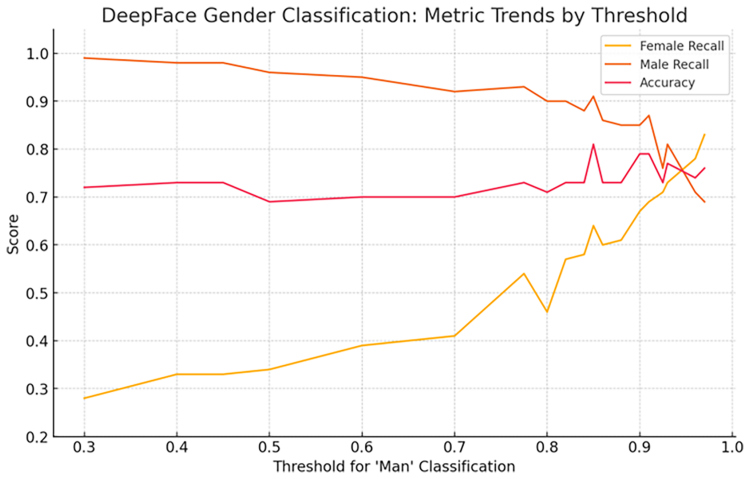

Figure 10: Trends of male recall, female recall and overall accuracy for DeepFace on the FairFace validation set.

3.7.3 System-Level Bias and Safety Margin

With ThresholdYOLOv8 = 0.07 and ThresholdDeepFace = 0.30, the combined pipeline achieves recall > 0.57 with precision ≈ 0.71 for person detection, while classifying ~99% of males correctly. Together, these

configurations bias ambiguous detections toward higher occupant counts and male weights. The net effect is a systematic overestimation of live load, consistent with the safety-first philosophy of structural engineering. This deliberate bias ensures that the occupancy monitoring system errors on the side of conservatism, reducing the likelihood of underestimating critical occupancy-induced loads.

4 Results

The proposed framework was applied to four monitored zones within the academic building: two hallways, a stairwell segment, and a computer laboratory. For each space, the total live load was estimated by integrating occupant detections from the YOLOv8–DeepFace pipeline with empirically measured static furniture weights.

Results are presented in two complementary formats: (i) load–time histories illustrating the temporal evolution of live loads, and (ii) histograms summarizing their distribution over the monitoring period. Together, these views capture both variability and distributional characteristics, supporting extreme-value modeling and design-value comparisons.

A brief description of the observations and data obtained in each zone follows.

4.1 Computer Laboratory (Camera Code: 13H-PA-LAB102)

The computer laboratory was monitored continuously during July 2024, including both weekdays and weekends to capture variations in occupancy associated with academic schedules. Furniture and equipment contributed a static load of 711 kg across a floor area of 61.28 m^2. The computed live loads ranged from a minimum of 0.12 kN/m^2 to a maximum of 0.35 kN/m^2, both well below the design values prescribed by Eurocode (3.0–4.0 kN/m^2) and IBC (4.8 kN/m^2). The histogram (Fig. 11) and time-series plot (Fig. 12) summarize these results, highlighting pronounced temporal fluctuations aligned with academic schedules and minimal loads during passive periods. The complete frame-level dataset for this zone is available as supplementary materials (LoadData_13H-PA-LAB.csv).

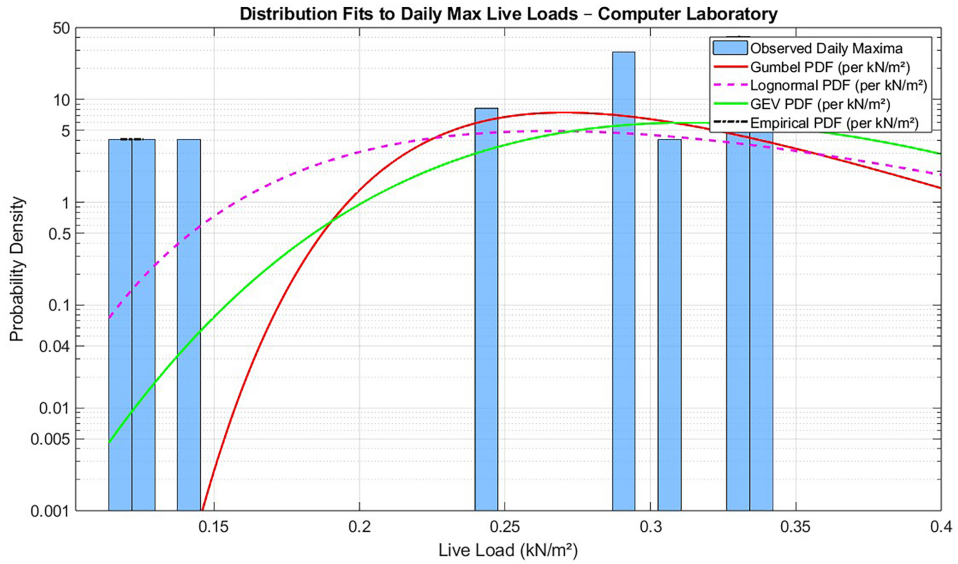

Figure 11: Histogram of daily maximum live loads in Computer Laboratory 2 (camera code: 13H-PA-LABORATORIO 102), showing measured peaks (≤0.35 kN/m^2) well below Eurocode/IBC design values.

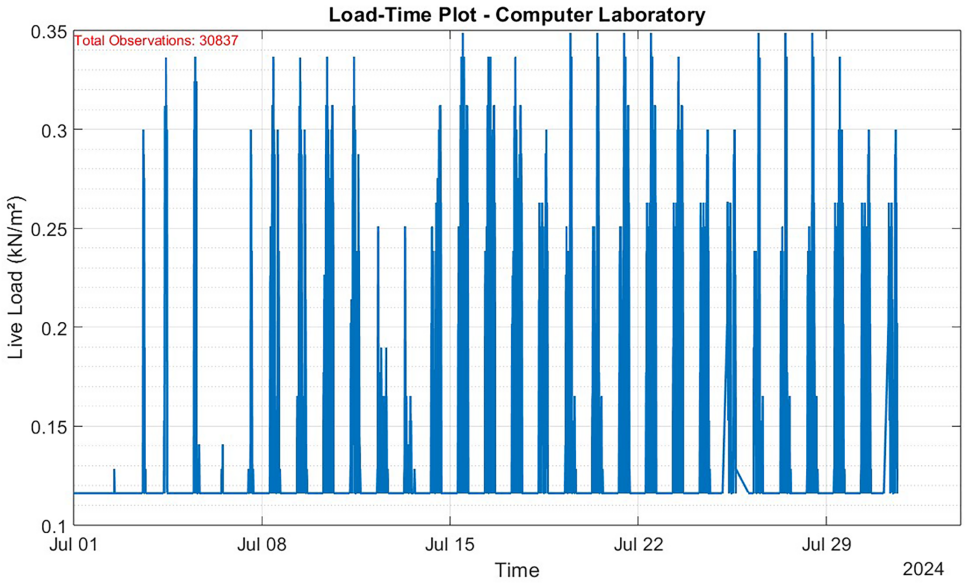

Figure 12: Time-series of computed live loads in Computer Laboratory 2 (camera code: 13H-PA-LABORATORIO 102), illustrating fluctuations aligned with active and passive occupancy periods.

To illustrate the detection and load-estimation process, Fig. 13 presents annotated frames corresponding to peak and low-occupancy conditions. Each detection is assigned a gender-specific weight (77 kg male, 61 kg female), with the summed occupant weight added to the measured furniture load. This example demonstrates how the YOLOv8 + DeepFace pipeline translates video data into quantifiable structural loads per frame.

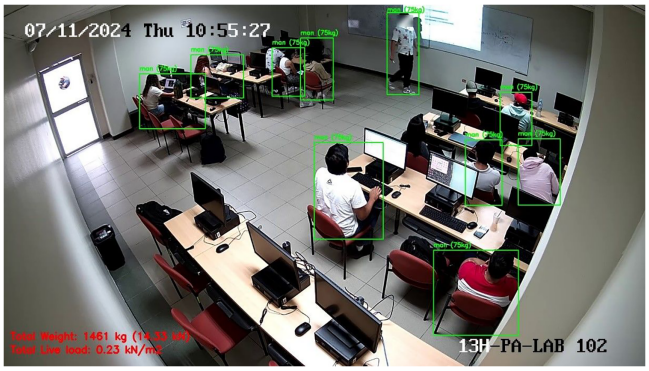

Figure 13: Annotated frame from Computer Laboratory 2 (camera code: 13H-PA-LABORATORIO 2), showing ten detected individuals with a computed total load of 1461 kg including furniture load (0.23 kN/m^2).

4.2 Hallway 1 (Camera Code: 13H-PA-PASILLO 1)

Hallway 1 was monitored for approximately one month. Static furniture contributed 38 kg (two metal seats) over a floor area of 30.58 m^2. The computed live loads ranged from a minimum of 0.012 kN/m^2 to a maximum of 0.158 kN/m^2, observed when 10 occupants were simultaneously present. Both values are far below the Eurocode (3.0–5.0 kN/m^2) and IBC (4.8 kN/m^2) design requirements. Higher loads occurred near the stairwell section, which also provides access to a faculty office and two computer laboratories, which explains observed localized clustering near high-traffic access points. Results are summarized

in Figs. 14 and 15. The complete frame-level dataset for this zone is available as supplementary materials (LoadData_13H-PA-PASILLO-1_extracted_frames.csv).

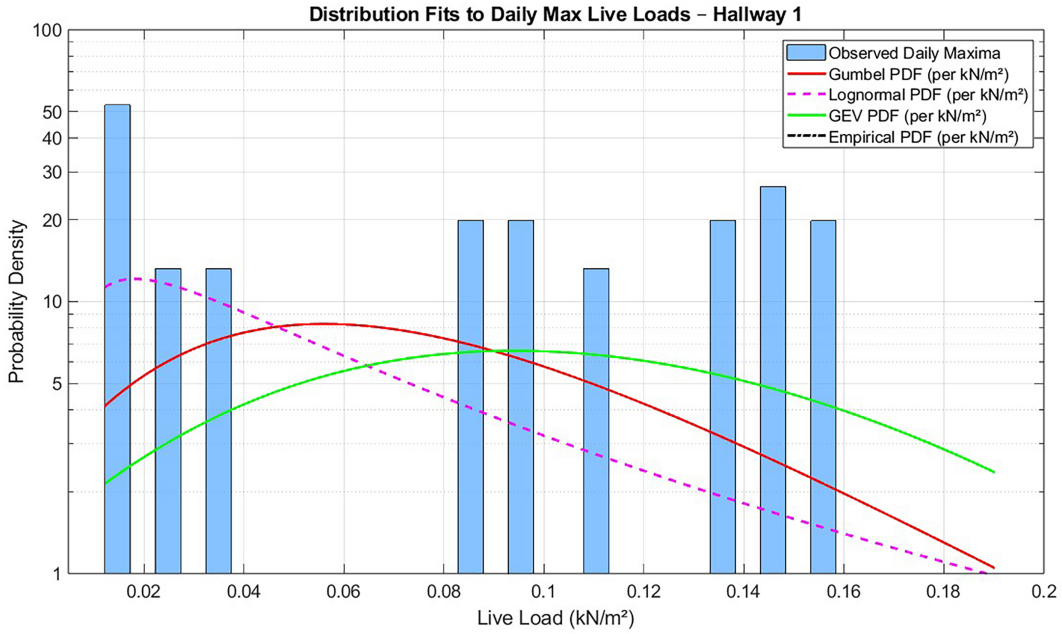

Figure 14: Histogram of daily maximum live loads in Hallway 1 (camera code: 13H-PA-PASILLO 1), with peaks ≤0.158 kN/m^2, well below Eurocode/IBC values.

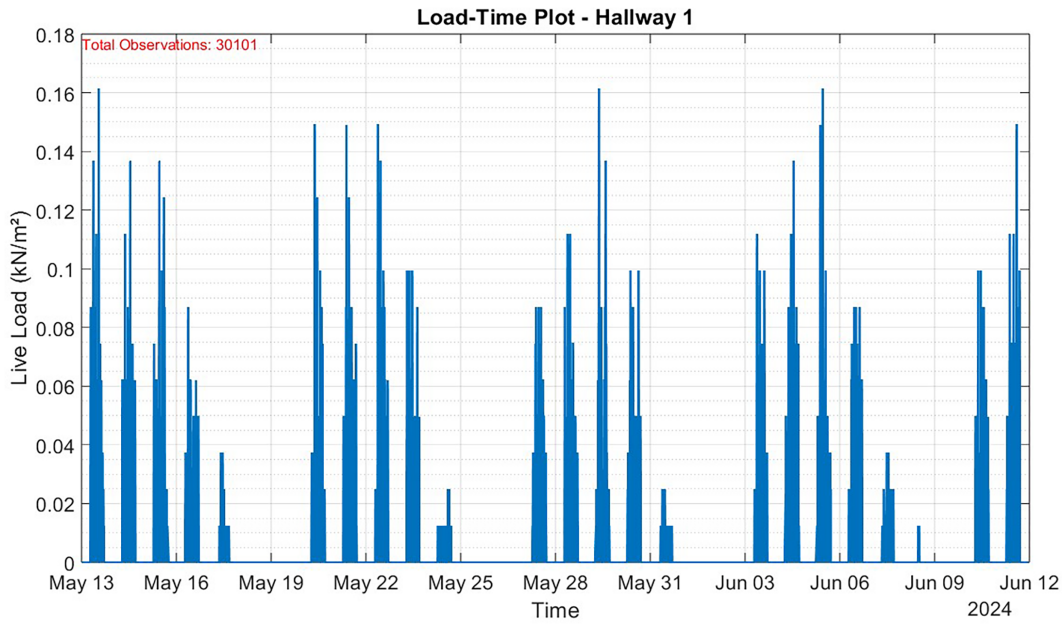

Figure 15: Time-series of computed live loads in Hallway 1 (camera code: 13H-PA-PASILLO 1), showing variability linked to circulation near the stairwell.

To illustrate the pipeline in operation, Fig. 16 presents an annotated frame from a weekday afternoon. Detected individuals are assigned gender-specific weights, which are then summed with the static furniture

load to compute an instantaneous live load of 0.08 kN/m^2. This demonstrates how transient peak events are captured in circulation areas, even when occupancy is brief.

Figure 16: Annotated frame from Hallway 1 (camera code: 13H-PA-PASILLO 1), with three occupants detected and a computed load of 0.08 kN/m^2.

4.3 Hallway 2 (13H-PA-PASILLO 2)

Hallway 2 was monitored for approximately 20 days. The area is 62.4 m^2, with no fixed seating or furniture. Computed live loads ranged from a minimum of 0.158 kN/m^2 to a maximum of 0.26 kN/m^2, both substantially below the Eurocode (3.0–5.0 kN/m^2) and IBC (4.8 kN/m^2) design values. As expected, peak values were lower than those observed in Hallway 1, consistent with the lack of furniture and shorter occupancy durations. Results are presented in Figs. 17 and 18. The complete frame-level dataset for this zone is available as supplementary materials (LoadData_13H-PA-PASILLO-2_extracted_frames.csv).

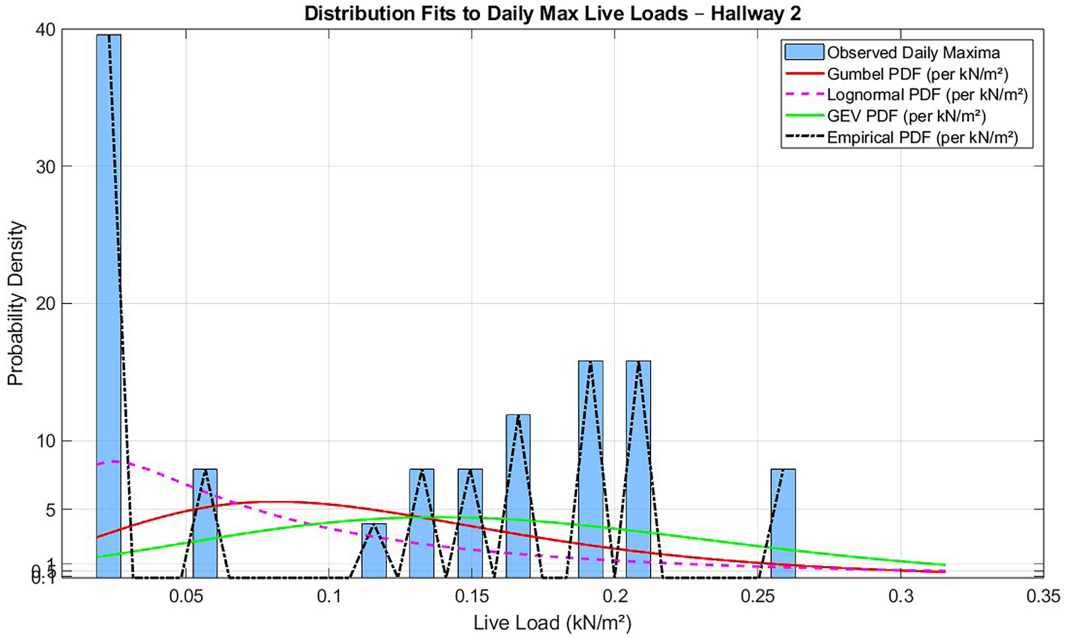

Figure 17: Histogram of daily maximum live loads in Hallway 2 (camera code: 13H-PA-PASILLO 2), with peaks ≤0.26 kN/m^2, far below Eurocode/IBC design thresholds.

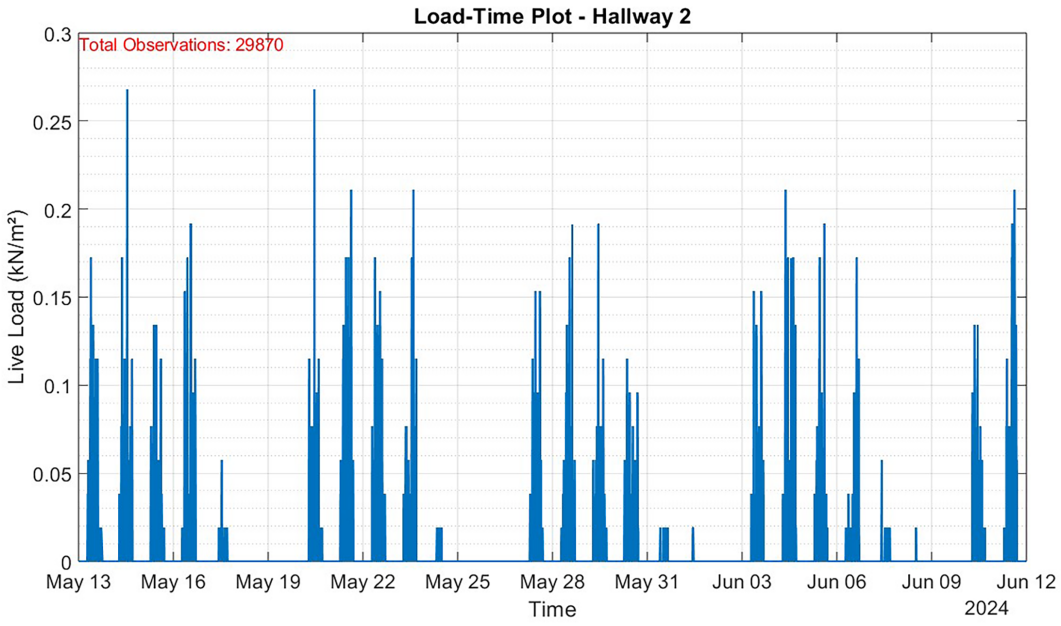

Figure 18: Time-series of computed live loads in Hallway 2 (camera code: 13H-PA-PASILLO 2), showing typical low-occupancy patterns.

Fig. 19 illustrates a representative frame during a low-occupancy period. Three individuals were detected, with DeepFace assigning gender-specific weights. The computed live load was 0.02 kN/m^2, demonstrating how the pipeline quantifies transient passage events in corridors.

Figure 19: Annotated frame from Hallway 2 (camera code: 13H-PA-PASILLO 2), with three detected individuals and a computed load of 0.02 kN/m^2.

4.4 Stairwell Connector (Camera Code: 13H-PA-ESCALERA)

The stairwell connector was monitored for 30 days, covering an area of 24.20 m^2. With no furniture present, the minimum live load recorded was 0 kN/m^2, while the peak reached 0.527 kN/m^2—the highest among all monitored zones. These peak loads correspond to high-circulation events, as the stairwell provides vertical access to the second floor, adjacent to two computer laboratories and a faculty office. The complete frame-level dataset for this zone is available as supplementary materials (LiveLoad_Data_Month_13H-PA-ESCALERA.csv).

The histogram in Fig. 20 exhibits a right-skewed distribution, characterizing short-duration congestion events typical of vertical circulation. Overlaid probabilistic fits (Gumbel, Lognormal, and GEV) highlight the probabilistic nature of transient loads, providing the basis for extreme-value projections. The time-series in Fig. 21 illustrates pronounced peaks during morning transitions and midday breaks, while night hours show near-zero loads.

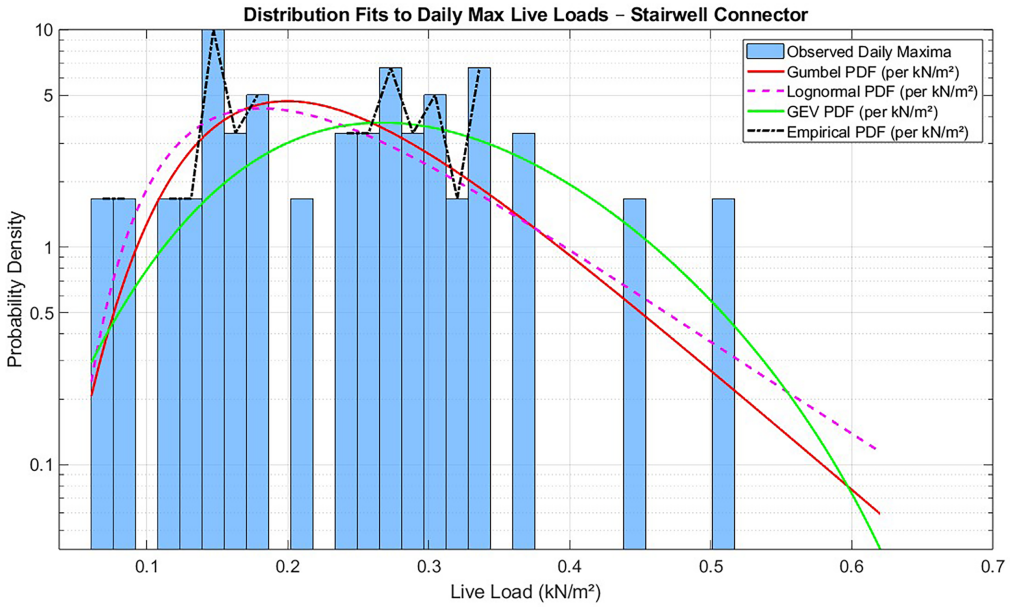

Figure 20: Histogram of daily maximum live loads in the Stairwell Connector (camera code: 13H-PA-ESCALERA), with peaks ≤0.527 kN/m^2, still well below Eurocode/IBC thresholds.

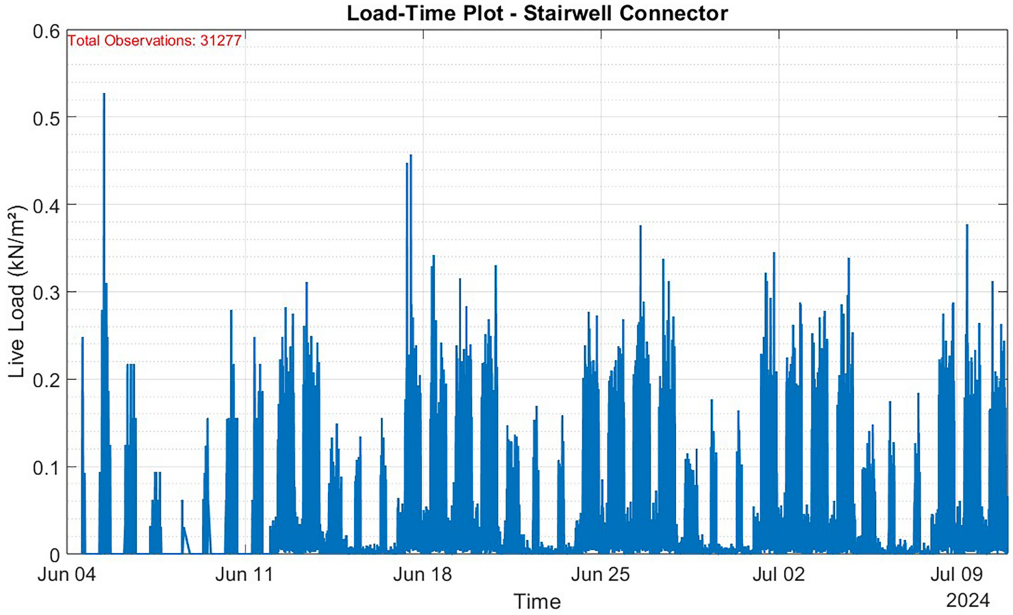

Figure 21: Time-series of computed live loads in the Stairwell Connector (camera code: 13H-PA-ESCALERA), showing peaks during class transitions and near-zero loads overnight.

Fig. 22 presents an annotated frame from a low-occupancy moment. Four individuals were detected, yielding a computed load of 0.12 kN/m^2. This demonstrates how the YOLOv8–DeepFace pipeline captures even short-term occupancy in transitional areas, supporting probabilistic modeling of peak loads.

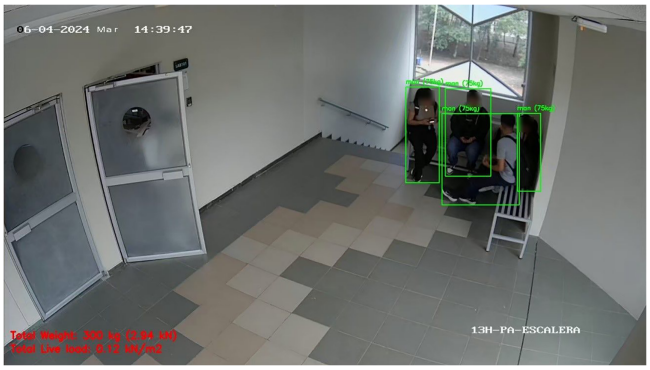

Figure 22: Annotated frame from the Stairwell Connector (camera code: 13H-PA-ESCALERA), showing four detected individuals and a computed load of 0.12 kN/m^2.

5 Discussion

5.1 Comparative Analysis with Prior Load Modeling Studies

Research on live load characterization has traditionally relied on probabilistic modeling frameworks supported by surveys, short-term inventories, or limited field campaigns. Foundational studies—such as those by Kumar [7,8] and Chalk and Corotis [10]—established key insights into occupancy-driven variability in commercial and institutional buildings, providing the basis for modern design code provisions and probabilistic live-load models.

More recent contributions, including the comprehensive review by Costa and Beck [11], have emphasized the need to revisit code assumptions in light of empirical evidence and have advanced statistical methodologies for recalibrating live-load models using reliability-based approaches.

These contributions have greatly advanced the state of knowledge. At the same time, the level of methodological detail reported in earlier studies—such as observation duration, sampling frequency, or the number and type of monitored buildings—was sometimes constrained by the practical limitations of data collection at the time. As a result, it can be difficult to directly compare synthetic or semi-empirical models with continuous, real-world occupancy measurements.

5.2 Novelty of the Study

The present work contributes a complementary and novel perspective to live-load research by providing a fully transparent empirical dataset derived from more than 100,000 annotated video frames collected over approximately one month of continuous monitoring (28–31 days depending on zone) in an operational academic building. Each of the four monitored zones contributed on the order of 25,000–30,000 usable observations, ensuring balanced temporal coverage and reproducibility. Every step of the framework—from sampling schedules to zone-specific monitoring strategies—is documented in detail.

At first glance, the consistently low live loads observed in these academic spaces may appear to limit the practical relevance of the study. However, from a structural engineering perspective, this outcome constitutes a pivotal contribution rather than a limitation. While existing design codes intentionally prescribe conservative live-load values to address uncertainty, rare events, and potential changes in use, continuous

empirical evidence quantifying the magnitude of these implicit safety margins under real service conditions remains scarce.

Previous studies have largely relied on short-term surveys, inventories, or probabilistic assumptions, often lacking continuous, zone-specific occupancy measurements. By integrating YOLOv8 person detection and DeepFace-based conservative weight assignment, the proposed framework enables direct conversion of occupancy observations into live-load histories using existing surveillance infrastructure. The results demonstrate that, even under deliberately conservative assumptions, observed peak live loads remain an order of magnitude below IBC and Eurocode provisions. Making this margin explicit through real operational data is a prerequisite for future developments in performance-based assessment, structural optimization, and digital-twin-assisted monitoring.

Importantly, the novelty of this work does not lie in proposing immediate reductions of code-mandated loads, but in providing a validated empirical baseline that complements traditional design approaches and supports informed decision-making in structural assessment, retrofitting, and sustainability-oriented design of institutional buildings.

5.3 Empirical Gaps in Prior Work

One frequently noted limitation in the live load modeling literature is the variability in empirical detail provided with probabilistic models. Foundational frameworks by Chalk and Corotis [10] have significantly shaped the field but may not always include specifics such as sampling frequency, monitoring duration, or building typologies. More recent contributions, including those by Márquez Peñaranda [19] and Costa and Beck [11], often emphasize statistical rigor while relying on secondary or aggregated data sources with limited procedural transparency.

These approaches remain invaluable for code calibration and reliability-based design; however, they highlight the inherent challenges in obtaining continuous, zone-level occupancy data. While such limitations were understandable given the technological constraints at the time of data collection, they underscore the value of modern, vision-based monitoring approaches capable of providing high-resolution empirical evidence.

5.4 Spatial Distribution and Clustering

One limitation of this pilot is that occupant positions within each zone were not explicitly modeled; instead, zone-average load densities were used. This simplifying assumption may under- or over-estimate local effects when occupants cluster near stairwells, doorways, or seating areas. A natural extension is to apply homography-based projection of detections onto floor-plan coordinates and to combine overlapping camera views for multi-angle triangulation. Preliminary internal tests confirm that such methods can generate high-resolution spatial load maps, enabling precise quantification of localized peak effects without necessitates modifications to the existing YOLOv8–DeepFace pipeline.

5.5 Live Load Projection and Model Comparison

The consistent observation of peak loads below 0.60 kN/m^2—in contrast to code requirements of 3.0 to 5.0 kN/m^2—highlights a significant material efficiency gap. While high code values are intended to cover extreme events, the magnitude of this discrepancy (over 500%) in a controlled university environment suggests that current classifications for ‘institutional’ use may be overly broad. For a structural engineer, these data indicate that lighter structural systems may be feasible for specific zones, provided that change-of-use is controlled. This supports the argument for ‘adaptive building codes’ that allow for load reduction

factors based on real-time monitoring history. Table 4 presents the projected maximum live loads for a 50-year and 100-year return periods, computed using Gumbel, GEV, and Lognormal distributions [20]. The live load histories, computed using Eq. (1), served as the basis for these projections, which were derived by fitting each distribution to the daily maximum observed live loads for each monitored space. All loads are presented in kN/m^2. These values enable a direct comparison with international building standards. Recent studies have demonstrated the potential of vision-based surveys to identify extraordinary or transient load conditions in buildings, extending occupancy monitoring beyond energy applications toward structural safety assessments [21].

Table 4: Projected maximum live loads (in kN/m^2) for 50-year and 100-year return periods.

Monitored Space	Distribution	50-Year Max Load (kN/m^2)	100-Year Max Load (kN/m^2)	Max Observed (kN/m^2)
13H-PA-LAB	Gumbel	0.633	0.668	
13H-PA-LAB	GEV	0.543	0.552	0.340
13H-PA-LAB	Lognormal	0.741	0.784	
13H-PA-ESCALERA	Gumbel	0.791	0.845	
13H-PA-ESCALERA	GEV	0.636	0.651	0.527
13H-PA-ESCALERA	Lognormal	0.993	1.083	
13H-PA-PASILLO 1	Gumbel	0.381	0.412	
13H-PA-PASILLO 1	GEV	0.301	0.309	0.158
13H-PA-PASILLO 1	Lognormal	1.575	1.930	
13H-PA-PASILLO 2	Gumbel	0.566	0.612	
13H-PA-PASILLO 2	GEV	0.446	0.459	0.240
13H-PA-PASILLO 2	Lognormal	2.474	3.050	

As observed, Gumbel and GEV models yield relatively consistent projections within expected ranges for hallways and laboratory spaces, while Lognormal models produce significantly higher values. This behavior is consistent with the heavier tails often observed in probabilistic modeling and reflects the higher sensitivity of the Lognormal distribution to extreme values, making it more reactive to rare events. See Table 5 for a clear comparison of observed live loads against IBC and Eurocode standards.

Table 5: Comparative analysis of observed live loads with design standards (IBC and Eurocode).

Monitored Space	Max Observed (kN/m^2)	IBC Standard (kN/m^2)	Eurocode Standard (kN/m^2)	Overestimation (%)
Laboratory	0.34	4.8	3.0–4.0	>85%
Hallways	0.16–0.26	4.8	3.0–5.0	>90%
Stairwell	0.527	4.8	3.0–5.0	>80%

These insights enable structural optimization: e.g., potential 80%–90% load reductions support lighter RC designs in academic corridors, cutting embodied carbon ~20%–30% without safety loss, informing code calibration for monitored buildings.

5.6 Limitations and Scope of Inference

This study provides a detailed empirical snapshot of live loads in a single academic facility, based on continuous video surveillance over approximately one month (28–31 days, depending on zone). The four monitored areas—two hallways, a computer laboratory, and a stairwell—represent distinct usage profiles, though the results may not fully generalize to other building types, occupancy behaviors, or regional contexts.

Detector-level errors were evaluated using a manually labeled subset of frames. The mean absolute deviation was approximately 0.31 persons on a per-frame basis, corresponding to an uncertainty of about ± 0.02 kN/m^2 in instantaneous load estimates under typical occupancy. False positives were rare (<2%) and primarily occurred under partial occlusion. These deviations are statistically negligible compared to the 80%–90% discrepancy between measured and code-prescribed loads. However, because high-density events (groups larger than eight occupants) represented less than 0.01% of observations, error quantification under clustered conditions remains incomplete and is identified as a key target for future validation. With sufficient labeled high-density frames (≈ 100 or more), occlusion-related amplification factors could be tuned to further refine load estimates. This improvement will be feasible once a full academic year of monitoring is completed—approximately nine months of continuous data collection across nine zones within Building 13H—which will provide the density and temporal coverage required for reliable calibration under diverse occupancy conditions.

Person-to-weight mapping was based on conservative gender-specific average body weights (77 kg for males and 61 kg for females), consistent with international anthropometric statistics and regional studies [17]: 77 kg for males and 61 kg for females. Sensitivity analysis showed that varying these weights by ±10 kg changes instantaneous live loads by up to ±13%. Even under conservative assumptions (+10 kg), the 95th-percentile daily peak load remained more than 40% below the minimum design threshold of 2.5 kN/m^2, and return-level estimates shifted by less than 5%, confirming the robustness of the results.

Strict privacy protocols were enforced. Raw video data were processed locally, and only anonymized metadata (bounding boxes and gender tags) were retained; no identifiable images were stored. Future work will implement edge-processing and in-camera anonymization (e.g., pixelation outside bounding boxes) to further enhance compliance with data-protection policies.

Although this pilot was limited to one academic building and a one-month observation window, its methodological framework is fully transferable. Ongoing work is extending the monitoring to all corridor-type spaces in Building 13H over a full academic year, capturing seasonal and typological variability and paving the way for multi-building studies across different occupancy types.

5.7 Operational Profile

All processing is performed in Python, combining the Ultralytics YOLOv8 API for person detection with DeepFace for gender attribution. On a mid-range desktop (Intel i5-9400 CPU, 16 GB RAM, Radeon RX 550 2 GB), the median inference time is approximately 0.8 s per 1080p frame (0.45–3.3 s across test images). At that speed, a month of four-camera footage (~2.8 TB) can be analyzed overnight on a single workstation. Faster GPUs would push processing well below real-time, but the present throughput is sufficient for daily batch updates. The main outstanding challenges are organizational—maintaining continuous GPU availability and enforcing CCTV-data privacy—rather than computational bottlenecks.

6 Conclusions

This study demonstrated the feasibility of using a YOLOv8–DeepFace pipeline combined with continuous video surveillance to generate empirical live-load data in an operational academic building. The proposed framework processed more than 100,000 annotated frames over approximately one month of monitoring, producing detailed load–time histories and probabilistic extreme-value projections. Observed live loads—peaking at 0.35 kN/m^2 in laboratories and 0.53 kN/m^2 in stairwells—remained consistently below the conservative values prescribed by Eurocode (3.0–5.0 kN/m^2) and the International Building Code (4.8 kN/m^2). Sensitivity analyses further confirmed that these results are robust with respect to detector thresholds and conservative person-weight assumptions. Together, these findings provide empirical evidence that occupancy-induced live loads in typical academic environments are substantially lower than those implicitly assumed in current design provisions.

Based on these outcomes, several directions for future research and practice are identified. Future research should prioritize: (i) extending monitoring to a full academic year to capture seasonal variability; (ii) implementing homography-based mapping to resolve spatial clustering; and (iii) integrating edge-processing for enhanced privacy compliance. Rather than advocating immediate reductions in code-prescribed live loads, the proposed framework establishes a validated empirical baseline that can support performance-based assessment, digital-twin-assisted monitoring, and future discussions on adaptive building codes aimed at improving material efficiency and reducing embodied carbon in monitored buildings.

Acknowledgement: The authors would like to acknowledge the institutional support provided by the authorities of the Faculty of Engineering in Earth Sciences (FICT), the Administrative Management Office (Gerencia Administrativa), and all personnel of the ESPOL-911 department of the Escuela Superior Politécnica del Litoral (ESPOL), Ecuador, for granting access to the video surveillance infrastructure used in this study under the corresponding institutional authorization and privacy protocols. During the preparation of this work, the authors utilized the Large Language Models "Chat GPT", "Gemini" and "Deepseek" to assist in the writing process and enhance the manuscript's readability and language. Following the use of these tools, the authors reviewed and edited the content as necessary and assume full responsibility for the published article's content.

Funding Statement: This research was supported by the Escuela Superior Polite´cnica del Litoral (ESPOL), Ecuador. No external funding was received for this work.

Author Contributions: Luis Sánchez Calderón: conceptualization, methodology, software, data curation, formal analysis, writing—original draft preparation. David Valverde Burneo: validation, investigation, writing—review & editing. Walter Hurtares Orrala: supervision, resources, writing—review & editing. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: All analyzed data supporting the extreme-value modeling results are included within the Supplementary Materials accompanying this article (four CSV files). However, restrictions apply to the raw video footage due to institutional privacy policies regarding video surveillance data. The raw data is therefore not publicly available, but non-anonymized data may be available from the corresponding author upon reasonable request and formal authorization from ESPOL's ESPOL-911 department.

Ethics Approval: This study involved the analysis of anonymized video surveillance data collected from publicly accessible academic areas. It did not involve any intervention, human subjects, or animal use. Therefore, Ethics Approval is not applicable. Data handling complied strictly with ESPOL's institutional privacy protocols, including local processing and the retention of only anonymized metadata.

Conflicts of Interest: The authors declare no conflicts of interest.

Supplementary Materials: The supplementary material is available online at https://www.techscience.com/doi/10.32604/sdhm.2026.077137/s1. The following data are submitted as supplementary files for this article. The data supporting the findings of this study are available as four CSV files with the computed live loads and its correspondent timestamp value for each zone monitored on this study:

- LiveLoad_Data_Month_13H-PA-ESCALERA.csv
- Ex_LoadData_13H-PA-LAB.csv
- LoadData_13H-PA-PASILLO 1_extracted_frames.csv
- LoadData_13H-PA-PASILLO 2_extracted_frames.csv

References

1. Cha YJ, Ali R, Lewis J, Büyüköztürk O. Deep learning-based structural health monitoring. Autom Constr. 2024;161(3):105328. doi:10.1016/j.autcon.2024.105328.
2. Washington State Building Code Council. International building code. Washington, DC, USA: Washington State Building Code Council; 2008.
3. EN 1991-1-1. Eurocode 1: actions on structures—Part 1-1: general actions: densities, self-weight, imposed loads for buildings. Brussels, Belgium: European Committee for Standardization; 2002.
4. JCSS-OSTL/DIA/VROU-10-11-2000. Probabilistic model code Part 1—basis of design. Zürich, Switzerland: Joint Committee on Structural Safety (JCSS); 2000.
5. Wu W, Dou K, Xu C, Li Y, Weng Q, Chen J. Floor live load survey and modelling by house VR viewing and public registration information. J Build Eng. 2024;98(5):111379. doi:10.1016/j.jobe.2024.111379.
6. Andam KA. Floor live loads for office buildings. Build Environ. 1986;21(3–4):211–9. doi:10.1016/0360-1323(86)90032-6.
7. Kumar S. Live loads in office buildings: lifetime maximum load. Build Environ. 2002;37(1):91–9. doi:10.1016/S0360-1323(00)00075-5.
8. Kumar S. Live loads in office buildings: point-in-time load intensity. Build Environ. 2002;37(1):79–89. doi:10.1016/S0360-1323(00)00074-3.
9. Tapia-Hernández E, Dominguez-Palacios AC, Martínez-Ruíz M. Live loads on floors of libraries and newspaper archive buildings. Int J Adv Struct Eng. 2019;11(2):285–96. doi:10.1007/s40091-019-0230-8.
10. Chalk PL, Corotis RB. Probability model for design live loads. J Struct Div. 1980;106(10):2017–33. doi:10.1061/jsdeag.0005542.
11. Costa LGL, Beck AT. A critical review of probabilistic live load models for buildings: models, surveys, Eurocode statistics and reliability-based calibration. Struct Saf. 2024;106(3):102411. doi:10.1016/j.strusafe.2023.102411.
12. Actions on structures—live loads in buildings. Ottawa, ON, Canada: International Council for Building Research Studies and Documentation; 1989. Report No.: 116.
13. Zhang W, Calautit J, Tien PW, Wu Y, Wei S. Deep learning models for vision-based occupancy detection in high occupancy buildings. J Build Eng. 2024;98:111355. doi:10.1016/j.jobe.2024.111355.
14. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016 Jun 27–30; Las Vegas, NV, USA. doi:10.1109/CVPR.2016.90.
15. Girshick R, Donahue J, Darrell T, Malik J. Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition; 2014 Jun 23–28; Columbus, OH, USA. doi:10.1109/CVPR.2014.81.
16. Girshick R. Fast R-CNN. In: Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV); 2015 Dec 7–13; Santiago, Chile. doi:10.1109/iccv.2015.169.

17. Sinchiguano BY, Sinchiguano YK, Vera EM, Peña SI. Estado nutricional en estudiantes de medicina de la Universidad San Gregorio de Portoviejo, Ecuador. Rev Española De Nutr Comunitaria. 2023;29(1):e41600. (In Spanish).
18. Tien PW, Wei S, Chow TW, Darkwa J, Wood C, Calautit JKS. Enhancing the detection performance of a vision-based occupancy detector for buildings. Proc Inst Civ Eng-Eng Sustain. 2023;176(6):301–14. doi:10.1680/jensu.22.00013.
19. Márquez Peñaranda JF. Variabilidad de las cargas vivas obtenidas del análisis de estructuras. Rev De Ing De Construcción. 2020;35(1):73–80. doi:10.22463/17948231.980 (In Spanish).
20. Coles S. An introduction to statistical modeling of extreme values. London, UK: Springer; 2001.
21. Li Y, Chen J, Wang P. Vision-based survey method for extraordinary loads on buildings. Front Struct Civ Eng. 2024;18(6):815–31. doi:10.1007/s11709-024-1029-7.