Designing a Real-World Telemetry Study of E-Rickshaw Batteries in Odisha: Duty Cycle, Thermal Exposure and Degradation from a Simulation-Calibrated Pre-Study

Logger specification, capacity observability and sample size for Bhubaneswar and Cuttack

[1]Sonia Mohapatra, [2]Saswat Swain, [3]Chandrasekhar Panda
[1]Researcher, [2]Ph.D Scholar, [3]Researcher
[1,3]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[2]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[1]smctc98@gmail.com, [2]saswat.mck@gmail.com, [3]contact@chandrasekharpanda.com

***Abstract*—E-rickshaws dominate short urban trips in Odisha, yet the public evidence on how their batteries are used and how fast they wear out rests on a 2019 energy test and on driver surveys from Delhi and Bihar that disagree with each other. This paper designs an instrumented telemetry study for Bhubaneswar and Cuttack and, before any hardware is deployed, tests that design with a coupled duty-cycle, electro-thermal and degradation simulation of a 24-vehicle fleet over 730 days. The simulator is calibrated to published field anchors and its outputs are labelled as synthetic throughout. It reproduces published energy use (median 49.1 Wh/km for lead-acid) and charging time, and attributes the conflicting range reports to pack size. Within the model, lead-acid packs run hottest during overnight absorption charging rather than on the road; stand-side top-up charging raises the distance delivered per pack by about 8 % without shortening its life in days; and lithium iron phosphate fade (7.4 % in year one) is dominated by calendar ageing at full charge overnight, which a 90 % charge ceiling or charge timer reduces. For instrumentation, the study shows that point sampling at 30 s gives a 95th-percentile throughput error of 5.2 %, that hardware charge integration with a recorded mean of current squared is needed, that capacity anchored on every pre-charge rest outperforms scheduled rest tests, and that rest-voltage capacity estimation is weak for lithium iron phosphate packs seldom discharged below 25 %. The planned fleet can track individual fade within three to six months but is under-powered for group comparisons; design changes are proposed. Code, synthetic data and logs are released.**

***Index Terms*—Battery degradation, duty cycle, e-rickshaw, field telemetry, lead-acid battery, LiFePO4, state of health, IoT data logger, electric three-wheeler, India.**

I. Introduction

Battery-operated e-rickshaws now carry a large share of short urban trips in eastern India. They are cheap to buy, cheap to run and, in cities such as Bhubaneswar and Cuttack, they have taken over feeder routes once served by cycle-rickshaws and diesel auto-rickshaws. The battery is the component that decides whether the vehicle remains economic. For a flooded lead-acid (PbA) pack the owner pays for a replacement roughly once a year; a lithium iron phosphate (LFP) pack costs more up front but is expected to last several years. Which of these is the better choice for a driver, a fleet financier or a state subsidy scheme depends almost entirely on how the battery is actually used and how quickly it wears out under that use.

The public evidence on that question is thin and dated. The most cited technical source on Indian e-rickshaw energy use is an SAE paper from 2019 comparing the on-road energy consumption of PbA and Li-ion vehicles [1]. The most quoted operational figures come from surveys rather than measurements: a 2014 Delhi study reported a mean range of 63 km per charge and a battery replacement interval of about 7.5 months [3], while a 2023 survey in Bihar found that 70 % of drivers charge overnight at unauthorised private stations, that a full charge takes 8–12 h, and that drivers cover 80–100 km per day [2]. These two field reports do not agree with each other on usable range, neither reports current, voltage or temperature, and neither was carried out in Odisha, whose coastal climate imposes a hot, humid pre-monsoon season on batteries that sit in unventilated under-seat trays.

Odisha's Electric Vehicle Policy 2021 set a target of 20 % electric share in new registrations and offers purchase incentives for three-wheelers [10]. Policy of that kind is calibrated against assumed duty cycles and assumed battery lives. Without a measured duty-cycle dataset, those assumptions cannot be checked, and battery-life claims made by vendors cannot be tested against local operating conditions.

This paper addresses the gap by designing an instrumented field study for Bhubaneswar and Cuttack and by carrying out the analysis that should precede any such deployment: a physics-based, simulation-calibrated pre-study that establishes what the logger must measure, how fast it must sample, how long it must run and how many vehicles it must cover before the resulting degradation estimates mean anything. The distinction matters. A field campaign that samples too slowly, cannot observe capacity, or covers too few vehicles will produce a dataset that looks authoritative and is statistically empty.

Scope statement. No telemetry has yet been collected. Every quantitative result in Sections V–VIII is produced by the simulation described in Section IV, calibrated to the published

anchors cited above. These results are design inputs and testable predictions, not field observations, and they are labelled as such throughout. The contributions are:

- A logger architecture and measurement protocol for e-rickshaw traction batteries, specified for both PbA and LFP packs and for the Odisha operating environment (Section IV-A).
- A coupled duty-cycle, electro-thermal and ageing simulator for a 24-vehicle, 730-day synthetic fleet, released with a dataset schema identical to the planned logger output (Sections IV-B to IV-F).
- A reconciliation of the conflicting published range figures through vehicle physics, and a ranked attribution of ageing stressors for each chemistry (Sections V–VII).
- Quantitative logger design rules: sampling interval, rest-anchor protocol for state-of-health (SoH) estimation, campaign duration and fleet sample size, derived by Monte-Carlo analysis (Section VIII).
- An explicit account of what the pre-study cannot establish, including a calibration step that must not be mistaken for validation (Section IX).

II. Related Work

Table 1 summarises the field evidence on e-rickshaw batteries in India together with the modelling literature this study draws on. Three features stand out. First, the Indian field studies are either energy-only [1] or questionnaire-based [2], [3]; none logs battery current, voltage and temperature together over months. Second, the ageing models with the strongest experimental basis were developed for stationary or automotive cells under controlled conditions [4]–[7], and their transfer to a hot, high-throughput, poorly maintained traction duty is untested. Third, the lead-acid lifetime literature itself warns that simple cycle-count and Ah-throughput models are systematically optimistic under partial-state-of-charge operation, which is why weighted-throughput formulations were developed [4], [11], [13].

Table 1 Prior Evidence and Its Limits for the Present Question

Ref.	Type / setting	What it provides	What it lacks here
[1] 2019	On-road test, Indian city	Wh/km, PbA vs Li-ion	No I, V, T logs; no ageing
[2] 2023	Driver survey, Bihar	80–100 km/day; 8–12 h charging; 70 % private overnight charging	Self-reported; no battery data
[3] 2014	Survey + vendor data, Delhi	48 V, 85 Ah; 63 km/charge; ~7.5-month battery life	Dated; self-reported life
[4] 2007	PbA lifetime model	Weighted Ah-throughput ageing	Stationary RE duty
[5], [6] 2018–20	LFP ageing study, lab	Calendar and cycle laws; SoC, DoD, C-rate forms	Small cylindrical cells; 40 °C cycling only
[11] 2021	Model comparison, PV	Shows simple models over-predict PbA life	No traction duty
This work	Field-study design + calibrated simulation	Logger spec, stress attribution, sample size	No measured telemetry (yet)

Table entries summarise the cited sources in the authors' own words.

III. Problem Formulation

Let a pack j in the fleet be described by its measured current $I_j(t)$, terminal voltage $V_j(t)$, pack temperature $T_j(t)$ and vehicle speed $v_j(t)$. The quantities of interest are the capacity fade $Q_j(t) = 1 - C_j(t)/C_{nom}$ and its dependence on stressors such as daily depth of discharge, minimum state of charge (SoC), time at high SoC, temperature and charge regime. A field study must therefore solve three linked problems:

- **Fidelity.** Choose a sampling scheme such that throughput, I^2R heat and cycle depth are recovered within a stated error.
- **Observability.** Construct an estimator of C_j from field signals alone, without removing the pack for a laboratory capacity test, and quantify its error.
- **Power.** Choose campaign duration T_c and fleet size N such that a fade rate, and differences in fade rate between groups, can be detected at $\alpha = 0.05$ with 80 % power.

Practical constraints bound the design: the logger must cost a small fraction of the battery, must not alter driver behaviour, must survive monsoon humidity and road shock, and must work with packs that have no accessible BMS (PbA) as well as packs that do (LFP).

IV. Methodology

A. Logger Architecture and Measurement Protocol

Fig. 1 shows the proposed instrumentation. Pack current is measured with a 200 A/75 mV manganin shunt in the negative leg, read by a 20-bit digital power monitor that integrates charge in hardware. Integration in hardware matters: as Section VIII-A shows, it removes the throughput error that point sampling introduces at slow logging rates. For PbA packs each 12 V block is tapped through a divider to a 16-bit ADC so that a weak block, the usual cause of early pack failure, can be identified. For LFP packs, cell voltages and BMS flags are read over the BMS serial link where available. Four NTC sensors record pack, hottest-block, compartment and ambient temperature. A GNSS receiver and a six-axis IMU provide speed, road grade and shock exposure. Acquisition runs at 10 Hz for current and voltage and is recorded at 1 Hz; five-minute aggregates are sent over LTE-M, and the full 1-Hz file is uploaded nightly while the vehicle is parked at its charging stand. Location is stored at full resolution internally but published only after snapping to a 500 m grid, and drivers' identities are replaced by pseudonymous vehicle codes, with written consent obtained before installation. The published dataset follows the FAIR principles [16].

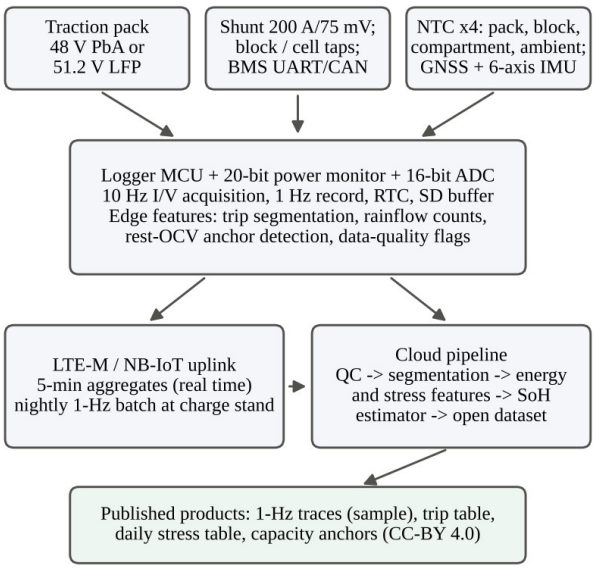


Fig. 1. Proposed logger architecture and data path. Component values are design choices for the field study; none has yet been built or tested.

Capacity cannot be read from a field logger directly; it must be inferred from a charge count between two points of known SoC. Two anchoring strategies were specified and are compared in Section VIII-B. In the *opportunistic* strategy the charger's termination event serves as the full-charge anchor and every end-of-shift rest before plug-in, typically 15–60 min, is used as the low-SoC anchor, subject to a quality gate requiring at least 30 % SoC between anchors. In the *scheduled* strategy drivers are additionally asked, twice a month, to leave the vehicle switched off for 2 h before charging, and only those well-relaxed rests are used. The scheduled strategy was the authors' initial design choice; the analysis below shows it to be the weaker of the two.

The planned deployment is 24 vehicles for 12 months, split evenly across chemistry (PbA, LFP) and city (Bhubaneswar, Cuttack), i.e. six vehicles per cell of the design. Section VIII tests whether this allocation is adequate.

B. Synthetic Fleet and Duty-Cycle Generator

To test the design before building it, a synthetic fleet with the same 2 × 2 structure was simulated for 730 days at 1-s resolution. Each driver has a personal mean daily distance drawn uniformly between 70 and 100 km, bracketing the Bihar survey figure [2]; daily targets scatter around that mean with a standard deviation of 10 km, and each day carries a 12 % probability of being a rest day. A shift starts between 06:30 and 08:00 and ends when the daily target is met, the battery runs low, or 21:30 is reached, whichever comes first. Each working day is built from passenger rides whose lengths are log-normal with a median of 2.8 km, consistent with the short-trip character reported in [1]; a quarter of movements are empty repositioning runs. Occupancy is 1 + Poisson(2.0), capped at six. Each ride is split into stop-to-stop micro-trips whose mean length and cruising speed differ by city: Bhubaneswar's wider arterials allow 18–30 km/h with micro-trips averaging 450 m, while Cuttack's dense old town limits speeds to 14–24 km/h with 300 m micro-trips and longer traffic stops. A lunch break of 45–90 min falls between 13:00 and 13:45.

Half of the PbA drivers are given access to stand-side charging; they top up at lunch if SoC is below 60 % and take a 1.5-h charge whenever SoC falls below a personal threshold between 15 % and 30 %. Drivers without access end their shift at that threshold, and the unmet distance is recorded as lost kilometres. This reproduces the two behaviours that the field reports imply but never measure directly.

C. Longitudinal Vehicle Model

Battery power follows from the road-load equation

$$F = mg(C_{rr}\cos\theta + \sin\theta) + \tfrac{1}{2}\rho C_d A\, v^2 + m(1+\delta)a, \qquad (1)$$

where m is the instantaneous vehicle mass (chassis 230 kg, driver 65 kg, 62 kg per passenger, plus the pack: 144 kg for four 130 Ah PbA blocks or 48 kg for a 105 Ah LFP pack), C_{rr} = 0.015 is rolling resistance on mixed urban surfaces, C_dA = 1.10 m^2 is the drag area of the box-shaped body, ρ = 1.16 kg/m^3, θ is road grade (zero-mean, σ = 0.8 %) and δ = 0.05 accounts for rotating inertia. Mechanical power $P_m = Fv$ is converted to battery power through a loss model

$$P_b = P_m + P_0 + k_q P_m^{\,2} + P_{aux}, \;\; P_m > 0, \qquad (2)$$

with P_0 = 40 W fixed loss, $k_q = 1.8 \times 10^{-4}$ W^{-1} (giving 82 % drive-train efficiency at 1 kW) and P_{aux} = 25 W for lights and accessories. Typical e-rickshaw hub-motor controllers do not regenerate, so negative P_m is dissipated in the brakes. The controller limits battery power to 2.2 kW and wheel force to 450 N; these limits, not the driver, set the acceleration of a loaded vehicle. Current is obtained from $P_b = (U_{oc} - IR)I$, i.e.

$$I = [U_{oc} - (U_{oc}^{\,2} - 4RP_b)^{1/2}] \,/\, 2R, \qquad (3)$$

where U_{oc}(SoC) is the open-circuit voltage and R(SoC, T, Q) the internal resistance, which rises with fade ($\times(1 + 2.5Q)$ for PbA and $\times(1 + 1.5Q)$ for LFP).

D. Battery Electrical and Thermal Models

The PbA pack is 24 flooded tubular cells with U_{oc} = 24(1.95 + 0.17 SoC) V. Rate-dependent capacity is represented by Peukert's law with exponent n = 1.15, so that charge is depleted at an effective rate

$$I_{eff} = I \cdot \max[1, (I/I_{20})^{n-1}], \qquad (4)$$

where $I_{20} = C_{20}/20$. Charging follows the constant-current/absorption behaviour of the 48 V, 15 A chargers common at roadside stands: bulk charge at 92 % Ah efficiency to 80 % SoC, then an exponential approach to full with a 1.8-h time constant and 75 % Ah efficiency, the remainder going to gassing. The LFP pack is 16 cells in series with a tabulated U_{oc}(SoC) whose plateau between 30 % and 80 % SoC varies by only 30 mV per cell; it is charged at 25 A by a BMS-terminated CC–CV charger [12], [14].

Each pack is treated as a lumped thermal mass:

$$m_b c_p \, dT/dt = I^2R + Q_{gas} - hA(T - T_{amb} - \Delta T_{sun}), \qquad (5)$$

with $m_b c_p$ = 130 kJ/K and hA = 2.6 W/K for the open PbA tray, and 50 kJ/K and 1.6 W/K for the sealed LFP box. Q_{gas} is the heat released by overcharge current during absorption. ΔT_{sun} = 2 K is added while driving in daylight to represent a sun-heated compartment. Ambient temperature is built from the WMO monthly normals for Bhubaneswar [9] with an assumed diurnal half-range of 3.5–6 K and a first-order autoregressive day-to-day anomaly (ϕ = 0.8, σ = 1 K).

E. Degradation Models

For PbA a reduced-order weighted-throughput model in the spirit of Schiffer *et al.* [4] is used. Daily capacity loss is the sum of a corrosion term and a throughput term:

$$\Delta Q = k_c 2^{(\bar{T}-25)/10}(1 + \kappa h_g) + k_a(A_{eff}/C_{nom})(1 + \beta\psi)(1 + \gamma n_{nf}), \quad (6)$$

where $\bar{T}$ is the day's mean pack temperature (°C), h_g the hours spent gassing, A_{eff} the Peukert-weighted discharge throughput, $\psi = \max(0, (0.5 - \mathrm{SoC}_{min})/0.5)$ penalises deep discharge, and n_{nf} counts consecutive days without a full charge (capped at 7), a proxy for sulphation and acid stratification. The values $k_c = 1.1 \times 10^{-4}$ day^{-1} (about five years to 20 % fade on float at 25 °C), $\kappa = 0.05$ h^{-1}, $\beta = 1.0$ and $\gamma = 0.1$ are stated assumptions and are varied in Section VII. The remaining constant k_a was calibrated by bisection so that a reference driver (85 km/day target, overnight charging only, Bhubaneswar climate) reaches 80 % capacity after 300 days, the geometric mean of the 7.5-month [3] and roughly 12-month [2] replacement intervals reported in the field. The result is $k_a = 3.73 \times 10^{-4}$ per weighted equivalent cycle.

For LFP the functional dependencies follow Naumann *et al.* [5], [6]: calendar loss grows with the square root of time, with an Arrhenius temperature factor and a cubic SoC factor; cycle loss grows with the square root of equivalent full cycles, with a cubic depth-of-discharge factor and a linear C-rate factor. Because the source cells are small cylindrical cells unlike the prismatic cells in e-rickshaw packs, the magnitudes are normalised to datasheet-type anchors rather than copied: 2.5 % calendar loss after one year at 25 °C and 50 % SoC with $E_a = 50$ kJ/mol, and 20 % cycle loss after 3000 cycles at 1C and full depth. Under time-varying stress both terms are integrated with the virtual-time method,

$$Q_{cal}(t + \Delta t) = k(T, \mathrm{SoC})\,[\,(Q_{cal}(t)/k)^2 + \Delta t\,]^{1/2}, \quad (7)$$

applied hourly, with the analogous update in equivalent-cycle space for cycle ageing. Cycles are extracted from the SoC trace with the ASTM E1049 rainflow algorithm [8]. The source study ran cycle tests at a single temperature, so the cycle term carries no temperature dependence; this omission is tested in Section VII. A pack is replaced when fade reaches 20 %.

F. Experimental Design of the Simulation

The baseline run uses a fixed global seed (20261023), common to all vehicles through seeded sub-streams, and produces 17,520 vehicle-days, of which 15,368 are working days containing 472,206 rides. Sensitivity analysis is one-at-a-time around a single reference vehicle with common random numbers; the spread across five independent seeds of the unperturbed case is reported alongside, so that effects smaller than stochastic noise are not over-read. One week in May and one day in January are stored at 1 Hz for every vehicle for the logger studies in Section VIII. Simulation used Python 3 with NumPy, SciPy and pandas; figures were produced with Matplotlib. Code, data and logs are released with the paper.

V. Model Plausibility Checks

Before the simulator is used to design anything, it should at least reproduce the coarse facts that are already published. These checks are labelled *plausibility checks* rather than validation because the published figures are themselves uncertain survey values, and in one case (battery life) the model was calibrated to the same number it is then compared with.

Energy intensity. The simulated PbA fleet consumes a median of 49.1 Wh/km at the battery terminals (5th–95th percentile 44.3–54.7 Wh/km). The Delhi study's 48 V, 85 Ah pack and 63 km range [3] imply 65 Wh/km if the full nominal 4.1 kWh were usable and about 52 Wh/km at a more realistic 80 % usable fraction. The two figures are compatible. The on-road measurements in [1] could not be used as a numerical check because their per-vehicle values were not accessible to the authors; this is a gap the field study is intended to close.

Charge duration. From the median end-of-shift SoC of 25 %, the bulk phase of the PbA charger lasts about 5 h and the absorption phase a further 6.6 h, giving roughly 11–12 h to termination. This sits inside the 8–12 h reported in Bihar [2] and above the 8.2 h mean reported in Delhi [3], where smaller packs were charged.

Range. The two field reports disagree: 63 km per charge in Delhi [3] against 80–100 km per day in Bihar [2]. The model suggests that the disagreement need not reflect different usage. With a 130 Ah pack, simulated PbA drivers cover a mean of 80 km per working day while their median daily minimum SoC is 25 %; scaling the same energy budget to the 85 Ah packs reported in 2014 gives a range in the low 60s. A larger pack, and in some cases a mid-day top-up, is sufficient to explain the Bihar figure. This is an inference from the model, not a finding, because neither survey reports the pack size of its respondents' vehicles.

Trip structure. Simulated rides average 2.7 km and 12 min with 2.2 passengers on average, consistent with the sub-10 km trip lengths reported in [1]. Table 2 lists the model parameters and their epistemic status.

TABLE 2 Model Parameters and Their Status

Quantity	Value	Status
Daily distance target	U(70, 100) km ± 10 km	Bracketed by [2]
Ride length (median)	2.8 km, log-normal	Assumed; cf. [1]
Vehicle mass (empty, PbA / LFP)	439 / 343 kg	Assumed
C_{rr}; C_dA	0.015; 1.10 m^2	Assumed, varied
Controller limit	2.2 kW; 450 N	Assumed
PbA pack	48 V, 130 Ah, n = 1.15	Assumed, varied
LFP pack	51.2 V, 105 Ah	Assumed
Ambient	WMO normals [9]	Sourced
PbA k_a	3.73×10^{-4}	Calibrated to [2], [3]
PbA k_c, β, γ, κ	1.1×10^{-4}; 1; 0.1; 0.05	Assumed, varied
LFP ageing forms	√t, √FEC; cubic SoC/DoD	Sourced [5], [6]
LFP magnitudes, E_a	2.5 %/yr; 3000 FEC; 50 kJ/mol	Assumed, varied

VI. Duty Cycle and Thermal Exposure

A. Energy and Current

Fig. 2 shows one simulated May day for a PbA and an LFP vehicle in Bhubaneswar, and Fig. 3 summarises the full fleet. Two effects separate the groups in Fig. 3(a). The LFP vehicles use about 14 % less energy per kilometre (median 42.1 against 49.1 Wh/km), almost entirely because the LFP pack is 96 kg lighter; rolling and inertial losses dominate at these speeds, so mass is the lever that matters. Cuttack vehicles use 8–9 % less energy per kilometre than Bhubaneswar vehicles (47.0 against 51.4 Wh/km for PbA), because lower cruising speed cuts aerodynamic and kinetic losses more than additional stops add. The saving is not free: Cuttack drivers spend 7.5 h on the road per day against 4.7 h in Bhubaneswar for similar distance, which lengthens exposure to daytime heat and shortens the evening rest before charging.

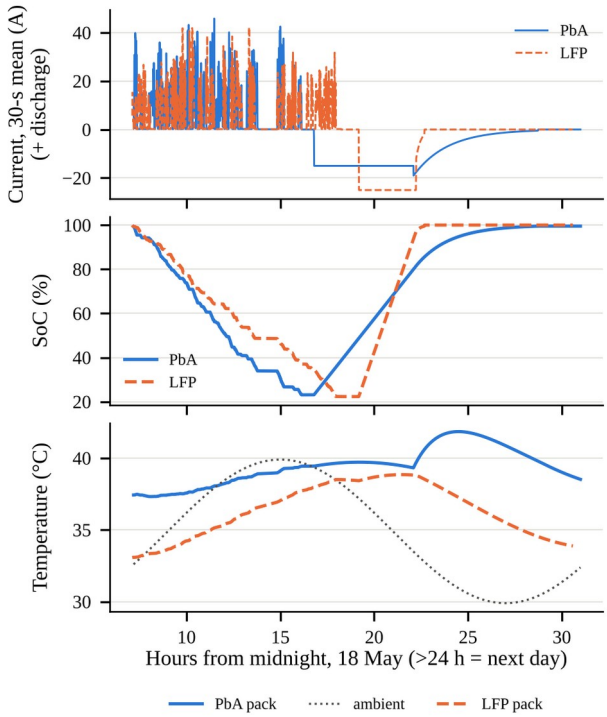


Fig. 2. Simulated 1-Hz traces for one PbA and one LFP vehicle on 18 May (synthetic data). The PbA driver stops at 17:00 on reaching the personal SoC threshold; the remaining target distance is recorded as lost. Note the PbA pack reaching its daily maximum temperature after midnight, during absorption charging.

Current is low on average and limited at the peak. Mean driving current is 13.6 A for PbA (RMS 19.0 A, about 0.15C) and 11.0 A for LFP (RMS 15.5 A). The 99th-percentile current, 47 A and 42 A respectively, sits at the controller limit on almost every day, so peak current carries little information about the driver or the route; it reflects the controller. Daily throughput is 0.74 Peukert-weighted equivalent full cycles for PbA and 0.64 for LFP (Fig. 3(b)). PbA packs are driven deeper (median daily minimum SoC 25 % against 31 %), and drivers without stand-side charging lose 13.6 % of their target distance to range limits, against 3.6 % for those with access and 0.3 % for LFP.

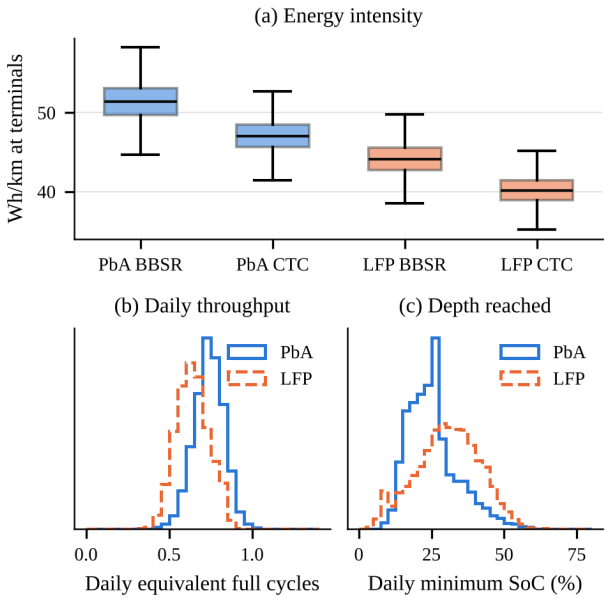


Fig. 3. Fleet duty-cycle statistics over 730 days (synthetic). (a) Energy intensity at the battery terminals by chemistry and city, days over 20 km; (b) daily equivalent full cycles (PbA values are Peukert-weighted); (c) daily minimum SoC.

TABLE 3 SIMULATED DUTY-CYCLE SUMMARY (WORKING DAYS)

Metric	PbA	LFP
Daily distance, mean (km)	79.9	85.2
Energy, median (Wh/km)	49.1	42.1
Drive current, mean / RMS (A)	13.6 / 19.0	11.0 / 15.5
Daily EFC, mean	0.74	0.64
Daily min. SoC, median (%)	25	31
Hours at SoC > 95 % idle, median	1.6	9.3
Pack T, annual mean (°C)	32.3	29.1
Pack T, May mean (°C)	37.7	34.6
Daily max T, P95 / max (°C)	40.1 / 44.0	37.3 / 41.0
Lost target distance (%)	8.7	0.3

All values are simulation outputs, not measurements.

B. Thermal Exposure

Fig. 4 shows that PbA packs run markedly hotter than LFP packs despite carrying similar currents. In May the mean PbA pack temperature is 37.7 °C against an ambient mean of 33.1 °C, and the 95th percentile of the daily maximum reaches 42.5 °C. The LFP pack runs only 1.4 K above ambient in the same month. The cause is not driving. As Fig. 2 shows, the PbA pack reaches its daily maximum after midnight, during the long absorption phase, when overcharge current that does not go into the plates is released as heat in an enclosed tray. Driving-related I^2R heating is modest for both chemistries at 0.1–0.15C.

This result has a direct consequence for instrumentation. A logger that records only while the vehicle is moving, or that sleeps once the ignition is off, would miss the hottest part of the PbA day. The field logger must therefore stay awake through charging, which is one reason it is supplied from the pack rather than from the vehicle's switched circuit.

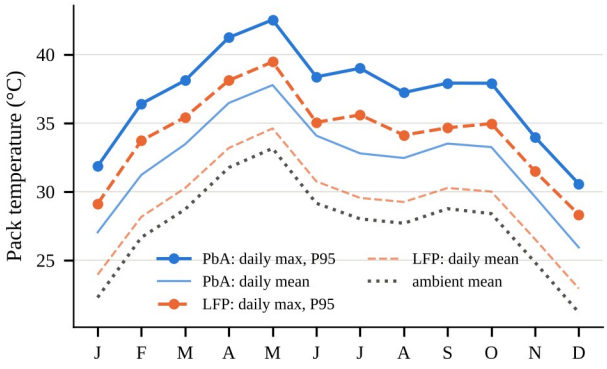


Fig. 4. Monthly pack temperature by chemistry against ambient (synthetic). Thick lines: 95th percentile of the daily maximum; thin lines: daily mean.

VII. Degradation and Its Drivers

A. Lead-Acid

Fig. 5(a) shows the twelve simulated PbA packs. First-life duration averages 327 days (5th–95th percentile 273–390 days), and the fleet needs 22 replacements over 24 months. That the mean lies close to the 300-day calibration target is expected and is not evidence that the model is right; the useful information lies in how life varies with the inputs. Corrosion accounts for 37 % of simulated fade and throughput-related mechanisms for the rest.

Mid-day charging access produces a result that survey data could not reveal. In calendar terms it makes almost no difference to pack life (328 against 326 days). In service terms it matters: each pack delivers 23,564 km before replacement with access against 21,855 km without, about 8 % more, because distance that would otherwise be lost to range limits is delivered at no extra cost in pack life. When only days-to-replacement are measured, this benefit is invisible. Battery life should therefore be reported per kilometre as well as per month in the field dataset.

The penalty term for consecutive days without full charge (γ) had no effect in any case. Under the simulated behaviour the overnight window is long enough for the pack to reach full charge on 90 % of days, so sulphation from chronic partial charge does not arise. This conclusion is only as good as the assumption that overnight sessions run uninterrupted; load-shedding or shared chargers at unauthorised stations [2] could change it, and the logger will show whether they do.

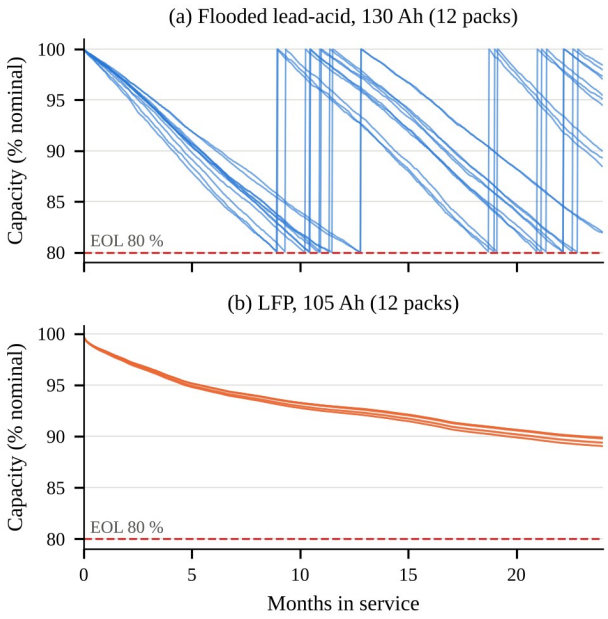


Fig. 5. Simulated capacity trajectories for all 24 packs over 24 months. PbA packs are replaced on reaching 80 % (vertical segments).

B. Lithium Iron Phosphate

The LFP packs lose 7.4 % of capacity in the first year and 10.5 % after two (Fig. 5(b)); the ratio of 1.415 is the √2 expected of square-root kinetics and confirms the integration is behaving as intended. Calendar ageing accounts for 62 % of first-year loss. The reason is visible in Table 3: an LFP pack charged at 25 A from the evening's typical SoC is full by around midnight and then sits at full charge for a median of 9.3 h each night, which is the most damaging storage condition in the calendar model, at night-time temperatures that remain near 30 °C for half the year.

Extrapolating the square-root law gives a projected life of 7.3 years to 80 %. That figure should be read as an upper bound. Traction LFP packs commonly depart from square-root behaviour as cells drift apart and a weakest cell limits the pack, and ageing within a cell can itself be non-uniform [15]; and the source model was fitted to cells far smaller than e-rickshaw prismatic cells [5], [6]. The field study is intended precisely to show whether real packs follow this curve.

C. Sensitivity and Interventions

Fig. 6 ranks the inputs by their effect on life. For PbA the two dominant terms, the deep-discharge weight β (247–408 days) and the corrosion rate k_c (212–386 days), are both assumptions that no published e-rickshaw data constrain. Because k_a was calibrated with them held fixed, the model can match the published life for many different combinations of β and k_c; it cannot tell how much of the observed wear is due to depth of discharge and how much to heat. Separating the two requires exactly the per-vehicle SoC and temperature histories the logger is designed to record. A 1.5 K warmer year shortens PbA life by about 11 days. Pack size between 120 and 140 Ah has no effect, because ageing in the model scales with throughput relative to capacity.

For LFP the calendar magnitude dominates (4.4–12.1 years), and adding Arrhenius temperature dependence to the cycle term, which the source study did not test, cuts projected life to 5.1 years. Two inexpensive charging interventions each move projected life by more than the seed spread (6.9–7.2 years):

capping charge at 90 % SoC extends it to 8.3 years, and a charge timer that finishes 30 min before the next shift extends it to 8.9 years, by shortening the time spent at full charge. Both are firmware settings rather than hardware changes, and both can be tested in the field.

For PbA the lunch-time top-up shortened reference life from 300 to 289 days, slightly outside the seed spread (305–311 days). Read together with the fleet result above, the top-up wears the pack slightly faster in calendar time because it is used more, while delivering more kilometres per pack.

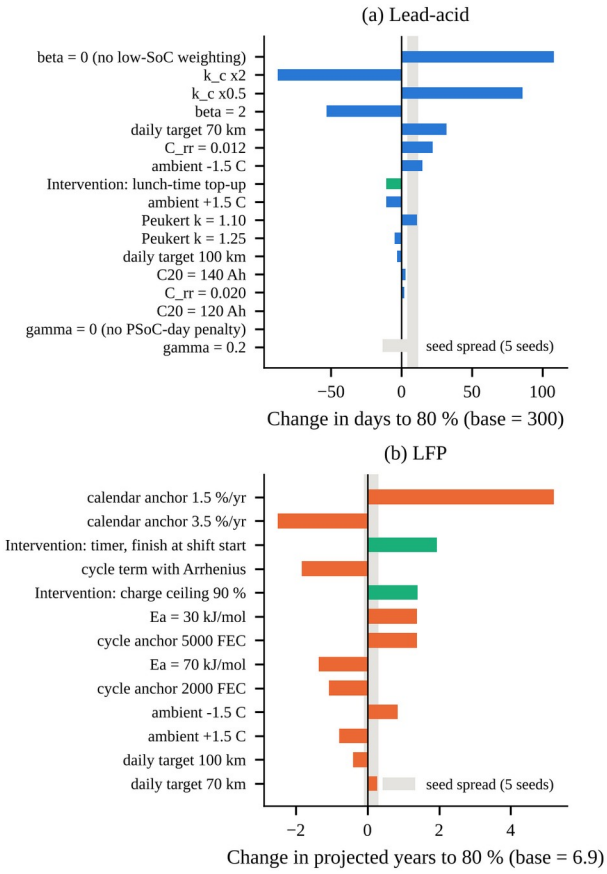


Fig. 6. One-at-a-time sensitivity of (a) PbA days to 80 % and (b) LFP projected years to 80 %. Grey band: range across five seeds of the base case. Green bars: practical interventions.

Cost per kilometre. Using the Bihar prices reported in [2] (Rs 40,000 for a new PbA set, Rs 8,000 scrap value), the simulated PbA battery cost is Rs 1.41 per kilometre. At the simulated 27,342 km per year, an LFP pack breaks even at a price of about Rs 2,81,814 if the projected 7.3-year life holds, and about Rs 1,40,907 if real life is half the projection. Odisha prices were not surveyed, and the ranking may change with them; the figures are offered as a breakeven threshold, not a recommendation.

VIII. Logger Design Results

A. Sampling Interval

Table 4 and Fig. 7 compare logging strategies against the 1-Hz synthetic reference. Point sampling, recording an instantaneous reading every Δt seconds, accumulates throughput error quickly: at 30 s the 95th-percentile daily Ah error is 5.2 % and at 60 s it is 8.7 %, larger than one month of LFP fade. Hardware charge integration with block averaging keeps the Ah error below 0.2 % at every interval tested. Averaging has its own cost: because heating scales with I^2, the square of a block-mean current understates the true mean of I^2 by 28 % at 30 s, and peak currents are flattened by 15 % at 60 s. The logger should therefore record, in each interval, the integrated charge, the mean of I^2 and the extreme values, rather than any single current reading. With those three quantities a 10-s record is sufficient for every metric in this study, and 1-Hz storage can be reserved for a rolling buffer and for sampled days. Because the synthetic reference has no sub-second content, these errors are lower bounds.

Table 4 Logging Error vs. 1-Hz Reference (%)

Δt (s)	Ah, point P95	I²R, point P95	I²R, averaged bias	Peak, averaged
5	0.51	1.2	−5.1	−0.0
10	1.41	3.2	−12.0	−0.4
30	5.16	9.7	−28.4	−5.3
60	8.74	15.8	−36.5	−14.7

Computed over 192 synthetic vehicle-days, five random sampling phases each.

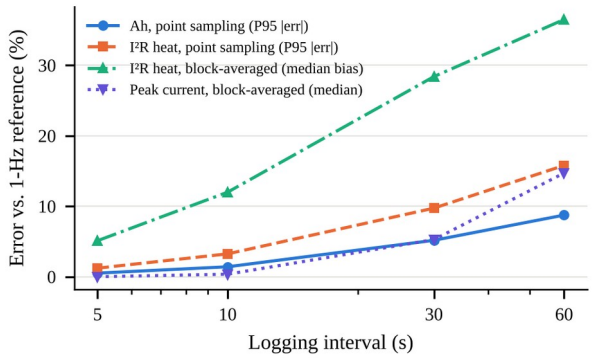


Fig. 7. Error of logging strategies against the 1-Hz synthetic reference as a function of logging interval.

B. Observability of Capacity

Fig. 8 gives the Monte-Carlo error of a single capacity estimate as a function of the SoC at the low anchor and of rest duration. The error model includes charger-termination ambiguity, voltage relaxation that is only partly corrected, sensor noise, and for LFP a residual ±5 mV per cell of the voltage hysteresis between charge and discharge branches [14]. Per-vehicle constant errors such as shunt gain are excluded because they cancel when fade *rate* is estimated.

The two chemistries behave in opposite ways. For PbA the open-circuit voltage is steep enough to read SoC anywhere, and the limit is relaxation: a 20 %-SoC anchor gives 3.5 % RMSE after 30 min of rest, 2.8 % after 2 h and 2.3 % after 8 h, with a persistent negative bias near −2 % from incomplete relaxation. For LFP rest duration hardly matters, because hysteresis rather than relaxation dominates, but the anchor SoC matters greatly: below 25 % SoC the error is 1.3–2.0 %, and on the voltage plateau between 40 % and 70 % it rises to 8–12 %. Only 30 % of simulated LFP working days end below 25 % SoC, against 53 % for PbA. An LFP SoH estimator built on rest voltage will therefore depend on the minority of days that end deep, and drivers who habitually charge early will be nearly unobservable by this method. For those vehicles the BMS's own capacity counter, or a periodic full discharge, will be needed.

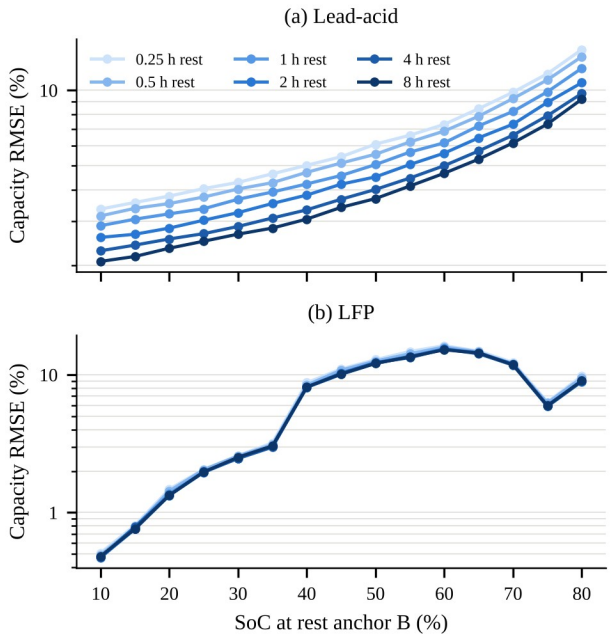


Fig. 8. RMSE of a single capacity estimate from a full-charge anchor and a rest-voltage anchor, by anchor SoC and rest duration (Monte Carlo, 4000 draws per point). In (b) the curves for different rest durations overlap.

C. Campaign Duration and Fleet Size

Combining the per-event errors with the anchors actually available in the synthetic fleet gives the monthly capacity-estimate noise in Table 5. The opportunistic strategy, using every end-of-shift rest that passes the depth gate, yields monthly noise of 0.58 % for PbA and 1.04 % for LFP. The scheduled strategy of two 2-h rests per month, which the authors had initially favoured, is worse: 1.55 % and 2.90 %. Better-relaxed anchors do not make up for having ten times fewer of them. This rests on the assumption that per-event errors are independent; if relaxation errors are correlated through temperature, the advantage of the opportunistic strategy will shrink, and the field data should be used to check this.

With opportunistic anchors, a single vehicle's fade becomes statistically detectable ($\alpha = 0.05$, power 0.8) after 3 months for PbA and 6 months for LFP. Comparing groups is far harder. Assuming a 20 % between-vehicle coefficient of variation in fade rate, which is itself an assumption to be measured, six LFP vehicles per group can detect only a 39 % relative difference in 12-month fade, and twelve per group a 27 % difference; detecting a 20 % difference needs 23 vehicles per group. The charging interventions that look most valuable in Section VII-C change 12-month fade by only about 12 %, which would require roughly 66 vehicles per group to confirm in a one-year field trial. For PbA, comparing first-life duration with the simulated between-vehicle spread of 39 days, a 30-day difference needs about 28 packs per group.

TABLE 5 DESIGN CONSEQUENCES OF THE PRE-STUDY

Quantity	PbA	LFP
Monthly SoH noise, opportunistic (%)	0.58	1.04
Monthly SoH noise, scheduled (%)	1.55	2.90
Months to detect fade, one vehicle	3	6
Min. detectable difference, 6/group	— *	39 %
Min. detectable difference, 12/group	— *	27 %
Vehicles/group for 20 % difference	—	23
Packs/group for 30-day life difference	28	—

** PbA packs reach end of life within 12 months, so groups are compared on life in days. $\alpha = 0.05$, power 0.8; between-vehicle CV of 20 % assumed.*

The planned 24-vehicle design is therefore adequate for its primary purpose, a published duty-cycle and stress dataset with per-vehicle fade curves, but under-powered for comparing groups. Three changes follow. First, the two cities should be pooled for chemistry comparisons unless a large city effect appears. Second, the LFP arm should be expanded to at least 12 vehicles, ideally by enrolling fleets that already have BMS telemetry. Third, the charging-intervention question should be moved off the road: a bench study holding spare packs at controlled SoC and temperature will answer it in months with a handful of packs, where a road trial would need dozens.

IX. DISCUSSION AND LIMITATIONS

The value of this work lies in what it allows a field team to avoid. Without the pre-study, a reasonable team might have logged current every 30 s, scheduled careful rest tests, relied on rest-voltage SoH for LFP packs, and planned to compare charging practices with six vehicles per group. Each of those choices would have weakened or defeated the dataset, and none would have been obvious until the money was spent.

The limitations are substantial and are stated plainly:

- **No measured data.** Every number in Sections V–VIII comes from a simulation. The ranking of stressors and the design rules are model-conditional.
- **Calibration is not validation.** PbA ageing was tuned so that a reference driver reaches end of life in 300 days, so agreement of the fleet mean with field-reported lives is circular.
- **Transferred ageing laws.** The LFP laws come from small cylindrical cells; their magnitudes were set to assumed datasheet anchors, and the square-root extrapolation to multi-year life ignores knee behaviour and cell imbalance.
- **Simplified vehicle and behaviour.** There is no regeneration, no explicit motor map, no driver-to-driver difference in aggressiveness, no monsoon flooding of roads, and no charger faults or grid outages. All PbA packs are treated as healthy flooded batteries with correct electrolyte levels, which field practice often violates.
- **Independent errors.** The sample-size results assume independent per-event SoH errors and a 20 % between-vehicle CV; both are guesses until the first months of field data are in.

The study design, code and synthetic dataset are released so that each of these assumptions can be replaced by measurement as the field data arrive, and the pre-study's predictions can be scored against them openly.

X. Conclusion

A field study of e-rickshaw batteries in Odisha was designed and stress-tested in simulation before deployment. The simulator reproduces the published energy-use and charging-time figures and reconciles the conflicting range reports through pack size. Within the model, PbA packs in Odisha's climate run hottest during overnight absorption charging, not on the road; stand-side top-up charging adds about 8 % to the distance delivered per pack without shortening its life in days; and LFP fade is dominated by calendar ageing at full charge overnight, which a 90 % charge ceiling or a charge timer could reduce materially. For the logger, the study shows that charge must be integrated in hardware and the mean of I^2 recorded alongside it; that every pre-charge rest should be used as a capacity anchor rather than a few scheduled ones; that rest-voltage SoH is weak for LFP packs that are seldom discharged below 25 %; and that the planned 24-vehicle fleet can track individual fade within three to six months but cannot compare groups unless it is enlarged or the comparison is moved to the bench. The next step is a pilot of four instrumented vehicles to test the logger hardware and the independence assumptions before the full campaign begins.

Acknowledgment

The authors thank the e-rickshaw drivers' associations of Bhubaneswar and Cuttack for discussions that shaped the field protocol. [Placeholder: confirm acknowledgments and any funding source before submission.]

References

[1] N. Jain and S. Gupta, “Real world energy efficiency calculation for e-rickshaws—A comparative study (lead acid vs lithium ion battery vehicles),” SAE Tech. Paper 2019-28-2486, NuGen Summit, Manesar, India, 2019, doi: 10.4271/2019-28-2486.

[2] N. Kataria, S. K. Gupta, and V. Tejaswi, “Understanding the electric transition in Bihar: A survey of 3W e-rickshaw operational practices,” Asian Development Research Institute (ADRI), Patna, India, 2023.

[3] S. Singh, “A study of the battery operated e-rickshaws in the State of Delhi,” Centre for Civil Society, New Delhi, India, Working Paper 323, 2014.

[4] J. Schiffer, D. U. Sauer, H. Bindner, T. Cronin, P. Lundsager, and R. Kaiser, “Model prediction for ranking lead-acid batteries according to expected lifetime in renewable energy systems and autonomous power-supply systems,” *J. Power Sources*, vol. 168, no. 1, pp. 66–78, 2007.

[5] M. Naumann, M. Schimpe, P. Keil, H. C. Hesse, and A. Jossen, “Analysis and modeling of calendar aging of a commercial $LiFePO_4$/graphite cell,” *J. Energy Storage*, vol. 17, pp. 153–169, 2018.

[6] M. Naumann, F. B. Spingler, and A. Jossen, “Analysis and modeling of cycle aging of a commercial $LiFePO_4$/graphite cell,” *J. Power Sources*, vol. 451, Art. no. 227666, 2020, doi: 10.1016/j.jpowsour.2019.227666.

[7] J. Wang *et al.*, “Cycle-life model for graphite-$LiFePO_4$ cells,” *J. Power Sources*, vol. 196, no. 8, pp. 3942–3948, 2011.

[8] *Standard Practices for Cycle Counting in Fatigue Analysis*, ASTM E1049-85 (Reapproved 2017), ASTM International, West Conshohocken, PA, USA.

[9] Japan Meteorological Agency, Tokyo Climate Center, “Monthly normals: Bhubaneshwar, India (WMO station 42971),” ClimatView/Normals data. Accessed: Oct. 2026.

[10] Commerce and Transport Department, Government of Odisha, *Odisha Electric Vehicle Policy, 2021*, Bhubaneswar, India, 2021.

[11] R. Dufo-López, T. Cortés-Arcos, J. S. Artal-Sevil, and J. L. Bernal-Agustín, “Comparison of lead-acid and Li-ion batteries lifetime prediction models in stand-alone photovoltaic systems,” *Appl. Sci.*, vol. 11, no. 3, Art. no. 1099, 2021.

[12] G. L. Plett, *Battery Management Systems, Vol. I: Battery Modeling*. Norwood, MA, USA: Artech House, 2015.

[13] D. U. Sauer and H. Wenzl, “Comparison of different approaches for lifetime prediction of electrochemical systems—Using lead-acid batteries as example,” *J. Power Sources*, vol. 176, no. 2, pp. 534–546, 2008.

[14] M. A. Roscher, J. Vetter, and D. U. Sauer, “Characterisation of charge and discharge behaviour of lithium ion batteries with olivine based cathode active material,” *J. Power Sources*, vol. 191, no. 2, pp. 582–590, 2009.

[15] M. Klett *et al.*, “Non-uniform aging of cycled commercial $LiFePO_4$//graphite cylindrical cells revealed by post-mortem analysis,” *J. Power Sources*, vol. 257, pp. 126–137, 2014.

[16] M. D. Wilkinson *et al.*, “The FAIR guiding principles for scientific data management and stewardship,” *Sci. Data*, vol. 3, Art. no. 160018, 2016.