

Highlights

Semi-physical Gamma-Process Degradation Modeling and Performance-Driven Opportunistic Maintenance Optimization for LED Lighting Systems

Haohao Shi, Huy Truong-Ba, Michael E. Cholette, Brenden Harris, Juan Montes, Tommy Chan

- A performance-driven, simulation-in-the-loop framework is developed for opportunistic maintenance optimization of large-scale LED lighting systems.
- A semi-physical non-homogeneous Gamma process with Bayesian calibration is proposed for LED package degradation under accelerated tests.
- Competing package degradation and driver outages are integrated to define luminaire operating states for system-level simulation.
- Static lighting indices are converted into a long-term dynamic deficiency-ratio metric using event-based performance-deficiency durations.
- A surrogate-based performance mapping replaces repeated ray-tracing, enabling scalable Monte Carlo evaluation and multi-policy optimization in a real case study.

Semi-physical Gamma-Process Degradation Modeling and Performance-Driven Opportunistic Maintenance Optimization for LED Lighting Systems

Haohao Shi^{a,*}, Huy Truong-Ba^a, Michael E. Cholette^a, Brenden Harris^b, Juan Montes^b and Tommy Chan^a

^a Faculty of Engineering, Queensland University of Technology, Brisbane, 4000, Queensland, Australia

^b Fredon Queensland, Brisbane, 4119, Queensland, Australia

ARTICLE INFO

Keywords:

LED lighting systems
Non-homogeneous Gamma process
Bayesian calibration
Performance deficiency ratio
Opportunistic maintenance
Simulation-based optimization
Surrogate modeling

ABSTRACT

Large-scale LED lighting systems degrade through gradual package degradation and abrupt driver outages, while acceptability is determined by spatio-temporal illuminance compliance rather than component reliability alone. This paper proposes a performance-driven, simulation-in-the-loop framework for opportunistic maintenance optimization of LED lighting systems. LED package degradation is modeled by a semi-physical non-homogeneous Gamma process whose mean follows an exponential lumen-maintenance trend, and driver outages are described by a Weibull lifetime model. Parameters are calibrated from LM-80 accelerated degradation data via Bayesian inference, enabling uncertainty propagation to operating conditions. System performance is evaluated using ray-tracing-based illuminance mapping, and static indices (average illuminance and uniformity) are converted into a long-term dynamic deficiency-ratio metric via performance-deficiency durations over event intervals. To enable scalable Monte Carlo policy evaluation and search, a surrogate-based performance mapping replaces repeated ray-tracing with negligible loss of fidelity. An opportunistic policy is optimized in a multi-objective setting to balance performance deficiency, site visits, and replacements. A case study demonstrates the practicality of the framework and the resulting Pareto trade-offs for maintenance decision support.

1. Introduction

LED lighting systems are ubiquitous in buildings because they offer high efficacy, long service life, and flexible control [2, 10, 29, 38]. The LED lighting systems inherently prioritize system performance as the primary criterion guiding operation and maintenance management, since it directly affects occupant well-being and cognitive ability [3]. In practice, lighting quality is not defined by the condition of any single luminaire, but by whether the *working-plane* illuminance field satisfies prescribed requirements, typically expressed in terms of maintained illuminance and illuminance uniformity [15, 17, 18]. These requirements are inherently spatial, and in large facilities they must be delivered by hundreds to thousands of luminaires distributed across multiple zones.

Long-term performance assessment of such systems is non-trivial. Although luminaire degradations may be statistically independent at the component level, their *collective* evolution drives strongly coupled, system-level performance changes because each luminaire contributes to many working-plane points simultaneously. Moreover, the spatio-temporal illuminance field is shaped not only by luminaire states, but also by room geometry, surface reflectances, mounting height, photometric distributions, and layout. As a result, system-level acceptability is fundamentally a *service* question defined on the working plane, rather than a simple aggregation of component reliabilities.

A further challenge is that lighting standards and common engineering practices are largely *commissioning-centric*. Compliance is typically verified for an as-designed, initial condition, using static indices such as average illuminance and uniformity [15, 17, 18]. Over long operating horizons, however, degradation and failures accumulate and can cause intermittent or persistent violations of these requirements. Despite the practical importance of long-term compliance, a widely adopted definition of *system-level failure* and a corresponding *dynamic* performance metric that quantifies how often and how long requirements are violated remains insufficiently developed for large-scale LED lighting fleets.

*Corresponding author: Haohao Shi

✉ haohao.shi@hdr.qut.edu.au (H. Shi)

ORCID(s):

Evaluating long-term system performance requires uncertainty-aware luminaire degradation and failure models that can be propagated through a lighting simulation model. LED luminaires are dominated by two competing component-level mechanisms: gradual degradation of LED packages [13] and abrupt outages caused by driver failures [7]. Due to long lifetimes (often up to 60,000 hours) [13], accelerated degradation tests (ADT) such as LM-80 are commonly used for packages [12]. The TM-21-style exponential lumen-maintenance trend is widely used in practice [13], but deterministic curves do not represent path uncertainty. Stochastic degradation processes (e.g., Gamma processes) have therefore been increasingly adopted to capture monotone degradation with uncertainty [5, 6, 20, 27, 39]. In contrast, driver reliability lacks standardized evaluation protocols; Weibull-based lifetime models and hybrid physics-of-failure/statistical approaches have been explored [7, 27, 28, 42]. Despite these advances, the two mechanisms are often modeled separately; however, for system-level analysis, these two mechanisms have to be integrated at the luminaire level to define an operational state that reflects both gradual degradation and catastrophic failures. In addition, because field observations are scarce over service lifetimes, accelerated degradation tests are typically required; a stress–acceleration relationship should be introduced to extrapolate degradation dynamics from multiple stress temperatures to operating conditions, and parameter uncertainty should be quantified for downstream performance and maintenance evaluation. However, the propagation of calibrated uncertainty to use conditions is not always made explicit in these studies.

Even with suitable degradation models, computational scalability becomes a significant barrier. Ray-tracing engines such as Radiance can model luminaire photometry, geometry, and reflectances with high fidelity, making them well-suited for mapping luminaire states to working-plane illuminance fields. However, traditional design workflows are not designed for repeated, automated evaluation over thousands of stochastic trajectories and many candidate policies. A programmatic *degradation–simulation–metric* pipeline is therefore required to enable Monte Carlo performance evaluation and simulation-in-the-loop optimization, and additional acceleration (e.g., surrogate modeling) becomes essential at practical scales.

Once long-term, system-level performance can be evaluated efficiently, maintenance policies can be formulated and optimized in a performance-driven manner. Existing research on lighting-system maintenance and optimization has predominantly focused on control- and retrofit-oriented formulations, where objectives are defined by energy consumption, retrofit cost, and payback period. In these studies, maintenance is typically represented through simplified assumptions or proxy terms (e.g., fixed-interval PM or run-to-failure CM, sometimes penalizing replacement actions within the controller), rather than being explicitly optimized against long-term lighting-performance compliance [14, 21, 23, 32–36, 40, 41]. In these studies, the maintenance policy is simply used as the input for retrofit project design, but the maintenance policy itself is not optimized.

In parallel, opportunistic maintenance has a mature methodological base for multi-unit systems, particularly in applications where maintenance actions can be coordinated across assets and system impact can be expressed through reliability/availability or cost models [8, 9, 25, 37]. However, applying OM to lighting requires policy evaluation against a different notion of “system performance”: lighting service quality is governed by a spatial illuminance field shaped by luminaire states, layout, and room characteristics, rather than by a single component-level failure event. Consequently, the system-level benefit of maintaining one luminaire is inherently coupled with the states and geometry of other luminaires. This motivates a performance-driven, simulation-in-the-loop framework that evaluates opportunistic policies using physics-based (or surrogate) performance mapping over long operating horizons.

Closed-form evaluation of long-term lighting-performance risk under competing failure modes and discrete-event OM rules is generally intractable, particularly when system acceptability depends on a spatial illuminance field obtained from ray tracing. Monte Carlo simulation therefore becomes the practical engine to propagate both parameter uncertainty (from Bayesian calibration) and process uncertainty (stochastic degradation/failures) to system-level performance metrics. This enables statistically robust comparison of maintenance policies and quantification of trade-offs beyond mean behavior, including variability and tail risk relevant to asset management.

Motivated by these gaps, a performance-driven modeling, evaluation, and optimization framework is developed for large-scale LED lighting systems. The main contributions are summarized as follows:

- *Dynamic system-level performance metric*: Commissioning indices are extended to a long-term dynamic deficiency metric that quantifies the duration of requirement violations over the operating horizon, enabling an explicit quantification of system-level performance failure.

- *Programmatic performance-mapping pipeline*: A scalable, automated degradation–simulation–metric workflow is established to propagate stochastic luminaire-state trajectories through a Radiance-based lighting model and parse working-plane outputs for Monte Carlo evaluation.
- *Uncertainty-aware luminaire modeling under competing failures*: A semi-physical non-homogeneous Gamma-process model is used for LED package degradation whose mean recovers a TM-21-style exponential lumen-maintenance trend under constant stress, while retaining path uncertainty; a Weibull model is used for driver failures; and both mechanisms are integrated through a luminaire-level competing-failure formulation.
- *Unified ADT-to-use inference with uncertainty*: An ADT–extrapolation–Bayesian inference workflow is employed to estimate use-condition degradation parameters jointly across multiple stress levels and quantify parameter uncertainty for downstream simulation.
- *Scalable policy evaluation and acceleration*: Performance-driven opportunistic maintenance is optimized via simulation-in-the-loop multi-objective evaluation, and a surrogate-based performance mapping is introduced to reduce computational overhead and enable large-scale policy search.

The remainder of this paper is organized as follows. Section 2 develops the semi-physical degradation and failure models and presents the ADT-based Bayesian calibration workflow. Section 3 introduces the system-level performance mapping and proposes a dynamic deficiency metric for long-term compliance assessment. Section 4 formulates the performance-driven maintenance-optimization framework and details the simulation-based policy evaluation and surrogate acceleration. Section 5 presents a realistic case study with calibration, optimization results, and sensitivity analyses. Section 6 concludes the paper and outlines future directions.

1.1. Problem statement

Consider an LED lighting system deployed in a building zone, where lighting service is specified on a working plane by maintained illuminance and uniformity requirements [15, 17, 18]. The zone geometry, surface reflectances, luminaire photometry, and luminaire layout define a physics-based lighting model (e.g., Radiance) that maps the system-level luminaire states to a working-plane illuminance field. Each luminaire evolves according to competing stochastic mechanisms: gradual package degradation and abrupt driver failure, whose parameters are inferred from accelerated tests and extrapolated to operating conditions with quantified uncertainty.

The objective is to (i) propagate stochastic luminaire-state trajectories through the lighting model to obtain time-indexed working-plane illuminance fields, (ii) convert static compliance indices into a dynamic performance metric that quantifies the duration of requirement violations over an operating horizon, and (iii) evaluate and optimize maintenance policies in a performance-driven manner. Because cost information is typically operator-specific and unavailable, policy performance is characterized by a multi-objective vector consisting of the long-term performance deficiency metric and two resource proxies: the total number of site visits and the total number of luminaire replacements. Decision variables parameterize an opportunistic maintenance policy (e.g., the preventive-maintenance interval and opportunistic threshold), and Pareto-optimal policy settings are sought via simulation-in-the-loop Monte Carlo evaluation with computational acceleration.

2. Semi-physical degradation and failure modeling of LED luminaires

This section develops component-level models for the two dominant failure mechanisms of an LED luminaire, namely gradual LED package degradation (manifested as light-output loss) and abrupt LED driver failures. Package degradation is modeled by a semi-physical non-homogeneous Gamma process whose mean evolution follows a physically-motivated exponential form. Driver lifetime is modeled by a Weibull distribution. These two mechanisms are integrated through a competing-failure formulation to define the overall luminaire state for subsequent system-level performance analysis. Since field data are scarce due to the long service life of LED packages, accelerated degradation tests (ADT) are employed; a stress–acceleration relationship is introduced to link model parameters across temperatures and to extrapolate them to operating conditions. Finally, Bayesian inference is employed to estimate model parameters jointly across multiple stress levels and to quantify parameter uncertainty for downstream Monte Carlo simulation and maintenance optimization.

Table 1

Key notation in the problem statement.

Symbol	Description
J	Number of luminaires in the lighting system
N	Number of working-plane calculation points
$s = 1, \dots, S$	Monte Carlo replication index
$k = 1, \dots, K$	Event index
T_{over}	Operational horizon
T_{PM}	Preventive-maintenance (PM) interval (decision variable)
H_{OM}	Opportunistic-maintenance (OM) threshold (decision variable)
T_{OM}	Equivalent OM age trigger, $T_{\text{PM}}(1 - H_{\text{OM}})$
$\mathbf{L}(t)$	Luminaire degradation-state vector at time t
$\mathbf{E}(t)$	Illuminance vector on the working plane at time t
$E_{\text{avg}}(t)$	Average illuminance at time t
$U(t)$	Illuminance uniformity at time t
S_E, S_U	Performance requirements (standards)
$T_{\text{defi}}(k)$	Deficiency duration within (t_{k-1}, t_k)
R_{DR}	Deficiency ratio over $[0, T_{\text{over}}]$
N_{tv}	Total number of site visits
N_{tr}	Total number of replacements

2.1. Component-level degradation and failure modes of LED packages and drivers

In an LED luminaire, the LED package (hereafter, *package*) is responsible for light generation, whereas the LED driver (hereafter, *driver*) provides electrical power conversion and regulation. Package degradation is gradual and typically observed as a loss of normalized light output over time, while driver failures are often abrupt and result in a complete loss of functionality. These two mechanisms are modeled below and then integrated to define the overall luminaire state.

LED package degradation: a semi-physical non-homogeneous Gamma process. Package degradation is quantified via the normalized light output. Following the widely adopted L_{70} criterion from the Illuminating Engineering Society of North America (IESNA), a degradation failure is defined as the condition in which the light output falls below 70% of its initial value [13].

Let $\mathbf{X}(t) = [X_1(t), X_2(t), \dots, X_J(t)]$ denote the package degradation states at time t , where

$$X_j(t) = 1 - P_{j,\text{out}}(t), \quad j = 1, 2, \dots, J, \quad (1)$$

and $0 \leq P_{j,\text{out}}(t) \leq 1$ represents the normalised light output of the j -th package (with $P_{j,\text{out}}(0) = 1$). Hence, L_{70} corresponds to $P_{j,\text{out}}(t) < 0.7$, equivalently $X_j(t) > 0.3$.

For each package j , $\{X_j(t), t \geq 0\}$ is modeled as a non-homogeneous Gamma process. Specifically, for any time instants $0 \leq t_1 < t_2 < \dots < t_K$, the increments $X_j(t_2) - X_j(t_1)$, $X_j(t_3) - X_j(t_2)$, \dots , $X_j(t_K) - X_j(t_{K-1})$ are independent and non-negative, with $X_j(0) = 0$ almost surely. Moreover, for $t > s$,

$$X_j(t) - X_j(s) \sim \text{Ga}(\alpha(t) - \alpha(s), \beta), \quad (2)$$

where $\alpha(\cdot)$ is a non-decreasing shape function and β is the rate parameter. Accordingly, the marginal distribution of $X_j(t)$ is Gamma with shape $\alpha(t)$ and rate β , and its probability density function (PDF) is

$$f_{X_j(t)}(x) = \frac{\beta^{\alpha(t)}}{\Gamma(\alpha(t))} x^{\alpha(t)-1} \exp(-\beta x), \quad x > 0, \quad (3)$$

where $\Gamma(\cdot)$ denotes the Gamma function.

Compared with the standard exponential model used in IES TM-21 [13], the Gamma process allows uncertain degradation paths to be represented. As discussed above, TM-21 assumes an exponential average light-output decay

under constant stress. This is physically motivated: thermal factors are a primary cause of package degradation, and thermal accumulation inside the package can exhibit an exponential trend dependent on temperature [30]. Therefore, a time-dependent exponential shape function is introduced to build a semi-physical non-homogeneous Gamma process:

$$\alpha(t) = A \exp(bt), \quad (4)$$

where A and b are constants (assumed invariant across stress levels under the linear acceleration assumption introduced later). This choice yields an exponential mean degradation trajectory:

$$\mathbb{E}[X_j(t)] = \frac{\alpha(t)}{\beta}, \quad (5)$$

and for $t > s$, the expected increment is

$$\mathbb{E}[X_j(t) - X_j(s)] = \frac{\alpha(t) - \alpha(s)}{\beta}. \quad (6)$$

LED driver failures: a Weibull lifetime model. Driver malfunctions are characterized by a sudden and complete loss of functionality. This behavior is appropriately modeled using a traditional two-state approach: operational versus failed, where the failure time of drivers is assumed to follow a Weibull distribution [42]. The PDF is:

$$f_{T_d}(t) = \frac{\eta}{\lambda} \left(\frac{t}{\lambda}\right)^{\eta-1} \exp\left(-\left(\frac{t}{\lambda}\right)^\eta\right), t \geq 0, \quad (7)$$

where λ is the scale parameter and η is the shape parameter.

Competing failure integration for luminaire state. To evaluate the overall state of an *LED luminaire* (comprising both the package and the driver), let $\mathbf{L}(t) = [L_1(t), L_2(t), \dots, L_j(t)]$ denote the luminaire degradation states in the lighting system. The degradation state of the j -th luminaire at time t is defined as

$$\begin{aligned} L_j(t) &= 1 - (1 - X_j(t)) \cdot (1 - S_j(t)) \\ &= S_j(t) \cdot (1 - X_j(t)) + X_j(t), \end{aligned} \quad (8)$$

where $S_j(t)$ is the driver failure indicator:

$$S_j(t) = \begin{cases} 1 & \text{if the driver has failed before } t \\ 0 & \text{otherwise} \end{cases}. \quad (9)$$

Equivalently,

$$L_j(t) = \begin{cases} 1 & \text{if the driver has failed before } t \\ X_j(t) & \text{otherwise} \end{cases}. \quad (10)$$

This competing-failure formulation captures both gradual package degradation and abrupt driver failures, enabling consistent luminaire-state trajectories to be generated for subsequent system-level performance evaluation and maintenance optimization.

2.2. Linear accelerated model and parameter extrapolation in ADT

Given that the typical LED package lifetime exceeds 60,000 hours [13], field-failure data are scarce, and natural aging proceeds slowly. Therefore, accelerated degradation tests (ADT) are widely employed to acquire sufficient degradation data in a practical time frame. The LM-80 method [12] is the industry-standard constant-stress ADT for LED packages and is routinely performed prior to commercial release.

Linear acceleration theory assumes that the underlying failure mechanism remains invariant across stress levels. Under this assumption, the Gamma-process shape function is taken as invariant, while the rate parameter varies with

stress (e.g., temperature) [4, 31]. Let $\beta(T)$ denote the rate parameter at case temperature T . The Arrhenius model is adopted to describe the temperature dependence [4, 16, 22, 31]:

$$\beta(T) = C \exp\left(\frac{E_a}{k_B T}\right), \quad (11)$$

where C is a pre-exponential factor, E_a is the activation energy (eV), k_B is the Boltzmann's constant (8.62×10^{-5} eV/K), and T is the absolute temperature (K). Taking logarithms yields

$$\ln \beta(T) = \ln C + \frac{E_a}{k_B T} \quad (12)$$

Following completion of the ADT, an acceleration model is fitted to the high-stress data and then extrapolated to normal operating conditions. The acceleration factor (AF) is defined as

$$AF = \frac{t(T_{\text{stress}})}{t(T_{\text{normal}})} = \frac{\beta(T_{\text{stress}})}{\beta(T_{\text{normal}})}, \quad (13)$$

where $t(T_{\text{stress}})$ and $t(T_{\text{normal}})$ are failure times under stress and normal temperatures, T_{stress} and T_{normal} , respectively.

However, AF-based extrapolation yields only a nominal (point) estimate of the use-condition rate parameter and does not quantify parameter uncertainty. To address this limitation, Bayesian parameter inference is adopted in the next subsection to capture epistemic uncertainty under normal operating conditions while retaining the linear acceleration relationship across stress levels.

2.3. Bayesian parameter inference under multiple stress conditions

A Bayesian inference framework is adopted to estimate parameters of the non-homogeneous Gamma-process model and to qualify uncertainties associated with parameter estimates and the degradation process. Observations collected under different case temperatures are flattened and combined into a unified dataset of the form:

$$\mathcal{D} = \{(x_{k-1}, x_k, t_{k-1}, t_k, T_k) \mid k = 1, 2, \dots, M\}, \quad (14)$$

where x_k is the observed degradation value at time t_k under (constant) stress temperature T_k . Under the non-homogeneous Gamma-process model,

$$X_k - X_{k-1} \sim \text{Ga}(\alpha(t_k) - \alpha(t_{k-1}), \beta(T_k)) \quad (15)$$

with the exponential shape function $\alpha(t) = \exp(\ln A + bt)$ and the temperature-dependent rate function $\beta(T) = \exp(\ln C + E_a/(k_B T))$.

In the Bayesian approach, the unknown model parameters are treated as random variables characterized by prior distributions, which encode prior engineering knowledge or physical constraints. These priors are updated with the observed data \mathcal{D} to yield the posterior distribution. Let $\theta = (\ln A, b, \ln C, E_a)$ denote the transformed parameters. Given the observations, the *likelihood* function is expressed as

$$f(\mathcal{D}|\theta) = \prod_{k=1}^M f_{\Gamma}(x_k - x_{k-1} | \alpha(t_k; \theta) - \alpha(t_{k-1}; \theta), \beta(T_k; \theta)) \quad (16)$$

where $f_{\Gamma}(\cdot)$ denotes the Gamma PDF. By Bayes' theorem, the posterior distribution is:

$$p(\theta|\mathcal{D}) = \frac{f(\mathcal{D}|\theta)p(\theta)}{\int f(\mathcal{D}|\theta)p(\theta)d\theta} \quad (17)$$

where the denominator represents the evidence.

To propagate both epistemic and aleatory uncertainties, a two-stage Monte Carlo scheme is employed. First, n_{path} parameter sets $\theta_l = (\ln A_l, b_l, \ln C_l, E_{a_l})_{l=1}^{n_{\text{path}}}$ are drawn from the posterior distribution using Markov chain Monte Carlo




Figure 1: An LED lighting system, working plane, and calculation grid

(MCMC). Then, for each draw θ_l , n_{ps} independent gamma-process trajectories are generated:

$$G(s)_{(l,m)} - G(t)_{(l,m)} \sim \text{Ga}(\alpha_l(s) - \alpha_l(t), \beta_l(T)), \quad m = 1, \dots, n_{ps}. \quad (18)$$

Consequently, a total of $n_{\text{path}} \times n_{ps}$ trajectories are generated to represent virtual luminaires accounting for both parameter uncertainty and process variability. Maintenance policies are applied to these luminaires, and the optimization objectives are estimated by averaging over trajectories. This approach ensures that the variability of degradation outcomes is adequately characterized for each plausible parameter, thereby reducing estimator noise and avoiding undue influence of extreme single-path realizations.

The inferred luminaire-state trajectories $L(t)$ provide the time-varying health states of individual luminaires and serve as inputs to the system-level lighting analysis in the next section, where dynamic illuminance distributions and performance indices are evaluated over the working plane.

3. System-level performance mapping and dynamic deficiency metrics for LED lighting systems

Conventional lighting performance assessments are predominantly static, focusing on whether commissioning-time metrics meet relevant standards. Such assessments do not capture performance fluctuations over the service life of the system caused by luminaire degradation and failures, which may lead to risks in occupant safety, visual comfort, and productivity. In this section, a system-level mapping from luminaire health states to working-plane illuminance is established using a physics-based simulation model, and two standard performance indices are adopted to quantify static performance. These static indices are then recast into a dynamic long-term metric through the performance deficiency duration (PDD) and the deficiency ratio (DR), enabling system-level lifetime performance evaluation for subsequent maintenance optimization.

3.1. System-level performance mapping and static performance indices

The illuminance distribution on a horizontal working plane (WP) provides a practical basis for assessing the system-level performance of an LED lighting system in indoor environments. The WP is defined as a virtual horizontal plane, typically located at a height of approximately 0.7–0.85 m above the floor to represent desk height for tasks such as writing and data processing. As illustrated in Figure 1, the gray plane represents the WP, red rectangles denote LED luminaires, and black markers denote calculation points used for quantitative performance assessment.

A physics-based ray-tracing simulator, Radiance, is employed to compute WP illuminance given the building geometry, material properties, luminaire photometry, and luminaire states. Radiance can process 3D models generated by Building Information Modeling (BIM) tools and extract geometric and material properties for lighting analysis. Its simulation engine combines Monte Carlo ray tracing with deterministic techniques [1, 24], and its command-line interfaces facilitate integration with Python for batch evaluation.

Accurate lighting source modeling is essential for reliable illuminance prediction. In this work, luminaire photometry is specified using IESNA LM-63 (.IES) files [11] and combined with a 3D BIM-based geometry/material model, luminaire layout, and a working-plane calculation grid to initialize a physics-based performance mapping in Radiance. Preparing these 3D inputs (geometry and material assignment, photometric specification, luminaire placement, and grid definition) is an engineering-intensive step; however, once the model is established, it can be

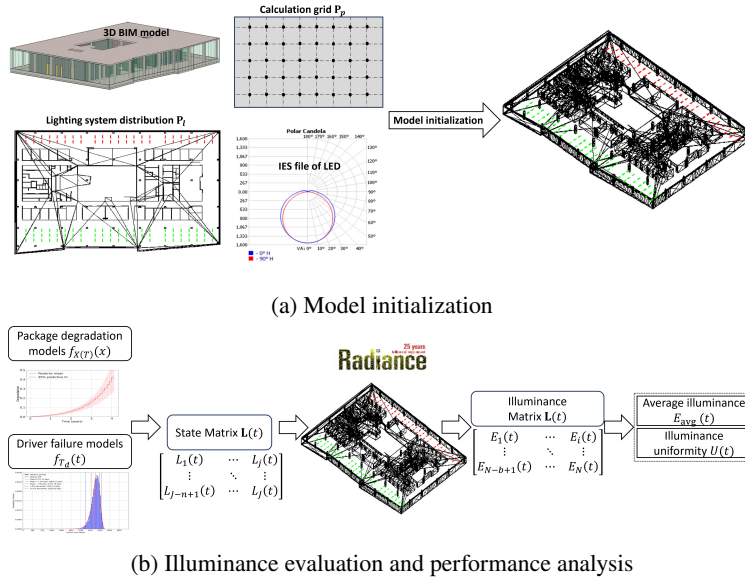


Figure 2: Workflow of Radiance-based working-plane illuminance evaluation

repeatedly evaluated at discrete time points as luminaire degradation trajectories evolve over the operational life. These inputs are fed into Radiance to initialize the performance mapping model, as depicted in Figure 2a.

As summarized in Figure 2b, the Radiance-based performance mapping is executed in a programmatic pipeline: stochastic luminaire state trajectories generated in Python are used to parameterize the model at discrete time points, and the resulting working-plane illuminance outputs are automatically parsed for performance evaluation. This automation is essential for large-scale Monte Carlo evaluation and optimization considered in later sections.

Let the luminaire degradation state vector be defined by Equation 1–10 as

$$\mathbf{L}(t) = [L_1(t), L_2(t), \dots, L_J(t)], \quad (19)$$

where $j = 1, 2, \dots, J$ denotes the j -th luminaire. Given $\mathbf{L}(t)$, the performance mapping model outputs the corresponding WP illuminance vector

$$\mathbf{E}(t) = [E_1(t), E_2(t), \dots, E_N(t)], \quad (20)$$

where $E_i(t)$ is the illuminance at the i -th calculation point, measured in lux (lx), i.e., luminous flux per unit area, and $i = 1, 2, \dots, N$.


To quantify static performance at time t , two standard illuminance-based indices are adopted [15, 17, 18]. The average illuminance is

$$E_{\text{avg}}(t) = \frac{1}{N} \sum_{i=1}^N E_i(t), \quad (21)$$

and the illuminance uniformity (hereafter, uniformity) is defined as

$$U(t) = \frac{E_{\min}(t)}{E_{\text{avg}}(t)}, \quad (22)$$

where $E_{\min}(t) = \min_{1 \leq i \leq N} E_i(t)$.



(a) The illuminance map for a brand-new lighting system (b) The illuminance map of day 1259 (c) The illuminance map of day 1882

Figure 3: Illustration of working-plane illuminance fluctuations over operation time

Let S_E and S_U denote the standard requirements (thresholds) for average illuminance and uniformity, respectively. Static compliance at time t is assessed by

$$E_{\text{avg}}(t) \geq S_E, \quad U(t) \geq S_U. \quad (23)$$

An illuminance map provides a visual representation of spatial performance. In Figure 3, warmer color indicate higher illuminance levels, and the progressive color and spatial-pattern changes over time reflect degradation- and failure-induced performance drift. This motivates the need for a long-term dynamic performance metric beyond commissioning-time indices.

3.2. Dynamic failure definition and deficiency-ratio metric

To evaluate long-term dynamic performance, the temporal evolution of system states over an operational horizon is represented by the luminaire-state trajectory matrix

$$\mathbf{L}_s = [\mathbf{L}(t_1), \mathbf{L}(t_2), \dots, \mathbf{L}(t_K)]^T, \quad (24)$$

where $\{t_k\}_{k=1}^K$ is a discrete time grid, $\mathbf{L}(t_k)$ is the system state vector at time t_k . Based on the performance mapping model, the corresponding illuminance sequence matrix is

$$\mathbf{E}_s = [\mathbf{E}(t_1), \mathbf{E}(t_2), \dots, \mathbf{E}(t_K)]^T. \quad (25)$$

Using Equation 21–23, the static indices at each time point are computed as $\{E_{\text{avg}}(t_k), U(t_k)\}_{k=1}^K$. For each interval (t_{k-1}, t_k) , define $\Delta t_k = t_k - t_{k-1}$ and introduce the deficiency indicators

$$I_E(k) = \mathbb{1}(E_{\text{avg}}(t_k) < S_E), \quad I_U(k) = \mathbb{1}(U(t_k) < S_U), \quad (26)$$

where $\mathbb{1}(\cdot)$ is the indicator function. The deficiency durations associated with the two indices over (t_{k-1}, t_k) are defined by

$$T_{E_{\text{defi}}}(k) = I_E(k) \Delta t_k, \quad T_{U_{\text{defi}}}(k) = I_U(k) \Delta t_k. \quad (27)$$

The performance deficiency duration (PDD) is then defined as the maximum of the two deficiency durations:

$$T_{\text{defi}}(k) = \max(T_{E_{\text{defi}}}(k), T_{U_{\text{defi}}}(k)), \quad (28)$$

which corresponds to the time within (t_{k-1}, t_k) , determined by the dominant (longer) violation duration among the two requirements. An illustration of index trajectories and requirements is shown in Figure 4.

Finally, the deficiency ratio (DR) is defined as the fraction of operational time during which the system violates at least one requirement:

$$R_{\text{DR}} = \frac{\sum_{k=1}^K T_{\text{defi}}(k)}{T_{\text{over}}}, \quad T_{\text{over}} = t_K - t_1. \quad (29)$$




Figure 4: Static performance indices and standard requirements over time

The DR recasts static illuminance-based indices into a long-term dynamic performance metric, which is subsequently used as a system-level objective for maintenance optimization.

Large-scale Monte Carlo evaluation under maintenance policies requires repeated mapping from $\mathbf{L}(t_k)$ to $\mathbf{E}(t_k)$. To enable efficient optimization, the Radiance-based mapping is accelerated using a surrogate model (see Section 4.4).

4. Performance-driven maintenance optimization for LED lighting systems

This section develops a simulation-in-the-loop framework to optimize opportunistic maintenance policies for LED lighting systems under stochastic luminaire degradation and failures. A unit-based time-window policy is formulated to exploit maintenance dependency across luminaires. Policy performance is evaluated by coupling a discrete-event maintenance simulation module with the system-level performance metric defined in Section 3.2, namely the deficiency ratio (DR), and by estimating trade-offs among DR, the total number of site visits, and the total number of replacements via Monte Carlo simulation. To make large-scale evaluation computationally feasible, illuminance evaluation within the loop is accelerated using a surrogate model that replaces repeated Radiance calls.

4.1. Opportunistic maintenance policy formulation

In an LED lighting system, performance degradation and failures are primarily driven by two component-level mechanisms: gradual LED package degradation (light-output loss) and abrupt LED driver failures (complete outage). Because packages and drivers are typically treated as non-repairable field units, maintenance is modeled as luminaire replacement.

Corrective maintenance (CM) is initiated when a luminaire is reported as failed or unacceptable in service, commonly due to a complete outage caused by driver failure or noticeable dimming caused by package degradation (with the degradation-failure definition following the L_{70} criterion in Section 2.1). Such reactive interventions can be costly because they require unscheduled site visits and access coordination, introduce safety and logistical constraints, and may allow degraded lighting performance to persist until intervention.

Preventive maintenance (PM) is therefore scheduled at fixed intervals to reduce disruptive CM events. Nevertheless, fixed-interval PM and purely reactive CM can be inefficient: PM may replace luminaires prematurely, whereas CM may occur only after unacceptable performance has already persisted, increasing the frequency of unscheduled visits and cumulative performance shortfalls. To exploit economic dependence among maintenance activities, an opportunistic maintenance (OM) strategy is adopted: when a maintenance visit occurs (triggered by either CM or PM), additional luminaires that are close to their next scheduled PM are replaced concurrently. By consolidating actions into fewer visits, OM reduces visit frequency and improves maintenance resource utilization.

Accordingly, CM is applied for driver failures (complete loss of function) and package degradation failures, while PM is performed periodically every T_{PM} . Both CM and PM events create an opportunity to execute OM. Following the condition-index concept in Shi et al. [25], a maintenance opportunity window (MOW) is defined for luminaire j as the ratio of the remaining time until its next scheduled PM to the PM interval:

$$H_{\text{mov},j} = \frac{T_{\text{re},j}}{T_{\text{PM}}}, \quad (30)$$




Figure 5: The schematic diagram of OM policy

where $T_{re,j}$ is the remaining time until luminaire j reaches its next scheduled PM time. The OM decision is determined by comparing $H_{mov,j}$ with a predefined threshold H_{OM} :

$$H_j(t) = \begin{cases} 1, & \text{if } H_{mov,j} \leq H_{OM}, \\ 0, & \text{if } H_{mov,j} > H_{OM}. \end{cases} \quad (31)$$

As illustrated in Figure 5, a CM event occurring between t_{k-1} and the next scheduled PM triggers the OM policy. For any luminaire j satisfying $H_{mov,j} \leq H_{OM}$, replacement is performed opportunistically during the same visit.

The total number of site visits over the operational horizon is defined as

$$N_{tv} = N_{cm,vis} + N_{pm,vis}, \quad (32)$$

and the total number of replacements is

$$N_{tr} = N_{cm} + N_{pm} + N_{om}, \quad (33)$$

where $N_{pm,vis}$ and $N_{cm,vis}$ denote the numbers of PM and CM visits, and N_{cm} , N_{pm} , and N_{om} denote the numbers of luminaires replaced via CM, PM, and OM, respectively.

Because detailed cost information (labor rates, access constraints, downtime penalties) is highly site- and operator-specific and unavailable in this study, we use the total number of site visits as a proxy for fixed/setup cost and the total number of replacements as a proxy for variable material-and-labor cost. This multi-objective formulation avoids imposing an arbitrary cost ratio; any specific cost model that is a weighted sum of these proxies corresponds to selecting a point on the Pareto front.

4.2. Multi-objective optimization formulation

The OM threshold and the PM interval jointly determine the aggressiveness of opportunistic replacement. We therefore formulate a multi-objective optimization problem over these two decision variables to balance long-term lighting performance against maintenance effort.

The maintenance policy aims to sustain high system performance while limiting maintenance workload. Three objectives are considered: the long-term deficiency ratio R_{DR} (Section 3.2), the total number of site visits N_{tv} (Eq. 32), and the total number of luminaire replacements N_{tr} (Eq. 33). Two decision variables govern these objectives:

1. *OM threshold* $H_{OM} \in [0, 1]$, which determines whether additional luminaires are replaced opportunistically during a CM or PM visit. A larger H_{OM} expands the set of luminaires eligible for opportunistic replacement, which can reduce future visit frequency and mitigate performance shortfalls, but may increase the number of replacements.
2. *PM interval* $T_{PM} \in (0, T_{over}]$, which specifies the periodic preventive-maintenance schedule. Shorter intervals generally improve performance by limiting degradation accumulation, but increase visit frequency and replacements; longer intervals reduce planned visits but may increase corrective interventions and performance deficiency.

Accordingly, the performance-driven maintenance optimization problem is formulated as

$$\begin{aligned} & \min_{H_{OM}, T_{PM}} (R_{DR}, N_{tv}, N_{tr}), \\ & \text{subject to } 0 \leq H_{OM} \leq 1, \quad 0 < T_{PM} \leq T_{over}, \end{aligned} \quad (34)$$

where T_{over} is the operational horizon and R_{DR} quantifies the proportion of time during which system performance falls below standard requirements.

To make the maintenance rule interpretable, the two decision variables can be expressed as an equivalent age-based trigger. From Eqs. 30–31, OM is executed when the remaining time to the next PM is within a fraction H_{OM} of the PM interval, which is equivalent to an “OM-eligible age” threshold

$$T_{OM} = T_{PM} (1 - H_{OM}). \quad (35)$$

Here, T_{PM} sets the maximum allowable age before the next scheduled PM action, and T_{OM} specifies how early a luminaire becomes eligible for opportunistic replacement within each PM cycle. A larger H_{OM} yields a smaller T_{OM} , corresponding to a more aggressive opportunistic strategy.

In principle, R_{DR} , N_{tv} , and N_{tr} could be mapped to monetary costs and aggregated into a single objective. However, the required cost coefficients (e.g., labor logistics, access constraints, contractor pricing, and service-level penalties) are highly site- and operator-specific and are unavailable in this study. Instead, a multi-objective formulation is adopted: N_{tv} serves as a proxy for fixed/setup effort associated with site access and visit organization, while N_{tr} reflects variable workload related to replacement actions. The set of Pareto-optimal (non-dominated) solutions [19] is then identified to characterize the trade-off among performance deficiency, the total number of site visits, and the total number of replacements.

4.3. Simulation-based policy evaluation via discrete-event Monte Carlo

A Monte Carlo simulation framework is developed to evaluate candidate maintenance policies under stochastic luminaire degradation and failures. For each policy setting (a (T_{PM}, H_{OM}) pair), S independent simulation runs are performed. Each run consists of two coupled modules: (i) a discrete-event maintenance simulator that generates the event timeline and the resulting luminaire-state trajectory, and (ii) a performance-evaluation module that maps the trajectory to the long-term deficiency metric defined in Section 3.2.

Discrete-event maintenance simulation. For a given simulation run s , the system evolution is represented by a sequence of events. Each luminaire j is associated with an event vector that stores candidate times of maintenance-related and performance-recording events. At event index k , the event vector of luminaire j is defined as

$$\mathbf{t}_{s,k,j} = \left[t_{s,k,j}^{\text{pm}}, t_{s,k,j}^{\text{cm,df}}, t_{s,k,j}^{\text{cm,pf}}, t_{s,k,j}^{\text{event-end}}, t_{s,k,j}^{\text{record}} \right]^T, \quad (36)$$

where

- $t_{s,k,j}^{\text{pm}}$ is the next scheduled PM time,
- $t_{s,k,j}^{\text{cm,df}}$ is the predicted CM time due to driver failure,
- $t_{s,k,j}^{\text{cm,pf}}$ is the predicted CM time due to package degradation failure (per the failure definition in Section 2.1),
- $t_{s,k,j}^{\text{event-end}}$ is the completion time of an ongoing maintenance action on luminaire j (if any); otherwise it is set to T_{over} ,
- $t_{s,k,j}^{\text{record}}$ is the next time at which system performance will be recorded for deficiency-duration evaluation.

Stacking all luminaires yields the system event matrix at event index k :

$$\mathbf{T}_{s,k}^{\text{evt}} = \left[\mathbf{t}_{s,k,1}, \mathbf{t}_{s,k,2}, \dots, \mathbf{t}_{s,k,J} \right]^T. \quad (37)$$




Figure 6: The schematic diagram of the discrete event simulation model

The next system event time is defined as the minimum over all entries of $\mathbf{T}_{s,k}^{\text{evt}}$:

$$t_{s,k}^{\text{ts-evt}} = \min \mathbf{T}_{s,k}^{\text{evt}} = \min_{j=1,2,\dots,J} \mathbf{t}_{s,k,j}, \quad (38)$$

where the minimization is taken over the five candidate event times for every luminaire.

Given $t_{s,k}^{\text{ts-evt}}$, the corresponding event type (PM, CM due to driver failure, CM due to package degradation failure, maintenance completion, or performance recording) is identified, and the relevant state variables and event times are updated. For example, in Figure 6, the first event is a CM triggered by driver failure on luminaire L_1 ; after replacement, the maintenance completion time is updated as $t_{s,k,1}^{\text{event-end}} = t_{s,k}^{\text{ts-evt}} + \text{service time}$. Once maintenance is completed, the event vector of the renewed luminaire is reinitialized and the system event matrix Equation 37 is updated for the next event index.

In addition to maintenance-triggered events, performance-recording events $t_{s,k,j}^{\text{record}}$ are inserted to ensure sufficient temporal resolution for computing performance deficiency durations (PDD), especially when maintenance events are sparse. The full discrete-event workflow is summarized in Algorithm 1.

The resulting strictly increasing event-time sequence for simulation s is

$$\mathbf{T}_s^{\text{ts-evt}} = [t_{s,1}^{\text{ts-evt}}, t_{s,2}^{\text{ts-evt}}, \dots, t_{s,K}^{\text{ts-evt}}]. \quad (39)$$

At each event time $t_{s,k}^{\text{ts-evt}}$, the system luminaire-state vector $\mathbf{L}(t_{s,k}^{\text{ts-evt}})$ is logged. Stacking these vectors yields the state trajectory

$$\mathbf{L}_s = [\mathbf{L}(t_{s,1}^{\text{ts-evt}}), \mathbf{L}(t_{s,2}^{\text{ts-evt}}), \dots, \mathbf{L}(t_{s,K}^{\text{ts-evt}})]^T. \quad (40)$$

Performance evaluation and objective estimation. Given \mathbf{L}_s , the performance-mapping model produces the corresponding illuminance sequence matrix \mathbf{E}_s (see Equation 25), from which the deficiency duration $T_{\text{defi},s}(k)$ is computed as described in Section 3.2. The resulting deficiency ratio for run s is

$$R_{\text{DR},s} = \frac{\sum_{k=1}^K T_{\text{defi},s}(k)}{T_{\text{over}}}. \quad (41)$$

The pseudocode of the performance-evaluation module is provided in Algorithm 2.

For each run s , the numbers of CM and PM visits ($N_{\text{cm,vis},s}$ and $N_{\text{pm,vis},s}$) and the counts of luminaire replacements via CM, PM, and OM ($N_{\text{cm},s}$, $N_{\text{pm},s}$, and $N_{\text{om},s}$) are recorded. The total numbers of site visits and replacements are

computed as

$$N_{tv,s} = N_{cm,vis,s} + N_{pm,vis,s}, \quad (42)$$

$$N_{tr,s} = N_{cm,s} + N_{pm,s} + N_{om,s}. \quad (43)$$

Averaging over S independent runs yields the estimated objectives for the policy:

$$N_{STV} = \frac{1}{S} \sum_{s=1}^S N_{tv,s}, \quad (44)$$

$$N_{STR} = \frac{1}{S} \sum_{s=1}^S N_{tr,s}, \quad (45)$$

$$R_{SDR} = \frac{1}{S} \sum_{s=1}^S R_{DR,s}. \quad (46)$$

For each policy setting, the objective triplet $(R_{SDR}, N_{STV}, N_{STR})$ is obtained and used for Pareto analysis in the optimization procedure. One-sided Welch's t-tests are applied across all three objectives to sort Pareto solutions, A significance level of α_{sv} is adopted. Solutions for which $p \geq \alpha_{sv}$ on all objectives are deemed statistically inferior and are removed.

4.4. Surrogate-based acceleration of performance evaluation

Direct illuminance evaluation via Radiance at every simulated event index (i.e., each $t_{s,k}^{ts-evt}$ in the event-time vector \mathbf{T}_s^{ts-evt} defined in Equation 39) is computationally prohibitive; we therefore introduce a surrogate-based performance mapping to enable efficient simulation-in-the-loop optimization.

During system operation, each Monte Carlo replication produces hundreds to thousands of luminaire-state vectors $\mathbf{L}(t_k)$ that must be mapped to working-plane illuminance for evaluating the static performance indices $E_{avg}(t_k)$ and $U(t_k)$. Repeated Radiance calls at all event times lead to an excessive computational burden. For example, for a working zone with $J = 76$ LED luminaires and $N = 265$ working-plane calculation points, a single Radiance evaluation takes on the order of one minute. When many policy settings and thousands of Monte Carlo replications are required, the overall runtime becomes prohibitive.

Physics-informed linear surrogate. For a fixed room geometry and material configuration, the illuminance at a given point is approximately an affine function of luminaire photometric output scaling, consistent with the superposition property of light transport. Since the luminaire degradation state $L_j(t) \in [0, 1]$ represents the loss-of-function level (Section 2.1), we define the corresponding normalized output-scaling vector as

$$\mathbf{Q}(t) = \mathbf{1}_J - \mathbf{L}(t), \quad (47)$$

where $\mathbf{1}_J$ is an all-ones vector. The working-plane illuminance vector $\mathbf{E}(t) \in \mathbb{R}^N$ is then modeled as

$$\mathbf{E}(t) = \mathbf{b}_E + \mathbf{C}_E \mathbf{Q}(t), \quad (48)$$

where $\mathbf{b}_E \in \mathbb{R}^N$ is an intercept term and $\mathbf{C}_E \in \mathbb{R}^{N \times J}$ is a coefficient matrix capturing the contribution of each luminaire to each calculation point under the fixed geometry/material setting.

Stacking the K event times in simulation run s yields the state trajectory matrix $\mathbf{L}_s \in \mathbb{R}^{K \times J}$ and the corresponding output-scaling matrix $\mathbf{Q}_s = \mathbf{1}_K \mathbf{1}_J^\top - \mathbf{L}_s$. The illuminance sequence matrix $\mathbf{E}_s \in \mathbb{R}^{K \times N}$ is

$$\mathbf{E}_s = \mathbf{1}_K \mathbf{b}_E^\top + \mathbf{Q}_s \mathbf{C}_E^\top, \quad (49)$$

where $\mathbf{1}_K$ is an all-ones vector of length K . This surrogate replaces Radiance in the performance evaluation module (Algorithm 2) to compute $\mathbf{E}(t_k)$ efficiently, after which $E_{avg}(t_k)$ and $U(t_k)$ are obtained via Equation 21 - 22.

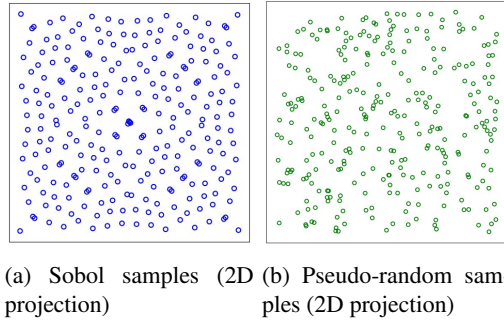


Figure 7: Distribution comparison of Sobol and pseudo-random samples in a 2D projection [26]

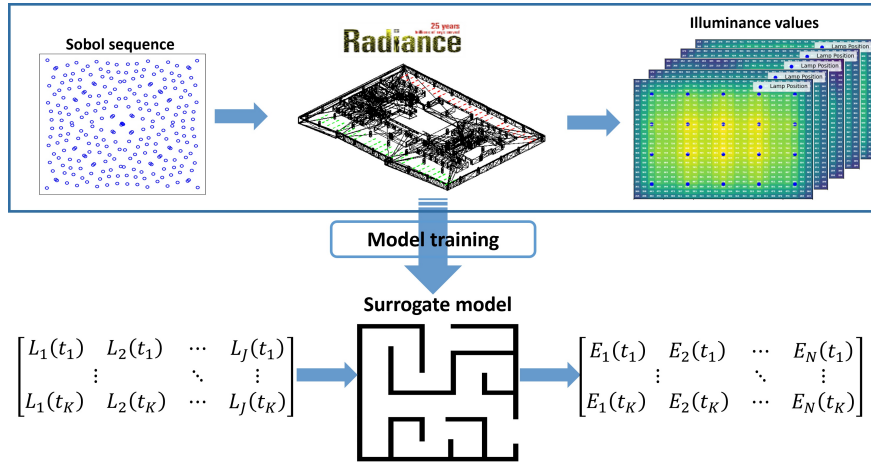


Figure 8: Surrogate-based acceleration of the performance evaluation pipeline

Training data generation and model fitting. The training dataset consists of paired samples $\{(\mathbf{L}^{(n)}, \mathbf{E}^{(n)})\}_{n=1}^{N_{\text{data}}}$, where each illuminance vector $\mathbf{E}^{(n)}$ is produced by a one-time Radiance simulation under the corresponding luminaire state $\mathbf{L}^{(n)}$. Because $\mathbf{L}(t) \in [0, 1]^J$ (Equation 10), a scrambled Sobol sequence is employed to generate quasi-random state samples that cover the high-dimensional state space efficiently. Figure 7 compares the distribution of Sobol-generated samples with conventional pseudo-random samples in a 2D projection.

The Sobol-generated luminaire states are input to Radiance to obtain illuminance outputs, and the resulting input–output pairs are used to fit the linear surrogate in Equation 48 (e.g., by least squares). Once trained, the surrogate provides a drop-in replacement for Radiance in large-scale simulations, as illustrated in Figure 8.

In this study, $N_{\text{data}} = 2016$ samples are generated for the case with $J = 76$ luminaires and $N = 265$ calculation points, producing an input matrix of size $[2016, 76]$ and an output matrix of size $[2016, 265]$. The dataset is split into 80% for training and 20% for validation. As shown in Figure 9, the predicted illuminance values closely match the Radiance outputs, yielding $R^2 = 0.9999$.

Computational benefit. As summarized in Table 2, the surrogate incurs a one-time offline cost for training-data generation and model fitting, but reduces the online cost of mapping luminaire states to illuminance from order minutes per snapshot (Radiance) to order milliseconds. This acceleration is essential for the large-scale Monte Carlo evaluation and multi-policy search required by the simulation-in-the-loop maintenance optimization framework.




Figure 9: Comparison of Radiance outputs and surrogate predictions ($R^2 = 0.9999$).

Table 2

Order-of-magnitude wall-clock estimates for performance evaluation (representative setting)

Task	Radiance (direct)	Surrogate-based
Offline training data generation	–	32 hours
Offline model fitting	–	1 minute
Online evaluation: 10,000 simulations [†]	~700 days	~20 seconds
Online policy search: 100 policy settings [‡]	~ decades	~33 minutes

[†] Estimates are based on a representative case with $J = 76$ luminaires and $N = 265$ calculation points, where one Radiance evaluation for one luminaire-state snapshot takes $\mathcal{O}(1)$ minute. Each simulation run evaluates the performance at K event time, and the reported time corresponds to evaluating 10,000 simulations for one policy setting. [‡] Policy-search time multiplies the per-policy evaluation by 100 policy settings; values are order-of-magnitude estimates intended to highlight scalability rather than exact hardware-dependent runtimes.

5. Case study: model calibration and maintenance-optimization results

The proposed performance-driven maintenance-optimization framework is instantiated on a representative office lighting system. A real building zone is modeled in Radiance, degradation and failure models are calibrated using LM-80 accelerated tests, and large-scale Monte Carlo evaluation is conducted to identify Pareto-optimal opportunistic maintenance policies.

5.1. Case description and simulation setup

To demonstrate the proposed framework, a working zone (Zone 1) in an office building at Queensland University of Technology (QUT) is considered. Zone 1 is designated for writing, reading, and data-processing tasks, as shown in Figure 10. The zone measures 65.38 m in length and 6.80 m in width. The working plane (WP) is defined at a height of 0.80 m above the finished floor and spans 57.73 m \times 4.80 m, consistent with AS/NZS 1680.1:2006 requirements [17]. The lighting system comprises $J = 76$ LED luminaires of two types: 46 B7 and 30 D13 luminaires (Figure 10b). Calculation points on the WP are placed on a 1 m grid, satisfying the maximum spacing requirement in [15]. Surface reflectances for the ceiling, walls, and floor are set to 0.7, 0.5, and 0.2, respectively, following AS/NZS 1680.1:2006 [17]. The maintained illuminance requirement is set to $S_E = 500$ lux and the uniformity requirement to $S_U = 0.6$, also consistent with AS/NZS 1680.1:2006 [17].

Package-degradation parameters are estimated using the Bayesian procedure in Section 2.3 and LM-80 test data. Since only B7 package test data are available, the resulting posterior is applied to both B7 and D13 packages to demonstrate the end-to-end framework; this simplifying assumption is relaxed in future work when D13-specific test data become available. In the absence of standardized lifetime-test data and maintenance records for LED drivers, baseline driver lifetimes are specified using Weibull distributions calibrated to match the mean time to failure (MTTF) of the package-degradation model (Scenario S1). Sensitivity analyses are then performed by perturbing the Weibull parameters to represent earlier/later driver-failure regimes (Scenarios S2-S3).

Corrective maintenance (CM) actions are assumed to have 3 days (package degradation failure) and 2 days (driver failure) service time to complete replacement on site. Each policy setting is evaluated using $S = 10,000$ Monte Carlo replications. The lighting system is assumed to operate 12 hours per day over an operational horizon of $T_{\text{over}} = 50$ years,

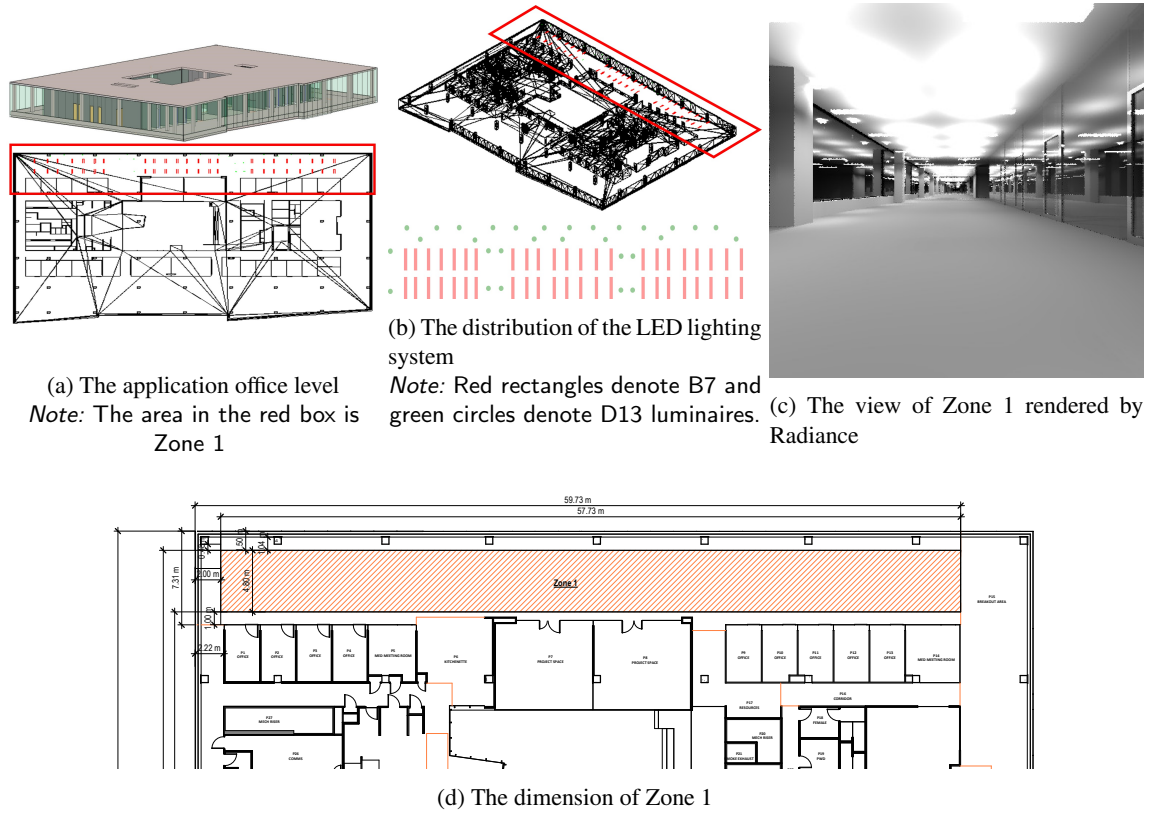


Figure 10: The application case and lighting system configuration.

corresponding to a typical building lifetime. Preventive-maintenance (PM) intervals are varied from 365 days to 18,250 days (50 years) in steps of 365 days, and the opportunistic-maintenance thresholds are swept from 0.05 to 1.00 in steps of 0.05. Statistical filtering of candidate policies is performed at significance level $\alpha_{sv} = 0.05$. Key simulation settings are summarized in Table 3.

Table 3

Parameters for the maintenance-optimization simulation study.

Parameter	Value
Zone 1 length, width, height (m)	65.38, 6.80, 3.55
Ceiling, wall, floor reflectance	0.7, 0.5, 0.2
WP length, width, height (m)	57.73, 4.80, 0.80
WP grid spacing (m)	1
Number of B7, D13 luminaires	46, 30
Maintained illuminance requirement S_E (lux)	500
Uniformity requirement S_U	0.6
CM service time (days): Package degradation (d_{CM-pf}), driver failure (d_{CM-df}) ,	3, 2
On-site replacement duration (days)	1
Monte Carlo replications per policy (S)	10,000
Statistical filtering level α_{sv}	0.05
Operational horizon T_{over} (years)	50
PM interval range T_{PM} (days)	[365, 18,250]
OM threshold range H_{OM}	[0.05, 1.00]




Figure 11: Posterior and extrapolated predictive degradation trajectories.

All simulations are parallelized on QUT's high-performance computing (HPC) system using *joblib*. For the full policy sweep, wall-clock runtimes are on the order of a few hours per PM-interval batch under multi-core execution.

5.2. Bayesian calibration using LM-80 accelerated degradation data

After defining the zone geometry, lighting requirements, and simulation settings, the package-degradation model is calibrated from LM-80 accelerated test data using Bayesian inference.

LM-80 accelerated degradation tests are conducted at three temperature levels, 55°C, 85°C, and 105°C, with 25 LED packages tested at each temperature. Each unit is inspected every 1,000 hours over a total duration of 10,000 hours. Inspection times are converted into years for model fitting.

Diffuse Normal priors are assigned to transformed parameters, $\ln A \sim \mathcal{N}(0, 10^2)$ and $\ln C \sim \mathcal{N}(0, 10^2)$, and half-Normal priors are assigned to nonnegative parameters to enforce physically consistent monotone degradation behavior, $b \sim \mathcal{N}^+(0, 10^6)$ and $E_a \sim \mathcal{N}^+(0, 10^6)$. Posterior inference is performed using *CmdStanPy* with four Markov chains, each run for 2,000 warm-up iterations and 2,000 sampling iterations.

To validate extrapolation capability across stress levels, only the 55°C and 105°C datasets are used to obtain the posterior, and posterior-predictive trajectories at 85°C are generated and compared against the observed 85°C data. As shown in Figure 11, the posterior mean (solid red line) and 95% predictive interval (shaded band) encompass the observed step trajectories, supporting both the non-homogeneous gamma-process model and the Bayesian calibration procedure.

Finally, all three temperature datasets are combined to re-estimate the model parameters. Table 4 summarizes posterior means and 95% credible intervals, with \hat{R} statistics indicating good convergence ($\hat{R} < 1.01$). Figure 12 further compares the posterior-predictive trajectories against observations at each temperature, showing consistent agreement across the full accelerated-test range.

Table 4
Posterior means and 95% credible intervals for LM-80 calibration.

Parameters	Posterior mean	95% CrI	\hat{R}
$\widehat{\ln A}$	2.2393	[1.9472, 2.5366]	1.0012
\hat{b}	0.8841	[0.7161, 1.0370]	1.0013
$\widehat{\ln C}$	3.7446	[2.7509, 4.7680]	1.0020
\widehat{E}_a	0.0815	[0.0505, 0.1112]	1.0020

5.3. Extrapolation to operating conditions and baseline scenario specification

The posterior over model parameters is propagated to operating temperatures to obtain use-condition degradation dynamics, and driver lifetime distributions are specified to define baseline component-reliability scenarios.

Based on measured operating conditions and practical relevance, 45°C is selected as the reference temperature for extrapolation to normal use. Using the fitted stress–acceleration relationship, the extrapolated rate parameter is $\hat{\beta}(T = 45^\circ\text{C}) = 828.68$ with a 95% credible interval of [700.77, 963.62].

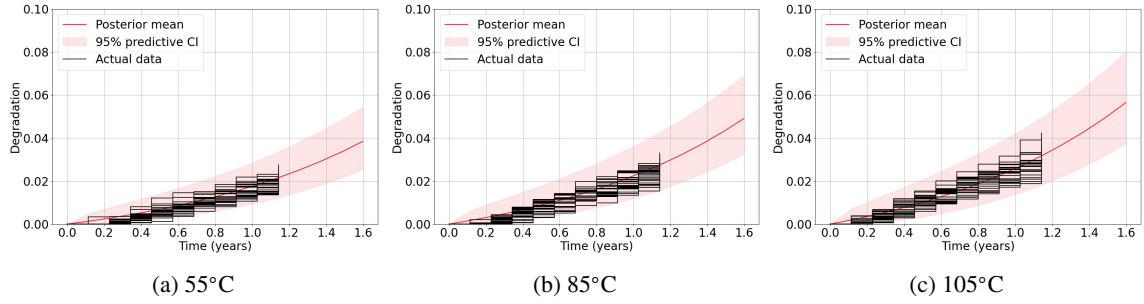


Figure 12: Observed degradation data (black steps) and posterior-predictive trajectories (red mean and 95% interval) at each temperature.

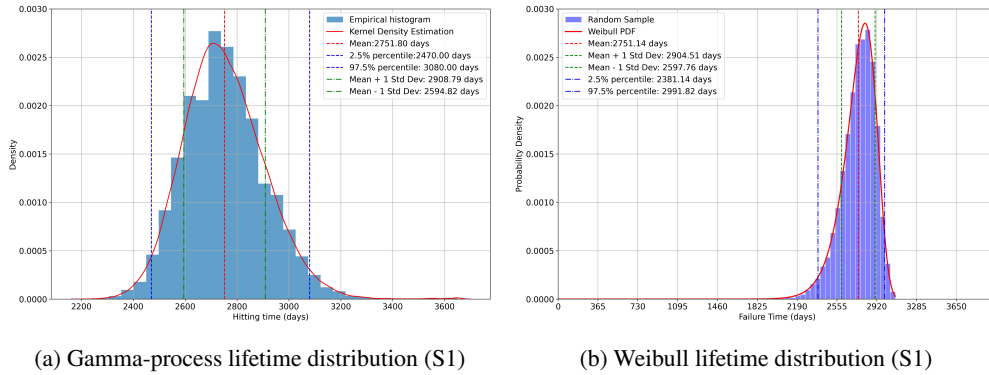


Figure 13: Lifetime distributions for LED packages and drivers under the baseline scenario (S1).

For LED drivers, a Weibull distribution is specified to match the package MTTF under the baseline setting (Scenario S1). Specifically, Weibull parameters (21.82, 2818.09) are calibrated such that the driver MTTF aligns with the package MTTF. The resulting MTTFs for packages and drivers (7.54 years and 7.53 years, respectively) are illustrated in Figure 13. Alternative Weibull parameterizations are introduced later to assess sensitivity to driver-reliability assumptions.

5.4. Maintenance-optimization results and policy interpretation

With the calibrated degradation and baseline driver-lifetime setting (S1), the proposed multi-objective optimization is evaluated over the grid of (T_{PM}, H_{OM}) policy settings. The full design includes 50 PM intervals and 20 OM thresholds (1,000 candidate policies). After statistical filtering using one-sided Welch's t -tests, 36 policies remain as statistically robust, non-dominated solutions.

Figure 14a depicts the resulting Pareto front in three-dimensional objective space defined by the deficiency ratio (f_1), the total number of site visits (f_2), and the total number of luminaire replacements (f_3). The background color illustrates variations in the deficiency ratio (f_1), where blue regions correspond to higher values of f_1 , highlighting scenarios with poorer lighting performance. In contrast, yellow regions indicate lower values of f_1 , signifying scenarios with superior system performance. Across the Pareto front, a reduction in total luminaire replacements (f_3) is generally accompanied by an increase in site visits (f_2), indicating that opportunistic maintenance activities become less aggressive and are more spread out over time, resulting in an increase in overall site visits.

As shown in Figure 14b, a useful interpretation is obtained by relating Pareto solutions to the effective OM aggressiveness captured by the OM age T_{OM} (Section 4.2). The Pareto solution can be classified into three main groups. In Group 1, as the T_{OM} increases, the total number of luminaire replacements (f_3) initially decreases significantly from approximately 759.5 to 532. During this transition, the total number of site visits (f_2) slightly decreases from approximately 9.99 to 7.81 before subsequently rising to approximately 13.68. This behavior occurs because, with an increasing T_{OM} , fewer OM replacements are executed, leading to lower replacement but poorer lighting performance.

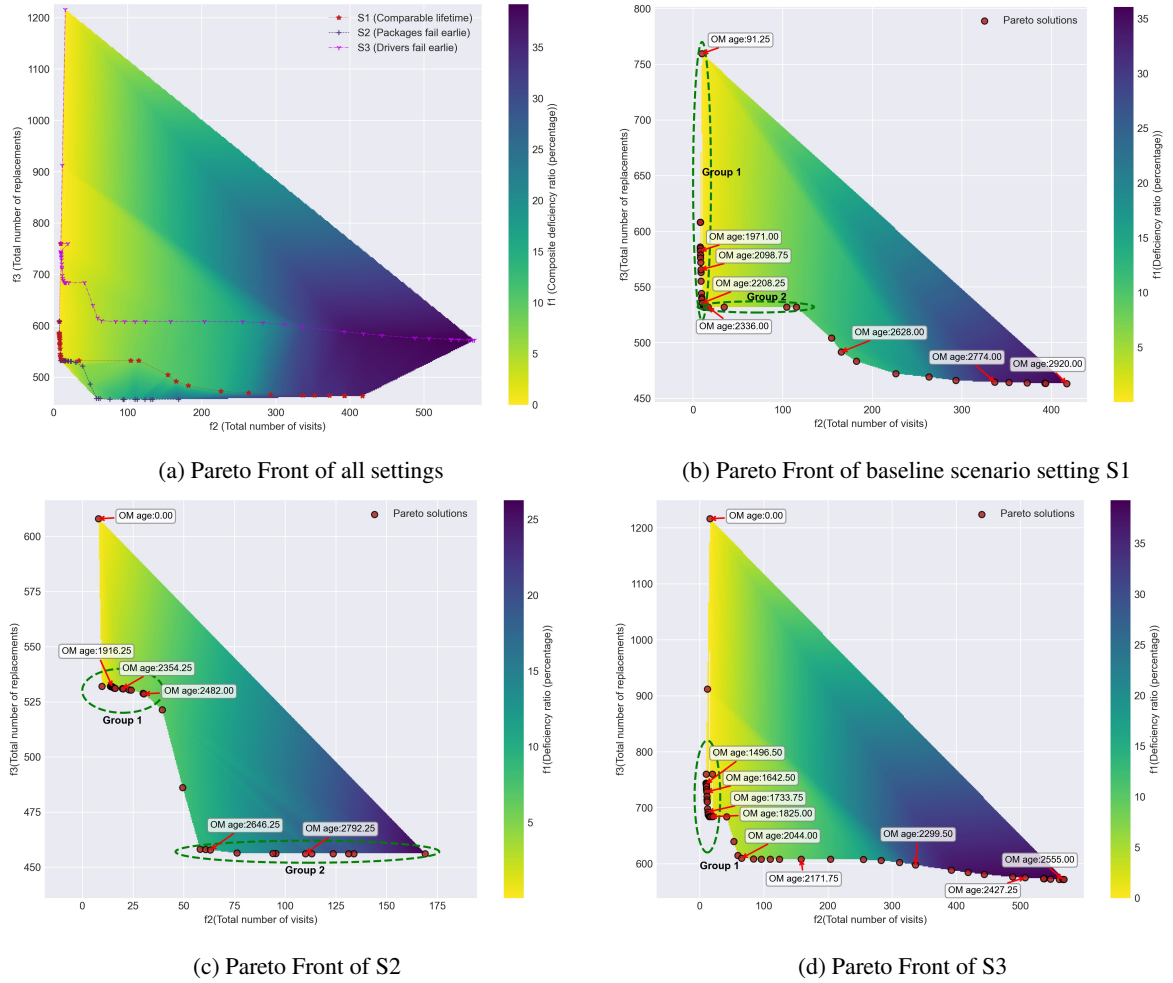


Figure 14: The results of baseline scenario S1, S2 (Packages fail earlier) and S3 (Drivers fail earlier)

However, since the T_{OM} within this range remains below the 2.5% percentile of the lifetime distributions of the packages and drivers, most LEDs are replaced by OM, and thus the increase in the number of site visits (f_2) remains moderate.

When the T_{OM} reaches the MTTF of the two critical components, solutions transition to Group 2. The replacements initiated by CM visits partially offset the reduction in OM replacements, yielding a marginal reduction in f_3 . As the OM policy adopts a more conservative stance, the system “waits” for enough CM visits to replace all LEDs. Once the T_{OM} exceeds the components’ MTTF, the solutions migrate to Group 3 (the remaining points) and the OM policy becomes markedly more conservative. Consequently, both the total number of site visits (f_2) and the deficiency ratio (f_1) begin to noticeably increase, indicating a clear trade-off: conservative policies reduce replacements but significantly compromise system performance.

Along the Pareto front, moving from the upper-left extreme toward the “corner” solutions yields a substantial reduction in luminaire replacements (f_3) with only modest increases in site visits (f_2). In practice, the corner solution represents an attractive compromise for facility managers, minimizing total maintenance interventions while maintaining acceptable system performance. Notably, the upper-left point corresponds to a PM policy. These results also suggest that a CM-driven OM policy might outperform a PM policy in practice, but the greater reliance on CM visits would place additional strain on spare-parts inventory management.

To illustrate the operational meaning of the performance objective, Figure 15 presents illuminance maps from one representative simulation trajectory. Warmer colors indicate higher illuminance values, while cooler colors indicate dimmer regions on the working plane. As degradation accumulates and failures occur, the illuminance distribution

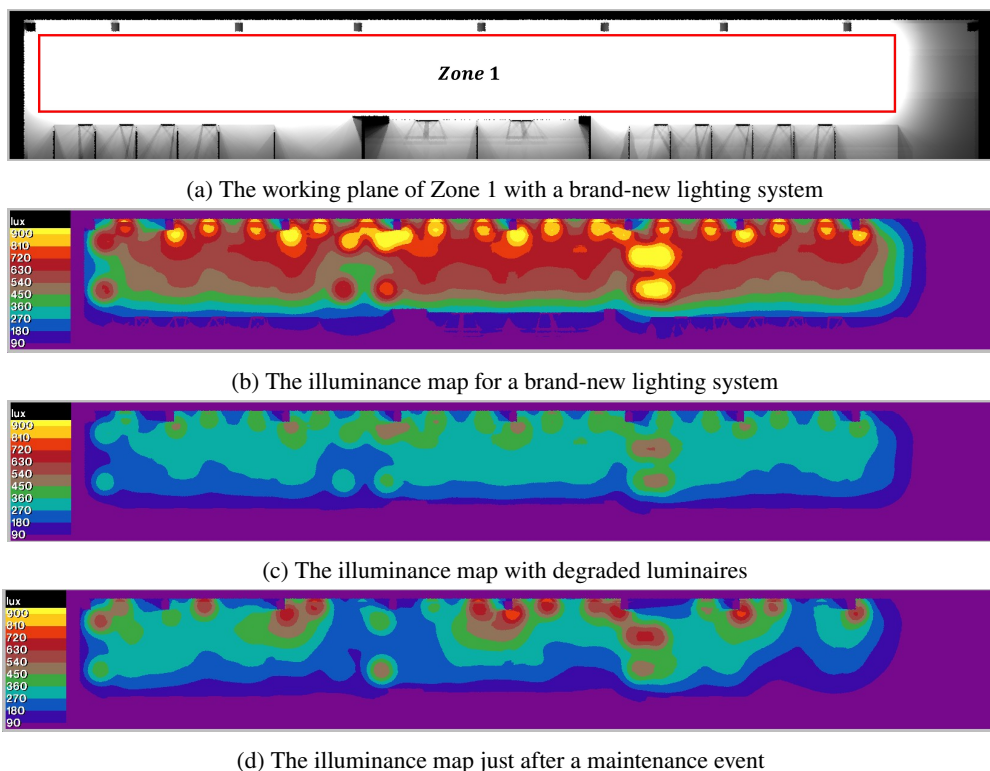


Figure 15: The lighting system's performance fluctuation under a maintenance scenario for Zone 1

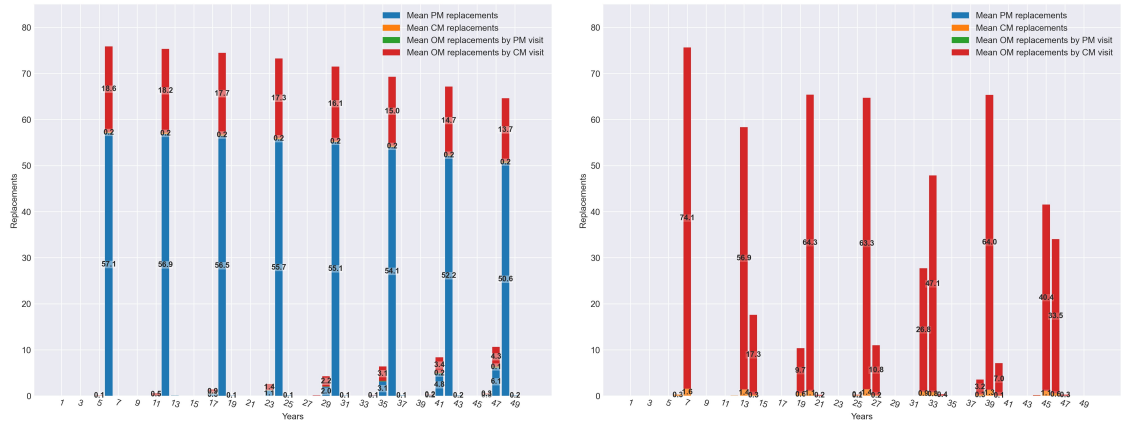
becomes less uniform and the average level decreases (e.g., Figure 15c), increasing the deficiency ratio. After a maintenance event (Figure 15d), the map moves closer to the reference state (Figure 15b), and the growth of performance deficiency is arrested.

From the Pareto set, two representative policies are selected to demonstrate how different parameter choices alter the composition of maintenance actions. Their decision variables and objective values are summarized in Table 5, and the corresponding annual mean replacement breakdowns are shown in Figure 16.

For the first scenario, the OM threshold is 0.2 and the PM interval is 2190 days – below the 2.5th percentile of the package/driver lifetime distribution – yielding a *PM-driven OM policy*. Although the OM age T_{OM} (1752 days) is relatively short – implying an aggressive OM policy – the low variance in component lifetimes means that most replacements still occur at scheduled PM visits. As shown in Figure 16a, a PM every 6 years (2190 days) replaces about 70% of the LEDs in the system, and any remaining units are opportunistically swapped out during CM visits in the same year. Over successive cycles, both the PM and OM-by-CM replacements shift to occur immediately before or after each 6-year PM cycle, reflecting the stochastic variability inherent in the degradation models. While this variability complicates precise maintenance planning, it provides facility managers with lead-time information for spare-parts provisioning and workforce allocation to better mitigate unexpected failures.

For the second scenario, the PM interval is 11315 days – above the 97.5th percentile of component lifetimes – and the OM threshold is 0.8 ($T_{OM} = 2263$ days). This produces a *CM-driven OM policy* dominated by failure-triggered replacements. As Figure 16b illustrates, the mean lifetimes are 7.4 years, and the T_{OM} is 6.2 years, so nearly the entire LED lighting fleet is renewed within a two-year window every 6-7 years via CM and OM-by-CM replacements. The inherent uncertainty in stochastic degradation models causes these renewal events to occur slightly earlier or later around the 6- to 7-year mark, yielding actionable insights for asset managers to optimize maintenance lead times, inventory strategies, and overall resource planning.

In summary, the Pareto front reveals a cluster of solutions in which the chosen maintenance parameters (T_{PM} and H_{OM}) achieve a balanced compromise among the conflicting objectives. This demarcation of viable strategies offers decision-makers multiple options. For example, if the primary concern is to minimize disruptions and maintenance



(a) The stacked bar chart under OM 0.2 and PM 2190 days (T_{OM} 1752 days) (b) The stacked bar chart under OM 0.8 and PM 11315 days (T_{OM} 2263 days)

Figure 16: The stacked bar charts of selected scenarios

Table 5
Selected Pareto solutions

Scenario	f_1 (Deficiency ratio (percentage))	f_2 (No. of visits)	f_3 (No. of replacements)	OM threshold	PM interval	OM age	PM replacements	CM replacements	OM (PM)	OM (CM)
1	0.1283	8.0917	608	0.2	2190	1752	457.2635	2.004	0.9075	147.8250
2	0.8719	12.0519	532.0176	0.8	11315	2263	0	12.2372	0	519.7804

^a OM (PM): number of OM replacements triggered by PM visits;

^b OM (CM): number of OM replacements triggered by CM visits.

resources, a strategy with fewer site visits (and a slightly higher deficiency ratio) may be preferred. Conversely, if maintaining optimal system performance is paramount, then a strategy that involves more frequent site visits and replacements (thereby lowering f_1) would be the logical choice.

5.5. Sensitivity analysis: impact of driver reliability assumptions

To investigate the robustness of the proposed framework against variations in LED driver failure timing, a sensitivity analysis was conducted through two additional settings: Scenario setting 2 (S2) with parameters (25.61, 3301.09) and Scenario setting 3 (S3) with parameters (17.91, 2342.09). In S2, LED drivers exhibit delayed failures, where the mean driver failure time is set at 3-standard-deviations later than the LED package MTTF, reflecting scenarios in which the driver reliability is significantly better than that of the packages, as shown in Figure 17a. In S3, drivers exhibit earlier failures, where the mean driver failure time is set at 3-standard-deviations earlier than the packages' MTTF, indicating scenarios with comparatively poorer driver reliability, as shown in Figure 17b.

The Pareto front under S2 (Figure 14c) exhibits trends broadly consistent with S1. Because drivers fail later, system renewal is primarily governed by package degradation, and many policies tolerate substantial degradation before replacement. Initially, the solution with the lowest deficiency ratio (f_1), situated in the upper-left corner, represents a PM-driven maintenance strategy. For these Group 1 solutions, increasing the OM age T_{OM} reduces the OM replacements; however this reduction is partially offset by a rise in CM replacements, resulting in a slight net decrease in f_3 . As the T_{OM} approaches and then surpasses the package MTTF, the policy becomes conservative (Group 2), resulting in decreased OM replacements but increased site visits (f_2). Additionally, replacements triggered by CM visits partially offset the reduction in replacements, causing a minor decrease in f_3 . Compared to S1, S2 typically demonstrates poorer system performance (higher f_1) since the failures become package-driven; thus, LEDs typically reach significant degradation before replacement.

For Scenario S3 (Figure 14d), the top-left solutions also represent PM-driven maintenance, similar to S2. In Group 1, the aggressive OM policy results in slightly increased site visits (f_2). As the T_{OM} increases, the deficiency ratio (f_1) worsens. When the OM policy shifts to conservative strategies, OM replacements decrease significantly, but

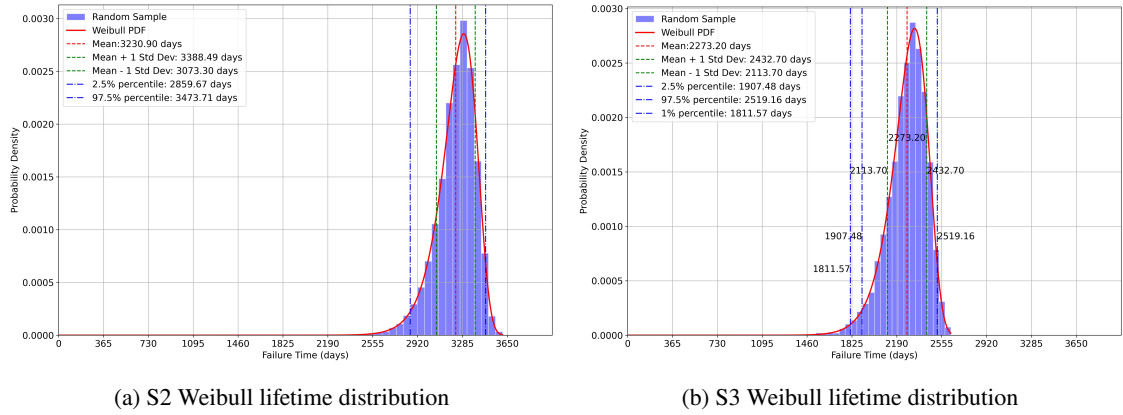


Figure 17: The setting of two scenarios for sensitive analysis

CM-driven replacements rapidly increase, indicating faster system deterioration. Notably, S3 reveals a pronounced upward shift in the total replacements (f_3). Given the earlier driver failures, the minimum number of replacements in S3 nearly equals the maximum replacements observed in S1 and S2, underscoring the critical impact of premature driver failure on maintenance planning. These shifts highlight the trade-off between failure-dominated schedules and lighting performance, and underscore how the model can guide policy selection under differing component lifetime distributions.

Taken together, S2–S3 confirm that the proposed framework is structurally robust (the same trade-offs persist), while also demonstrating that relative component reliability strongly influences where the Pareto front lies.

6. Conclusion

This paper proposed an integrated, performance-driven, simulation-in-the-loop framework for long-term maintenance optimization of large-scale LED lighting systems. At the component level, gradual degradation of LED packages was modeled using a semi-physical non-homogeneous Gamma process whose mean evolution follows an exponential lumen-maintenance trend, while abrupt driver outages were represented by a Weibull lifetime model. The two mechanisms were combined through a competing-failure formulation to define luminaire operating states. Model parameters were calibrated from LM-80 accelerated degradation data via Bayesian inference, enabling both parameter uncertainty and stochastic process variability to be propagated to use conditions for downstream Monte Carlo evaluation. Sensitivity studies further examined how driver reliability assumptions affect maintenance outcomes.

At the system level, spatial illuminance simulations were linked to operational requirements through two standard commissioning indices (average illuminance and uniformity). These static indices were then converted into a long-term dynamic performance metric by defining performance-deficiency durations across event intervals and aggregating them into a deficiency ratio over the operating horizon. This dynamic metric provides an interpretable measure of spatio-temporal compliance that is directly actionable for maintenance decision making. To make large-scale policy evaluation computationally feasible, a surrogate-based illuminance mapping was developed to replace repeated Radiance evaluations, yielding orders-of-magnitude acceleration while preserving high predictive fidelity.

Building on these elements, a performance-driven opportunistic maintenance policy was evaluated and optimized under three competing objectives: minimizing the deficiency ratio, the total number of site visits, and the total number of luminaire replacements. Preventive-maintenance intervals and opportunistic thresholds were used as decision variables, and policies were assessed through a discrete-event Monte Carlo simulator coupled with the performance-evaluation module. A real office-zone case study demonstrated that the proposed framework can reveal clear Pareto trade-offs and provide decision support for selecting maintenance policies aligned with different performance and resource priorities.

Several extensions merit future investigation. First, spatial heterogeneity could be incorporated more explicitly by accounting for location-dependent criticality (e.g., zones with stricter task requirements) and by developing location-aware maintenance actions. Second, practical resource constraints such as spare-parts inventory, crew capacity, and access scheduling could be integrated to produce more implementable policies. Finally, richer cost models and real

maintenance records would allow the current resource proxies (visits and replacements) to be mapped into monetary objectives and further strengthen field applicability.

Acknowledgments

I would like to express my sincere gratitude to the reviewers for their valuable feedback and constructive comments. This research was funded by the Australian Research Council through the ARC Research Hub for Resilient and Intelligent Infrastructure Systems (RIIS) (IH210100048). Additionally, I appreciate the collaboration and support from my industry partner, Fredon, and Queensland University of Technology (QUT). Computational (and/or data visualization) resources and services used in this work were provided by the eResearch Office, Queensland University of Technology, Brisbane, Australia.

A. Appendix

A.1. Algorithms of the Simulation Model

Algorithm 1 Discrete-event maintenance simulation for a given policy

1: **Input:**

- Degradation/failure models for luminaires: Gamma-process parameters (or posterior draws) and Weibull parameters;
- Policy variables: OM threshold H_{OM} and PM interval T_{PM} ;
- Time parameters: operational horizon T_{over} , maintenance service times (e.g., d_{CM-df} , d_{CM-pf});
- Recording schedule for performance evaluation (e.g., next record time rule for t^{record});
- Simulation size: number of Monte Carlo replications S .

2: **Output:** For each run s : event-time vector \mathbf{T}_s^{ts-evt} (Equation 39), state trajectory \mathbf{L}_s (Equation 40), and counts ($N_{tv,s}$, $N_{tr,s}$) (Equation 42 - 43); and averages (N_{STV} , N_{STR}) (Equation 44 - 45).

3: Initialize accumulators for N_{STV} and N_{STR} .

4: **for** $s = 1$ to S **do**

5: Initialize time $t \leftarrow 0$, event index $k \leftarrow 1$.

6: Initialize luminaire states $\mathbf{L}(0)$ and schedule initial event vectors $\mathbf{t}_{s,1,j}$ for all j using Equation 36.

7: Initialize $N_{cm,vis,s} \leftarrow 0$, $N_{pm,vis,s} \leftarrow 0$, $N_{cm,s} \leftarrow 0$, $N_{pm,s} \leftarrow 0$, $N_{om,s} \leftarrow 0$.

8: **while** $t < T_{over}$ **do**

9: Form the system event matrix $\mathbf{T}_{s,k}^{ts-evt}$ using Equation 37.

10: Identify the next event time and its source: $(t_{s,k}^{ts-evt}, j^*, \text{type}) \leftarrow \arg \min \mathbf{T}_{s,k}^{ts-evt}$

11: Set $t \leftarrow t_{s,k}^{ts-evt}$ and append t to \mathbf{T}_s^{ts-evt} .

12: **if** type = PM (i.e., $t = t_{s,k,j^*}^{pm}$) **then**

13: $N_{pm,vis,s} \leftarrow N_{pm,vis,s} + 1$; replace luminaire j^* ($N_{pm,s} \leftarrow N_{pm,s} + 1$).

14: Set maintenance completion time $t_{s,k,j^*}^{event-end} \leftarrow t + d_{CM-df}$.

15: **OM selection:** For each luminaire $\ell \neq j^*$, compute $H_{mov,\ell}$ via Equation 30.

16: **if** $H_{mov,\ell} \leq H_{OM}$ **then**

17: Replace luminaire ℓ opportunistically ($N_{om,s} \leftarrow N_{om,s} + 1$).

18: Update/reset its future event times in $\mathbf{t}_{s,k,\ell}$ (e.g., next PM and predicted CM times).

19: **end if**

20: Update/reset future event times for luminaire j^* (e.g., resample failure times after replacement, schedule next PM).

21: **else if** type = CM-df (i.e., $t = t_{s,k,j^*}^{cm-df}$) **then**

22: $N_{cm,vis,s} \leftarrow N_{cm,vis,s} + 1$; replace luminaire j^* ($N_{cm,s} \leftarrow N_{cm,s} + 1$).

23: Set $t_{s,k,j^*}^{event-end} \leftarrow t + d_{CM-pf}$.

24: **OM selection:** For each luminaire $\ell \neq j^*$, compute $H_{mov,\ell}$ via Equation 30; if $H_{mov,\ell} \leq H_{OM}$, apply OM and update/reset its event times.

25: Update/reset future event times for luminaire j^* after replacement.

26: **else if** type = CM-pf (i.e., $t = t_{s,k,j^*}^{cm-pf}$) **then**

27: $N_{cm,vis,s} \leftarrow N_{cm,vis,s} + 1$; replace luminaire j^* ($N_{cm,s} \leftarrow N_{cm,s} + 1$).

28: Set $t_{s,k,j^*}^{event-end} \leftarrow t + d_{CM-pf}$.

29: **OM selection:** For each luminaire $\ell \neq j^*$, compute $H_{mov,\ell}$ via Equation 30; if $H_{mov,\ell} \leq H_{OM}$, apply OM and update/reset its event times.

30: Update/reset future event times for luminaire j^* after replacement.

31: **else if** type = END (i.e., $t = t_{s,k,j^*}^{event-end}$) **then**

32: Mark maintenance complete for luminaire j^* and set $t_{s,k,j^*}^{event-end} \leftarrow T_{over}$.

33: Update \mathbf{t}_{s,k,j^*} (e.g., enable future PM/CM events as applicable).

34: **else if** type = RECORD (i.e., $t = t_{s,k,j^*}^{record}$) **then**

35: No maintenance action; update the next record time(s) in the event vectors (e.g., advance t^{record}).

36: **end if**

37: Record system state $\mathbf{L}(t)$ at the current event index and append to the trajectory \mathbf{L}_s .

38: Increment event index $k \leftarrow k + 1$.

39: **end while**

40: Compute totals for run s : $N_{tv,s}$ and $N_{tr,s}$ using Equation 42 - 43.

41: **end for**

42: Average over runs: N_{STV} and N_{STR} using Equation 44–Equation 45.

Algorithm 2 Performance evaluation model (PEM) for one simulation run1: **Input:**

- Performance standards: S_E and S_U ;
- Event-time vector $\mathbf{T}_s^{\text{is-evt}} = [t_{s,1}, \dots, t_{s,K}]$ (Equation 39);
- Luminaire-state trajectory \mathbf{L}_s (Equation 40);
- Performance mapping model (Radiance or surrogate; Section 4.4).

2: **Output:** Deficiency ratio $R_{DR,s}$ for run s (Equation 41) (and optionally illuminance sequence \mathbf{E}_s).

▷ Step 1: map states to static performance indices at event times

3: **for** $k = 1$ to K **do**4: Retrieve luminaire state vector $\mathbf{L}(t_{s,k})$ from \mathbf{L}_s .5: Evaluate the working-plane illuminance vector $\mathbf{E}(t_{s,k})$ using the performance mapping model.6: Compute $E_{\text{avg}}(t_{s,k})$ via Equation 21 and $U(t_{s,k})$ via Equation 22.7: **end for**

▷ Step 2: compute interval-wise deficiency durations and PDD

8: Initialize $\sum T_{\text{def},s} \leftarrow 0$.9: **for** $k = 2$ to K **do**10: Using the index trajectories on $(t_{s,k-1}, t_{s,k})$, compute deficiency durations: $T_{E,\text{def},s}(k)$ is the total time in $(t_{s,k-1}, t_{s,k})$ for which $E_{\text{avg}}(t) < S_E$, $T_{U,\text{def},s}(k)$ is the total time in $(t_{s,k-1}, t_{s,k})$ for which $U(t) < S_U$.

11: Compute the performance deficiency duration (PDD) using Equation 28:

$$T_{\text{def},s}(k) \leftarrow \max(T_{E,\text{def},s}(k), T_{U,\text{def},s}(k)).$$

12: Accumulate $\sum T_{\text{def},s} \leftarrow \sum T_{\text{def},s} + T_{\text{def},s}(k)$.13: **end for**

▷ Step 3: compute deficiency ratio for the run

14: Compute $R_{DR,s} = \frac{\sum T_{\text{def},s}}{T_{\text{over}}}$ using Equation 41.**CRedit authorship contribution statement**

Haohao Shi: Writing - original draft, Visualization, Software, Methodology, Conceptualization.. **Huy Truong-Ba:** Writing - review & editing, Supervision, Funding acquisition, Methodology, Conceptualization. **Michael E. Cholette:** Writing - review & editing, Supervision, Funding acquisition, Methodology, Conceptualization. **Brenden Harris:** Writing - review, Funding acquisition, Data support. **Juan Montes:** Writing - review, Data support. **Tommy Chan:** Writing - review, Supervision, Funding acquisition.

References

- [1] Bean, R., 2014. Lighting: Interior and Exterior. Taylor & Francis Group, London, UNITED KINGDOM. URL: <http://ebookcentral.proquest.com/lib/qut/detail.action?docID=1600511>.
- [2] Cui, J., Zhao, Y., Xie, Q., Huang, L., Hua, Y., 2026. Distributed control of large-scale light-emitting-diodes for providing operating reserve based on adaptive comfort-oriented consensus algorithm. Expert Systems with Applications 303, 130458. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417425040734>, doi:10.1016/j.eswa.2025.130458.
- [3] Elhami, M., Goodarzi, S.S., Maleki, S., Sajadi, B., 2025. Three-objective optimization of the HVAC system control strategy in an educational building to reduce energy consumption and enhance indoor environmental quality (IEQ) using machine learning techniques. Journal of Building Engineering 105, 112444. URL: <https://www.sciencedirect.com/science/article/pii/S2352710225006813>, doi:10.1016/j.jobe.2025.112444.
- [4] Elsayed, E.A., 2020. Reliability Engineering. Third ed., John Wiley & Sons, Inc. doi:<https://doi.org/10.1002/9781119665946.fmatter>.
- [5] Fan, J., Chen, Y., Jing, Z., Ibrahim, M.S., Cai, M., 2021a. A Gamma process-based degradation testing of silicone encapsulant used in LED packaging. Polymer Testing 96, 107090. URL: <https://www.sciencedirect.com/science/article/pii/S0142941821000404>, doi:10.1016/j.polymertesting.2021.107090.
- [6] Fan, J., Jing, Z., Cao, Y., Ibrahim, M.S., Li, M., Fan, X., Zhang, G., 2021b. Prognostics of radiation power degradation lifetime for ultraviolet light-emitting diodes using stochastic data-driven models. Energy and AI 4, 100066. URL: <https://www.sciencedirect.com/science/article/pii/S2666546821000203>, doi:10.1016/j.egyai.2021.100066.
- [7] Fan, L., Wang, K., Fan, D., 2021c. A combined universal generating function and physics of failure Reliability Prediction Method for an LED driver. EKSPLOATACJA I NIEZAWODNOSC-MAINTENANCE AND RELIABILITY 23, 74–83. URL: <https://www.webofscience.com/api/gateway?GWVersion=2&SrcAuth=DynamicDOI&SrcApp=WOS&KeyAID=10.17531%2Fuin.2021.1.8&DestApp=DOI&SrcAppSID=EUW1EDOF2ALmBooTwxjT19p2IAcdv&SrcJTitle=EKSPLOATAACJA+I+NIEZAWODNOSC-MAINTENANCE+AND+RELIABILITY&DestDOIRegistrantName=Polskie+Naukowo-Techniczne+Towarzystwo+Eksploatacyjne>, doi:10.17531/ein.2021.1.8. num Pages: 10 Place: Lublin Publisher: Polish Maintenance Soc Web of Science ID: WOS:000605643900008.

- [8] Gorenstein, A., Kalech, M., 2022. Predictive maintenance for critical infrastructure. *Expert Systems with Applications* 210, 118413. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417422015196>, doi:10.1016/j.eswa.2022.118413.
- [9] Hu, J., Zhang, L., 2014. Risk based opportunistic maintenance model for complex mechanical systems. *Expert Systems with Applications* 41, 3105–3115. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417413008580>, doi:10.1016/j.eswa.2013.10.041.
- [10] Ibrahim, M.S., Jing, Z., Yung, W.K., Fan, J., 2021. Bayesian based lifetime prediction for high-power white LEDs. *Expert Systems with Applications* 185, 115627. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417421010216>, doi:10.1016/j.eswa.2021.115627.
- [11] IESNA, 2002. IESNA+LM-63-02. URL: https://webstore.ansi.org/preview-pages/IESNA/preview_ANSI+IESNA+LM-63-02.pdf.
- [12] IESNA, 2008. Approved Method Measuring Lumen Maintenance of LED Light Sources_ies-LM-80-08-4-pdf-free.pdf. URL: https://webstore.ansi.org/preview-pages/IESNA/preview_IESNA+LM-80-08.pdf.
- [13] IESNA, 2021. ANSI IES TM-21-21 Technical Memorandum Projecting Long-Term Lumen, Photon, and Radiant Flux Maintenance Of LED Light Sources .pdf. URL: <https://store.ies.org/product/tm-21-21-projecting-long-term-luminous-photon-and-radiant-flux-maintenance-of-led-light-sources/>.
- [14] Ikuzwe, A., Ye, X., Xia, X., 2020. Energy-maintenance optimization for retrofitted lighting system incorporating luminous flux degradation to enhance visual comfort. *Applied Energy* 261, 114379. URL: <https://www.sciencedirect.com/science/article/pii/S0306261919320665>, doi:10.1016/j.apenergy.2019.114379.
- [15] Institution, T.B.S., 2021. Light and lighting. Lighting of work places. Part 1 Indoor work places. OCLC: 1282861987.
- [16] Jiang, P.H., Wang, B.X., Wu, F.T., 2019. Inference for constant-stress accelerated degradation test based on Gamma process. *Applied Mathematical Modelling* 67, 123–134. URL: <https://www.sciencedirect.com/science/article/pii/S0307904X18305080>, doi:10.1016/j.apm.2018.10.017.
- [17] Joint Technical Committee LG-001, 2006. AS NZS 1680.1 2006 Interior and workplace lighting.pdf.
- [18] Judith, B. (Ed.), 2000. The IESNA lighting handbook: reference & application. 9. ed ed., Illuminating Engineering Society of North America, New York, NY.
- [19] Keller, A.A., 2019. Multi-objective optimization in theory and practice II: metaheuristic algorithms. 1st ed. ed., Bentham Science Publishers, Sharjah, UAE.
- [20] Ling, M.H., Ng, H.K.T., Tsui, K.L., 2019. Bayesian and likelihood inferences on remaining useful life in two-phase degradation models under gamma process. *Reliability Engineering & System Safety* 184, 77–85. URL: <https://www.sciencedirect.com/science/article/pii/S0951832017303083>, doi:10.1016/j.res.2017.11.017.
- [21] Malatji, E.M., Zhang, J., Xia, X., 2013. A multiple objective optimisation model for building energy efficiency investment decision. *Energy and Buildings* 61, 81–87. URL: <https://www.sciencedirect.com/science/article/pii/S0378778813000856>, doi:10.1016/j.enbuild.2013.01.042.
- [22] Rocchetta, R., Zhan, Z., van Driel, W.D., Di Bucchianico, A., 2024. Uncertainty analysis and interval prediction of LEDs lifetimes. *Reliability Engineering & System Safety* 242, 109715. URL: <https://www.sciencedirect.com/science/article/pii/S0951832023006294>, doi:10.1016/j.res.2023.109715.
- [23] Salata, F., Golasi, I., Falanga, G., Allegri, M., De Lieto Vollaro, E., Nardecchia, F., Pagliaro, F., Gugliermetti, F., Vollaro, A.D.L., 2015. Maintenance and Energy Optimization of Lighting Systems for the Improvement of Historic Buildings: A Case Study. *Sustainability* 7, 10770–10788. URL: <https://www.mdpi.com/2071-1050/7/8/10770>, doi:10.3390/su70810770. number: 8 Publisher: Multidisciplinary Digital Publishing Institute.
- [24] Shakespeare, R., Ward Larson, G., 2004. Rendering with Radiance: The Art and Science of Lighting Visualization, Revised Edition. BookSurge, LLC.
- [25] Shi, H., Zhang, J., Zio, E., Zhao, X., 2023. Opportunistic maintenance policies for multi-machine production systems with quality and availability improvement. *Reliability Engineering & System Safety* 234, 109183. URL: <https://www.sciencedirect.com/science/article/pii/S0951832023000984>, doi:10.1016/j.res.2023.109183.
- [26] Sobol', I.M., Asotsky, D., Kreinin, A., Kucherenko, S., 2011. Construction and Comparison of High-Dimensional Sobol' Generators. *Wilmott* 2011, 64–79. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/wilm.10056>, doi:10.1002/wilm.10056. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/wilm.10056>.
- [27] Sun, B., Fan, X., van Driel, W., Cui, C., Zhang, G., 2018. A stochastic process based reliability prediction method for LED driver. *Reliability Engineering and System Safety* 178, 140–146. doi:10.1016/j.res.2018.06.001.
- [28] Sun, B., Fan, X., Ye, H., Fan, J., Qian, C., van Driel, W., Zhang, G., 2017. A novel lifetime prediction for integrated LED lamps by electronic-thermal simulation. *Reliability Engineering & System Safety* 163, 14–21. URL: <https://www.sciencedirect.com/science/article/pii/S0951832017301205>, doi:10.1016/j.res.2017.01.017.
- [29] Tan, C.M., Singh, P., 2023a. Chapter 1 - Introduction, in: Tan, C.M., Singh, P. (Eds.), *Reliability and Failure Analysis of High-Power LED Packaging*. Woodhead Publishing. Woodhead Publishing Series in Electronic and Optical Materials, pp. 1–17. URL: <https://www.sciencedirect.com/science/article/pii/B9780128224083000046>, doi:10.1016/B978-0-12-822408-3.00004-6.
- [30] Tan, C.M., Singh, P., 2023b. Chapter 2 - General failure analysis techniques for packaged LED, in: Tan, C.M., Singh, P. (Eds.), *Reliability and Failure Analysis of High-Power LED Packaging*. Woodhead Publishing. Woodhead Publishing Series in Electronic and Optical Materials, pp. 19–48. URL: <https://www.sciencedirect.com/science/article/pii/B9780128224083000083>, doi:10.1016/B978-0-12-822408-3.00008-3.
- [31] Tobias, P.A., 2012. Applied reliability. Third edition. ed., CRC Press, Boca Raton, Florida.
- [32] Wang, B., Wu, Z., Xia, X., 2017. A Multistate-Based Control System Approach Toward Optimal Maintenance Planning. *IEEE Transactions on Control Systems Technology* 25, 374–381. URL: <https://ieeexplore.ieee.org/document/7460255>, doi:10.1109/TCST.2016.

2550505. conference Name: IEEE Transactions on Control Systems Technology.
- [33] Wang, B., Wu, Z., Zhu, B., Xia, X., 2015. Optimal control of maintenance instants and intensities in building energy efficiency retrofitting project, in: 2015 54th IEEE Conference on Decision and Control (CDC), pp. 2643–2648. URL: <https://ieeexplore.ieee.org/document/7402614>, doi:10.1109/CDC.2015.7402614.
- [34] Wang, B., Xia, X., 2014. A Control System Approach to Corrective Maintenance Planning of Building Retrofitted Facilities. IFAC Proceedings Volumes 47, 12056–12061. URL: <https://www.sciencedirect.com/science/article/pii/S1474667016435342>, doi:10.3182/20140824-6-ZA-1003.00964.
- [35] Wang, B., Xia, X., Cheng, Z., Liu, L., 2020. Optimal maintenance planning in building retrofitting with interacting energy effects. Optimal Control Applications and Methods 41, 2023–2036. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/oca.2593>, doi:10.1002/oca.2593. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/oca.2593>.
- [36] Wang, B., Xia, X., Zhang, J., 2014. A multi-objective optimization model for the life-cycle cost analysis and retrofitting planning of buildings. Energy and Buildings 77, 227–235. URL: <https://www.sciencedirect.com/science/article/pii/S037877881400245X>, doi:10.1016/j.enbuild.2014.03.025.
- [37] Wang, J., Luo, L., Mu, G., Ma, Y., Ni, C., 2025. Joint optimization of quality control and maintenance policy for a production system with quality-dependent failures. Expert Systems with Applications 272, 126800. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417425004221>, doi:10.1016/j.eswa.2025.126800.
- [38] Wen, M., Ibrahim, M.S., Meda, A.H., Zhang, G., Fan, J., 2024. In-Situ early anomaly detection and remaining useful lifetime prediction for high-power white LEDs with distance and entropy-based long short-term memory recurrent neural networks. Expert Systems with Applications 238, 121832. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417423023345>, doi:10.1016/j.eswa.2023.121832.
- [39] Wen, M., Jing, Z., Ibrahim, M.S., Fan, J., Zhang, G., 2021. A hybrid degradation modeling of light-emitting diode using permutation entropy and data-driven methods, in: 2021 22nd International Conference on Electronic Packaging Technology (ICEPT), pp. 1–6. doi:10.1109/ICEPT52650.2021.9568181.
- [40] Ye, X., Xia, X., 2014. Optimal lighting project maintenance planning by a control system approach. IFAC Proceedings Volumes 47, 3152–3157. URL: <https://www.sciencedirect.com/science/article/pii/S1474667016420926>, doi:10.3182/20140824-6-ZA-1003.02030.
- [41] Ye, X., Xia, X., Zhang, L., Zhu, B., 2015. Optimal maintenance planning for sustainable energy efficiency lighting retrofit projects by a control system approach. Control Engineering Practice 37, 1–10. URL: <https://www.sciencedirect.com/science/article/pii/S0967066114002871>, doi:10.1016/j.conengprac.2014.12.014.
- [42] Zhang, H., 2018. A Viable Nontesting Method to Predict the Lifetime of LED Drivers. IEEE Journal of Emerging and Selected Topics in Power Electronics 6, 1246–1251. URL: <https://ieeexplore.ieee.org/document/8336874/>, doi:10.1109/JESTPE.2018.2826364.