

MATERIALS CHALLENGES AND ML-DRIVEN INNOVATION IN NEXT GENERATION CLEAN ENERGY - REVIEW



Feedback, physics, and forecasts: The emerging paradigm of machine learning-driven battery research

Sreya Vangara , Department of Mechanical Engineering, Stanford University, Stanford, CA, USA; SLAC National Accelerator Laboratory, Applied Energy Division, SLAC-Stanford Battery Research Center, Menlo Park, CA, USA

Jagjit Nanda and **Yan-Kai Tzeng**, SLAC National Accelerator Laboratory, Applied Energy Division, SLAC-Stanford Battery Research Center, Menlo Park, CA, USA

Address all correspondence to Yan-Kai Tzeng at ytzeng@slac.stanford.edu

(Received: 8 December 2025; accepted: 9 February 2026)

ABSTRACT

- *Machine learning has reshaped battery research from post hoc analysis toward closed-loop prediction, control, and design.*
- *Physics-informed and hybrid models enable earlier failure forecasting and improved generalization across operating conditions.*
- *Integrating machine learning with economic, manufacturing, and sustainability constraints is essential for real-world battery deployment.*

Machine learning (ML) is reshaping how we understand, predict, and optimize electrochemical systems. In batteries, ML accelerates discovery across chemistry, design, and operation by transforming massive experimental and simulated datasets into predictive, interpretable models. This review consolidates a decade of progress in ML-driven battery innovation, from early-cycle feature extraction to operando image analysis and physics-informed modeling. We categorize approaches by data domain and physical fidelity, emphasizing interpretable ML for diagnostics, reinforcement learning for control, and multi-objective optimization for lifetime extension strategies. Additionally, we demonstrate how integrated models accelerate discovery, reduce testing time, and guide sustainable design. Economic analyses furthermore illustrate how these advances can lower cost per cycle and improve circularity. Together, these developments chart a path toward self-optimizing, sustainable battery technologies.

Keywords energy storage · machine learning · sustainability · artificial intelligence · autonomous research

Discussion

As machine learning-driven battery optimization becomes increasingly autonomous, how should responsibility be assigned when algorithmic decisions trade off lifetime, safety, and cost in real-world systems?

Physics-informed machine learning promises interpretability, yet overly simplified physical assumptions may encode systematic bias, raising the question of when “less physics” may yield more reliable predictions. Open battery datasets accelerate innovation, but without standards for data provenance and protocol diversity, they may unintentionally narrow the design space explored by the community.

Jagjit Nanda was an Editorial Advisory Board member of this journal during the review and decision stage. For the MRS Energy & Sustainability policy on review and publication of manuscripts authored by editors, please refer to mrs.org/editormanuscripts.

Introduction

The global push for decarbonization has elevated electrochemical energy storage to center stage. From enabling electric vehicles (EVs) and fast-charging infrastructure to buffering intermittent renewable energy sources on the electric grid, the demand for electrochemical energy storage is actively accelerating. Currently, lithium-ion batteries (LiBs) power these green energy transitions due to their high energy density and rapidly dropping costs. Worldwide deployment of LiBs is skyrocketing; annual LiB production for EVs alone reached over 116 GWh in 2019 and is projected to surpass 600 GWh by 2025.¹ Yet, critical challenges remain. Consumers demand batteries that charge in minutes and last for thousands of cycles, but hastening performance often accelerates degradation or failure. Meanwhile, novel battery chemistries like sodium-ion and solid-state lithium promise improvements in cost, safety, or energy density, but face steep development hurdles toward commercialization. A major bottleneck is time: evaluating a new battery formulation or usage protocol can take months or years of cycling tests.² This delayed feedback stifles innovation and inflates research costs. As a result, battery advancement has historically been iterative and slow, failing to match the pace of the overarching green energy sector.

Machine learning (ML) has recently emerged as a powerful tool to break this bottleneck by predicting battery failures and lifetime in early cycles and by optimizing materials and operating protocols with far fewer experiments than brute-force approaches. Instead of treating battery development as an uninformed trial-and-error process, ML algorithms can identify subtle patterns in initial performance that foreshadow long-term degradation.³ These data-driven models, especially when combined with physical domain knowledge, offer a paradigm shift: the ability to rapidly forecast cycle life, detect emergent failure modes, and design experiments in a closed-loop manner that accelerates discovery.² The past few years have seen an explosion of battery ML research, from pioneering early-life prognostics and closed-loop fast-charging optimization to physics-informed models that reveal why batteries fail.^{2,4} In parallel, techno-economic analysis enables an understanding of how these technological advances translate to real-world impact, such as assessing how extending battery lifetime or improving fast-charge capability affects EV adoption.⁵

In this review, we provide a comprehensive overview of machine learning-driven failure prediction and characterization in battery materials. We survey key developments in data-driven battery life prediction (“[Data-driven battery life prediction](#)” section) and closed-loop optimization of charging protocols (“[Closed-loop optimization of protocols](#)” section) that reduce testing times by orders of magnitude. We then examine how physics-informed and mechanistic ML approaches (“[Physics-informed and mechanistic ML](#)” section) integrate domain knowledge, from atomic-scale simulations to operando imaging, to improve model interpretability and accuracy. An economic and societal perspective is provided in “[Economic and societal perspectives](#)” section, highlighting how techno-economic

analysis frames the value and challenges of deploying these ML-accelerated innovations at scale. Finally, we outline open science efforts in “[Open science efforts and community resources](#)” section, such as public battery datasets and codes that are fueling community-wide progress.

Data-driven battery life prediction

One of the earliest and most impactful applications of machine learning in the battery field has been early-life prognostics: using data from a cell’s initial cycles or formation process to accurately predict its eventual lifetime or failure modes. If successful, this enables high-throughput screening in which battery materials or designs can be rapidly triaged and assessed without waiting for full aging tests. Early-life prognostics can thus quickly eliminate inferior designs, optimize formation protocols, sort cells by quality, and save tremendous time and cost. Traditional approaches to battery lifetime estimation have relied on semi-empirical aging models or diagnostic metrics that require significant degradation to occur before extrapolation. In contrast, data-driven prognostics strive to make accurate predictions before significant capacity fade, instead teasing out subtle variations in initial performance that preclude long-term behavior. Indicators and impact of different sections in the battery lifetime are explored in Fig. 1, and an overview of various early-life prognosis methods is provided in Table 1.

A landmark study in this area was the 2019 Stanford-MIT-Toyota collaboration by Severson et al., who demonstrated that lithium-ion cycle life can be predicted with high accuracy using only the first 100 charge/discharge cycles of data.⁶ They generated a comprehensive dataset of 124 commercial LFP/graphite cells aged under various fast-charging protocols (cycle lives ranged from 150 to 2300 cycles) and used the resulting time-series of electrochemical measurements (voltage, current, capacity, impedance, and temperature) to train ML models. By extracting engineered features from the early-cycle voltage curves (*i.e.*, ΔQ_{100-10}) and applying regularized regression, their model achieved a mean absolute error (MAE) of 9% in cycle life prediction on hold-out test cells using only the first 100 cycles of each cell’s data. Remarkably, even using the first five cycles of data, the algorithm could classify cells into long-life vs. short-life groups with ~95% accuracy. These results were among the first to show that subtle early signals, invisible to the typical researcher, could foretell failure. This study defined a paradigm for data-driven battery prognostics, inspiring many subsequent efforts. A summary of representative early-life prediction studies and their approaches is available in Fig. 2. Commonly-used benchmark datasets in these cycle life prediction studies are outlined in Table 2.

Recent improvements

Following the above breakthrough, researchers have expanded data-driven life prediction along multiple fronts. One direction has been to improve the interpretability and feature




Figure 1. Battery cycle lifetime from pre-formation rest to end of life. Five key phases define a cell's usable life. Pre-formation rest (OCV/wetting) ensures uniform electrolyte wetting and reduces later impedance growth. Formation ($\approx 1\text{--}3$ cycles) establishes the SEI and electrode utilization window that governs aging. Early-life stabilization ($\approx 1\text{--}100$ cycles) shows flat capacity but predictive voltage–capacity shifts enabling ML-based lifetime forecasting. Mid-life to EOL involves gradual fade, a “knee” acceleration, and rising polarization that broadens dQ/dV peaks and limits usable capacity.

engineering behind such models. For instance, Cui et al. recently evaluated dozens of candidate early-life features to determine which ones most impactfully drive predictive performance.⁷ By systematically ranking features (*i.e.*, voltage plateau metrics, discharge capacity trends, impedance trends, *etc.*) using Shapley values and similar statistical techniques, they identified a minimal subset of predictive features and provided insight into why these may correlate with longevity. Such interpretability studies help focus attention on the physical mechanisms underlying statistical predictions; for example, certain voltage curve distortions might correspond to lithium plating onset or the loss of active material, which guides ML predictors to link particular electrochemical measurement patterns to failure modes.

Another thrust has been the application of advanced machine learning architectures to improve prediction accuracy and generality. Whereas the initial Severson et al.⁶ model used simple linear regression and decision trees, subsequent works have incorporated more complex techniques, such as deep learning. Zhou et al. combined a deep residual network (ResNet) with a genetic algorithm (GA) feature selection to predict remaining life, achieving over 91% accuracy using the first 100 cycles of a 40-cell dataset.⁹ The ResNet + GA model outperformed simpler models (plain ResNet, SVR) in tests. Rout et al. also harnessed deep learning by combining boosted trees with recurrent architectures to model capacity fade; in their experiments, an LSTM network achieved a coefficient of determination of $R^2 \approx 0.998$ on test data, effectively learning long-horizon state of health (SoH) trajectories from cycling histories.¹¹ Similarly, Li et al. developed

particle-swarm-optimized (PSO) deep models, including LSTMs and CNNs, and found that a PSO-tuned LSTM outperformed alternative models such as SVR in SoH prediction accuracy.¹⁰ These results illustrate how hyperparameter optimization and architecture search can materially improve deep learning battery health predictors.

Hybrid models have also been explored, with great success. Fan et al. proposed a hybrid approach that first applied signal decomposition (CEEDMAN) to the capacity versus cycle data, then used Support Vector Regression (SVR) and smoothing (SSA) to extract features, finally fed the features into a Long Short-Term Memory (LSTM) neural network optimized for remaining useful life (RUL) prediction.⁸ This multi-step pipeline achieved near-perfect RUL predictions on a public NASA battery dataset. Zhang et al. adopted a hybrid, attention-based “Informer” architecture to estimate health status from long-horizon time series, demonstrating how transformer-style sequence models from broader AI can be adapted to battery diagnostics.¹⁴

These hybrid approaches leverage domain knowledge, such as known aging modes and identifiable cycling landmarks, to constrain ML models and improve generalization across different usage profiles.


Other researchers have explored probabilistic ML (*i.e.*, Gaussian process regression and Bayesian neural networks) to quantify uncertainty in life predictions, which is crucial for applications requiring uncertainty-aware control and risk-informed decision-making. An early example by Richardson

Table 1. Comparison of early-life prognostics methods.


Model type	Input features (early cycle)	Dataset/chemistry	Reported accuracy/error	Highlights
Elastic-net regression (linear ML) ⁶	Variance of ΔQ_{100-10} (V); discharge plateau shape; temperature + resistance features	124 LFP/graphite cells (A123 APR18650M1A) under 72 fast-charge conditions	9.1% RMSE (quantitative) using first 100 cycles; 4.9% error (classification with first 5 cycles)	First mechanism-agnostic early prediction demonstration; established $\Delta Q(V)$ variance as key feature
Bayesian optimization + linear early Predictor (elastic-net) ²	$\Delta Q(V)$ and capacity fade trend features from first 100 cycles	224 six-step fast-charge protocols on LFP/graphite cells (A123 systems)	RMSE ~ 80 cycles (~8–10%) for early-life model	First closed-loop optimization linking early ML prediction to experimental design
Feature screening + random forest regression ⁷	dV/dQ slope variance; plateau shift rate; charge-transfer resistance proxy from formation data	186 SC-NMC532/graphite pouch cells	MAE $\approx$ 7–9% of true lifetime (5-cycle input)	Demonstrated formation stage prognostics without cycling to 100 cycles
CEEMDAN decomposition + SVR and LSTM ⁸	Statistical fusion of $\Delta Q(V)$, IR, temperature rise, Coulombic efficiency	NASA Dataset (18,650 Li-ion)	$R^2 = 0.963$ (MAE $\approx$ 0.04 SOH fraction)	Improved robustness to sensor noise via multi-domain fusion
Deep residual network (ResNet) ⁹	Raw voltage–capacity curves (first 50 cycles) + derived spectral features	124 LFP/graphite cells (A123 APR18650M1A) under 72 fast-charge conditions (Severson et al. dataset)	MAE $\approx$ 8.97% using the first 100 cycles	End-to-end deep learning approach for early-cycle prediction
PSO-optimized LSTM/CNN/SVR ¹⁰	Raw cycling data + engineered temporal features	NASA datasets (multiple Li-ion chemistries)	LSTM (PSO tuned) best accuracy; significantly > CNN/SVR	Demonstrates automated hyperparameter optimization for SoH prediction
Boosted trees + LSTM ¹¹	Time-series capacity fade trajectories	Li-ion 18,650 NMC cells arranged in packs of 520	$R^2 \approx 0.998$	High-accuracy long-horizon SoH prediction

A comparison of major early-life prognostic studies showing model types, input features, datasets, and reported predictive performance.

FEATURE SELECTION



MODEL ARCHITECTURES



PREDICTIONS




Figure 2. From raw cycling data to lifetime prediction. A three-stage workflow common to recent battery-prognostics studies. Feature extraction: transforms raw voltage–current–capacity–time data into statistical, derivative, interpolated, or spectral predictors. Model architectures: features feed into linear, classical ML, hybrid, or deep sequence models. Prediction targets: models output cycle life or remaining useful life (RUL) to defined state-of-health thresholds. Letter–reference key: A = Ref. 12; B = Ref. 13; C = Ref. 7; D = Ref. 9; E = Ref. 8; F = Ref. 2; and G = Ref. 6.

Table 2. Benchmark datasets for cycle-life prediction.

Dataset name	Description/chemistry	Size (# cells, cycles)	Key attributes/measurements
Severson–Chueh dataset (TRI 2019) ^{6, 9}	LFP/graphite cells under 72 fast-charging protocols	124 cells, 96,700 cycles	Voltage, current, temperature, internal resistance
Attia–Chueh closed-loop set (2020) ²	224 fast-charge protocols (LFP/graphite) used for Bayesian Optimization	185 cells, 100 cycles each	Current profile, predicted lifetime labels
NASA battery prognostics dataset (B0005–B0018) ^{8, 10}	18,650 Li-ion cells under fixed current charge/discharge regimes	34 cells, full lifetime	Voltage, current, temperature, capacity vs cycle
Real-world EV field dataset ⁵	Real mid-size electric SUV with a 95 kWh Li-ion battery pack, 36 modules of 12 pouch cells each, driven regularly over one year	> 10,000 vehicles aggregated	Current, power, temp, drive-cycle usage

A summary of widely used benchmark battery-life datasets, including chemistries, size, measurement types, and representative studies.

et al. used Gaussian processes on partial cycling data to predict capacity fade trajectories with uncertainty bounds, enabling estimation of not just when a battery might fail but also confidence in that estimate.¹⁵ Recent reviews highlight that incorporating uncertainty quantification and probabilistic predictions is becoming standard in battery prognostics to account for variability and rare events.

Challenges for early lifetime prediction

Notably, a key challenge for data-driven life prediction is ensuring generalizability across diverse operating conditions and cell types. Models trained on one cycling regime often falter if the usage profile changes. To address this, Zhang et al. introduced an “inter-cell transfer learning framework” called BatLiNet that learns from pairs of cells under different conditions to predict lifetime across a variety of cycling protocols.¹² By combining conventional single-cell modeling with an inter-cell comparison mechanism, their deep learning model achieved higher accuracy and robustness when tested on cells aged under conditions not seen in the training data. BatLiNet even demonstrated transferability across different Li-ion chemistries, hinting at more universal prediction models that can leverage cross-dataset insights to perform chemistry-agnostic inference. In parallel, researchers are seeking ways to incorporate physical knowledge into lifetime predictions, as well as methods to handle real-world field data produced under variable conditions.

To contextualize the diverse modeling approaches discussed above, Table 3 provides a qualitative comparison of major machine-learning paradigms used in battery research, highlighting tradeoffs in data requirements, generalization risk, interpretability, and deployment maturity. As summarized in Table 3, no single modeling paradigm consistently

dominates across all battery tasks. Physics-informed and hybrid approaches increasingly offer the best balance between predictive accuracy, data efficiency, and generalization. In contrast, purely data-driven deep models excel under controlled conditions but remain sensitive to protocol and chemistry shifts, while mechanistically constrained models trade some expressiveness for improved robustness. This comparison highlights that future state-of-the-art performance is likely to emerge from adaptive hybrid frameworks rather than monolithic architectures.

An additional challenge is the generalizability of datasets and the transfer of lab-borne algorithms to real-world conditions. Widely used benchmark datasets are typically generated under controlled laboratory conditions with limited protocol diversity and uniform cell formats; examples include all but the final row in Table 2. While invaluable for model development, such datasets may underrepresent real-world variability arising from manufacturing tolerances, environmental fluctuations, and user behavior. As a result, models trained exclusively on these datasets may exhibit optimistic early-life performance while failing to generalize under field conditions, underscoring the need for protocol-diverse and deployment-informed data collection.

Overall data-driven battery life prediction has evolved from a niche demonstration into a broad research thrust. Continued advances in model interpretability, handling of heterogeneous data, and integration of domain knowledge are rapidly improving the ability to forecast battery lifetimes. Early-life prognostics not only accelerate R&D by forgoing long aging tests, but also open up new opportunities in battery fleet management, such as dynamically adjusting operating conditions for cells showing early failure signs. These capabilities form a foundation for many subsequent ML applications in the battery field, including the closed-loop optimization techniques discussed next.

Table 3. Representative machine learning paradigms for battery modeling.

Model class	Representative tasks			Data requirements		Generalization risk		Interpretability		Deployment maturity
	Lifetime prediction, capacity fade forecasting	Voltage/SoH prediction, degradation trajectories	State estimation, degradation modeling, failure prediction	Charging protocol optimization, experimental design	Low–moderate; curated cycling datasets	High sensitivity to dataset bias	High	High	High	
Early-life regression models ^{2, 6}										High
Deep sequence models (<i>i.e.</i> , RNNs, Transformers) ^{8, 10}					High; diverse cycling protocols	Risk of overfitting to protocols	Low	Low	Medium–high	Medium
Physics-informed/hybrid ML ^{16–18}					Moderate; augmented by physical constraints	Reduced, but dependent on model fidelity	Medium–high	Medium–high	Medium	Medium
Closed-loop optimization and RL ^{2, 14}					Moderate–high; online or simulated interaction	Risk under distribution shift	Low–medium	Low–medium	Low	Low

Comparative overview of major machine-learning paradigms for battery modeling across tasks, data requirements, generalization, interpretability, and deployment readiness.

Closed-loop optimization of protocols

Predicting battery behavior is powerful, but controlling and optimizing that behavior is the logical next frontier. Closed-loop optimization refers to using machine learning in the experiment loop itself: iteratively designing, testing, and refining battery designs or protocols to rapidly achieve target outcomes. This approach treats experimentation itself as an optimization problem, for which machine learning helps intelligently navigate the vast space of possible material formulations and operational protocols. Two areas of recent significant strides are in optimizing fast-charging protocols and initial conditioning formation protocols.

Fast charging optimization

Charging Li-ion cells faster is highly desirable, but aggressive charging can induce lithium dendrite growth and other failure mechanisms that drastically shorten life. Traditionally, engineers would test a handful of charging profiles, such as various combinations of constant current, pulses, and rest steps, and observe which protocols best balance speed and longevity. This was a costly and lengthy process which impeded rapid iterative battery R&D. Attia et al. revolutionized this lethargic cycle by applying Bayesian Optimization (BO) and early-life prediction.² In their closed-loop scheme, an initial set of fast-charge protocols was tried on LFP cells and then used to train a model predicting cycle life from early-cycle data (following the Severson et al. approach). That model then guided the selection of new protocols to test via BO acquisition functions aiming to maximize lifespan. By cycling only a small fraction of candidate protocols to failure and using ML to infer outcomes for the rest, the algorithm identified an optima fast-charge protocol in just 40 cycles of testing per candidate: a span of a few weeks, in stark comparison to the conventional brute-force exploration of the large protocol space that could have taken years. The optimized protocol found was twice as fast as the conventional baseline, yet still yielded comparable cycle life (~92% capacity retention after 600 cycles vs. ~90% for the baseline). This was a remarkable result; machine learning had discovered a non-intuitive multi-step charging waveform that mitigated degradation by reducing time spent at a high SOC and other pertinent factors, all while charging in ~15 min. Beyond the novel protocol, the study demonstrated a workflow now being adopted widely: using data-driven models as surrogates to evaluate many conditions rapidly, then experimentally confirming only the most promising ones. A conceptual illustration of this closed-loop optimization process combining predictive modeling with intelligent experiment selection can be found in Fig. 3, and a summary of ML-optimized charging protocols can be found in Table 4.

Building on that work, subsequent studies have applied various ML techniques for optimizing both charging and discharging strategies; the major paradigms of control algorithms for this battery management are outlined in Table 5. For example, Thelen et al. employed a probabilistic ML model to design fast-charge protocols that avoid lithium plating by mapping out




Figure 3. Closed-loop Bayesian optimization for fast-charging protocol discovery. A six-stage feedback cycle adapted from Attia et al. and Jiang et al.^{2,19} Candidate current profiles are generated and tested experimentally, producing voltage–capacity data used for feature extraction (*i.e.*, $\Delta Q(V)$ variance, internal resistance growth). These early-cycle descriptors update a Gaussian process surrogate model that estimates lifetime or charge time performance. The Bayesian optimizer evaluates the acquisition function (expected improvement) to select the next protocol, closing the loop between experiment and model. This iterative cycle reduces experimental search time by more than an order of magnitude while systematically converging on fast-charge conditions that maximize cell life.

plating probabilities across the charge curve.²³ By focusing on the portions of the charge process where plating is most thermodynamically likely, they devised charging profiles that significantly reduced plating while only slightly lengthening charge time. Zhang et al. for example, implemented a feature-driven closed-loop optimization framework for multistage fast-charge protocols; an ML surrogate mapped early-cycle features to projected cycle life, and a search algorithm iteratively adjusted current steps and durations.¹⁴ This approach achieved up to ~16% reduction in charge time relative to baseline protocols, with negligible loss in lifetime. Notably, the optimized protocols were structurally different from standard CCCV schemes, illustrating how ML can uncover non-intuitive charging waveforms that human designers might overlook. Furthermore, Jiang et al. extended these ideas by integrating Bayesian optimization with recurrent neural networks to identify fast-charge profiles that

satisfy electrochemical constraints.²⁰ Their RNN surrogate predicted aging outcomes from early cycles, and the Bayesian optimizer explored the protocol design space under constraints designed to avoid lithium plating. The resulting framework identified high-current protocols that substantially shortened charge times while meeting health and safety limits, demonstrating that constrained global optimization can efficiently navigate large protocol spaces.

Other researchers have explored reinforcement learning (RL); Zhang et al. formulated fast-charging as a continuous control problem and solved it with a Deep Deterministic Policy Gradient (DDPG) algorithm, which was able to learn an optimal charging current profile that adapted in real time to the cell's state.²¹ Their DDPG-based charger achieved higher efficiency and extended life compared to standard constant current charging. He et al. trained an RL agent directly on extensive

Table 4. Summary of ML-optimized charging protocols.

Study/method	Main objective	Data/setting	Performance gain
Attia et al. closed-loop optimization (early prediction + Bayesian optimization) ²	Maximize cycle life of 10 min fast-charge protocols	224 candidate 6-step protocols (5xCC + 1xCV); CLO in four rounds	Identified top-lifetime protocols in 16 days versus ~560 days by exhaustive testing (~35 × fewer test days for discovery)
Jiang et al. Bayesian optimization (BO) with constraints ²⁰	Minimize charging time under degradation constraints	Physics-based simulator; compares acquisition functions and # of steps	3-step CC protocol charges 126.6 s faster than 1 step and 14.0 s faster than 2-step baseline in their tests
Zhang et al. deep RL (DDPG) for pack-level charging + balancing ²¹	Minimize total charge time and balancing time with thermal/voltage/SOH constraints	Simulated multi-cell pack, electro-thermal model	Basic DDPG: 650 s (110 s slower than benchmark) while satisfying safety constraints; CC-CV is faster but violates thermal limit (~51 °C)
He et al. RL-enabled fast charging using real-world experimental data ²²	Accelerate charging while ensuring safety and reducing aging	Real experimental dataset: 600 charge/discharge cycles per cell; Q-learning with ϵ -greedy + DFN model validation	RL policy reduces voltage overshoot and temperature excursions versus CC-CV; achieves faster time-to-SOC while lowering aging after 500 cycles (capacity loss 0.0493 vs. 0.0576 Ah for CC-CV)
Thelen et al., probabilistic ML plating-avoidance control ²³	Avoid lithium plating while enabling fast charging	SOC-dependent plating probability model; probabilistic ML maps plating risk versus SOC and operating conditions	Achieves lower plated-Li probability while maintaining comparable charge time to CC-CV; safety envelopes allow higher allowable currents in low-risk SOC regions
Zhang et al. (2023), feature-driven closed-loop optimization ¹⁴	Identify actionable feature drivers for fast-charge optimization	Lab cycling data; feature ranking + CLO iterations	Discovers robust feature subsets that correlate with optimal fast-charge protocols; improves search efficiency over random/naive optimization and guides interpretable protocol updates

A comparison of machine learning-optimized fast-charging strategies, their objectives, and experimentally validated performance gains.

Table 5. Control algorithms for battery management.

Decision variables (examples)		Typical objective(s)	Constraints handled	Pros	Cons
Control paradigm					
Bayesian Optimization (BO) (experiment in the loop or sim in the loop) ^{2,19}	Step currents/voltages, step durations (protocol design)	Minimize charge time or maximize cycle life	Can explicitly include degradation/temperature constraints in acquisition/objective	Sample-efficient search in large spaces; works without first-principles model	Requires a surrogate; outcomes depend on feature choice/early predictors
Deep reinforcement learning (e.g., DDPG) ²¹	Continuous current profile per cell; balancing currents	Minimize charge time and balancing cost; protect SOH	Voltage, temperature, SOC, SOH limits; thermal constraints	Learns adaptive, cell-aware policies; reduces total time and balancing time versus CC-CV	Policy training stability and safety constraints must be engineered carefully
Rule-based (CC-CV) (baseline) ^{19,21}	Fixed CC to voltage limit, then CV taper	Simplicity/repeatability; defacto baseline	Voltage/temperature limits via fixed set-points	Simple and robust; easy to implement	Can hit voltage limit early at high currents → slower overall charge; risk of polarization-induced plating in fast charge
Probabilistic ML controllers ²³	SOC ranges; current ramps	Avoid plating/thermal runaway	Plating probability, thermal gradients	Incorporates uncertainty; safer control	Requires calibrated probabilistic models; computational overhead for real-time inference
Hybrid BO + physics models ²⁰	Multi-step waveform parameters	Trade-off charge speed vs plating	Physical simulators + BO	High interpretability; better extrapolation to unseen protocols;	Accuracy limited by physics-model fidelity; BO performance still tied to surrogate modeling choices

A high-level comparison of control paradigms for battery charging and management, including decision variables, constraints, and advantages.

experimental datasets to adapt charging current in real time.²² Their policy accelerated charging to 80% state of charge compared with standard methods while respecting voltage and temperature bounds. Crucially, this study used real-world experimental data rather than purely simulated environments. It is important to note the practical consideration that any learned policy must be vetted for safety and would typically need to be simplified for implementation on battery management system hardware, which may involve model compression or simplified rule extraction from the ML policy. Regardless, the trend is clear: ML is enabling a shift from fixed, conservative charging rules to adaptive, optimized strategies tailored to each battery or use-case.

Formation protocol optimization

The formation step, or the initial charging cycles performed to “form” the solid-electrolyte interphase (SEI) on anode surfaces, is another critical process that influences cell longevity and performance. Formation can take weeks per cell under carefully tuned conditions and suboptimal protocols can lead to poor SEI and early failure. Recognizing this, Cui et al. undertook a data-driven structure of how varying the formation charge rate, temperature, and other factors affects long-term cycle life.³ They systematically tested dozens of formation protocols on 18,650 Li-ion cells and gathered an extensive dataset of post-formation characteristics and eventual cycle lives. Using statistical learning, they discovered an intriguing phenomenon: fast formation (at higher currents and temperature) caused a shift in electrode utilization that actually increased cycle life by up to 70%. Specifically, fast formation drove lithium inventory loss early, which shifted the cathodes’ usable voltage window away from the range that normally induced lithium plating on the anode during later cycling. In essence, a harsher formation sacrificed some capacity early, but pre-conditioned the cells to avoid the most damaging voltage regime during operation, enabling an extended lifetime. This counterintuitive insight, which was validated by differential voltage analysis and impedance measurements, reverts the conventional wisdom that gentler formation is always better. It showcases how data-driven analysis can reveal hidden levers to improve lifetime. As a result, one can envision entire factories adopting adaptive formation protocols informed by machine learning models, which would quickly screen a given cell and then choose a formation protocol optimized for that cell’s unique response.

Others are also merging ML with electrochemical knowledge to optimize battery usage. Jiang et al. applied Bayesian Optimization to discharge protocols and operating set-points for identifying cycling patterns that minimized degradation for grid storage applications.¹⁹ Likewise, Lu et al. showed that when future usage conditions are uncertain, a Recurrent Neural Network (RNN)-based algorithm can recommend a robust charging policy that hedges against worst-case degradation scenarios.²⁴ These examples illustrate the broader theme: ML is enabling a proactive approach to battery longevity in which we no longer only predict failure, but also actively extend time to failure through

intelligent control. Coupling life prediction models (“[Data-driven battery life prediction](#)” section) with optimization closes the loop. Predictions inform better protocols, and data from new protocols improve predictions.

Physics-informed and mechanistic ML

While purely data-driven models have proven their worth, there is increasing recognition that combining ML with physical knowledge can yield the best of both worlds: accurate predictions that respect known electrochemical behavior, and models that provide insight into why a battery fails, not just when. Physics-informed machine learning encompasses a variety of strategies, from adding physical constraints to learning algorithms, to using ML to analyze and interpret physical experiments. A hierarchy of these strategies and their requisite data and modeling requirements can be found in Fig. 4, and the main types of physics-informed constraints used in battery modeling can be found in Table 6.

Augmenting data-driven models

One class of work augments data-driven models with theoretical constraints or simulation data. For instance, Wang et al. demonstrated a physics-informed neural network (PINN) for battery degradation forecasting, where parts of the network architecture enforced known degradation kinetics.¹⁶ By training the neural network simultaneously on experimental cycling data and the governing differential equations for capacity fade, the model achieved stable long-horizon predictions and remained consistent with known physical laws such as enforcing a non-negative capacity, and reasonable SEI growth rates. Similarly, hybrid modeling approaches use output from physics-based battery models to inform or regularize ML models. Nascimento et al. built a hybrid model that combined a simplified electrochemical model with a neural network, improving extrapolation to unseen regimes.¹⁷ These approaches address a key issue with black-box ML: adding physical knowledge prevents models from violating conservation laws or providing unphysical results when pushed beyond the training domain. By injecting physics, whether as hard constraints, soft penalties, or extra training data from simulation, the ML models become more robust and interpretable. In fact, Navidi et al. recently compared several physics-informed ML methods, from differential equation informed networks to electrolyte feature-based models, for state of health prediction and found that these consistently outperformed purely empirical models, especially when training data were sparse.¹⁸ Table 7 outlines the specific physics-based loss terms and constraints commonly used to regularize these battery ML models.

Despite their advantages, though, it is important to note that physics-informed models may underperform purely data-driven approaches when the underlying physical assumptions are incomplete, poorly parameterized, or mismatched to the operating regime. In such cases, enforcing simplified physics can constrain model expressiveness and introduce systematic bias,




Figure 4. Mechanistic machine-learning hierarchy spanning atomic to electrode scales. At the top, atom-resolved models trained on quantum mechanical data predict lithium-ion diffusion barriers and migration pathways.²⁵ At the middle, imaging-based ML fits operando microscopy to reaction–diffusion physics, mapping local reaction rates and phase-boundary movement within individual particles.^{4,14} At the base, electrode-scale physics-informed neural networks learn degradation kinetics and lifetime trends from long-term cycling data.¹⁶

particularly under extreme conditions or emerging chemistries where mechanistic understanding remains limited.

High-dimensional experimental data

Another exciting avenue is using ML to glean insights from high-dimensional experimental data that traditional analysis struggles with. Modern battery diagnostics produce massive datasets, including high-resolution microscopy, spectroscopy, and tomography, that contain hidden information about failure mechanisms. Machine learning and computer vision (CV) are proving adept at extracting those insights. A striking example is the joint Stanford–MIT work by Zhao et al. in which researchers used computer vision techniques to analyze operando nanoscale X-ray microscopy videos of battery particles, learning the underlying reaction kinetics pixel by pixel.²⁶ In this study, researchers took 62 video sequences of LFP cathode particles charging and discharging, and applied a form of physics-constrained machine learning to infer local reaction rate constants across each particle. By training a model to fit the observed pixel intensity changes while obeying a phase-field reaction diffusion model, they could quantify how kinetics varied at different locations on a single particle. The learned kinetic heterogeneity map correlated strongly with the non-uniform carbon coating thickness on the particle, for which a thicker coating resulted in slower lithium insertion kinetics. This ML-driven analysis thus pinpointed a key reason for intraparticle heterogeneity and suggested that engineering more uniform electrode coatings could improve performance.

Researchers have similarly applied ML to connect micro-scale observation with macro-scale battery failure. Deng et al. used correlative image learning on X-ray and electron microscopy data to understand chemo-mechanical degradation

in battery electrodes.²⁷ By training a model on images of particle cracking and phase transformation, they linked certain microstructural patterns to faster capacity fade. In a different system, Park et al. combined theory and experimentation to explain anomalous “fictitious” phase separation in Ni-rich cathodes, identifying through data-driven modeling that an autocatalytic surface reaction was driving the apparent phase coexistence.⁴ The ability of ML to sift through complex datasets, whether they be images, spectroscopy, or multi-sensor time-series, and highlight meaningful correlations is accelerating the discovery of which physical processes dominate battery aging under various conditions. Table 8 summarizes these mechanistic ML frameworks across atomic, particle, electrode, and pack scales, and Table 9 outlines the diverse datasets used across physics-informed ML.

Materials discovery and design

Another domain where physics-informed ML is making strides is in materials discovery and design for improved battery longevity. Here, generative models and data-mining tools leverage large databases of known materials and hypothetical options to identify promising candidates. The Materials Project, a massive open database of computed properties for > 200,000 materials, exemplifies how data driven approaches can screen battery materials computationally.²⁵ The Materials Project team highlighted successes in using ML to recommend new cathode compositions and solid electrolytes by learning from Density Functional Theory (DFT) data and past experiments, significantly narrowing the search space for lab testing. Generative AI is also being explored; for instance, Datta et al. developed a generative adversarial model to design novel porous oxide structures for next-generation battery electrodes.²⁹ Such models can

Table 6. Types of physics-informed ML constraints.

Constraint type	Explanation	Typical implementation in batteries	Pros/caveats
Hard constraints ¹⁷	Physics is embedded directly into the model structure or state update; violations are (essentially) impossible	Governing ODEs embedded inside an RNN cell; battery voltage computed from Nernst + Butler–Volmer blocks with learnable non-ideal terms	High fidelity and interpretability; training more delicate
Soft constraints ¹⁸	Physics appears as penalty terms in the loss; the network is free to deviate but is nudged toward physical behavior	PINN loss combines data MSE with physics terms derived from a half-cell model (capacity Q, LLI consistency; peak position penalties on dQ/dV)	Flexible; needs weighting/regularization; may drift if penalties are weak
Hybrid (hard + soft) constraints ¹⁶	Part of the physics is baked into the model; additional physics appears in the loss or via surrogate mappings	Empirical state-space degradation equation enforced in a PINN (monotonicity + PDE loss) while learning the feature → SOH mapping	Good stability with limited data; still needs careful design

A classification of physics-informed constraints used in battery ML models, highlighting implementation styles and representative studies.

propose material compositions or nanostructures optimized for certain performance metrics, like high ionic conductivity with stability, which researchers can then synthesize and evaluate. While early in application to batteries, these AI-driven materials design efforts hint at a future in which ML not only interprets data from existing materials, but also suggests new materials less prone to failure.

Economic and societal perspectives

Advances in battery failure prediction and life extension have important economic and societal implications. After all, a battery that lasts twice as long can dramatically improve the total cost of ownership of an electric vehicle (EV) or the economics of grid storage, and vice versa, overly aggressive use that leads to early failure can negate sustainability gains. A key area to study is how ML-accelerated improvements will translate to the real world to ensure that technological breakthroughs align with market and user needs.

Charging optimization and cost

Techniques discussed in “Closed-loop optimization of protocols” and “Physics-informed and mechanistic ML” section, particularly closed-loop charging optimization and early failure prediction, have direct economic implications by reducing overdesign, extending usable lifetime, and enabling adaptive operation. Machine-learning-guided charging protocols can shift batteries away from conservative one-size-fits-all use toward condition-aware strategies, improving both performance and economic efficiency across electric vehicles and grid-scale storage systems.

One major topic is the second-life use of EV batteries. After an EV battery’s capacity fades to ~ 80% of its original value, it is typically retired from vehicular service, but it may still have years of useful life for less demanding applications, like stationary storage and backup power.³⁰ The crucial question is: Which used batteries are good candidates for this second-life reutilization, and under what conditions? Machine learning is helping answer this by analyzing field data and degradation patterns to both predict remaining useful life in second-life scenarios and optimize repurposing decisions. Takahashi et al. conducted a comprehensive techno-economic analysis using ML to evaluate the feasibility of various second-life strategies; by inputting battery health metrics and usage history into ML models, they could forecast whether a given spent EV battery would survive long enough in second-life applications to be profitable.³¹ For instance, their analysis indicated that batteries retired with relatively higher remaining capacity and those from less aggressive climates had a much better chance at economically viable second-life deployment. Other reviews echo these findings: Iqbal et al. surveyed second-life battery projects and noted that accurate state-of-health diagnostics, a

Table 7. Physics-informed loss functions and constraints used in battery ML.

Constraint category	Example mathematical form	Physical law enforced	Use-case
Mass conservation term ^{16, 18}	$\ \nabla \cdot (J) + \partial C / \partial t\ ^2$ penalty	Charge/species conservation	Electrode-scale degradation, diffusion
Monotonic capacity constraint ¹⁶	$\lambda \cdot \text{ReLU}(-dQ/dt)$	Ensures non-negative SEI growth/ no spontaneous capacity increase	SOH trajectory prediction
Thermodynamic consistency ¹⁷	$\ \Delta \mu_{\text{pred}} - \Delta \mu_{\text{phys}}\ ^2$	Aligns learned chemical potentials with Nernst relation	Voltage prediction, reaction energetics
Smoothness/continuity ^{26, 27}	$\ \nabla^2 Q_{\text{pred}}\ ^2$ regularizer	Suppresses unphysical discontinuities in learned fields	Operando mapping, imaging inversion
Hybrid empirical + PDE loss ^{16, 18}	$L = \alpha L_{\text{data}} + \beta L_{\text{phys}} + \gamma L_{\text{reg}}$	General composite objective combining experiment + theory	Physics-informed neural networks (PINNs)

A summary of common physics-based loss terms used in battery ML models and the physical laws they enforce.

Table 8. Multiscale mechanistic ML frameworks.

Scale	Typical model/data input	ML integration	Key physical process captured
Atomic/DFT ²⁵	Local charge density $\rho(r)$, potential $V(r)$	Graph NNs trained on DFT data to predict Li-diffusion barriers	Ion migration, defect formation
Particle/mesoscale ^{4, 26}	Phase-field variables (ϕ , c) from STXM or TEM	Physics-constrained CNNs estimating local k_{react} and D_{eff}	Phase transformation, autocatalytic delithiation
Electrode/cell ^{16, 17}	$dQ/dV(V)$, impedance $Z(\omega)$, $T(t)$	PINNs/hybrid RNNs with electrochemical ODEs	Capacity fade, SEI growth
Pack/system ⁵	BMS histogram features (voltage, current distributions)	ML models on aggregated field telemetry	ML models on aggregated field telemetry

Mechanistic machine-learning approaches across atomic, particle, electrode, and pack scales.

common application area of battery ML, are crucial for identifying which batteries can be reused safely and economically.³² Ultimately, by predicting failure and remaining life more precisely, machine learning can better enable warranty pricing or performance guarantees for repurposed batteries, incentivizing and reducing risk for second-life business models.

Failure prediction and risk

Early-life prognostics and failure prediction methods reviewed in “[Data-driven battery life prediction](#)” section enable condition-based valuation of batteries across both first- and second-life applications. By identifying high-risk cells before

catastrophic failure, these approaches reduce safety-related costs, support more accurate warranty pricing, and better align economic incentives with actual degradation behavior.

ML-driven insights on battery longevity also feed into techno-economic models for new battery technologies. For example, Hasselwander et al. compared the cost and performance of different Li-ion chemistries (NMC vs. LFP vs. others) for passenger EVs, incorporating not just upfront cost and energy density, but also cycle life and fast-charge capability.³³ Such analyses help quantify how much value is unlocked by extending lifespan or enabling faster charging; if an ML-optimized charging protocol projects a 20% increase on lifetime for a particular EV battery, techno-economic models can translate that into dollar

Table 9. Data sources and physics integration pathways.

Data type	Example datasets	Physical law integration mode	Typical ML architecture	Benefit/risk
Electrochemical cycling data ^{6, 16, 17}	MIT-Stanford fast-charge sets, NASA B0005–B0018	Differential equations for diffusion and reaction kinetics as hard/soft constraints	PINN/hybrid RNN	Better extrapolation and physical interpretability; requires parameter identifiability
Operando imaging data ^{4, 26, 27}	STXM, TEM, ptychography datasets	Phase-field or reaction–diffusion constraints in CNN training	Physics-guided UNet/CNN	Converts images to quantitative kinetics maps; data volume large
Spectroscopic data ¹⁸	Raman, XANES, EXAFS operando sets	Beer–Lambert/energy-shift laws as loss terms	Spectral PINN/autoencoder	Denser spectra → noise robust fitting; still needs cross-calibration
Simulated physics data ^{16–18, 28}	DFT/PyBaMM/phase-field outputs	Multi-fidelity training (co-kriging, transfer learning)	Hybrid surrogates and physics-guided NNS	Combines low- and high-fidelity data; domain-gap risk

A summary of the primary data sources used in physics-integrated ML and the modes by which physical laws are incorporated.

savings per kWh, reduction in carbon emissions per vehicle, or other quantitative figures that better incentivize the adoption of research findings. Roy et al. took this further by examining hybrid renewable energy systems and modeling how battery life improvements alter the optimal system configurations and cost, finding that in scenarios like solar-plus-battery microgrids, using batteries with longer cycle life significantly lowered the leveled cost of energy and reduced the carbon footprint of the system.³⁴ These studies underscore that the impact of machine learning in batteries is not just technical: it propagates to energy economics, climate outcomes, and the society we live in.

Open science efforts and community resources

Battery research has historically been hampered by proprietary data and partitioned research efforts. In the past several years, however, there has been a strong push toward open science in the battery community: sharing data, models, and software to accelerate processes for all. Machine learning thrives on data, and community-accessible datasets have been key to recent breakthroughs. Here, we survey notable open science resources and initiatives that are propelling battery ML.

Open datasets

Large, well-annotated battery cycling datasets are invaluable for developing and benchmarking ML models. One landmark was the publication by Harlow et al. of an extensive dataset of over 20,000 cycles of high-quality NMC/graphite cells, which have been used in numerous ML studies as a training or validation set.³⁵ Likewise, the NASA battery data repositories, which contain life tests on cells under various conditions, have enabled comparative studies. More recently, Luh and Blank released a comprehensive aging dataset for a modern NMC/SiO anode cell, including both capacity fade and impedance evolution over time, which opens the door for ML models that incorporate electrical signatures alongside capacity curves.³⁶ Another effort, the Battery Data Genome initiative, is attempting to standardize data formats to make shared data more usable.³⁷ The Chueh Lab at Stanford/SLAC has also made multiple datasets public *via* its website, including the complete datasets from the Severson and Attia studies. These include formation datasets, fast-charging cycling data, and more, often accompanied by code to load and analyze these datasets. Such sharing has encouraged reproducibility and spurred new work: for instance, researchers in Europe used the Stanford open dataset to test novel algorithms for knee-point detection in capacity curves.³⁸ Overarchingly, a key lesson learned from the open-data efforts thus far is the need for context: a cycle number is only useful if the exact testing conditions, cell specifications, etc. are also known, so it is recommended that metadata and even raw time-series records also be included in such datasets.

For emerging battery chemistries, though, data scarcity and noise present fundamental challenges to model generalization.

Promising strategies include transfer learning from well-characterized lithium-ion systems, physics-guided regularization to constrain extrapolation, and uncertainty-aware learning frameworks that explicitly encode confidence under sparse data regimes.

Open software and tools

Alongside data, having open code accelerates progress by allowing researchers to build on each other's methods. One prominent tool is Python Battery Mathematical Modeling (PyBaMM), an open-source library for simulating electrochemical battery models.²⁸ PyBaMM provides a flexible platform to run physics-based models, like the Doyle-Fuller-Newman model, and even perform parameter calibration and sensitivity analysis. It is becoming a common baseline to compare against when integrating ML with physics; for instance, one might use PyBaMM to generate synthetic data to train an ML model, or vice versa, use ML to accelerate parts of the simulation. On the purely ML side, the open-source BatteryML toolkit was released to provide standardized routines for common tasks like feature extraction, dataset splitting, and model evaluation specific to battery data.³⁹ This toolkit lowers the entry barrier for new researchers in battery ML, handling data preprocessing and even including some published models ready to apply on new data.

Community competition

Competitions and community challenges have also played a role in open science. The Materials Research Society (MRS) Spring 2023 Meeting featured a public Kaggle competition on battery life prediction, where participants worldwide built models to predict cycle life from voltage curve data.⁴⁰ Such challenges not only yield innovative solutions but also result in open-source code and discussion that benefit the field at large. There has also been the emergence of open-access benchmarks: Schauser et al. introduced an open tool to systematically evaluate different ML model types on a common battery dataset, providing a fairness comparison and illuminating which models tend to perform best under various training set sizes.⁴¹ This kind of work helps avoid “researcher bias” where each group reports only their model's performance on their data; an open benchmark enables apples-to-apples comparisons that can better guide newcomers on selecting an approach.

Conclusion

The marriage of machine learning and battery science is still in its early chapters, but progress thus far demonstrates that with data and algorithms we can predict and even prevent failures that once seemed unpredictable. The societal implications of this interdisciplinary approach are profound. Predictive lifetime modeling informs not only design but also circular economy

logistics: when and how a cell should be repurposed or recycled. Coupling these insights with policy incentives for open data will ensure that advances in intelligence align with sustainability. The next chapters will involve broadening these successes to all kinds of batteries and embedding them in the devices and grids that rely on this technology.

A pressing challenge, however, is generalization. Models trained on narrowly defined chemistries often fail when confronted with new materials, structures, protocols, or environmental conditions. To address this, transfer learning and multimodal datasets, such as combining electrochemical, imaging, and spectroscopic data, are essential. Equally transformative will be the integration of ML into experimental control. Bayesian optimization and reinforcement learning frameworks now guide closed-loop cycling experiments, autonomously tuning current densities, temperatures, and cut-off voltages to maximize lifetime or minimize formation time. Such systems dramatically reduce the empirical trial space, democratizing advanced research across institutions. Adapting these complex models to perform on edge computing devices will clear the hurdles for truly automated closed-loop systems.

Looking forward, several research priorities must be addressed for machine learning-driven battery science to transition from promise to practice. First, robustness to real-world noise and protocol variability must become a primary evaluation metric, rather than an afterthought. Second, standardized data schemas and interoperable machine learning pipelines are essential to enable reproducibility and cross-institutional collaboration. Third, generalization across chemistries will require hybrid approaches that combine transferable physical constraints with data-driven adaptability.

Ultimately, the success of this emerging paradigm will depend not only on algorithmic advances, but also on the co-design of experiments, models, and deployment pathways that align scientific insight with economic and societal impact.

Author contributions

Conceptualization: All authors contributed to manuscript conceptualization. Literature review and methodology design: S.V. and J.N. Data curation and analysis: S.V. Writing—original draft preparation: S.V. Writing—review and editing: all authors contributed to manuscript review and editing. Supervision: YK.T. and J.N. All authors have read and approved the final manuscript.

Funding

The authors would like to acknowledge support from the Assistant Secretary for Energy Efficiency and Renewable Energy (EERE), Office of Vehicle Technologies (VTO) of the U.S. Department of Energy (DOE), under the Battery Materials Research (BMR) Program and the Battery 500 Consortium. This work was also supported by the Knight-Hennessy Scholars Program, the Quad Fellowship, and the Stanford EDGE Fellowship.

Declarations

Conflict of interest

On behalf of all authors, the corresponding authors state that there is no conflict of interest.

Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

REFERENCES

1. M. Dean, S. Man, T. Yoshida, J. Chen, G. Reghelin, P. Lau, G. Davis, C. McKerracher, *Global Battery Electric Vehicle Outlook 2024: Survival of the Fittest* (Bloomberg Intelligence, Bloomberg, 2024)
2. P.M. Attia, A. Grover, N. Jin, K.A. Severson, T.M. Markov, Y.-H. Liao, M.H. Chen, B. Cheong, N. Perkins, Z. Yang, P.K. Herring, M. Aykol, S.J. Harris, R.D. Braatz, S. Ermon, W.C. Chueh, Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature* **578**, 397–402 (2020). <https://doi.org/10.1038/s41586-020-1994-5>
3. X. Cui, S.D. Kang, S. Wang, J.A. Rose, H. Lian, A. Geslin, S.B. Torrisi, M.Z. Bazant, S. Sun, W.C. Chueh, Data-driven analysis of battery formation reveals the role of electrode utilization in extending cycle life. *Joule* **8**, 3072–3087 (2024)
4. J. Park, H. Zhao, S.D. Kang, K. Lim, C.-C. Chen, Y.-S. Yu, R.D. Braatz, D.A. Shapiro, J. Hong, M.F. Toney, Fictitious phase separation in Li layered oxides driven by electro-autocatalysis. *Nat. Mater.* **20**, 991–999 (2021)
5. G. Pozzato, A. Allam, L. Pulvirenti, G.A. Negoita, W.A. Paxton, S. Onori, Analysis and key findings from real-world electric vehicle field data. *Joule* **7**, 2035–2053 (2023)
6. K.A. Severson, P.M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M.H. Chen, M. Aykol, P.K. Herring, D. Fraggadakis, M.Z. Bazant, S.J. Harris, W.C. Chueh, R.D. Braatz, Data-driven prediction of battery cycle life before capacity degradation. *Nat. Energy* **4**, 383–391 (2019). <https://doi.org/10.1038/s41560-019-0356-8>
7. X. Cui, S. Sun, W.C. Chueh, Evaluating and interpreting the predictive power of features in battery lifetime prediction. *Electrochem. Soc. Interface* **34**, 49 (2025). <https://doi.org/10.1149/2.F102521F>
8. Y. Fan, Z. Lin, F. Wang, J. Zhang, A hybrid approach for lithium-ion battery remaining useful life prediction using signal decomposition and machine learning. *Sci. Rep.* **15**, 8161 (2025)
9. J. Zhou, W. Huang, H. Dai, C. Wang, Y. Zhong, Data-driven battery remaining life prediction based on ResNet with GA optimization. *World Electr. Veh. J.* **16**, 267 (2025)
10. K. Li, X. Chen, Machine learning-based lithium battery state of health prediction research. *Appl. Sci.* **15**, 516 (2025)
11. S. Rout, S.K. Samal, D.J. Gelmecha, S. Mishra, Estimation of state of health for lithium-ion batteries using advanced data-driven techniques. *Sci. Rep.* **15**, 30438 (2025)
12. H. Zhang, Y. Li, S. Zheng, Z. Lu, X. Gui, W. Xu, J. Bian, Battery lifetime prediction across diverse ageing conditions with inter-cell deep learning. *Nat. Mach. Intell.* **7**, 270–277 (2025)
13. A. Geslin, L. Xu, D. Ganapathi, K. Moy, W.C. Chueh, S. Onori, Dynamic cycling enhances battery lifetime. *Nat. Energy* **10**, 172–180 (2025). <https://doi.org/10.1038/s41560-024-01675-8>
14. Y. Zhang, D. Han, R. Xiong, Feature-driven closed-loop optimization for battery fast charging design with machine learning. *J. Electrochem. Soc.* **170**, 060508 (2023)
15. R.R. Richardson, M.A. Osborne, D.A. Howey, Gaussian process regression for forecasting battery state of health. *J. Power. Sources* **357**, 209–219 (2017). <https://doi.org/10.1016/j.jpowsour.2017.05.004>
16. F. Wang, Z. Zhai, Z. Zhao, Y. Di, X. Chen, Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis. *Nat. Commun.* **15**, 4332 (2024)
17. R.G. Nascimento, M. Corbetta, C.S. Kulkarni, F.A. Viana, Hybrid physics-informed neural networks for lithium-ion battery modeling and prognosis. *J. Power. Sources* **513**, 230526 (2021)
18. S. Navidi, A. Thelen, T. Li, C. Hu, Physics-informed machine learning for battery degradation diagnostics: a comparison of state-of-the-art methods. *Energy Storage Mater.* **68**, 103343 (2024)
19. B. Jiang, M.D. Berliner, K. Lai, P.A. Asinger, H. Zhao, P.K. Herring, M.Z. Bazant, R.D. Braatz, Fast charging design for Lithium-ion batteries via Bayesian optimization. *Appl. Energy* **307**, 118244 (2022)
20. B. Jiang, Y. Wang, Z. Ma, Q. Lu, Fast charging of lithium-ion batteries using deep Bayesian optimization with recurrent neural network. *IEEE Trans. Autom. Sci. Eng.* **22**, 13150–13160 (2025)
21. Z. Zhang, T. Guo, Y. Liu, X. Pang, Z. Zheng, Fast-charging optimization method for Lithium-Ion battery packs based on deep deterministic policy gradient algorithm. *Batteries* **11**, 199 (2025)
22. J. He, T. Yang, L. Xie, Y. Yang, C. Chen, J. Wei, A data-driven reinforcement learning enabled battery fast charging optimization using real-world experimental data. *IEEE Trans. Ind. Electron.* **72**, 430–438 (2024)
23. A. Thelen, X. Huan, N. Paulson, S. Onori, Z. Hu, C. Hu, Probabilistic machine learning for battery health diagnostics and prognostics—review and perspectives. *NPJ Mater. Sustain.* **2**, 14 (2024)
24. J. Lu, R. Xiong, J. Tian, C. Wang, C.-W. Hsu, N.-T. Tsou, F. Sun, J. Li, Battery degradation prediction against uncertain future conditions with recurrent neural network enabled deep learning. *Energy Storage Mater.* **50**, 139–151 (2022)
25. M.K. Horton, P. Huck, R.X. Yang, J.M. Munro, S. Dwaraknath, A.M. Ganose, R.S. Kingsbury, M. Wen, J.X. Shen, T.S. Mathis, Accelerated data-driven materials science with the materials project. *Nat. Mater.* **2025**, 1–11 (2025)
26. H. Zhao, H.D. Deng, A.E. Cohen, J. Lim, Y. Li, D. Fraggadakis, B. Jiang, B.D. Storey, W.C. Chueh, R.D. Braatz, Learning heterogeneous reaction kinetics from X-ray videos pixel by pixel. *Nature* **621**, 289–294 (2023)
27. H.D. Deng, H. Zhao, N. Jin, L. Hughes, B.H. Savitzky, C. Ophus, D. Fraggadakis, A. Borbély, Y.-S. Yu, E.G. Lomeli, Correlative image learning of chemo-mechanics in phase-transforming solids. *Nat. Mater.* **21**, 547–554 (2022)
28. V. Sulzer, S.G. Marquis, R. Timms, M. Robinson, S.J. Chapman, Python battery mathematical modelling (PyBaMM). *J. Open Res. Softw.* (2021). <https://doi.org/10.5334/jors.309>
29. J. Datta, A. Nadimpally, N. Koratkar, D. Datta, Generative AI for discovering porous oxide materials for next-generation energy storage. *Cell Rep. Phys. Sci.* (2025). <https://doi.org/10.1016/j.xcrp.2025.102665>

30. M.K. Al-Alawi, J. Cugley, H. Hassanin, Techno-economic feasibility of retired electric-vehicle batteries repurpose/reuse in second-life applications: a systematic review. *Energy Clim. Change* **3**, 100086 (2022)
31. A. Takahashi, A. Allam, S. Onori, Evaluating the feasibility of batteries for second-life applications using machine learning. *Iscience* (2023). <https://doi.org/10.1016/j.isci.2023.106547>
32. H. Iqbal, S. Sarwar, D. Kirli, J.K. Shek, A.E. Kiprakis, A survey of second-life batteries based on techno-economic perspective and applications-based analysis. *Carbon Neutral*. **2**, 8 (2023)
33. S. Hasselwander, M. Meyer, I. Österle, Techno-economic analysis of different battery cell chemistries for the passenger vehicle market. *Batteries* **9**, 379 (2023)
34. D. Roy, S. Zhu, R. Wang, P. Mondal, J. Ling-Chin, A.P. Roskilly, Techno-economic and environmental analyses of hybrid renewable energy systems for a remote location employing machine learning models. *Appl. Energy* **361**, 122884 (2024)
35. J.E. Harlow, X. Ma, J. Li, E. Logan, Y. Liu, N. Zhang, L. Ma, S.L. Glazier, M.M. Cormier, M. Genovese, A wide range of testing results on an excellent lithium-ion cell chemistry to be used as benchmarks for new battery technologies. *J. Electrochem. Soc.* **166**, A3031 (2019)
36. M. Luh, T. Blank, Comprehensive battery aging dataset: capacity and impedance fade measurements of a lithium-ion NMC/C-SiO cell. *Sci. Data* **11**, 1004 (2024)
37. L. Ward, S. Babinec, E.J. Dufek, D.A. Howey, V. Viswanathan, M. Aykol, D.A. Beck, B. Blaiszik, B.-R. Chen, G. Crabtree, Principles of the battery data genome. *Joule* **6**, 2253–2271 (2022)
38. P. Fermín-Cueto, E. McTurk, M. Allerhand, E. Medina-Lopez, M.F. Anjos, J. Sylvester, G. Dos Reis, Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells. *Energy AI* **1**, 100006 (2020)
39. H. Zhang, X. Gui, S. Zheng, Z. Lu, Y. Li, J. Bian, BatteryML: An open-source platform for machine learning on battery degradation (2023), Preprint at <https://arxiv.org/abs/2310.14714>
40. S. Sun, *MRS Spring 2023—Battery Prediction Challenge* (Kaggle, 2023)
41. N.S. Schausser, C.N. Lininger, E.S. Leland, T.Z. Sholklapper, An open access tool for exploring machine learning model choice for battery life cycle prediction. *Front. Energy Res.* **10**, 1058999 (2022)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.