

A Unified Continuous Time Framework for Adaptive Smartphone Battery Drain Prediction with Finite Horizon Error Analysis

Supplementary Information

Chuhan Wang^{1,*}, Xinxin Li¹, Yiran Sun¹, Shunxiang Liu¹, Haifeng Zhang^{1,*}

¹Harbin Institute of Technology, Harbin, Heilongjiang Province, 150006, China

*Correspondence: Chuhan Wang (2024110886@stu.hit.edu.cn); Haifeng Zhang (zhanghf@hit.edu.cn).

These materials document the data provenance, fixed evaluation procedures and supplementary results supporting the manuscript. Power, capacity and numerical verification use distinct targets and experimental units. Source measurements are retained separately from model-generated predictions.

S1 Data provenance and capacity extraction

UCL: DOI 10.5522/04/12159462, CC0 in publisher metadata. The original collection contains one electrochemical cycling CSV; its other EIL entries are tomography records. The downloaded EIL-MJ1-015.csv matches the original local record by SHA-256. NASA: all four physical cells in BatteryAgingARC-FY08Q4 are included. The official repository links the PHM-hosted archive at <https://phm-datasets.s3.amazonaws.com/NASA/5.+Battery+Data+Set.zip>. Its SHA-256 is 82302a7db4fc1b34e0b6676326610438d43b816bdf11a69d1d012a464ef2f92e. The current NASA catalog specifies no named license; its acknowledgement requirements are retained.

NASA MAT files contain charge, discharge and impedance operations. Only discharge operations supply capacity targets, and the ordinal index follows original discharge order. All 636 values are finite and positive; none are smoothed or excluded. The original operation index, date vector and ambient temperature are retained in derived CSVs. Charging uses 1.5 A to 4.2 V and constant voltage to 20 mA; discharge uses 2 A to operating cutoffs of 2.7, 2.5, 2.2 and 2.5 V for B0005, B0006, B0007 and B0018, respectively. The documented Capacity field is used without redefining its 2.7 V basis. The released files do not identify a specific electrode chemistry for these cells.

AndroWatts: dataset DOI 10.5281/zenodo.14314943, CC-BY-4.0. The aggregate CSV and source analysis script were retrieved from the original ZIP using HTTP byte ranges and ZIP CRC verification. The complete archive MD5 was not independently verified. The source's fixed 4.4 V conversion defines a power proxy; no per-profile terminal voltage is available in the aggregate file. External source code was inspected but not executed. Locally authored scripts perform the new analysis.

S2 Phone preprocessing and temporal protocol


The analytical archive combines CSV and raw JSONL records by same-acquisition nearest timestamp, with a 100 ms tolerance and one-to-one matching. File-rotation suffixes do not define new acquisitions. The raw current source is held fixed. A mislabeled charge_uah field has approximately the scale of mAh, and many raw charge-counter entries equal 1; neither is used to infer usable phone capacity. Charging-state screening also excludes plugged/full

fallback-current anomalies. Source flags and original timestamps are preserved in the local audit.


The additional protocol is exploratory because the original session holdout was already inspected. Models comprise identity, training mean, scale NNLS, affine NNLS, CPU NNLS, CPU-plus-screen NNLS and affine FFRLS. CPU utilization is divided by 100. FFRLS uses $[1, x]$, ridge 10^{-6} and $\lambda = 0.995$ per scored observation, predicting before updating with the current target. No test target is used for initialization. The update is not reset at automatic file-part boundaries.

Time weights are half of each eligible edge duration allocated to both endpoints. Edges require original-row adjacency, a shared acquisition and $0 < \Delta t \leq 5$ s. We do not bridge excluded observations or evaluation boundaries. In the later acquisition, 308 continuous segments include 200 isolated rows with zero time support. Total support is 18,679.055 s. Of 618 candidate segment/minute bins, 297 meet at least 20 observations and 30 s support; they contain 17,067 raw observations and 17,416.024 s support. Minute metrics are equally weighted across bins after within-bin time averaging.

Forward checks use 30 min UTC-clock blocks. A block requires at least 300 eligible test rows and 600 rows strictly earlier than its start minus 60 s. Seventeen of 23 candidate blocks qualify. All earlier eligible observations initialize a block; future blocks are excluded. FFRLS then updates sequentially within the block. Every skipped-block reason, model coefficient and prediction is saved. These blocks are repeated periods on one phone, not additional independent experiments.



Supplementary Figure S1. Recording coverage and field availability in the new ZTE archive. Acquisition duration is distinct from observed temporal support. Battery percentage is device telemetry.



Supplementary Figure S2. All fixed model candidates across 17 forward temporal blocks. Larger aggregate improvements coexist with individual blocks where simpler methods have lower error. Blocks are correlated periods on the same device.


S3 Phone condition strata and dependence

Fixed strata use screen state and battery temperature <30 , $30-<35$, $35-<40$ and $\geq 40^\circ\text{C}$. Counts are 13,331 screen-off and 4304 screen-on observations, and 36, 14,635, 2335 and 629 observations by ascending temperature stratum. Scale MAEs in those temperature groups are 0.4846, 0.1901, 0.3365 and 0.8172 W. FFRLS MAEs are 0.2808, 0.0972, 0.2003 and 0.2460 W. Screen-on/off scale MAEs are 0.3320/0.2003 W and FFRLS MAEs are 0.2065/0.0875 W. These are observational comparisons with co-varying workload and time.

Raw neighboring residual correlations use 17,327 eligible pairs; minute-neighbor correlations use 291 consecutive-minute pairs in shared segments. Scale/FFRLS correlations are 0.6357/0.5262 at raw resolution and 0.6067/0.5588 at minute resolution. No independent-sample p values, device-population confidence intervals or effective-sample-size estimates are inferred from these serial correlations.

S4 Capacity fits and identifiability

The fixed split is 300/100 for UCL, 126/42 for B0005–B0007 and 99/33 for B0018. Single-exponential starts number four and dual-exponential starts 27, with the exact start arrays and optimizer tolerances retained in `analyze_multicell.py`. All selected fits report numerical convergence, which is not a guarantee of global optimality or physical parameter identification. No forecast clipping is applied. The recent linear baseline uses 60 UCL calibration points, 26 for the three 168-cycle cells and 20 for B0018.



Supplementary Figure S3. Capacity trajectories and fixed-split predictions for the UCL reference and all four NASA cells. Shading marks held-out observations. Overlap of single- and dual-exponential forecasts reflects parameter degeneration; panel f uses initial-capacity-normalized errors.

Table S1. NASA equal-cell capacity error summaries

Model	Mean MAE Ah	Mean RMSE Ah	Mean normalized RMSE %
Dual exponential	0.02597	0.03099	1.641
Single exponential	0.02598	0.03100	1.641
Global linear	0.05854	0.06489	3.372
Persistence	0.05526	0.06406	3.310
Recent linear	0.02563	0.03087	1.620

S5 Public workload analysis

The eleven inputs are RGB levels, brightness, CPU_LITTLE_FREQ_KHz, CPU_MID_FREQ_KHz, CPU_BIG_FREQ_KHz, GPU0_FREQ, GPU_1FREQ, GPU_MEM_AVG and TOTAL_DATA_WIFI_BYTES. CPU-mid err and zero Wi-Fi values are handled as missing following the source's coding. The Wi-Fi sentinel rule was corrected after the first execution by inspecting the source, without altering candidates or test groups. Imputation uses training means; source-code full-dataset median imputation is not reused. Wi-Fi is fixed-duration profile traffic, not an independently measured instantaneous throughput.

Five blocks contain successive groups of 200 sorted IDs. CPU-only OLS uses an intercept and three CPU inputs. All-feature OLS uses all eleven inputs with an intercept. NNLS constrains the intercept and all coefficients to be nonnegative and uses uncentered, standard-deviation-scaled inputs. No hyperparameters are selected on the evaluation profiles. Pooled out-of-fold MAEs are 0.9242, 0.3469, 0.2959 and 0.4061 W for mean, CPU OLS, full OLS and NNLS, respectively. Equivalent current MAEs divide by 4.4 V and equal 210.0, 78.8, 67.2 and 92.3 mA. This conversion adds no independent voltage measurement.

S6 Numerical verification and reproducibility

The complete numerical protocol, all 18 solver-step outcomes, full uncertainty trajectories and an independent analytical check are supplied with `run_verification.py`. The current input is right-continuous for voltage evaluation at switching events; each preceding integration interval retains its own current. The scalar RC exponential reference is independent of the tested time-step recurrence. Errors are expressed as fractions for SOC in machine-readable files and percentage points in the propagation plot.

The uncertainty bound becomes an equality in this example because current bias and capacity underestimation act in the same direction under positive discharge. This designed equality is an implementation check, not an empirical uncertainty-coverage estimate or a proof of uniform-in-time stability. No synthetic trajectory is merged with the physical cell or phone datasets.

Independent verification recomputes saved predictions and metrics through separate implementations. Capacity metrics for all 25 cell/model combinations agree within 3.4×10^{-14} . Phone FFRLS is checked using exponential-weighted 2×2 normal equations rather than the covariance recursion, with a maximum prediction difference of 1.33×10^{-12} W. Forward training-time and 60 s embargo checks pass. Supplementary Data 1 and Supplementary Code 1 contain the analytical inputs, scripts, candidate predictions and source manifests. Extract both ZIP files into the same parent directory to combine their SR_Reproducibility folder, then follow its README.

S7 Relationship to recent smartphone battery models

Table S2 compares the focus of recent related studies with the present framework. Reference numbers [10–13] refer to the main manuscript. The comparison concerns objectives, evidence and the connection between modules; reported accuracies from different datasets and reference quantities are not ranked.

Table S2. Related approaches and the present contribution

Study	Objective and evidence	Focus of the present framework
Wu et al. [10]	Sensor based continuous state and lifetime prediction; GreenHub observations.	Sequential load calibration on newly collected phone sessions, linked to current and capacity error propagation.
Long et al. [11]	Modular electrothermal SOC and TTE modeling; PinePhone scenarios.	Empirical testing of telemetry driven adaptation alongside finite-horizon charge-balance error accounting.
Zheng et al. [12]	Aging aware SOC estimation; cell characterization and behavior based loads.	Repeated forward capacity extrapolation across five cells and an explicit capacity-error term.
Xu et al. [13]	Instantaneous and delayed drain decomposition; device logs and battery-level proxies.	Online battery-side power calibration connected to continuous charge balance and polarization dynamics.
Present study	Self-collected ZTE sessions, public workload profiles, five cell trajectories and numerical verification.	An observation-to-state interface with prediction-before-update evaluation and separate current, capacity and numerical error terms.

The distinctions identify the emphasis of this study and do not assert that related studies contain no overlapping components. Established circuit, recursive-estimation and integral-inequality methods remain the building blocks of the proposed architecture.