Sustainable Energy Technologies and Assessments 93 (2026) 105315

Contents lists available at ScienceDirect

Sustainable Energy Technologies and Assessments

journal homepage: www.elsevier.com/locate/seta

Energy reliability and feasibility analysis for battery-constrained off-grid solar street lighting: an uncertainty-aware dimming control framework

Salaheddine Chouikh [a], Nouh Elmitwally [a,b,*], Faisal Saeed [a]

[a] *School of Architecture, Built Environment, Computing, and Engineering, Birmingham City University, UK*
[b] *Faculty of Computers and Artificial Intelligence, Cairo University, Giza 12613, Egypt*

A R T I C L E I N F O

Keywords:
Off grid solar street lighting
Dimming control
Battery aging
Uncertainty quantification
Residual quantile calibration
Feasibility analysis
Energy reliability
Coverage validation
Public lighting

A B S T R A C T

Off grid solar street lighting serves millions of people in energy-scarce communities where grid access is unavailable, yet battery aging and uncertain nightly discharge endurance make reliable service difficult to guarantee. Strict blackout budgets can be physically unattainable under minimum comfort floor constraints, a feasibility limit that existing control frameworks do not explicitly diagnose. We present a simulation-based benchmarking framework that links cycle-level battery discharge forecasting to phase-specific brightness control, and introduces two contributions to the energy reliability analysis of off grid lighting systems: (i) comfort-floor feasibility bounds that distinguish physically infeasible blackout targets from controller shortcomings, and (ii) blackout severity evaluation via unmet minutes conditional on blackout (UM|BO p90), which exposes tail-risk differences invisible to frequency-only metrics. Forecast uncertainty is handled through a model agnostic residual quantile safety buffer with battery-ID grouped splits to reduce leakage. Across repeated grouped battery splits, the floors-only feasibility estimate was 0.86% [95% CI: 0.00, 3.26] for H = 0.75 h, 6.06% [2.07, 10.04] for H = 1.0 h, and 26.47% [18.03, 34.91] for H = 1.5 h. Results are simulation-based benchmarks on NASA laboratory battery discharge data; field validation remains future work. Empirical coverage validation on held-out grouped test batteries confirmed that the residual-quantile buffer achieved 93.03% pooled empirical coverage (mean 91.61%, 95% CI: [88.69%, 94.53%] across 10 repeated runs) against a nominal 80% target, supporting its conservative empirical robustness properties.

Introduction

Off grid solar street lighting provides essential night-time public lighting in energy-scarce communities where grid access is unavailable. Battery aging and uncertain discharge endurance make reliable service difficult to guarantee, blackouts are a first-class reliability outcome, not a rare edge case. The practical challenge is to conserve energy while keeping blackout risk within an operator-defined budget, under storage constraints that existing grid-oriented approaches do not address. Despite growing interest in adaptive operation and data driven control for smart lighting [1–3], most approaches treat prediction as an end rather than converting forecasts into reliability bounded actions. In reliability sensitive infrastructure, the tails of the error distribution matter more than average accuracy, this motivates control designs that explicitly incorporate forecast uncertainty rather than relying on point predictions alone.

This paper proposes an uncertainty aware budget–constrained dimming pipeline that forecasts battery discharge runtime from compact cycle level descriptors and translates these forecasts into brightness commands under operational constraints. The residual quantile buffer is inspired by conformal calibration concepts and is used here as a model agnostic mechanism to reduce optimistic risk in downstream control [4–6]. This study focuses on the battery-runtime and control-feasibility layer; PV generation, charging dynamics, seasonal irradiance variation, and photometric lux mapping are outside the scope of the present simulation benchmark and are treated as deployment-specific extensions.

Contributions: This work presents a constraint aware integration framework for adaptive street lighting that translates battery runtime forecasts and their uncertainty into blackout rate–bounded brightness control under comfort floor constraints. While battery runtime forecasting and uncertainty calibration are established techniques, their integration into an energy reliability framework with explicit feasibility bounds and severity aware evaluation constitutes, to our knowledge, a

* Corresponding author.
E-mail address: nouh.elmitwally@bcu.ac.uk (N. Elmitwally).

https://doi.org/10.1016/j.seta.2026.105315
Received 12 June 2026; Received in revised form 31 July 2026; Accepted 8 August 2026
Available online 12 August 2026

novel contribution to off grid lighting system analysis. Specifically, we:

1. Distinguish algorithmic failure from physical infeasibility through comfort floor feasibility analysis.
2. Evaluate reliability using both blackout frequency and severity (UM|BO p90) alongside energy savings.
3. Formulate a budget aware dimming strategy that maps uncertainty adjusted runtime predictions into brightness decisions under explicit blackout constraints.
4. Propose a leakage resistant forecasting and calibration protocol using grouped splits and residual quantile safety buffering.

Organization: Sections II–VIII cover background and methods; Section IX describes the experimental setup; Section X presents results; Sections XI–XIII discuss implications, limitations, and conclude.

Related work

Smart street lighting spans static schedules (e.g., time–of–night dimming), sensor/traffic–driven adaptation, and networked "smart city" deployments. These approaches show substantial energy saving potential but are commonly designed for grid–powered luminaires or assume uninterrupted availability, whereas off grid PV systems face storage constraints that can induce hard shutoffs. As a result, blackout behavior becomes a first–class reliability outcome rather than a rare edge case [1–3]. A second relevant thread is battery prognostics and cycle level prediction. Battery health monitoring and remaining–life modelling are mature, supported by public repositories including NASA's battery aging datasets [9,10]. However, most battery forecasting work optimizes predictive accuracy (e.g., MAE/RMSE) without explicitly analyzing how forecast errors propagate into control failures such as blackouts. In off grid lighting, that coupling is strongly nonlinear: small optimistic errors near feasibility margins can trigger late–night service loss. This motivates an evaluation design that connects the predictor directly to dimming decisions and reports reliability outcomes rather than forecasting error alone [9–11] Uncertainty quantification provides a natural bridge from forecasting to safety aware control. Conformal prediction and related calibration methods provide model agnostic uncertainty tools; in regression settings, they are often operationalized through residual quantiles or calibrated prediction intervals [4–6]. In this work we use a practical residual quantile safety buffer inspired by conformal calibration, while avoiding strong theoretical claims under battery aging and distribution shift [4–6]. This distinction is important: conformal methods speak to prediction–set coverage under assumptions, whereas meeting an operational blackout rate budget is an empirical property of the closed loop controller and constraints. Finally, prior evaluations of dimming policies often emphasize energy saving and (sometimes) blackout frequency [1–3]. Frequency alone can obscure human impact: two policies with similar blackout rates may differ substantially in how long outages last. We therefore adopt an explicit severity view using unmet–minutes conditional on blackout and its tail statistic (UM|BO p90), which reveals differences among policies that appear similar under BR–only reporting. We also introduce feasibility bounds that separate "controller weakness" from "budget impossibility," especially under minimum brightness floors and long effective night horizons. Prior off grid adaptive lighting work evaluates energy neutrality or LPSP–style blackout probability [25,26], and broader smart lighting surveys focus on energy saving without reliability constraints [1–3], but none jointly addresses comfort floor feasibility bounds, UM|BO p90 tail risk, or the separation of algorithmic failure from physical infeasibility. This work fills that gap (Table 1).

Existing adaptive street-lighting controllers, including solar-predictive approaches [25,26], SOC-threshold dimmers [22,23], and energy-neutral designs [25], mainly optimize energy efficiency or battery life. They do not explicitly test whether a blackout budget is physically achievable under minimum brightness floors, and they rarely quantify blackout severity beyond frequency. Although conformal prediction and residual-quantile methods have been used in time-series forecasting [4–6] and battery state estimation [13], their integration into a reliability-bounded off-grid lighting control pipeline remains limited. This work addresses that gap by introducing comfort-floor feasibility bounds, BR_min, and blackout tail-severity metrics, UM|BO p90, as explicit outputs of a simulation-based benchmarking framework.

Table 1
Notation and evaluation metrics.

Item	Symbol	Definition
Cycle index	t	Battery discharge cycle index
Endurance target	y_t	Discharge endurance (cycle duration), seconds
Features	x_t	Cycle level feature vector
Point forecast	$\widehat{y}_t$	$\widehat{y}_t = f(x_t)$
Horizon	H	Effective night horizon (hours)
Brightness schedule	$b_t(\tau)$	Commanded brightness (%) at time τ
Phase floors	$b_\Phi^{\min}$	Minimum brightness (%) per phase Φ
Blackout indicator	I_t	1 if blackouts occur, else 0
Unmet minutes	U_t	Shortfall duration (min), 0 if no blackout
Budget	B	Target blackout rate (%)
Blackout rate	BR	Mean blackout frequency, defined in Eq. (1)
Energy saving	ES	Relative energy saving, defined in Eq. (2)
Severity tail	UM|BO p90	90th percentile of U_t conditional on blackout, defined in Eq. (3)
Feasibility bound	$BR_{\min}(H)$	Minimum achievable BR under comfort floors
Block Length	L	Safety factor update interval

Problem formulation and metrics

We consider an off grid solar luminaire powered by a PV–battery system. Each battery discharge cycle t yields a discharge endurance target y_t (cycle duration in seconds) and a feature vector x_t extracted from cycle measurements. A supervised regression model produces a point forecast $\widehat{y}_t = f(x_t)$. This forecast is then mapped to a dimming decision for an effective night horizon H(hours), which is divided into phases $\phi \in \Phi$ (e.g., Active, Quiet, Predawn) with phase specific minimum brightness floors $b_\phi^{\min}$. Such phase–based dimming is consistent with practical part night strategies used in smart lighting deployments, where brightness is typically scheduled by time–of–night or activity level rather than held constant at 100% [1–3].

Control Objective. The controller maximizes achievable brightness subject to (i) a blackout rate constraint not exceeding a specified budget and (ii) minimum comfort floor brightness requirements. Energy savings are reported as a secondary outcome of satisfying these constraints.

Brightness schedule and constraints

Let $b_t(\tau) \in [0, 100]$ denote the commanded brightness percentage at time index τ within the horizon of night t. Phase floors impose:

$$b_t(\tau) \geq b_{\phi(\tau)}^{\min}, \forall \tau.$$

In this paper, floors are operational brightness thresholds (percentages) rather than direct lux constraints; mapping a dimming percentage to maintained illuminance depends on luminaire photometrics and installation geometry and is deployment specific.

Energy proxy and blackout event

We evaluate energy relative to a full brightness baseline. Following standard practice in LED dimming evaluations, we use brightness percentage as a proxy for relative power consumption, assuming approximately linear driver efficiency across the operating dimming range [27–29]. This is a simplifying assumption; linearity may deviate at very low dimming levels or with certain driver architectures. Accordingly, ES values reported in this paper should be interpreted as relative comparisons under this proxy, not as absolute energy measurements. Let $E_t(b_t)$ denote the total energy proxy consumed by schedule $b_t(\cdot)$ over the horizon; $E_t(100\%)$ is the proxy under constant 100% brightness.

A blackout occurs when the available endurance for cycle t is insufficient to complete the scheduled horizon. Let:

- $I_t \in \{0,1\}$ indicates whether blackouts occur
- $\text{unmet\_minutes}_t \geq 0$ denote the shortfall duration in minutes (0 if there is no blackout).

Reliability budget

An operator specifies a blackout rate budget B(e.g., 2%, 5%, 10%). Policies are compared under identical constraints and budgets.

Evaluation metrics

We report three primary metrics: blackout rate (BR), energy saving (ES), and UM|BO p90.

1. Blackout rate

$$\mathrm{BR} = \frac{1}{\mathrm{N}} \sum_{\mathrm{t}=1}^{\mathrm{N}} 1\{\mathrm{I_t} = 1\} \times 100\%$$

2. **Energy saving** relative to the 100% baseline:

$$\mathrm{ES} = \left(1 - \frac{\sum_{\mathrm{t}=1}^{\mathrm{N}} \mathrm{E_t}(\mathrm{b_t})}{\sum_{\mathrm{t}=1}^{\mathrm{N}} \mathrm{E_t}(100\%)}\right) \times 100\%$$

3. UM|BO p90 (blackout severity)

$$\mathrm{UM|BO\ p90} = \mathrm{Q}_{0.9}(\{\mathrm{U_t} : \mathrm{I_t} = 1\})$$

This statistic captures tail risk: two policies with similar BR can differ substantially in the length of their worst outages. To test sensitivity to this choice, Section X.F additionally reports UM|BO p75 and UM|BO p95.

Feasibility under comfort floors

For a fixed horizon H and fixed floors $\left\{b_\phi^{\min}\right\}$, some budgets may be infeasible regardless of policy design. We define the feasibility bound:

$$BR_{\min}(H) = BR\left(\pi^{\text{floors}}; H\right), \text{where } b_t = b_{\phi(t)}^{\min} \forall t$$

where π^{floors} is the floors only policy that holds each phase at its minimum brightness threshold. This policy is the most energy conservative admissible schedule in the evaluated set and serves as an empirical lower bound within the evaluated policy set on the achievable blackout rate under comfort constraints. If a target budget $B < \mathrm{BR}_{\min}(H)$, the target is infeasible under nominal comfort floors and must be relaxed (e.g., by adjusting floors, horizon assumptions, storage sizing, or PV capacity) rather than by controller tuning.

Dataset and preprocessing

Dataset source and scope

We use the publicly available NASA lithium–ion battery aging dataset from the NASA Ames Prognostics Center of Excellence (PCoE) repository [10], accessed via the Kaggle community conversion that provides a convenient two–tier CSV layout (metadata plus per–cycle trace files) for reproducible cycle level feature extraction [11]. The original dataset contains multiple cycle types (charge, discharge, and impedance). Because the control objective is to forecast usable energy for the forthcoming lighting period, only discharge cycles are modelled; charge cycles are retained only for contextual checks and impedance cycles are excluded [10,11]. (This modelling choice is consistent with battery PHM practice where discharge behavior is directly tied to usable runtime [9], and with data driven battery prediction work showing that discharge–curve descriptors carry strong prognostic signal [12]).

Data are publicly available via Kaggle: https://www.kaggle.com/datasets/patrickfleith/nasa–battery–dataset [11].

The NASA dataset is used here as a controlled simulation environment to evaluate the proposed framework under known battery aging trajectories. The dataset consists of laboratory discharge–to–cutoff profiles and does not represent full field–night profiles (no PV charging, no seasonal variation, no lamp–load variability). Absolute performance numbers should therefore be interpreted as simulation benchmarks rather than deployment predictions. The framework is designed to generalize deployment specific data when paired with appropriate field measurements and photometric mapping.

Cycle definition and target

A “cycle” corresponds to one contiguous discharge service event from activation at approximately nominal output to end–of–service. The target variable is the per–cycle runtime at ≈100% brightness recorded as cycle_duration (seconds) [10,11]. Throughout, we treat brightness percentage as proportional to electrical power in the operating range of interest (a standard approximation when driver efficiency is approximately flat across dimming levels), enabling relative energy comparisons across dimming schedules [1–3].

Terminology note: We refer to each discharge event as a “night cycle” for readability, although typical discharge durations in the dataset are on the order of $\sim$ 0.8 h rather than full calendar–night lengths [11]. Accordingly, H = 0.75 h is dataset–typical, H = 1.0 h is upper–tail, and H = 1.5 h is a beyond–data stress test; none of these correspond to literal calendar–night lengths. The framework generalizes to longer horizons (e.g., 6–10 h) when paired with deployment specific data and photometric mapping.

Inclusion screens and basic cleaning

To ensure physically meaningful targets and stable downstream control, discharge cycles must pass basic quality screens prior to modelling: (i) valid timestamps and non–missing telemetry, (ii) physically plausible discharge trajectories (e.g., excluding impossible voltage behavior), and (iii) nonnegative durations [11]. Where raw traces are irregularly sampled, they are resampled to a fixed cadence and re–aggregated to per–cycle descriptors so feature definitions remain consistent across cycles [11].

Feature engineering

The forecaster maps a compact per–cycle descriptor to the next cycle’s runtime. Features summarize voltage and current trajectories (proxies for usable energy) together with temperature dynamics (a proxy for resistance and thermal effects) [9,12]. The 13 active model input features span four families: voltage trajectory descriptors (avg_voltage, voltage_drop, voltage_gradient, voltage_std, start_voltage, end_voltage), current trajectory descriptors (avg_current, max_current, current_std), temperature descriptors (initial_temp, final_temp, temp_rise), and an energy proxy (energy_approx). Full feature definitions, units, and ordering are provided in Appendix Table 4.

Leakage resistant split and preprocessing protocol

Battery aging data exhibit strong within–battery dependence across cycles; naive row–level random splitting can leak information and inflate test performance. We therefore split data by battery identifier, ensuring train and test batteries are disjoint. Hyperparameter tuning and uncertainty calibration follow the same grouping principle to preserve out of sample integrity [14,15]. More generally, out of sample evaluation guidance in forecasting emphasizes strict separation of training decisions from test behavior, especially under temporal dependence and system drift [14,15].

Target outlier handling

Extreme runtime values can disproportionately affect both regression fit and downstream control stability. We therefore evaluate an optional target–only outlier filter using Tukey's inner–fence rule (bounds $Q_1 \pm 1.5\,\mathrm{IQR}$; optionally outer fences at $3\,\mathrm{IQR}$) [16,17]. To avoid preprocessing leakage, IQR bounds are computed on the training set only and applied unchanged to the test set [14,15]. We report results with and without IQR cleaning to make sensitivity visible.

Forecasting model and training protocol

Forecasting task

We learn a regression model $f(\cdot)$ that maps cycle level descriptors to discharge endurance y_t (seconds) [9,10,12].

Leakage resistant split

Following the battery ID grouped split described in Section IV, we evaluate generalization on disjoint test batteries [14,15].

Model candidates and selection

We screened several regression families (tree ensembles and nonlinear baselines). In the main pipeline we use Random Forest (RF) due to stable performance and ease of calibration/interpretation [8]. CatBoost was also competitive in screening and is reported as a strong comparator [7].

Group–aware hyperparameter tuning

Hyperparameters are tuned using GroupKFold cross–validation on training batteries, minimizing MAE across grouped folds; held out test batteries are never used in tuning [14,15,18,19]. The Random Forest grid search swept n_estimators ∈ {400, 800}, max_depth ∈ {None, 20}, max_features ∈ {sqrt, 0.5}, and min_samples_leaf ∈ {1, 2, 4}, with MAE as the tuning objective. The residual-quantile buffer was evaluated with α ∈ {0.05, 0.10, 0.20} (main configuration); α = 0.20 was selected for primary results, corresponding to a nominal 80% coverage target. A post-hoc sensitivity sweep over α ∈ {0.10, 0.20, 0.30} is reported in Section X. Results are aggregated across 10 runs (5 battery split seeds × 2 RF model seeds). Full hyperparameter and configuration details are provided in Appendix Table 5. Computational Complexity. Single-cycle RF inference completed with a median latency of 170.30 ms (p95: 208.35 ms). Batch prediction over 351 held-out cycles completed in 208.17 ms median. Residual quantile computation was negligible at 0.12 ms median. Since inference is invoked once per discharge cycle at night onset rather than continuously, the measured latency is operationally negligible for the simulated control setting. Training and calibration are performed offline; embedded deployment would require hardware-specific validation.

Forecast error metrics

We report MAE (primary), and optionally RMSE and R^2, on held out batteries. MAE is preferred for tuning due to robustness to heavy–tailed errors, which is important because rare optimistic errors can have outsized downstream impact on blackout behavior.

Output constraints for runtime prediction

To avoid physically invalid runtime outputs, predictions are constrained to nonnegative values. We screened ML/Deep Learning forecasters (rolling training) but treat the predictor as modular. The resulting point forecasts $\hat{y}_t$ feed the uncertainty buffer stage described next (Section VI), where grouped out–of–fold residuals are used to compute a conservative safety margin inspired by conformal calibration ideas [4–6].

Uncertainty buffer via residual quantile calibration

Motivation

Point forecasts of discharge endurance can be optimistic, and rare optimistic errors can disproportionately increase blackout risk in downstream control. We therefore apply a conservative, model agnostic safety margin based on the empirical distribution of prediction residuals. This is inspired by conformal–style calibration ideas for regression uncertainty quantification, where quantiles of residuals are used to form distribution–free uncertainty buffers under appropriate assumptions [4–6].

Grouped out–of–fold residuals (train only)

To avoid optimistic residual estimates due to within–battery dependence, we compute calibration residuals using GroupKFold out–of–fold (OOF) predictions on the training batteries (groups = battery IDs). This ensures each residual is generated from a model that did not train on the same battery as the evaluated point, consistent with strict out of sample evaluation guidance [14,15,18,19].

Let $\hat{y}_t^{\mathrm{OOF}}$ denote the OOF prediction for cycle t in the training set. The absolute residual is

$$r_t = \left|y_t - \hat{y}_t^{\mathrm{OOF}}\right|.$$

Residual quantile safety buffer

For a user–specified mis coverage level $\alpha \in (0,1)$, define the empirical $(1-\alpha)$–quantile of training residuals:

$$q_{1-\alpha} = \mathrm{Quantile}_{1-\alpha}(\{r_t\}).$$

Given a test–time point prediction $\hat{y}_t$, we form a conservative endurance estimate:

$$\hat{y}_t^{\mathrm{safe}} = \max(0, \hat{y}_t - q_{1-\alpha}).$$

This corresponds to a lower safety adjustment that reduces the impact of optimistic point forecasts using a single scalar buffer. The buffer is model agnostic and can be applied to any regressor used in Section V (e.g., RF [8], CatBoost [7]).

Coverage Validation. Empirical coverage validation was performed on held-out grouped test batteries using the conservative lower-runtime estimate $\hat{y}_t^{\mathrm{safe}} = \max\left(0, \hat{y}_t - q_{1-\alpha}\right)$. A cycle was considered covered when y_true ≥ ŷ_safe. For α = 0.20, corresponding to a nominal 80% target coverage level, the buffer achieved 93.03% pooled empirical coverage across 6,270 held-out cycle predictions. Across 10 repeated grouped split/model-seed runs, mean empirical coverage was 91.61% with an approximate 95% CI of [88.69%, 94.53%]. The lowest run coverage was 84.05%, remaining above the nominal target. These results support the buffer as a conservative empirical robustness mechanism, but not as a formal distribution-free conformal guarantee under battery aging or deployment shift. Fig. 2 shows per-run empirical coverage against the 80% nominal target.

Notes on interpretation and scope

The above procedure is conformal–inspired: residual quantiles are a

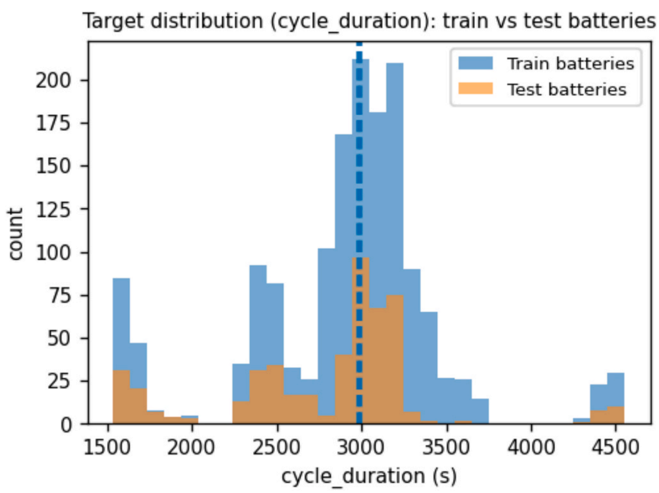

Fig. 1. Distribution of discharge–cycle runtime (cycle_duration) for train vs. test batteries under a battery ID split. H = 0.75 h (dataset–typical), H = 1.0 h (upper–tail), and H = 1.5 h (beyond–data stress test).

standard mechanism for constructing calibrated uncertainty margins in regression [4–6]. However, battery aging induces temporal dependence and distribution shift across cycles and batteries, and because the controller is evaluated in closed loop under constraints, The buffer is a practical robustness heuristic: it reduces the probability of optimistic–forecast driven blackouts but does not provide a formal calibration guarantee under battery aging and distribution shift. Budget satisfaction is an empirical property of the full controller–policy loop, not a consequence of the buffer construction alone (Sections VII–IX).

Interface to the controller

All dimming policies in Section VII operate on $\hat{y}_t^{\mathrm{safe}}$ (not $\hat{y}_t$). This keeps the control logic unchanged while making it more conservative under uncertainty. The effect of α on the energy reliability tradeoff is evaluated in Section X.F.

Dimming Policies

Policy set and notation

All results use the proposed pipeline (grouped battery split; residual quantile buffering). We evaluate two rule–based baselines and forecast driven policies that convert $\hat{y}_{safe}$ into brightness commands, enabling consistent comparison across controllers [4–6].

Let H be the effective night horizon (hours) and $T = 3600H$ its duration in seconds. Each policy outputs a commanded brightness trajectory $b_t(\tau) \in [0, 100]$ over $\tau \in [0, T]$, subject to phase specific minimum brightness floors (comfort floors). Practical smart lighting deployments commonly use time–of–night segmentation ("part night dimming"), motivating our use of phase–based schedules and baselines [1–3].

Rule based baselines

- Floors only baseline. The luminaire is held at the minimum comfort floors for each phase. This establishes a lower bound "comfort preserving" energy schedule and is useful for diagnosing feasibility under strict floors.
- Part night schedule baseline (100/50/70). A fixed three stage schedule is applied independent of battery state (a standard style of part night dimming), providing a realistic non–ML comparator [1–3].

Forecast driven brightness mapping

All ML policies use the same mapping from conservative endurance to brightness. Given $\hat{y}_t^{\mathrm{safe}}$, we compute an "energy headroom ratio" $\hat{y}_t^{\mathrm{safe}}/T$ and translate it to a scalar brightness set point, then enforce floors and actuator bounds:

$$b_t = \mathrm{clip}\left(100 \cdot s_t \cdot \frac{\max\left(0, \hat{y}_t^{\mathrm{safe}}\right)}{T},\ \mathrm{floors},\ 100\right)$$

where s_t is a safety factor (policy–dependent) and "clip(·, floors, 100)" applies phase floors and saturates at 100%. The residual quantile buffer used to produce $\hat{y}_t^{\mathrm{safe}}$ follows Section VI and is conformal–inspired [4–6].

ML fixed safety policy

ML fixed uses a train–selected constant safety factor s. For each (H, B) condition, s is chosen once offline using training batteries only (via out–of–fold predictions) and then held fixed for all test batteries/cycles (no test labels are used for tuning). This yields a simple, interpretable forecast–to–action policy that is easy to deploy and serves as a strong baseline for budgeted control.

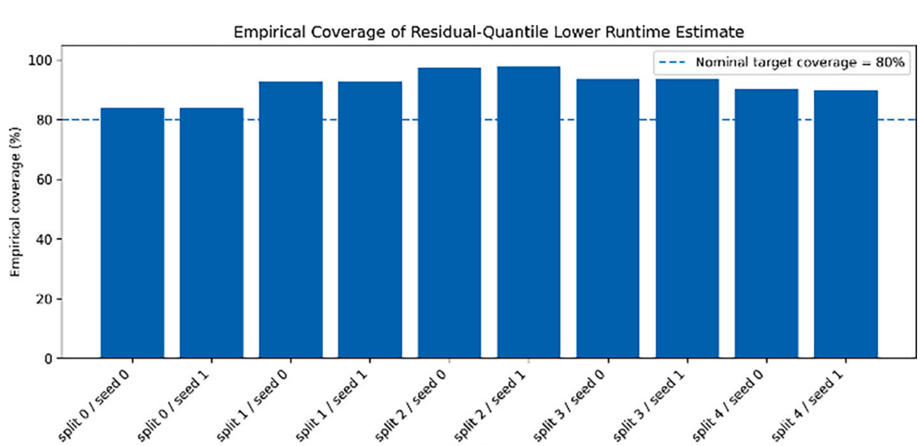

Fig. 2. Per-run empirical coverage of the residual-quantile buffer (α = 0.20) across 10 repeated grouped split/model-seed combinations. All runs exceed the 80% nominal target (dashed line). Pooled coverage: 93.03%; mean: 91.61% [95% CI: 88.69%, 94.53%].

ML block fixed policy

ML block fixed (L = 30) updates the safety factor only every L cycles (block length). Within each block, a constant s is applied; at block boundaries, s is re–selected (using the same tuning rule) to reduce sensitivity to transient shocks while remaining more adaptive than a fully fixed strategy.

ML online adaptive policy

ML adaptive updates s_t online to track a target blackout rate budget B%. Let $\widehat{\mathrm{BR}}_t$ denote a smoothed blackout rate estimate computed over a recent window (block) of cycles. The controller uses a multiplicative feedback update:

$$s_{t+1} = \mathrm{clip}(s_t \exp(\kappa(\widehat{\mathrm{BR}}_t - B)),\ s_{\min},\ s_{\max})$$

so s_t increases when observed blackout exceeds budget (becoming more conservative) and decreases when the system is comfortably below budget (allowing higher brightness). This controller is designed to be deployable (no hindsight) and to adapt to changing difficulty across batteries and aging. The paper's focus is the evaluation framework rather than controller optimization. The gain κ governs adaptation speed: higher κ produces faster correction but risks oscillation when the blackout rate estimate is noisy; lower κ yields smoother trajectories at the cost of slower recovery. The clip bounds [s_min, s_max] constrain divergence. Across the κ values tested, the qualitative ES–BR tradeoff was consistent, though the exact operating point shifted.

Oracle upper bounds and emergency floor relaxation

For context, we report oracle fixed and block fixed policies that select safety factors retrospectively using realized outcomes to meet the target budget while maximizing brightness. These are not deployable controllers; they are performance ceilings that use hindsight to establish how close deployable policies come to an upper bound. Floor relaxation is reported as an offline diagnostic only. Evaluated policies are ranked directly on BR, ES, and UM|BO p90, independently of the oracle.

Feasibility under comfort floors

Under comfort floors, some blackout rate targets may be physically infeasible regardless of controller design (see Section III for the BR_min (H) definition). Table 3 reports the bound estimated on held out test batteries using the floor–only baseline.

Remark: The NASA discharge cycles are laboratory discharge–to–cutoff runs (Fig. 1) [10,11], so H is used as an effective demand horizon parameter rather than a geographic night length. We therefore label $H = 0.75$ h as typical (dataset–scale), $H = 1.0$ h as upper–tail/long, and $H = 1.5$ h as a beyond–data stress test.

Experimental setup

We evaluate effective night horizons $H \in \{0.75, 1.0, 1.5\}$ hours and target blackout budgets $B \in \{2\%, 5\%, 10\%\}$ across held out test batteries [10,11,14,15,18,19]. For each (H, B) setting, if the target is infeasible under strict comfort floors (per Section VII), an emergency floor relaxation factor floor_scale is computed offline via bisection to identify the minimum comfort scaling required for the Oracle–Fixed upper bound to meet the budget. The same floor_scale is then applied across policies to compare behavior under identical comfort relaxation.

Results

This section reports end–to–end outcomes of the proposed forecast–buffer–control pipeline on held out test batteries under a battery ID split (Fig. 1). We evaluate performance under effective night horizons H ∈ {0.75, 1.0, 1.5} h and blackout budgets B (reported per figure/table), using the metrics in Table 2: blackout rate (BR), energy saving (ES), and UM|BO p90. Unless stated otherwise, the main results use the residual quantile buffer with α = 0.20 and compare deployable policies (ML fixed, ML adaptive) against oracle ceilings/diagnostics for context only. All metrics for model–driven policies are reported as mean [95% CI] over n = 10 repeated runs (five battery–split seeds × two RF model seeds); baselines, which do not depend on model initialization, are reported over n = 5 split seeds. Feasibility–bound and severity findings are expected to generalize; specific BR and ES values are dataset–specific (see Section XII for scope).

Feasibility dominates under strict comfort floors

A key result is that minimum brightness floors impose a feasibility limit on the blackout rate that cannot be overcome by controller tuning. Table 3 reports the minimum achievable blackout rate under nominal comfort floors on held out test batteries. This feasibility bound increases sharply with horizon, consistent with the runtime distribution in Fig. 1 (most observed discharge runtimes lie near the 0.75–1.0 h range; H = 1.5 h exceeds all observed runtimes and is a beyond–data stress test, not an in–distribution result). Practically, this means any target budget $B < BR_{\min}(H)$ is infeasible under nominal floors, so violations in that regime should be interpreted as budget impossibility, not algorithmic failure.

Strict–floor outcomes: budgets can be unattainable even when the system is conservative

Under strict floors, the longer–horizon regimes saturate near the feasibility bound. For $H = 1.0$ h, both deployable policies produce blackout rates close to the 11% –13% range across budgets, aligning with mean BR_min(1.0 h) = 6.06% [2.07, 10.04]. Under the beyond-data stress test (H = 1.5 h), BR saturates near the feasibility bound of mean BR_min(1.5 h) = 26.47% [18.03, 34.91]; this result confirms the bound's tightness in the stress regime but should not be interpreted as an in-distribution finding. These results reinforce the role of Table 3 as a diagnostic: once H is sufficiently demanding, comfort floors (not forecasting method choice) set the attainable reliability region. For the shorter horizon $H = 0.75$ h (typical dataset–scale horizon), budgets are achievable in a looser regime, but tight BR budgets remain difficult to track out of sample: BR improves with increasing budget, but the deployable policies do not reliably operate near the $y = x$ budget line. This behavior is consistent with two structural properties of the system: (i) the closed loop mapping from forecast error to blackout is nonlinear near the endurance margin, and (ii) the train–selected safety factor is intentionally chosen without test tuning.

Emergency–mode feasibility analysis via minimum floor scaling (offline diagnostic)

When budgets are infeasible under strict floors, we compute the minimum comfort scaling factor floor_scale required for the oracle–fixed diagnostic to meet the budget (floor_scale = 1.0 = nominal floors;

Table 2
NASA battery dataset summary.

Aspect	Details
Number of batteries	34 (NASA dataset collection)
Total cycles (after preprocessing)	7,565
Cycle types in original dataset	Charge, Discharge, Impedance
Cycle types retained	Charge, Discharge
Cycle types dropped	Impedance
Features before engineering	10
Features after engineering	31
Batteries used for runtime prediction	All available batteries

Table 3
$BR_{min}(H)$values mean [95% CI] across 5 repeated splits.

H	Minimum blackout rate under comfort (%)	Scenario label
0.75	0.86% [0.00, 3.26]	Typical(dataset–scale) night
1.0	6.06% [2.07, 10.04]	Long (upper tail) night
1.5	26.47% [18.03, 34.91]	Stress test (beyond observed runtimes)

smaller values = stronger reduction).

Fig. 3 summarizes the minimum floor_scale required for feasibility under relaxed_oracle. Two consistent patterns emerge:

1. More demanding horizons require larger comfort sacrifice.
2. At $H = 1.0$ h, feasibility is recovered with modest relaxation (approximately floor_scale ≈ 0.87 for $B = 1\%$, rising toward ≈ 0.96 for $B = 10\%$).
3. At $H = 1.5$ h, feasibility requires substantially stronger relaxation (approximately floor_scale ≈ 0.58 at $B = 1\%$, rising toward ≈ 0.64 at $B = 10\%$).
4. As budgets loosen, required relaxation decreases.
5. floor_scale increases monotonically with B, because a larger allowable blackout budget demands less comfort floor relaxation to become feasible.

This analysis is an offline diagnostic: it separates controller limitations from feasibility constraints. Deployment analogues would require forecast– or feedback driven floor adjustments rather than hindsight.

Energy reliability tradeoffs under survivability floors (relaxed_oracle)

Given the feasibility normalized setting of relaxed_oracle (Fig. 3), we compare policies under identical relaxed floors and report the ES–BR tradeoff (Fig. 4).

The main text reports results for H = 1.5 h (Figs. 4–5); the corresponding H = 1.0 h results are provided in Appendix Figs. C1–C2. Across horizons the oracle curve is shown as a performance ceiling, it uses hindsight and is not deployable, but it provides a useful upper bound confirming how close the best deployable policies come to the optimum. The ranking of deployable policies (ML fixed, ML adaptive, baselines) is determined directly from their BR, ES, and UM|BO p90 values, independently of the oracle. Within deployable policies, the following trends are consistent:

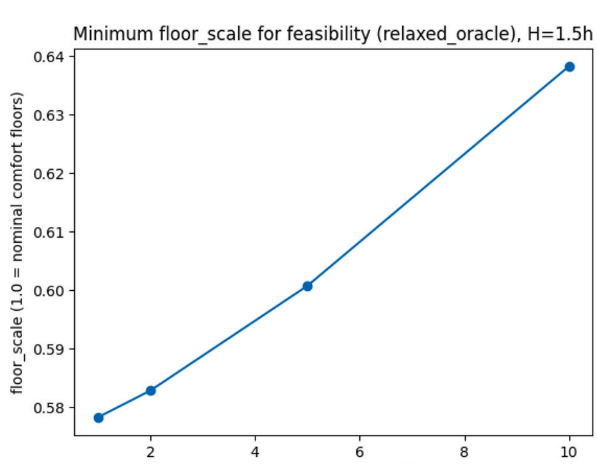

Fig. 3. Feasibility bound under nominal comfort floors: $BR_{min}(H)$ versus horizon H (held out test batteries).

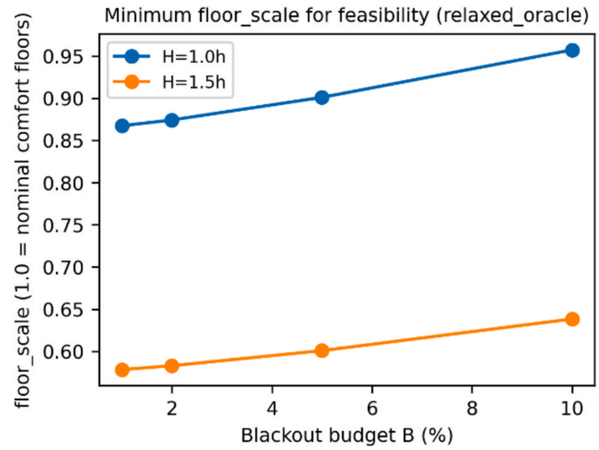

Fig. 4. Offline survivability diagnostic (relaxed_oracle): minimum comfort floor scaling (floor_scale) required for feasibility versus blackout budget B (shown for $H = 1.0$ h and $H = 1.5$ h).

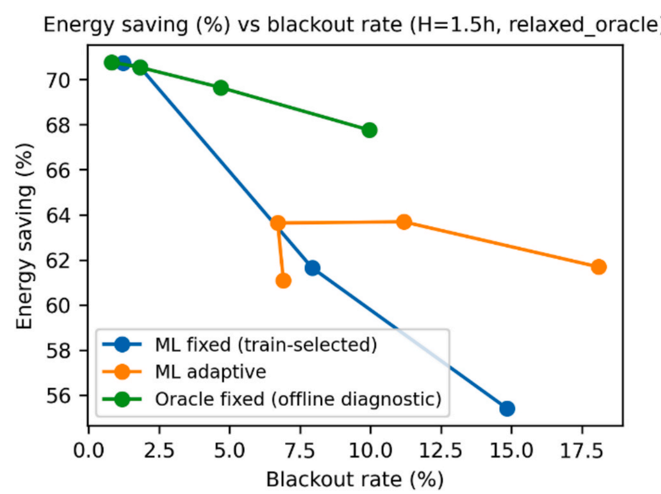

Fig. 5. Energy saving versus blackout rate under survivability floors (relaxed_oracle).

- **ML fixed** achieves strong energy savings at low–to–moderate blackout rates, indicating that a single train–chosen safety factor can translate conservative endurance forecasts into substantial energy reduction while maintaining survivability floors.
- **ML adaptive** often improves the energy–reliability operating point relative to static schedules in several settings, but it does not consistently track strict BR targets and can drift to higher BR at larger budgets, reflecting the empirical nature of closed loop budget targeting under distribution shift.

For the stressed horizon $H = 1.5$ h, the ES–BR frontier becomes particularly informative: at very low BR (near 1–2%), high energy savings remain attainable under survivability floors, whereas moving toward higher BR corresponds to diminished ES and/or increased outage risk. This supports that feasibility analysis with tradeoff reporting is more informative than reporting BR alone (Fig. 6).

Blackout severity reveals tail risk differences invisible to BR alone

Fig. 5 reports UM|BO p90 versus blackout rate under relaxed_oracle. The results show that frequency and severity can decouple: Policies with comparable BR can exhibit markedly different UM|BO p90: some show higher BR but lower UM|BO p90, indicating more frequent but shorter

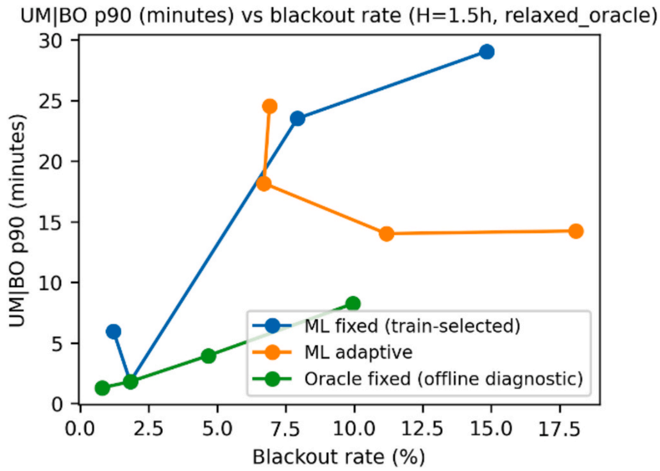

Fig. 6. Blackout severity versus blackout rate under survivability floors (relaxed_oracle), $H = 1.5$ h.

interruptions versus rarer but longer outages.

This validates severity aware reporting as a complementary reliability lens: BR alone can understate practical impact, while UM|BO p90 exposes tail risk differences that matter for public lighting acceptability and safety.

Robustness to buffer quantile and buffer ablation (Appendix)

- **Buffer ablation (n = 10 paired runs).**

Enabling α = 0.20 reduced mean BR at H = 0.75 h from 3.27% to 2.90% at B = 2% (ΔBR = − 0.37 pp, 95% CI [-1.05, +0.31]) and from 3.74% to 3.00% at B = 5% (ΔBR = − 0.75 pp, 95% CI [−1.89, +0.40]). At H = 1.0 h the reduction was smaller (ΔBR = − 0.11 pp at B = 2%); at H = 1.5 h the effect was zero (feasibility limit binding). All ΔBR confidence intervals include zero (paired t–test $p > 0.17$; Wilcoxon $p \geq 0.25$): the buffer provides a conservative shift in non–saturated regimes but not a guaranteed BR reduction. Budget satisfaction is an empirical property of the full closed loop pipeline.

- **Severity–percentile sensitivity.**

UM|BO p75, p90, and p95 increase monotonically across all 18 evaluated settings. At H = 0.75 h the p90–p95 gap is < 0.15 min (p95 adds no discrimination). At H = 1.5 h, B = 2%, ML adaptive: p75 = 29.17, p90 = 49.09, p95 = 57.27 min. Policy ordering is stable across percentiles except one cell (H = 1.0 h, B = 10%: ML fixed p95 = 28.41 vs. ML adaptive p95 = 29.41 min), which does not change any qualitative conclusion. UM|BO p90 is retained: more tail focused than p75, less sensitive to isolated extremes than p95. Sensitivity to Residual-Buffer Quantile Level (α). A post-hoc sensitivity sweep over α ∈ {0.10, 0.20, 0.30} was performed at H = 0.75 h, B = 5%. Energy saving was effectively unchanged across all values (40.42% at α = 0.10; 40.40% at α = 0.20 and α = 0.30). Blackout rate varied modestly (3.81% at α = 0.10; 4.00% at α = 0.20 and α = 0.30). UM|BO p90 was lowest or effectively tied within rounding at α = 0.20 and α = 0.30 (both 12.55 min), with α = 0.10 marginally higher (12.67 min). This supports α = 0.20 as a reasonable operating choice; the appropriate value for a given deployment should be selected based on the operator's coverage requirements and local battery population. Fig. 7 shows the UM|BO p90 profile across α values.

Discussion

Three findings from this study are expected to generalize beyond the NASA simulation. First, feasibility bounds imposed by comfort floors are a structural property of any budget constrained dimming system: missed blackout targets in infeasible regimes should be interpreted as budget impossibility rather than algorithmic failure, pointing to comfort policy or storage/PV sizing changes rather than further controller tuning [22–24]. Second, severity aware evaluation via UM|BO p90 captures tail risk differences among policies that are invisible to blackout rate alone, and this diagnostic value is independent of the specific dataset [1–3]. Third, the separation of controller–induced blackouts from physically infeasible targets is a structural property of the framework and holds for any system where comfort floors constrain the feasible policy set.

The specific BR and ES values are simulation specific and will shift under field deployment conditions; scope constraints and extension requirements are detailed in Section XII.

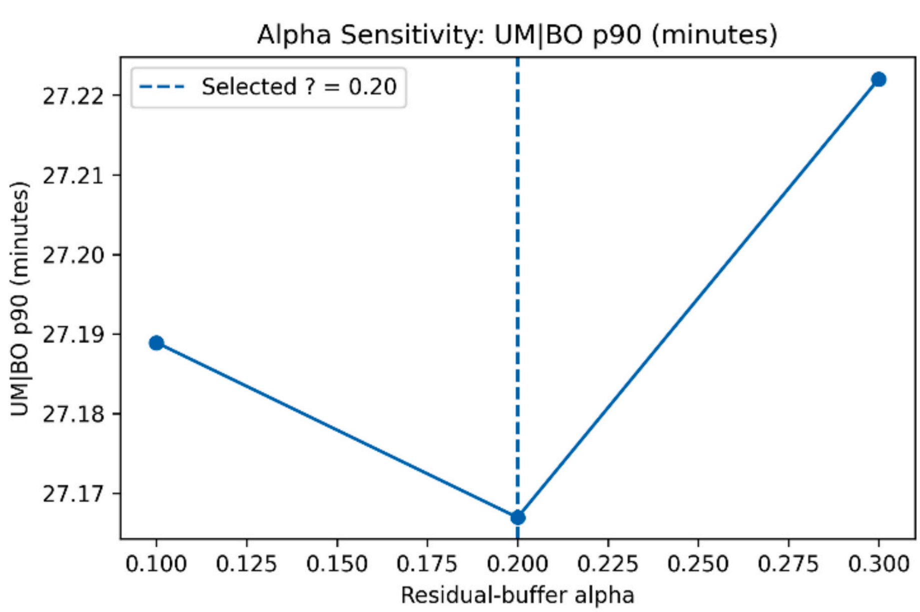

Fig. 7. UM|BO p90 as a function of residual-buffer quantile level α (H = 0.75 h, B = 5%). α = 0.20 (dashed line) yields the lowest observed UM|BO p90; α = 0.30 is effectively tied within rounding (both 12.55 min).

Limitations

Off-grid solar street lighting is widely deployed in energy-scarce communities where grid access is unavailable or unreliable. By providing a principled framework for diagnosing physically infeasible blackout targets and quantifying tail-risk severity, this work offers a transparent basis for evidence-based infrastructure planning and controller benchmarking in resource-constrained settings. The framework is model-agnostic and computationally lightweight, supporting accessibility for practitioners without specialized hardware. Open publication of the methodology and evaluation code supports reproducibility. This study uses publicly available laboratory battery data (NASA PCoE dataset) and involves no human subjects. Potential deployment risks include unequal service quality for batteries with accelerated aging or installations in harsher thermal environments than those represented in the benchmark; these are noted as motivations for field validation.

Several constraints define the scope and generalisability of this study. First, the analysis is focused on the battery-runtime and control-feasibility layer of off-grid solar street lighting. PV generation, charging dynamics, seasonal irradiance variation, and photometric lux mapping are not modelled here; they are deployment-specific extensions of the present simulation benchmark. Second, the results are based on NASA laboratory discharge-cycle data, covering 34 batteries and 2,794 discharge cycles after preprocessing. These controlled laboratory profiles do not include PV charging, seasonal effects, ambient temperature variation, or lamp-load variability. As a result, the absolute BR and ES values should be expected to change under field conditions, although the feasibility-bound and severity trade-off structure may still transfer [10,11]. The NASA data should therefore be viewed as a controlled benchmark proxy for battery ageing behaviour, not as field-night discharge data. The framework is intended as a design and analysis tool for benchmarking, rather than evidence of field-ready deployment performance.

Third, each split uses 9 held-out batteries, corresponding to a 25% held-out fraction. This is suitable for benchmarking within the evaluated dataset, but it still limits broader generalisation; larger and more diverse battery datasets would strengthen the evidence. Fourth, brightness percentage is used as a proxy for illuminance. In practice, the relationship between dimming percentage and maintained road-surface lux depends on luminaire photometrics, installation geometry, ageing, and environmental conditions. Fifth, the phase-based dimming template is configurable, but deployment-specific choices of phase durations and minimum brightness floors are outside the scope of this study. Sixth, the oracle variants are included only as hindsight diagnostic ceilings. They should not be interpreted as deployable controllers, and the adaptive controller is evaluated under baseline tuning rather than through an exhaustive control-design study [20,21]. Finally, the residual-quantile buffer is used as an empirical robustness mechanism, not as a formal distribution-free guarantee under battery ageing or distribution drift. Coverage under deployment shift remains an open question [4–6,14–17]. Future work will include field trials, PV generation forecasting, multi-node lighting networks, and formal uncertainty analysis under non-stationary battery degradation.

Conclusion

This paper presented an uncertainty aware budget targeting dimming framework for off grid solar street lighting that links cycle level battery runtime forecasting to constrained brightness schedules. Two findings are central. First, feasibility bounds under comfort floors, derived analytically from the floors only policy, separate physically infeasible blackout targets from targets that are achievable but not met by the evaluated controllers, providing a diagnostic that pure blackout rate reporting cannot offer. Second, reporting UM|BO p90 alongside blackout rate and energy saving exposes tail risk differences among policies that are invisible to frequency only metrics, enabling more informative benchmarking of survivability–oriented control under uncertainty. All results are simulation–based benchmarks evaluated on NASA laboratory battery discharge data; absolute BR and ES values are dataset–specific and field validation under realistic PV charging, seasonal irradiance variation, and deployment specific loads remains future work. The framework is intended as a design and analysis tool for off grid lighting system benchmarking, not as field–deployment evidence.

CRediT authorship contribution statement

Salaheddine Chouikh: Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Nouh Elmitwally:** Validation, Supervision, Methodology, Conceptualization. **Faisal Saeed:** Resources, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The author would like to thank Khalid Cherifi for early-stage discussions that provided practical context on real-world street lighting systems and helped inform the problem framing.

Appendix

..

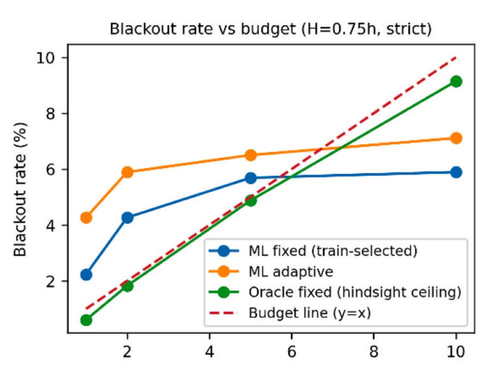

Fig. A1. Blackout rate versus blackout budget (H = 0.75 h, strict floors); dashed line is .

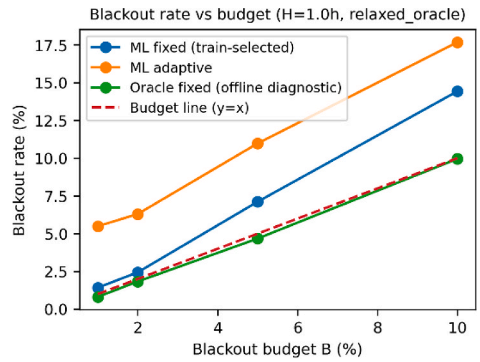

Fig. A2. Blackout rate versus blackout budget (H = 1.0 h, relaxed_oracle survivability floors); dashed line is .

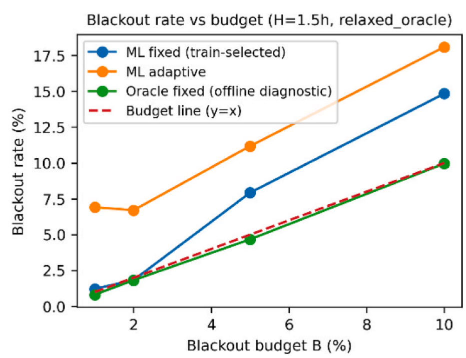

Fig. A3. Blackout rate versus blackout budget (H = 1.5 h, relaxed_oracle survivability floors); dashed line is .

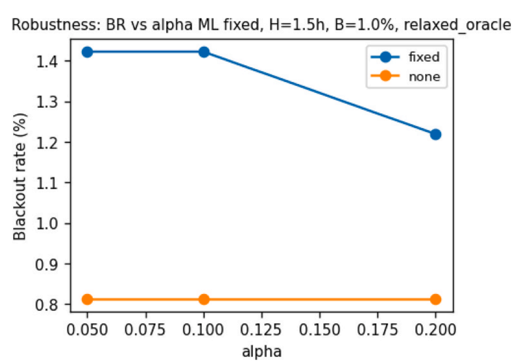

Fig. B. Sensitivity of blackout rate to residual quantile level and buffer ablation (H = 1.5 h, B = 1%, relaxed_oracle; ML fixed).

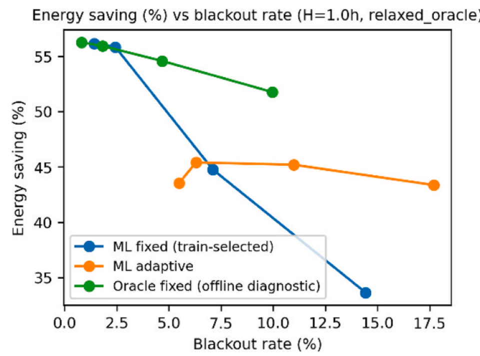

Fig. C1. Energy saving versus blackout rate under survivability floors (relaxed_oracle), H = 1.0 h.

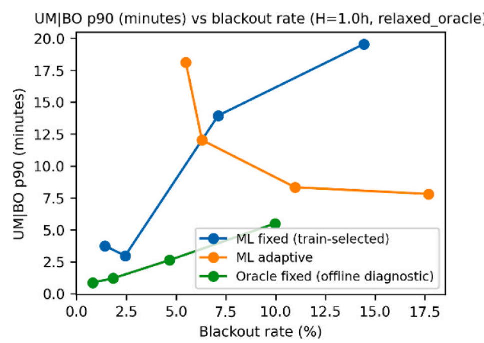

Fig. C2. Blackout severity (UM|BO p90) versus blackout rate under survivability floors (relaxed_oracle), H = 1.0 h.

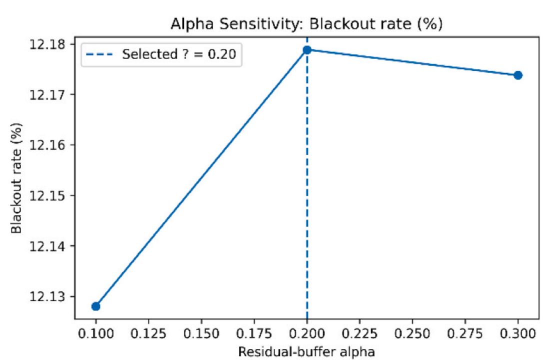

Fig. D. Blackout rate as a function of α (H = 0.75 h, B = 5%). BR variation across the sweep is modest (3.81%–4.00%). α = 0.20 shown as dashed line.

Table 4
Cycle-level feature set used for RF discharge runtime forecasting.

Feature name	Description	Unit	Feature family
avg_voltage	Mean discharge voltage over the cycle	V	Voltage
voltage_drop	Start-to-end voltage decrease over the cycle	V	Voltage
voltage_gradient	Estimated voltage change rate over the cycle	V/s	Voltage
voltage_std	Standard deviation of discharge voltage over the cycle	V	Voltage
start_voltage	Voltage at the start of the cycle	V	Voltage
end_voltage	Voltage at the end of the cycle	V	Voltage
avg_current	Mean discharge current over the cycle	A	Current
max_current	Maximum observed current magnitude over the cycle	A	Current
current_std	Standard deviation of discharge current over the cycle	A	Current
initial_temp	Initial measured temperature for the cycle	°C	Temperature

(continued on next page)

Table 4 (*continued*)

Feature name	Description	Unit	Feature family
final_temp	Final measured temperature for the cycle	°C	Temperature
temp_rise	Final temperature minus initial temperature	°C	Temperature
energy_approx	Approximate cycle energy derived from voltage, current, and duration, using the Riemann-sum proxy $\Sigma_k V_k I_k \Delta t$	approx. W·s	Energy

Table 5
Model and framework hyperparameter/configuration summary.

Component	Parameter	Value
Run config	Mode	review_fast
Run config	Runs	5 split seeds × 2 model seeds = 10 runs
RF grid	n_estimators	{400, 800}
RF grid	max_depth	{None, 20}
RF grid	max_features	{sqrt, 0.5}
RF grid	min_samples_leaf	{1, 2, 4}
RF grid	Tuning objective	MAE, GroupKFold CV
RF grid	Group CV splits	5
Buffer – main config	α values, ROBUST_ALPHAS	{0.05, 0.10, 0.20}
Buffer – main config	α selected for primary results	0.20, corresponding to 80% nominal target
Buffer – main config	Buffer modes	{none, fixed}
Buffer – main config	Conservative runtime estimate	ŷ_safe = max(0, ŷ − q_α)
Buffer – sensitivity sweep	α values swept, post-hoc only	{0.10, 0.20, 0.30}
Grouped splits	Split seeds	{0, 1, 2, 3, 4}
Grouped splits	Model seeds	{0, 1}
Grouped splits	Test battery fraction	0.25
Dimming controller	Horizons H	{0.75, 1.0, 1.5} h
Dimming controller	Blackout budgets B	{2%, 5%, 10%}
Phase floors	Active / Quiet / Pre-dawn	70% / 30% / 50%
Phase split	Active / Quiet / Pre-dawn	40% / 40% / 20%
Adaptive controller	Update rule	s_{t + 1} = clip(s_t · exp(κ(BR_t − B)), s_min, s_max)
Adaptive controller	κ, adaptive update gain	0.02
Adaptive controller	s_min / s_max	0.2 / 1.2
Adaptive controller	Window length	200 cycles
Adaptive controller	Initial safety factor	0.85
Adaptive controller	Hysteresis	5.0
Adaptive controller	Ramp limit	10.0
Outlier filter	Method	Tukey inner fence, 1.5 × IQR
Outlier filter	Applied to test set	Yes
Outlier filter	Bounds fitted on	Training set only

Data availability

The code and experiment configuration are available at: https://github.com/Galuyoo/offgrid-streetlight-dimming-control. The NASA battery dataset is not redistributed; instructions for obtaining and preprocessing it are included.

References

[1] Shahzad G, et al. Energy–efficient intelligent street lighting system using traffic–adaptive control. IEEE Sens J 2016.
[2] Putrada AG, et al. Machine learning methods in smart lighting toward achieving user comfort: a survey. IEEE Access 2022.
[3] Yang YS, et al. An implementation of high efficient smart street light management system for smart city. IEEE Access 2020.
[4] Shafer G, Vovk V. A tutorial on conformal prediction. JMLR 2008.
[5] Angelopoulos AN, and S. Bates, "A gentle introduction to conformal prediction…," arXiv, 2021.
[6] Romano Y, et al., "Conformalized quantile regression," arXiv, 2019.
[7] Prokhorenkova L, et al., "CatBoost…," NeurIPS, 2018.
[8] Breiman L. Random forests. Mach Learn 2001.
[9] Zhang J, Lee J. A review on prognostics and health monitoring of Li–ion battery. J Power Sources 2011.
[10] NASA Ames PCoE, "Prognostics data repository (battery data sets)," 2007.
[11] Fleith P. NASA battery dataset (Kaggle distribution). Kaggle 2021.
[12] Severson KA, et al. Data driven prediction of battery cycle life before capacity degradation. Nat Energy 2019;4(5):383–91. https://doi.org/10.1038/s41560-019-0356-8.
[13] Richardson RR, Osborne MA, Howey DA. Gaussian process regression for forecasting battery state of health. J Power Sources 2017;357:209–19. https://doi.org/10.1016/j.jpowsour.2017.05.004.
[14] Tashman LJ. Out of sample tests of forecasting accuracy: an analysis and review. Int J Forecast 2000;16(4):437–50. https://doi.org/10.1016/S0169-2070(00)00065-0.
[15] Cerqueira V, Torgo L, Mozetič I. Evaluating time series forecasting models: an empirical study on performance estimation methods. Mach Learn 2020;109:1997–2028. https://doi.org/10.1007/s10994-020-05910-7.
[16] Tukey JW. Exploratory data analysis. Addison-Wesley; 1977.
[17] NIST/SEMATECH. (2024). e–Handbook of Statistical Methods: Outliers. https://www.itl.nist.gov/div898/handbook/.
[18] Hyndman RJ, Athanasopoulos G. Forecasting: Principles and Practice (3rd ed.). OTexts. https://otexts.com/fpp3/.
[19] Bergmeir C, Benítez JM. On the use of cross–validation for time series predictor evaluation. Inf Sci 2012;191:192–213. https://doi.org/10.1016/j.ins.2011.12.028.
[20] DALI Alliance, "Technical guidance: IEC 62386 (DALI) fade time and fade rate," https://www.dali–alliance.org.
[21] Sofrony J, Turner MC, Postlethwaite I. Anti–windup synthesis for systems with rate limits using Riccati equations. Int J Control 2010;83(2):233–45. https://doi.org/10.1080/00207170903137026.
[22] Panguloori RB, Mishra PR, Kumar S. Power distribution architectures to improve system efficiency of centralized medium scale PV street lighting system. Sol Energy 2013. https://www.sciencedirect.com/science/article/pii/S0038092X13003447.
[23] Panguloori RB, Mishra PR (2014). Analysis on system sizing and secondary benefits of centralized PV street lighting system. IEEE Conference Publication. https://ieeexplore.ieee.org/abstract/document/6805322/.
[24] Bouroussis CA, Topalis FV (2004). Optimization of potential and autonomy of a photovoltaic system for street lighting. (PDF) http://147.102.23.61/files/Publications/Journals/Optimization%20of%20Potential%20and%20Autonomy%20of%20a%20Photovoltaic%20System%20for%20Street%20Lighting.pdf.
[25] Lau SP, Weddell AS, Merrett GV, White NM, "Energy–neutral solar–powered street lighting with predictive and adaptive behaviour," in Proc. 2nd Int. Workshop Energy Neutral Sensing Systems (ENSSys '14), Memphis, TN, USA, Nov. 2014, pp. 13–18. DOI: 10.1145/2675683.2675690.

[26] Ramadhani F, Bakar KA, Shafer MG, "Optimization of standalone street light system with consideration of lighting control," in Proc. Int. Conf. Technological Advances in Electrical, Electronics and Computer Engineering (TAEECE), Konya, Turkey, May 2013, pp. 583–588. DOI: 10.1109/TAEECE.2013.6557340.
[27] Hsia, Lin, Wu. A high–efficiency dimmable LED driver for low–power lighting applications. IEEE Trans Ind Electron 2013.
[28] Long, Ozpineci, Khwan–on "Analysis and design of a high power factor, single–stage LED driver," IEEE Trans Power Electronics, 2009.
[29] Lau, Luk, Lam. LED luminaire with improved dimming performance for street lighting applications. IET Power Electron 2015.