11-2021

Using role play to develop an empathetic mindset in executive education

Anuradha RAVI
Singapore Management University

Kasun GAMLATH
Singapore Management University

Siyan HU
Singapore Management University

Archan MISRA
Singapore Management University

Citation

RAVI, Anuradha; GAMLATH, Kasun; HU, Siyan; and MISRA, Archan. Using role play to develop an empathetic mindset in executive education. (2021). *BuildSys '21: Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 61-70.
Available at: https://ink.library.smu.edu.sg/sis_research/6267

CS-Light : Camera Sensing Based Occupancy-Aware Robust Smart Building Lighting Control

Anuradha Ravi
School of Computing and Information Systems
Singapore Management University
anuradhar@smu.edu.sg

Kasun Galmath
School of Computing and Information Systems
Singapore Management University
kasung@smu.edu.sg

Siyan Hu
School of Computing and Information Systems
Singapore Management University
siyanhu@smu.edu.sg

Archan Misra
School of Computing and Information Systems
Singapore Management University
archanm@smu.edu.sg

Abstract

We describe the practical development of a smart lighting control system, *CS-Light*, that uses a preexisting surveillance camera infrastructure as the sole sensing substrate. At a high level, the camera feeds are used to both (a) estimate the illuminance of individual, fine-grained (roughly $12m^2$) sub-regions, and (b) identify sub-regions that have non-transient human occupancy. Subsequently, these estimates are used to perform fine-grained (non-binary) power optimization of a set of LED luminaires, collectively minimizing energy consumption while assuring comfort to human occupants. The key to our approach is the ability to tackle the challenging problem of translating the luminance (pixel intensity) of image frames into accurate estimates of the illuminance (LUX) of the various sub-regions, under variations in ambient lighting and layouts. To overcome this challenge, we develop a novel technique that (a) classifies image pixels as corresponding to light vs. dark-colored surfaces, and (b) uses unsupervised ML-based color-specific, pixel-to-LUX classifiers and statistical aggregation to provide robust LUX estimates. Experimental studies, conducted over a collaborative work area in an operational ZEB, demonstrate *CS-Light's* efficacy: it supports accurate pixel-to-LUX estimation (median error= 8.5%), and its real-time multi-LED adaptation results in appreciable energy savings (63.5% in low occupancy situations), while ensuring negligible perceptual discomfort to human occupants.

CCS Concepts

• **Computer systems organization → Embedded and cyber-physical systems**.

Keywords

Smart building, Smart lighting, LED Lighting, cyber physical system

BuildSys '21, November 17–18, 2021, Coimbra, Portugal

ACM ISBN 978-1-4503-9114-6/21/11…$15.00
https://doi.org/10.1145/3486611.3486657

ACM Reference Format:
Anuradha Ravi, Kasun Galmath, Siyan Hu, and Archan Misra. 2021. *CS-Light* : Camera Sensing Based Occupancy-Aware Robust Smart Building Lighting Control. In *The 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '21), November 17–18, 2021, Coimbra, Portugal.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3486611.3486657

1 Introduction

Despite the adoption of power-efficient LED luminaires, lighting remains the second-most dominant component of energy consumption (after HVAC) in smart buildings. To reduce lighting energy consumption without adversely affecting occupant safety and comfort, many commercial buildings have adopted occupancy-based adaptive lighting strategies. Broadly speaking, most such strategies involve binary and coarse-grained (on/off) control, whereby motion sensors are used to detect the presence vs. absence of humans and the lights are dimmed when an *entire area* is estimated to be unoccupied. Such binary strategies are often sub-optimal for open-plan office layouts (such as the collaborative work area illustrated in Figure 1), where only part of the space may be occupied and where a specific luminaire's impact diffuses over a wider area [8]).

In this work, we present *CS-Light*, an operationally-deployed system that performs fine-grained, occupancy-driven, **non-binary** lighting intensity adaptation, over such collaborative open indoor spaces. One of *CS-Light's* key and attractive features is that the entire adaptation process is based on inputs provided by an existing infrastructure of CCTV cameras, deployed previously to ensure campus safety and security–i.e., it does not require any additional custom sensor (e.g., lux sensors) deployment or external knowledge (e.g., level of ambient lighting). In addition, *CS-Light* also supports daylight harvesting [5, 22], automatically taking advantage of external daylight that enters via glass windows/doors. Most importantly, *CS-Light* is *robust* to a variety of real-world artifacts (e.g., changes in layout, occlusion effects under medium-to-high occupancy density) that previously-proposed approaches do not effectively tackle. At a high-level, *CS-Light* operates as follows. The image frames captured by such cameras are used to perform both (a) *occupancy sensing*: state-of-the-art AI-based vision algorithms are used to identify the sub-regions currently experiencing non-transient human occupancy, and (b) *illuminance estimation:* novel algorithms are used to translate the intensity of pixel clusters to the current

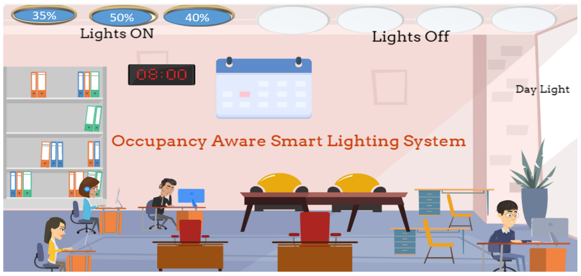

Figure 1: Occupancy-Aware, Fine-Grained Smart Lighting

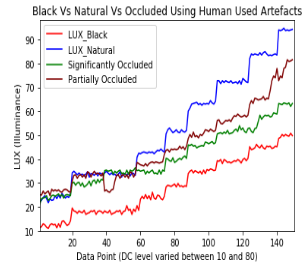

Figure 2: LUX-vs.Pixel (Day)

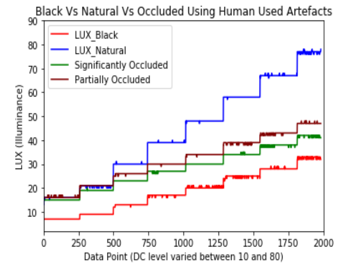

Figure 3: LUX-vs.Pixel (Night)

illuminance (LUX) level of corresponding sub-regions. These inputs are then used to dynamically control LED luminaire intensity level, so as to minimize the total power consumption while ensuring adequate illuminance at all human-occupied sub-regions.

The Challenge of Accurate LUX Estimation: Accurate illuminance estimation is the most challenging of the various *CS-Light* components. In commercial buildings, such illuminance is often measured by using ceiling-mounted light sensors (LUX meters). However, our experimental studies, conducted in an indoor 100m^2 open-collaboration space (detailed in Section 3) reveal that such measurements are sensitive to variations in the reflectivity of commonplace objects placed by humans, on the tables and floor, during regular usage. Figures 2 & 3 plot the LUX readings measured by multiple ceiling-mounted LUX sensors, during daylight and night conditions respectively, as the intensity of the LED lights are varied (by varying the luminaire's duty cycle (*DC*) values). The different lines in the plots correspond to different ambient conditions, when the tables are covered either by black-colored objects or are completely uncovered (thereby revealing the furniture's natural white color), as well as partially covered by objects (such as laptops, bags and books) during regular usage. We can see that the LUX readings differ by $\sim$ 25 lumen ($\approx$ 33%) when the table is completely uncovered vs. covered with black objects, and $\sim$ 12 lumen, on average, under different levels of naturally-occurring human occupancy and usage. As a consequence, the ceiling-mounted LUX sensor paradigm not only imposes additional operational overhead (e.g., for replacing batteries) and measures ceiling-level illuminance (occupant comfort is known to be higher [13] when lighting adjustment is based on table-level measurements), but is also vulnerable to natural variations in ambient reflectivity. An alternative approach, of embedding the LUX sensors on the work tables, also generates erroneous readings, when being obscured by objects that occupants place on the tables. In contrast, *CS-Light* attempts to estimate the table-top illuminance levels based on the images captured by the CCTV cameras.

While image pixel-to-LUX estimation is not conceptually new [18], we shall experimentally demonstrate that such past approaches fail to consider the high variability in pixel-to-LUX mappings observed naturally, due to factors such as changes in ambient lighting levels, changes in the color (light vs. dark), type and reflectivity of objects present in the visual scene and possible saturation of the image's pixel values. *Practical* and *robust* illuminance estimation is thus a challenging, non-trivial problem. Figures 4 & 5 illustrate such variability observed at two distinct times of the day over 2 regions within our experimental testbed: for region (a), we observe that the illuminance has changed dramatically, due to changes in the LED & ambient lighting, while the pixel intensity remains roughly the same; in contrast, for region (b), while the illuminance remains unchanged, the image pixel intensity changes with the color of foreground objects. We shall show *CS-Light's* individual components are carefully designed to overcome these challenges.

Key Contributions: Our key contributions are:

- *Robust & Stable Pixel⟶LUX Mapping:* We introduce a practical and robust approach, for mapping pixel values of camera images to illuminance (LUX) levels at the corresponding sub-regions. The approach first partitions the large collaborative space into multiple sub-regions, with each sub-region then subject to a two-step process: first, classifying each micro-region within the sub-region into a *dark vs. light* colored surface using a state-of-the-art classifier [27], followed by the use of a color-specific ML-based model (built during an offline, fingerprinting phase) to map the corresponding pixel values to a presumptive *luminance range*. The resulting ranges from each micro-region are then statistically aggregated to estimate the overall illuminance (LUX) level of each sub-region. We show that this approach is significantly more accurate, achieving a median LUX estimation error of 8.5% across a diverse real-world variety of environmental and usage artifacts.
- *Practical & Accurate LED Lighting Control:* We develop an optimization based approach for LED lighting control that implements *non-binary* LED adaptation. This approach utilizes a fingerprinting-based strategy [18] that effectively captures the non-uniform "channel" between a cluster of LEDs and each sub-region, and then uses this channel matrix to iteratively modify the intensity (duty cycle) of individual LED clusters and ensure that we satisfy each sub-region's illuminance constraint. Using diverse real-world deployment conditions, we show that the the illuminance stays within ±10% of the ideal levels.
- *Practical Validation of* CS-Light*:* We deploy and evaluate *CS-Light's* performance in an open-layout, shared collaboration space of ~100 m^2 within a university ZEB (Zero Energy Building). We demonstrate that *CS-Light* is effective under diverse occupancy regimes, (a) achieving energy savings of 63.5% (low occupancy) and 37% (high occupancy), compared to a non-adaptive baseline, and (b) saving ~26% energy savings vs. a binary control mechanism. Equally importantly, explicit occupant feedback demonstrates that such energy savings are achieved without any degradation in occupant-perceived comfort. Additionally, results from a longer duration, operational deployment show that

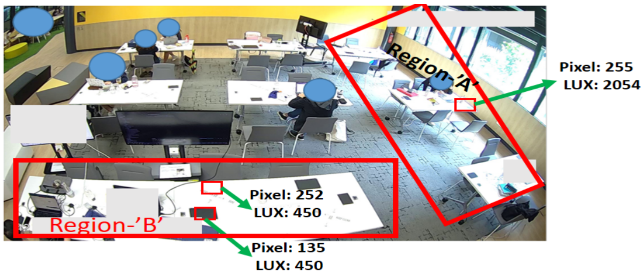

Figure 4: (Pixel,LUX) Values: CollabZone Daytime View

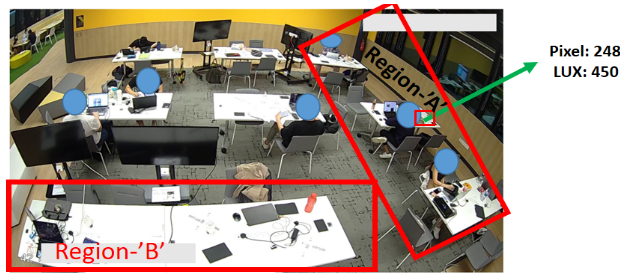

Figure 5: (Pixel,LUX) Values: CollabZone Nighttime View

CS-Light continues to provide ∼40% energy savings in spite of significant, uncontrolled variations in ambient layout and usage.

2 Related Work

The past work on automated, occupancy-aware lighting control in smart spaces can be characterized along multiple dimensions, including the choice of sensors, the type of environment targeted and the techniques used for camera image-based LUX estimation.
Occupancy & Illuminance Sensing: Researchers have explored different strategies (e.g., [19]) to optimally tune the intensity of lights to reduce energy consumption, while maintaining an acceptable level of comfort for the occupants. Recognizing the diffusive properties of open spaces, researchers have recently moved beyond classical binary (ON/OFF) control strategies to perform finer-grained illumination control. Lee et. al. [10] proposed the usage of desk-mounted photometry sensors (measuring illuminance) and RFID based user tags (providing location) to perform occupancy-aware adaptation of ceiling luminaires, while taking into account the natural ambient illumination (daylight harvesting). Meugheuvel et.al. [21] utilize occupancy and light sensors mounted alongside each LED luminary to measure illuminance and occupancy, and communicate such status across LEDs to perform distributed lighting control. More recently, the ReViCEE framework [7] utilized similar luminaire-attached sensors to measure desk-level illuminance and occupancy, and proposed a centralized, joint control of multiple LED luminaires while incorporating explicit user preferences. A variety of past work (e.g., [4, 25] have explored the deployment of other sensors (e.g., RFID user tags, microwave sensors) to measure occupancy together with photosensor-based illuminance estimation. To support stable indoor illuminance via daylight harvesting, SunCast [11] uses historical patterns of sunlight variation to predict time-varying sunlight intensity, and subsequently adjust the transparency level of external windows. Pandharipande et. al. [15] explored the use of ultrasound sensors and photosensors to estimate illuminance and the trajectory of a person walking in an office-space to uniformly distribute LED lighting across the occupied and unoccupied regions. Such past work has largely ignored the significant errors in LUX estimation, caused by real-world artifacts such as the presence of objects with varying color/reflectivity or desktop occlusion.
Implicit & Image-based Illuminance Estimation: Several researchers have utilized theoretical models, correlating illuminance with LED lighting intensity. For example, Zevgolis et. al., [26] placed image sensors on the ceiling to estimate the luminance maps via idealized, isotropic Lambertain reflectance models [9]. Similarly, lighting control algorithms in [2, 3] utilize LUX estimates derived from the distortion of illuminance patterns (captured by photometric sensors) across different locations in an open office. Researchers have previously used HDR images [14], as well as RGBD sensors that provide depth information [20], to estimate the luminance intensity and further estimate the illuminance of a given scene. Our proposed approach of estimating the lux from the pixel intensity of images is conceptually similar to techniques described in [12, 18]. However, the relatively-simple mathematical models proposed for pixel-to-LUX estimation are inadequate in capturing a variety of real-world challenges, such as changes in surface reflectance caused by commonplace objects and day-vs.-nighttime variations in ambient lighting.
Neural Models: Recent work (e.g., [6, 23]) has also explored the possibility of using DNN models to develop strategies for effective lighting control and building management. Wang et. al.[24] used a 2-layer neural network to model the non-linear relationship between the DC level and table-top measured LUX values within an open-plan office space. While these DNN models effectively replace idealized models (such as Lambertian) on the spatial diffusion of lighting, they do not tackle the problem of image→LUX estimation.

3 Real-World Experimental Testbeds

To build and test our lighting control strategies, we conduct experimental studies on two separate physical locations, described in chronological order of use.

3.1 Micro-Study Lab Testbed

We first utilized a $25m^2$ area research lab-based testbed (Figure 6), purely to perform *controlled* micro-studies and understand the LUX→pixel behavior. This research lab was subject to varying levels of ambient daytime lighting, arriving by a window located on one outer wall, and included 4 work cubicles with *birch* table-tops and additional workplace furniture. We installed 6 clusters of 8-LED puck lights that could collectively generate a maximum intensity of ∼400 lm, and that were controlled by a set of Raspberry Pi devices. In addition, to carefully understand how colored object affected the (lux,pixel) relationship, we placed a variety of colored "patches" (paper strips) and commonplace objects (such as bags and laptops) to emulate human artifacts, at multiple locations on the cubicles and furniture. The research lab was also instrumented with 2 Dahua ceiling-mounted cameras (same versions as our deployment testbed) with auto-tuning switched-off.

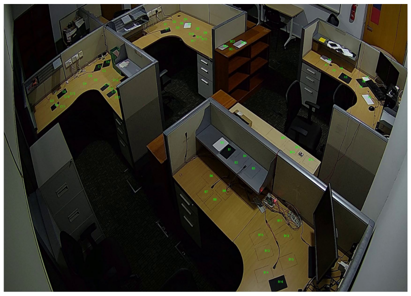

Figure 6: Research Lab TestBed

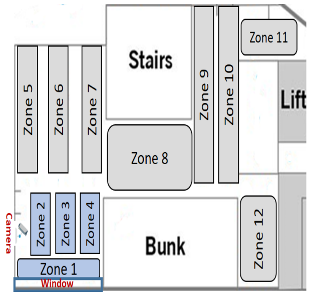

Figure 7: Net-ZEB CollabZone

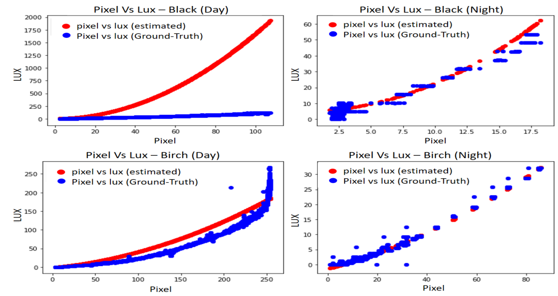

Figure 8: LUX vs Pixel (Black & Birch - MicroTestbed

3.2 Net-ZEB Collaborative Zone (CollabZone)

The *CS-Light* system was operationally deployed on one floor of a 5-story Net-ZEB equipped with Building Management System (BMS), which permits DALI (Digital Addressable Lighting Interface)-based control of individual LED luminaires. This floor houses an open-floor space of $\sim 750m^2$ that university students and staff can normally access 24x7 for collaborative study and work. The overall collaborative area is partitioned into a set of smaller sub-zones (as illustrated in Fig. 7): each subzone has a cluster of 6 LED luminaires controlled jointly by a single DALI end-point. The CollabZone is also visually monitored by ceiling-installed Dahua cameras. We turned OFF the camera's auto-tuning features (gain, exposure and white balance), keeping them fixed at all times, to provide reliable and reproducible image frames under varying illuminance and weather conditions. We empirically verified that disabling the auto-tuning did not materially affect the quality of the captured images used for conventional security/safety monitoring.

For our final measurement studies and analysis, we principally focus on Zones 1-4 (an area of $\sim 70m^2$ with *white* tabletops, as illustrated in Figures 4 & 5), as these areas typically exhibited the highest occupancy during the observation period. While Zones 2-4 have a cluster of 6 horizontally-spaced LED luminaires, Zone 1 has 6 vertically-spaced LEDs and is adjacent to a window (as indicated in the figure) that lets in ambient light and is equipped with manually-controlled blinds (whose status is not explicitly known to *CS-Light*). Under normal pre-*CS-Light* operation, the LED lights were lowered to 5% intensity during the early morning hours and operated with DC =75% otherwise. Additionally, the CollabZone is also equipped with standard motion sensors (3-4 sensors in each zone), to (a) gradually decrease the intensity of the LED luminaires in the sustained absence of motion, and (b) instantaneously increase the LED intensity to 75% if motion is detected.

4 Challenges in Pixel - Lux Mapping

The ability to map the pixel content of camera images to illuminance intensity is a core requirement for *CS-Light*. However, in contrast to past work (e.g., [12] that employs fairly straightforward Lambertian reflectance models, our real-world studies (demonstrated in Figures 4 & 5) show that the pixel-to-LUX relationship is non-trivial and varies both across time and space due to a variety of natural artifacts.

Through careful controlled experiments (involving deliberate placement of different-colored patches on tables) on our micro-study testbed, we established that this relationship (a) was non-linear (with a strong quadratic component), and (b) depended strongly on the nature/color of the object being captured within each pixel cluster. In particular, object surfaces can be broadly categorized as: (a) *Low-Intensity*: corresponding to dark-colored (such as black or dark brown) surfaces, where the pixel intensity is significantly lower even when the illuminance is very high, and (b) *High-Intensity*: corresponding to lighter-colored (white or birch) surfaces, where the pixel intensity is significantly higher even under low illuminance. As illustrated in Figure 8, under varying DC levels of LEDs, we observe that: (a) Black surfaces have significantly lower pixel intensity values, reaching at most 130+ even under extremely high illuminance (LUX~=1200 - under daylight conditions); (b) in contrast, Birch (natural veneer) surfaces can reach pixel intensities close to saturation (250+)[1] under similar illuminance conditions; and (c) a single (quadratic) regressor does not work well for both daylight and night lighting conditions.

Table 1: LUX vs. Pixel: High Variability

Color	Time-of-Day	Pixel	LUX (Est)	LUX (GT)	TestBed
Birch	Night	135.77	32.62	32.08	Micro
Birch	Day	255	121.6	98	Micro
Black	Night	11.9	30.0	31.6	Micro
Black	Day	135.24	121.6	1930	Micro
White	Night	255	270	384.2	Collab
White	Day	255	584	384.2	Collab
Black	Night	18.6	70.8	84	Collab
Black	Night	18.6	70.8	102	Collab

These observations are further reinforced by the values plotted in Table 1, corresponding to measurements on both the Micro (Research) and the CollabZone sites. We see that, for the birch and white colored patches, the pixel intensity is high, even under relatively low illuminance during the day, with the pixel values saturating at 255 under high illuminance. For the Black colored patch, the changes in pixel intensity are relatively muted even when the illuminance changes significantly. **Dependency on Light Source Location:** To further confound matters, we additionally established that this lux-vs.-pixel relationship is not just surface color and lighting quality (natural daylight vs. LED) dependent, it also varies based

[1]With each pixel represented by a 24-bit RGB value, the dynamic range for camera pixel intensity varies from 0-255.

on the physical location of the patch relative to its neighboring LED luminaires. In other words, due to the natural directionality of the light reflected from different LED sources, the pixel→LUX mapping is also *location-specific*. To illustrate this phenomenon, Table 2 tabulates the (pixel, LUX) readings observed in our testbeds, for both Black and White/Natural patches at representative locations. We see that, on the Micro-Lab testbed, a pixel intensity=109 maps to both LUX=17.0 (when illuminated solely by LED with DC=100% placed on top of cubicle B) and LUX= 9.3 (when illuminated solely by an LED with DC=40% on top of cubicle F).

Table 2: Lux-vs.-Pixel: Different Lights

LED DC	LUX)	NaturalPixel	BlackPixel	TestBed
LED-B (100%)	17.08	109.8	35.2	Micro
LED-F (40%)	9.3	109.8	32.6	Micro

Key Takeaway: Our experimental studies demonstrate the need to (a) first, during an offline phase, create distinct **location and color/surface** (Black vs. Natural/Light) specific pixel-to-LUX models, and then (b) during runtime, perform location-specific classification of the color/surface.

5 CS-Light System & Key Components

We now describe the overall design of the *CS-Light* system and its key components, explicitly differentiating between: (a) the offline training phase used to build the various illuminance estimation models, and (b) the online phase where *CS-Light* estimates LUX levels from camera images and issues commands to modify the LED brightness (DC) values.

5.1 Offline Phase - Fingerprinting and Modelling

Fingerprinting Process and Patch Definition: To first create the color-specific pixel-to-LUX mappings, we adopt the fingerprinting-based approach: we place multiple {black, white} paper strips at table-level in a zone, and measure the intensity of camera-captured pixels as the LED light intensity (DC) is varied between 5-100%. (Ground-truth LUX readings for the annotated patches are obtained using collocated Texas Instrument BLE-equipped SensorTags equipped with a light sensor.) We further manually annotate the images captured during the fingerprinting process with each annotation referring to a cluster of pixels , called a "patch", with size ≈25*25 pixels. The number of patches (per zone) varies between 8–25, depending on the size of a zone and the distance of the zone from the camera. Such patches correspond to an 10*10cm^2 area if the patch is near to the camera, an an 70*70cm^2 are when farther away. A larger number of patches provides greater robustness against occlusion (caused as the occupants fill up the region with artifacts such as laptops and books). From an empirical 3.5 hour study, 78% of patches gave correct LUX estimations during real-time, while 22% were erroneous. Accordingly, to improve accuracy, we utilize the median of the lux estimates (Algorithm 1) across all the patches. Given our observation that a majority of occupant-specific objects (laptops, phones and bags) are either dark or light in color, we focus purely on building a two-class model: one for high-intensity (white, birch and similar) surfaces and another for low-intensity (black, dark brown and similar) surfaces. Figure 9 outlines the (offline) fingerprinting process. While fingerprinting is performed using only the LED luminaires at night-time (to avoid ambient lighting fluctuations), we shall see that this approach works well under daylight conditions as well.

Color/Patch Classifier: As the first step in the runtime execution of *CS-Light*, we will need to classify the type of surface (light-colored vs. dark) corresponding to each pixel micro-cluster. Based on similar prior work for light-vs.-dark classification [16], we utilize an off-the-shelf color-recognizer [27], that computes the histogram of "R,G,B" values for a given micro-cluster/patch and uses a KNN (K-Nearest Neighbour) classifier to classify the corresponding patch. In empirical testing, we discovered that due to the absence of sufficient daylight-based training data, such a model could incorrectly classify white patches as black under low-light daytime conditions. Hence, we incorporate additional statistical spatial smoothing techniques (described in Algorithm 1) for robust LUX estimation.

Pixel-to-LUX Mapping: To create a pixel→LUX map, we first extract the 'mean' pixel values for the pre-defined patches in the fingerprinted images and the corresponding measured LUX values. To build this map, we utilize only the night-time data (i.e., only under LED illumination) as the daylight LUX values are often well beyond the specified illuminance level (300 − 500 lumens, as per local regulations) for human-occupied spaces. As mentioned earlier, and illustrated in Figure 10, the (pixel, lux) ranges also depended on the source of the illumination–i.e., whether the light incident on a particular patch was primarily from an overhead LED luminaire or from LEDs located in nearby sub-regions. Accordingly, instead of embracing the regressor approach that we investigated initially, we adopted an *unsupervised clustering approach*, using the well-known DBSCAN algorithm to perform clustering of the observed (lux, pixel) tuple, for *each* (patch, location) combination. Once the clusters are determined, we represent each cluster by a 4-element vector that includes: (a) (min,max) of pixel values, and (b) the (min, max) of LUX values observed across cluster elements. We shall shortly see how the runtime illuminance predictor utilizes such multiple, possibly overlapping clusters.

Weight Matrix (Transfer Function): As detailed in Section 6, the adaptive control of LED lights requires knowledge of the relationship between an LED-cluster's intensity and the LUX level of the zone and it's nearby zones. This relationship is defined via a *weight matrix* $\mathcal{W}$, such that W_{ij} represents the contribution from LED-cluster i to the sub-region j. Mathematically, if DC_i is the current duty cycle of luminaire i, it contributes a value $DC_i * W_{ij}$ to the LUX value at sub-region j. We empirically verified the linear LUX-vs.-DC relationship implicitly assumed in this model–e.g., Figure 11, which plots the LUX-vs.-DC values across zones (sub-regions) in CollabZone confirms this hypothesis. The matrix elements are determined offline by independently varying the intensity of each zone's LED and observing the slope of the change in the LUX value.

5.2 Online Phase - LUX Estimation

The online phase involves the following steps:

- Capturing the images from the existing surveillance cameras (no additional cameras/equipment deployed for *CS-Light*)
- Localize occupants to zones (sub-regions)
- Estimating the luminance from images (using micro-regions).

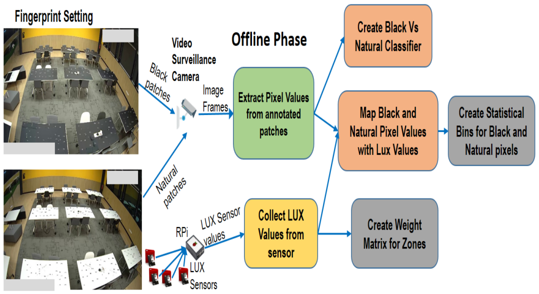

Figure 9: CS-Light Fingerprinting Process (Offline)

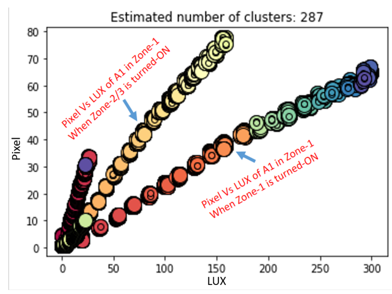

Figure 10: Clusters defined over (Pixel,LUX) Values for Black Patch

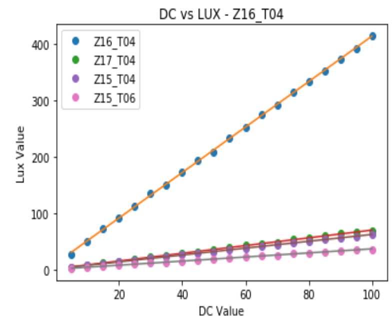

Figure 11: LUX vs. DC Variation (LED-1)

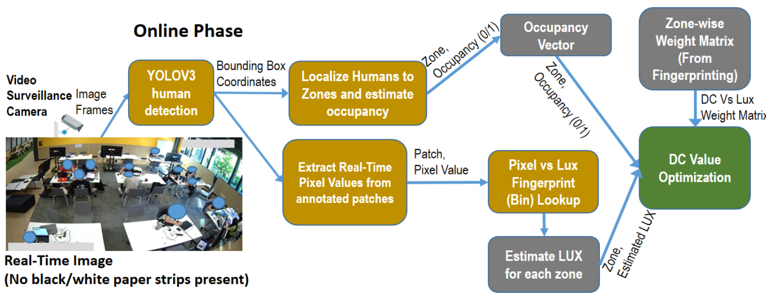

Figure 12: CS-Light Real-Time Estimation

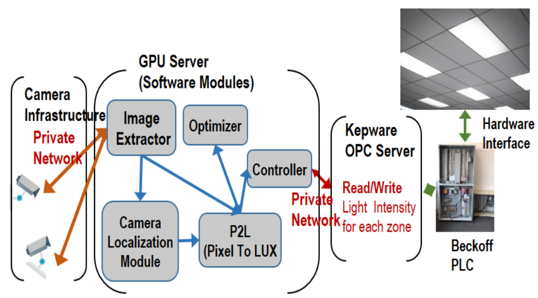

Figure 13: *CS-Light* System Implementation

- Map luminance to illuminance (LUX) values
- Aggregate occupancy and LUX values
- Estimate the ambient light
- With the estimated LUX values, ambient light and localized occupants, determine the new optimal LED lighting levels.

We now describe the four key steps, deferring the last step (determining the optimal LED lighting levels) to Section 6. Algorithm 1 describes the logical steps involved in classifying a patch and estimating the LUX.

Image Capture: We capture real time images from the surveillance cameras installed in the building. Note that the environment no longer contains any black/white paper strips (those were used only for offline model building). Our *CS-Light* solution also merely extracts 'anonymous' human object bounding box coordinates from the camera image stream (and does not store any images), satisfying the privacy requirements of campus facilities management.

Localizing Humans: The captured images from the surveillance cameras are fed through an object detector (Yolo-V3 [17]) to extract bounding box coordinates of detected human objects. These bounding boxes coordinates are further passed through a localization module, which effectively uses a classifier (trained using site-specific data) to map the object's bounding box pixel coordinates to the corresponding sub-zone. We use such localization data to calculate the binary-valued *occupancy vector OV*, with the i^{th} element equal to '0' if the sub-zone i is unoccupied, and '1' otherwise. Occupants can move around the table and often place objects (e.g., monitors) that can partially occlude the table surfaces (sub-regions).

LUX Estimation: For a robust LUX estimation, we first extract the mean pixel intensities (Y-Values) of the pre-defined patches (pixel clusters) and use binary KNN-based patch classifier to distinguish between {high,low} intensity surface classes. We then use the class-specific clusters (bins) to determine the feasible range of associated LUX values per patch. To overcome the errors in patch classification (specifically during low-light conditions), we "separately estimate" the LUX values for dark-colored and light-colored surfaces of each zone; note that each such sub-zone may have a mix of dark and light sub-regions (surfaces). We then compute the median of the LUX values for each of the two patch classes, and finally compute the average of these two median values. However, if the median values obtained for black and white patches vary significantly (empirically set to $\geq 100 lumen$), we assign the cumulative LUX to be equal to the median LUX of white colored patches. This is based on the observation that LUX values estimated for white patches are more accurate than black patches at low illuminance (which suffer from low sensitivity).

Aggregation: To eliminate sporadic estimation errors, we average each zone's LUX estimate over each consecutive 2-min window. Also, because the object detector might miss human objects in a single frame (especially in low light and partially occluded regions) or capture transient movement (false positives), we first aggregate the occupancy vector over 2 minutes, classifying a zone as occupied if it was declared to have human objects 4 or more times (i.e., $\geq$ 50% of the frames sampled at $\frac{1}{15}$ fps) within the 2 minute interval.

Algorithm 1 Patch Classification and LUX Estimation Process

```
for Each Zone in List of Zones do
    For Each Patch:
    Calculate Mean of Y values for each patch (P_i)
    Classify patch color using KNN Classifier
    if Patch Color = Black then
        Select the closest LUX bin corresponding to P_i for black
        Add estimated LUX to C_lux_black
    else if Patch Color = White then
        Select the closest LUX bin corresponding to P_i for white
        Add estimated LUX to C_lux_white
    end if
end for
if MEDIAN (C_lux_black) - MEDIAN (C_lux_black) > Threshold then
    Estimated LUX = C_lux_white
else
    Estimated LUX = MEAN(C_lux_black, C_lux_white)
end if
```

6 Optimization Control Logic

The real-time aggregated LUX values and occupancy vector, together with the pre-computed weight matrix and current DC values of the LED luminaires (obtained from the BMS), are passed as input to the optimizer. As the first step, the optimizer calculates the required illuminance levels for each zone depending on the occupancy vector and the ambient light present in each zone. The ambient light AL is calculated by subtracting the *anticipated* LUX value, DCL, of zone i (obtained via the product of the weight matrix and current LED levels) from the image-derived LUX estimate IL–i.e., for any zone i, $AL(i) = IL(i) - \sum_j DC_j * W_{ji}$. If the image LUX estimation is less than the DC LUX value (due to estimation errors), we consider the ambient LUX as 0; in practice, AL is computed only during daylight hours.

Subsequently, the optimizer determines the required LED contributed LUX value for each sub-zone i, designated by $RL(i)$, needed to ensure adequate illuminance. If the zone is unoccupied, the target illuminance is set to the mandated minimum safety value $SafetyL$; else, it is set to the comfort level $ComfortL$. RL is thus obtained as:

$$RL(i) = max(0, (OV[i]*ComfortL+(1-OV[i])*SafetyL)-AL(i)) \quad (1)$$

Given these objectives, the optimization problem can then be expressed (Equation 2) as one of minimizing the total power of all the N LED luminaires, across all $\mathcal{J}$ zones, while satisfying the required LUX values in each individual zone–i.e.,

$$\min \sum_{i=1}^{N} DC(i) \text{ such that} \sum_{i=1}^{N} W_{ij} * DC(i) \geq RL(j) \quad \forall j \in Zone; \quad (2)$$

We solve the above optimization equation (which implicitly achieves non-binary lighting control) using the least squares method to find the best possible solution (DC values) for all the zones.

7 Practical CS-Light Implementation

We now describe the implementation of *CS-Light* (illustrated in Figure 13) that is currently deployed in the Net-ZEB Collaboration Zone (Section 3.2). The *CS-Light* software components are executed on an Nvidia Tesla T4-equipped GPU server, with the GPU cores used to run the object detector and pixel extraction tasks. The coordination between the individual *CS-Light* components is implemented via the Kafka middleware [1], which supports high throughput, low-latency processing of event streams. The hardware infrastructure includes the already-deployed Dahua[2] IP camera infrastructure and the Beckoff PLC system[3] used to control the LED intensity values via standard DALI interfaces.

The overall system consists of the following modules:

- **Kepware OPC Server:** The Kepware OPC-UA server acts as a middleware between *CS-Light* and Beckoff PLC. We subscribe to the OPC tags (published by the Kepware OPC server whenever the tag value changes) via Python-3's OPC-UA package. These OPC tag values are stored in a Postgres SQL DB and also published via the Kafka middleware to be utilized by the other *CS-Light* components.
- ***CS-Light* Software:** The *CS-Light* software consists of 5 modules– *Image Extractor, Camera Localizer, Smartlight Optimizer, Pixel2LuxMapper (P2L) and the Smartlight Controller.* The Image Extractor component uses an RTSP (Real-time Streaming Protocol) based private API to retrieve the image feeds from 6 different cameras, with a nominal rate of 1/15 FPS. Each such image is then streamed to both (a) the P2L module to perform LUX estimation, and (b) the Camera Localizer to identify the location coordinates of the human objects. The Localizer uses a state-of-the-art 800*800 YOLOV3 object detector [17] to first identify each bounding box, with a logistic regressor then mapping each such box to one of the 4 Collaboration Area sub-zones (Z1-Z4, as seen in Fig 7). The bounding boxes coordinates are fed (with appropriate timestamps) to the P2L estimator. The P2L module uses these inputs to estimate each subzone's LUX value, averaged over 2 mins. These estimates are then fed into the SmartLight Optimizer, which uses the Python *numpy* linear solver to determine the 'optimal' values (the duty cycle) for each of the LED luminaires. These optimal values are then provided to the SmartLight Controller, which uses an OPC-Client to adjust the operation of each LED luminaire. To avoid rapid, jarring changes to the lighting levels, the Controller performs a linear adjustment of each luminaire's intensity, modifying the DC values gradually once every 30 secs.

8 Experimental Results

We now present results based on extensive evaluation of *CS-Light* in the ~$100m^2$ CollabZone space. Our evaluation focuses on the following key performance metrics:

- *LUX Estimation Accuracy:* this helps establish the estimation error of our image-based pixel→LUX estimation approach.
- *Patch Classification & Occupancy Estimation Accuracy:* these establish our efficacy in estimating *CS-Light* intermediate variables.
- *Energy Savings:* obviously the most important "application-specific metric", this quantifies the amount of energy saved by *CS-Light's* fine-grained, collective LED adaptation approach.

8.1 LUX Estimation Accuracy

We study the errors in estimating the LUX for a single-frame on a per-zone basis, observed over a period of 2 hours and also evaluate how the sub-zone level aggregation (over a 2 minute interval) helps improve such LUX estimation. The errors are analyzed only during

[2] https://www.dahuasecurity.com/asia
[3] https://www.beckhoff.com/en-us/products/automation/

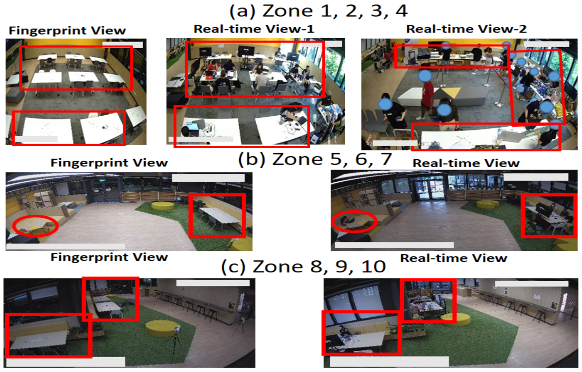

Figure 14: Fingerprint Vs Real-Time View of CollabZone

the low-occupancy period (one occupant per zone), due to errors in ground truth computation when the lux meters are occluded by the occupants and their artefacts. Fig 15 illustrates the estimated LUX compared to the measured Ground-truth LUX values. Over a representative study period (**120 mins**), we observed that the LUX estimator failed to produce any valid estimate ≤8.5% of the time. Overall, *CS-Light* achieved a median LUX Estimation error of 12.3% on a per-frame basis, with the error dropping to 8.5% (max=18%, min=5.5%) after temporal and spatial aggregation (in contrast to the ∼ 30% variance observed in Section 1).

Comparative Performance: One of *CS-Light's* strengths is its ability to deal with changes in the environment, such as rearrangement of tables or occlusion of a patch due to artefacts such as laptop and books. To underline the importance of patch-based color classification and LUX estimation, especially when the layout and ambient lighting vary over time and space (as shown in Fig 14), we compare *CS-Light* against a baseline that mimics [18]–i.e., excludes *CS-Light's* patch classification and statistical aggregation techniques. We observe that, in such dynamically changing environments, *CS-Light* achieves 47.9% lower LUX estimation error compared to [18].

8.2 Patch Classification Accuracy

Our patch classifier achieved an accuracy of 84.6%, across a variety of ambient lighting and occupancy conditions. This classification error is certainly not negligible and further validates our observation that correct and robust classification of patch types is a challenging problem. However, our results (Figure 15) show that *CS-Light* is still able to achieve low LUX estimation error via suitable spatiotemporal aggregation, thus demonstrating the practical utility of our approach. On closer inspection, we noted that our classifier achieved a precision of 80.2% and 96.4% for 'Black/dark-colored' patches and Natural/light-colored' patches, respectively.

8.3 Occupancy Estimation

CS-Light uses a logistic regressor to localize a human bounding box to a sub-zone. Inspection of the camera-recorded ground truth revealed an average instantaneous occupancy estimation error of 12.3%, primary due to 2 seats being occluded in Zone-2. However, this error is eliminated entirely via aggregation and majority-voting over 2-minute intervals.

8.4 Energy Savings and User Comfort

We calculate the total energy saving of *CS-Light*, as well as the user-report opinion on comfort levels, under three distinct occupancy conditions: (a) *Low* (4 occupants), (b) *Medium* (8 occupants) and (b) *High* (15 occupants), measured over observation periods of 30 mins each. In addition, via a survey[4] of subjects working in CollabZone, we ascertain their perceptual response to *CS-Light*-based lighting control. To quantify *CS-Light's* energy savings in the CollabZone deployment, we calculate the energy consumption using Equation 3:

$$P_{avg} = \frac{1}{T}\left(\int_0^{dc-T}\left(\frac{V_{on}^2}{R}dt\right) + \int_{dc-T}^{T}\left(\frac{V_{off}^2}{R}dt\right)\right) \quad (3)$$

Shorter Duration Studies: We first conduct 2-hour long user studies, in zones 1-4, and make the following key observations:

- During daylight hours, when the zone is occupied, *CS-Light* performs implicit daylight harvesting to achieve **60.2**% reduction in lighting energy, compared to the current operational baseline (where all the lights are set to dc = 75% by the campus building managers). Even compared to a more conservative baseline of dc = 50% (which is experimentally seen to satisfy the minimum LUX requirements), *CS-Light* reduces energy consumption by 27.2%. Moreover, *CS-Light's* energy consumption is 56.4% lower than a state-of-the-art occupancy-aware "binary control" strategy (where a sub-zone's LEDs are set to DC = 75% if the zone is occupied and 30% otherwise). More specifically, for unoccupied zones, *CS-Light* achieves an energy saving of 78.7% compared to the binary-occupancy controller during daylight.
- At night, *CS-Light's* savings were considerably lower, due to the absence of any *daylight harvesting* possibilities. When the area was occupied, *CS-Light* achieved energy savings of ≈29.7% compared to a baseline binary control strategy (DC = 75%) and binary-occupancy, with the savings being lower (12.7%) if the baseline used a lower illumination threshold (DC = 50%). When the area was unoccupied, the energy saving increases to 26.3%, compared to a binary control strategy and is obviously superior (66.9% reduction) to a non-adaptive lighting regime (DC = 75%).
- Overall, as expected, *CS-Light's* energy savings are higher at lower occupancy. Figures 16 & 17 plot the energy consumption for two different zones, one that can harvest daylight (Zone-1) and the other that does not receive adequate daylight (Zone-4), and show how *CS-Light* effectively harvests daylight to outperform binary and fixed lighting control.
- While *CS-Light* currently changes DC values *gradually* (to avoid distracting occupants), we noted that *instantaneous* lighting adaptation could provide an additional ≈5.3% savings in energy.

Occupant Comfort: To study human responses to *CS-Light*-based control, we recruited 15 human occupants[5], who sat in different areas of CollabZone while performing their natural life/work-style based activities. (They were, however, additionally asked to perform a few explicit activities– reading, writing and laptop usage– to help elicit their perception of lighting conditions during specific activities.) We list below the user responses for *CS-Light*.

- During daylight with high occupancy, 100% of the occupants felt comfortable working under *CS-Light*-controlled illumination.

[4]Questionnaire available at: https://bit.ly/3itaIxi
[5]All studies were performed with appropriate IRB approvals from our University

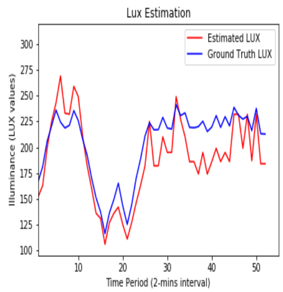

Figure 15: Ground-Truth Vs LUX Estimation

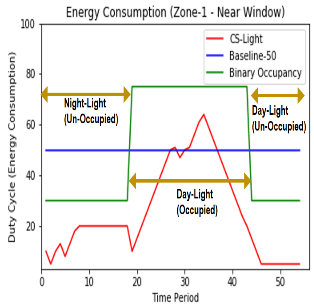

Figure 16: DC Variation (Low Occupancy)

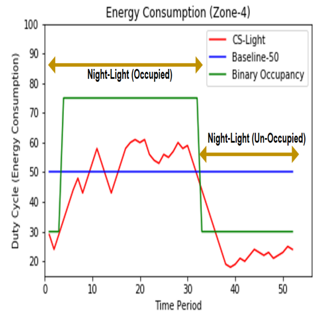

Figure 17: DC Variation (Medium Occupancy)

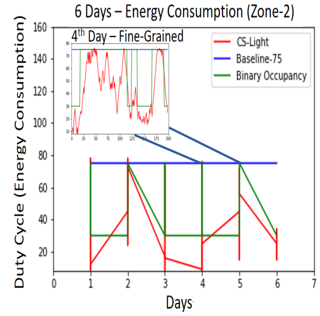

Figure 18: DC Variation (Varying Occupancy - 6 Days)

In low and medium occupancy, 100% of the occupants felt the area was perfectly lit. Across all occupancy states, 10% of the occupants *noticed* the fluctuation when lights were being adjusted, but everyone indicated that they were *comfortable* with the adaptation process.

- In the absence of daylight (i.e., with only LED-provided illumination), 86.6% of the occupants felt the area was perfectly lit under high occupancy condition, while 13.4% felt the area could have been *slightly* brighter (by 10%). In medium occupancy, 70% of the occupants felt the area was perfectly lit and 30% felt the area could have been a bit more bright. In low occupancy, 90% felt the area was perfectly lit, while 10% felt the area could have been brighter. Overall, 17.2% of the occupants *noticed* the LED adjustment process: while 2 occupants felt a bit distracted by the continual changes, they indicated that they were accepting of the illumination dynamics if it helped reduced energy overheads.

Robustness of *CS-Light* over Longer Periods: Going beyond our controlled studies, we deployed *CS-Light* operationally over a period of 6 days (during the period 8.30am-11.30pm) of regular occupancy across Zones 1-10 in CollabZone. As visualized in Figure 14, while the layout does change significantly from the initial fingerprinting stage, *CS-Light* was able to adapt to the changes. Quantitatively, *CS-Light* achieved an average energy savings of 39% (max=53.5% and min=22.6%, depending on the occupancy) compared to baseline-75 (Figure 18), 18.2% compared to baseline-50 and 22% compared to binary occupancy, respectively, over our primary testbed area comprising zones 1-4.

8.5 Variable LED Intensity Control

We emphasize that *CS-Light's* novelty lies in collectively adjusting the intensity of multiple LEDs, so as to take advantage of the indirect illuminance benefits from LEDs not directly overhead an occupied area. Fig 19 illustrates one such adaptation, both before and after occupancy, for different zones and lighting conditions. We see that *CS-Light* is able to use non-binary lighting control to adapt to both occupancy and lighting variations: during daylight, Zone-1 is lit with DC=5% only in the absence of occupancy.

9 Discussion

***CS-Light* Sensitivity and Scalability:** *CS-Light* relies on the offline creation of (sub-region, color) specific pixel-to-LUX mappings, and involves (a) the manual effort in fingerprinting each region, and (b) the computational overhead in model training. The fingerprinting effort grows proportionately to the density of LEDs and the number of patches per zone. In our deployed system, the fingerprinting cost of laying out the paper strips on tables and changing the DC levels is $\approx$ 60 mins/zone/colored strip (zone $\approx$= $12m^2$). We empirically verified that reducing the number of patches by 25% (in effect, defining coarser sub-regions) leads to a 15-20% increase in LUX estimation error under different ambient conditions. Deploying *CS-Light* over the entire ZEB (total area= $750m^2$) is thus anticipated to require a one-time cumulative fingerprinting duration of 22 hours, which remains robust to significant subsequent variations in ambient lighting and layout. In addition, the computational time for training the models (using a Windows 10, 32 GB machine) was 30 mins/zone, implying a total time of 171 mins for the entire CollabZone. In addition, scaling *CS-Light* to our entire campus (10 buildings, approx. 20,000 m^2 of shared-occupancy spaces) will involve the continuous processing of ~120 camera feeds–i.e., avg. total network bandwidth of 500 Mbps (assuming a camera deployment density similar to CollabZone). Our implementation experience suggests campus-wide execution of *CS-Light* can be achieved using 5 (246GB, 4GPU) servers.

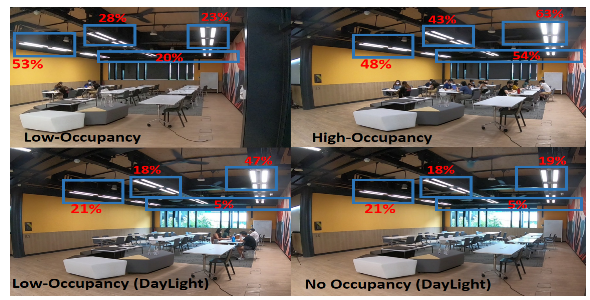

Figure 19: DC Levels of LED luminaries under different occupancy conditions

System Limitations: While *CS-Light* is robust to changes in ambient conditions, it makes certain implicit assumptions: (a) the LED lights may exhibit dynamic intensity changes, but not variations in color. Accommodating such color would require additional human effort in constructing of LED color-specific models; (b) the cameras themselves are static, such that specific pixel coordinates of each camera's image always map to a specific spatial sub-region. To accommodate non-static cameras (e.g., ones that continually sweep a larger area), *CS-Light* would need explicit knowledge of the camera's pose at any time instant; (c) the level of visual occlusion is relatively modest: if the camera's view is significantly obstructed, the occupancy estimates based on visual human object detection will be severely degraded; (d) the distance of the camera from the monitored region is not too large: empirical analysis of CollabZone data reveals that as the distance of the patches/sub-region from the camera increases beyond 10m, the LUX estimation error increases by 12%. Accordingly, *CS-Light's* camera-based LUX estimation may not be suitable for certain public spaces (e.g., convention centers with very high-ceilings and ceiling-mounted cameras); (e) while the *CS-Light* deployment currently disables camera auto-tuning (to support a static pixel-to-LUX mapping), certain environments may require auto-tuning enabled–e.g., for performance spaces where the camera feeds are streamed for live viewing. For such scenarios, we believe that a continuous pre-calibration step, involving a fixed set of colored strips mounted on walls and serving as visual markers, can preserve *CS-Light's* LUX estimation accuracy, by automatically modifying (performing "gain control") on the pixel intensity inputs. This is an area of ongoing work.

10 Conclusion

We've introduced *CS-Light*, a novel system for fine-grained, occupancy-aware lighting control that utilizes only an existing surveillance camera infrastructure and tackles several hitherto-unaddressed practical challenges for pixel-to-LUX estimation. Via extensive experimental and real-world studies on an open-plan collaboration space, we establish that *CS-Light's* approach, of using unsupervised clustering to build multiple such *color-specific* pixel-to-LUX models, is able to dramatically reduce this estimation error to ∼8.5%. In addition, *CS-Light* can reduce lighting energy consumption (without sacrificing occupant comfort) by as much as ∼53%, especially during daylight hours under normal occupancy. When zones are unoccupied (a phenomenon accentuated during Covid), *CS-Light* achieves ∼79% energy savings.

Acknowledgments

This research is supported by the National Research Foundation, Singapore under its Strategic Capabilities Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

[1] Apache. 2020. *Apache Kafka Streams.* https://kafka.apache.org/documentation/
[2] D. Caicedo and A Pandharipande. 2015. Daylight and Occupancy Adaptive Lighting Control System: An iterative optimization approach. *Lighting Research and Technology* 48, 6 (2015), 661–675.
[3] D. Caicedo, A. Pandharipande, and G. Leus. 2011. Occupancy-based illumination control of LED lighting systems. *Lighting Research and Technology* 43, 2 (2011), 217–234. https://doi.org/10.1177/1477153510374703
[4] Yusi Cheng, Chen Fang, Jingfeng Yuan, and Lei Zhu. 2020. Design and application of a smart lighting system based on distributed wireless sensor networks. *Applied Sciences (Switzerland)* 10, 23 (2020), 1–21. https://doi.org/10.3390/app10238545
[5] Francesco De Luca, Raimo Simson, Hendrik Voll, and Jarek Kurnitski. 2018. Daylighting and energy performance design for single floor commercial hall buildings. *Management of Environmental Quality: An International Journal* 29, 4 (2018), 722–739. https://doi.org/10.1108/MEQ-10-2017-0110
[6] Nandha Kumar Kandasamy, Giridharan Karunagaran, Costas Spanos, King Jet Tseng, and Boon Hee Soong. 2018. Smart lighting system using ANN-IMC for personalized lighting control and daylight harvesting. *Building and Environment* 139 (2018), 170–180. https://doi.org/10.1016/j.buildenv.2018.05.005
[7] Pushpendu Kar, Arish Shareef, Arun Kumar, Koh Tsyr Harn, Balaji Kalluri, and Sanjib Kumar Panda. 2019. ReViCEE: A recommendation based approach for personalized control, visual comfort & energy efficiency in buildings. *Building and Environment* 152, January (2019), 135–144. https://doi.org/10.1016/j.buildenv.2019.01.035
[8] Srinivasan Keshav. 2017. *Solar + Storage + IoT + LED = $30 Trillion.* https://youtu.be/YzLXNmHh_QA?list=PL7C97C3FD61272714/
[9] Johann Heinrich Lambert. 2020. *Lambartian Reflectance.* https://en.wikipedia.org/wiki/Lambertian_reflectance
[10] Heangwoo Lee, Chang ho Choi, and Minki Sung. 2018. Development of a dimming lighting control system using general illumination and location-awareness technology. *Energies* 11, 11 (2018). https://doi.org/10.3390/en11112999
[11] Jiakang Lu and Kamin Whitehouse. 2012. SunCast: Fine-Grained Prediction of Natural Sunlight Levels for Improved Daylight Harvesting. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks (IPSN '12).* Association for Computing Machinery, New York, NY, USA, 245–256. https://doi.org/10.1145/2185677.2185738
[12] A. R. Mead and K. M. Mosalam. 2017. Ubiquitous luminance sensing using the Raspberry Pi and Camera Module system. *Lighting Research and Technology* 49, 7 (2017), 904–921. https://doi.org/10.1177/1477153516649229
[13] Guy R. Newsham, M. B.C. Aries, S. Mancini, and G. Faye. 2008. Individual control of electric lighting in a daylit space. *Lighting Research and Technology* 40, 1 (2008), 25–41. https://doi.org/10.1177/1477153507081560
[14] Guy R. Newsham and C. Arsenault. 2009. A camera as a sensor for lighting and shading control. *Lighting Research and Technology* 41, 2 (2009), 143–160. https://doi.org/10.1177/1477153508099889
[15] Ashish Pandharipande and David Caicedo. 2011. Daylight integrated illumination control of LED systems based on enhanced presence sensing. *Energy and Buildings* 43, 4 (2011), 944–950. https://doi.org/10.1016/j.enbuild.2010.12.018
[16] Reza Fuad Rachmadi and I. Ketut Eddy Purnama. 2015. Vehicle Color Recognition using Convolutional Neural Network. *CoRR* abs/1510.07391 (2015). arXiv:1510.07391 http://arxiv.org/abs/1510.07391
[17] Joseph Redmon and Ali Farhadi. 2018. YOLOv3: An Incremental Improvement. *CoRR* abs/1804.02767 (2018). http://arxiv.org/abs/1804.02767
[18] Abhijit Sarkar and Richard G. Mistrick. 2006. A novel lighting control system integrating high dynamic range imaging and DALI. *LEUKOS - Journal of Illuminating Engineering Society of North America* 2, 4 (2006), 307–322. https://doi.org/10.1080/15502724.2006.10747642
[19] Abdul Salam Shah, Haidawati Nasir, Muhammad Fayaz, Adidah Lajis, and Asadullah Shah. 2019. A review on energy consumption optimization techniques in IoT based smart building environments. *Information (Switzerland)* 10, 3 (2019). https://doi.org/10.3390/info10030108
[20] Theodore Tsesmelis, Irtiza Hasan, Marco Cristani, Fabio Galasso, and Alessio Del Bue. 2019. RGBD2lux: Dense light intensity estimation with an RGBD sensor. *Proceedings - 2019 IEEE Winter Conference on Applications of Computer Vision, WACV 2019* (2019), 501–510. https://doi.org/10.1109/WACV.2019.00059 arXiv:1809.07558
[21] Niels Van De Meugheuvel, Ashish Pandharipande, David Caicedo, and P. P.J. Van Den Hof. 2014. Distributed lighting control with daylight and occupancy adaptation. *Energy and Buildings* 75, 2014 (2014), 321–329. https://doi.org/10.1016/j.enbuild.2014.02.016
[22] Kevin Van Den Wymelenberg, Mehlika Inanici, and Peter Johnson. 2010. The effect of luminance distribution patterns on occupant preference in a daylit office environment. *LEUKOS - Journal of Illuminating Engineering Society of North America* 7, 2 (2010), 103–122. https://doi.org/10.1582/LEUKOS.2010.07.02003
[23] Zhe Wang and Tianzhen Hong. 2020. Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy* 269, July (2020). https://doi.org/10.1016/j.apenergy.2020.115036
[24] Zizhen Wang and Yen Kheng Tan. 2013. Illumination control of LED systems based on neural network model and energy optimization algorithm. *Energy and Buildings* 62 (2013), 514–521. https://doi.org/10.1016/j.enbuild.2013.03.029
[25] Y. J. Wen and A. M. Agogino. 2011. Control of wireless-networked lighting in open-plan offices. *Lighting Research and Technology* 43, 2 (2011), 235–248. https://doi.org/10.1177/1477153510382954
[26] D Zevgolis and F V Topalis. 2014. Evaluation of Image Sensors for Lighting Control Applications. In *9th Mediterranean Conference on Power Generation, Transmission Distribution and Energy Conversion (MEDPOWER).* 1–6.
[27] Ahmet Özlü. 2018. *Color Recognition.* https://github.com/ahmetozlu/color_recognition