# The Evolution of Battery Lifecycle Prediction: From Physics to AI

Stanley Xie
Grunuss R&D Team

## A brief history of how we learned to predict when batteries die — and why we're still not quite there yet

*"Every battery is a ticking clock. You just don't know when it's going to stop."*

If you've ever owned a smartphone, you know the feeling. The battery that used to last all day now barely makes it through the afternoon. You check the settings, you see the percentage, and you wonder: *how much longer?* For most of us, the answer is a shrug. We just replace the phone, or the battery, when it gets too annoying.

Now imagine that phone is a $100,000 electric vehicle. Or a grid-scale battery system the size of a shipping container storing energy for an entire neighbourhood. Suddenly, *"it's getting a bit low"* isn't a minor inconvenience. It's a multi-million dollar question. *When will this battery fail? How many cycles do we have left? Can we sell this for second-life grid storage, or is it ready for recycling?*

This is the problem of **battery lifecycle prediction**. And it's been surprisingly difficult to solve.

In this article, I want to walk you through how we've tried to answer this question over the past few decades — from the early days of physics-based models, to the explosion of machine learning, to where we are today. And I want to explain why neither approach alone is enough.

### A Quick Note on Terms

Before we dive in, let's define a few key terms that will appear throughout this article.

| Term | Definition |
|---|---|
| RUL | Remaining Useful Life — the estimated number of cycles or time until the battery reaches end-of-life |
| SOH | State of Health — a metric indicating the battery's current capacity relative to its original capacity |
| DFT | Density Functional Theory — a quantum mechanical method for calculating electronic structure of materials |
| DFN | Doyle-Fuller-Newman model — a pseudo-2D electrochemical model for lithium-ion batteries |
| PINN | Physics-Informed Neural Network — a neural network with physical equations embedded in the loss function |

| Term | Definition |
|---|---|
| ECM | Equivalent Circuit Model — a simplified electrical representation of battery behavior |
| SEI | Solid-Electrolyte Interphase — a passivation layer that forms on the anode and consumes lithium over time |

## A Brief History of Batteries: From Curiosity to Cornerstone

Before we dive into the technical challenge of battery lifecycle prediction, it's worth taking a step back. Because the story of how we got here — how batteries transformed from a laboratory curiosity into the foundation of modern energy systems — helps explain why we care so much about predicting when they'll stop working.

The history of the battery begins long before lithium-ion. In 1800, Alessandro Volta stacked copper and zinc discs separated by brine-soaked cardboard to create the first device that could produce a continuous electrical current [14]. The "voltaic pile" was a sensation — it provided the first practical source of steady electricity and established the electrochemical principles that underpin all batteries to this day [15].

What followed was a series of incremental but important advances. In 1859, Gaston Planté invented the lead-acid battery, the first rechargeable storage battery — a technology still used in cars today [10]. In the 1860s, Georges Leclanché developed the zinc-manganese dioxide system that would evolve into the modern dry cell [21]. By 1896, the Columbia Dry Cell had become the first battery intended for widespread consumer use [13], and in 1899, Waldemar Jungner invented the nickel-cadmium rechargeable battery — the first truly portable rechargeable power source [10].

For most of the 20th century, batteries were the quiet workhorses of the economy: essential, but hardly transformative [18]. They started cars, powered radios, and provided backup when the grid failed. But they were also seen as a constraint rather than an enabler — devices were designed around their limitations: short runtimes, heavy weight, and slow charging [19].

That changed in 1991 with the commercial introduction of the lithium-ion battery [12]. Building on decades of research, including the foundational work of John Goodenough, Stanley Whittingham, and Akira Yoshino (who would later share the 2019 Nobel Prize in Chemistry), Li-ion offered a step-change in energy density, rechargeability, and flexibility [14]. Suddenly, portable electronics were no longer tethered to power outlets.

Over the following decade, lithium-ion batteries quietly powered the rise of laptops, mobile phones, and smartphones. By the early 2000s, batteries were no longer just supporting devices — they were enabling entirely new product categories [20]. And as global demand for consumer electronics surged, battery manufacturing scaled dramatically, particularly in East Asia. With scale came learning: manufacturing processes improved, yields increased, and costs began to fall — over roughly a decade, lithium-ion battery prices dropped by close to 80–90%, one of the steepest cost declines of any major industrial technology [16].

That cost curve proved decisive. Cheaper batteries didn't just expand existing markets; they created entirely new ones [17]. Electric vehicles turned the battery from a component into the defining feature of the product, effectively replacing the internal combustion engine as the heart of the car. Governments introduced incentives and emissions rules that accelerated adoption. Demand for raw materials surged, pulling batteries into the realm of geopolitics — supply chains became

strategic assets, and control over battery materials began to carry implications for energy security and industrial competitiveness [10].

The next frontier has been the power sector. As wind and solar generation expand, the variability of renewable energy has created a growing need for flexibility. Batteries have emerged as one of the most effective tools to provide it [18]. Utility-scale battery storage projects are now being deployed to balance supply and demand, provide grid stability, and store excess renewable generation for later use. In many markets, battery systems are now cost-competitive with traditional gas-fired peaking plants — a shift that would have been unimaginable a decade ago [16].

Today, batteries are at the center of a global industrial shift, underpinning everything from smartphones and electric vehicles to power grids increasingly reliant on renewable energy [11]. They have moved from niche component to strategic foundation — one of the defining technologies of the transition away from fossil fuels.

But with that transformation comes an urgent new question: as batteries become more central to our energy infrastructure, how do we know when they're going to fail?

That's where our story begins.

## The Early Days: Physics-Based Models

Before machine learning was a thing, if you wanted to understand a battery, you had to model it from first principles. You wrote down the physics. You solved the equations. You tried to describe *exactly* what was happening inside the cell.

This gave us models like the **pseudo-2D Newman model** (Doyle-Fuller-Newman or DFN) — a detailed electrochemical framework that simulates ion transport, intercalation kinetics, and degradation mechanisms [1, 2]. It's beautiful physics. It's also extremely complicated.

The DFN model solves a system of coupled partial differential equations:

- Conservation of charge in solid and electrolyte phases
- Conservation of species in solid and electrolyte phases
- Butler-Volmer kinetics for lithium intercalation
- Volume-averaged transport in porous electrodes

$$\frac{\partial c_s}{\partial t} = D_s \left( \frac{\partial^2 c_s}{\partial r^2} + \frac{2}{r} \frac{\partial c_s}{\partial r} \right) \tag{1}$$

where $c_s$ is the solid-phase lithium concentration, $D_s$ is the solid-phase diffusion coefficient, and $r$ is the radial coordinate.

These models work well in theory. But in practice, they require information you rarely have. What's the diffusion coefficient of that specific electrolyte? What's the exchange current density? What's the exact SEI growth rate? These parameters are difficult to measure, and they change with temperature, age, and chemistry.

You end up with a model that's physically accurate *in principle* — but practically limited by the data you can feed it. It's like having a perfect recipe that requires ingredients you can't find.

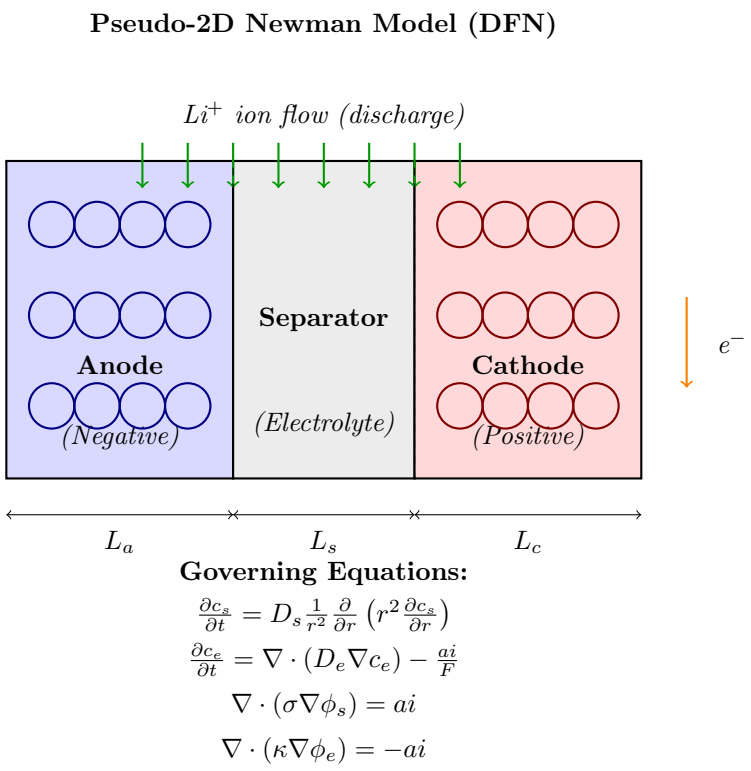


Figure 1: Simplified schematic of the pseudo-2D Newman (DFN) model showing electrode domains, separator, lithium-ion flow, and governing equations

## The Data Revolution: Machine Learning Arrives

Around the time that public battery datasets started becoming available (CALCE [3], NASA [4], Oxford [5]), something shifted. If we couldn't model the physics perfectly, maybe we could just learn the patterns from data.

Machine learning models — LSTM [6], CNN [7], Transformer architectures [8] — started appearing in battery research. The idea was simple: feed the model voltage, current, and temperature measurements, and let it figure out the relationship to cycle life. No physics required. Just data.

And it worked, up to a point. These models could sometimes match, or even exceed, the accuracy of physics-based models. They were faster, too. You didn't need to know the diffusion coefficient. You just needed a lot of cycling data.

But there was a catch.

Machine learning models are black boxes. They don't explain *why* they make a particular prediction. And when you show them a battery chemistry they haven't seen before, or a new cycling protocol, they often fail spectacularly. They've memorized patterns from training data, but they haven't learned the underlying physics. It's like teaching someone to drive by memorizing a map — they'll navigate that city just fine, but drop them in a new city and they're lost.

| Model | Strengths | Limitations |
|---|---|---|
| Equivalent Circuit Models (ECM) | Fast, interpretable, low-dimensional | Limited accuracy, empirical, extrapolates poorly |
| Pseudo-2D Newman (DFN) | Physically accurate, interpretable | Computationally expensive, requires detailed parameters |
| LSTM / RNN | Captures temporal dependencies | Requires large datasets, black-box |
| CNN / Transformer | Learns complex patterns | Black-box, requires extensive data |
| Gaussian Process Regression | Provides uncertainty estimates | Computationally expensive for large datasets |

### The Gap: Why Neither Approach Alone Is Enough

So where does that leave us?

On one hand, we have physics models that are *accurate but impractical*. They require too many parameters, take too long to run, and struggle with real-world variability.

On the other hand, we have data-driven models that are *practical but fragile*. They learn patterns from data but can't generalize beyond what they've seen.

The gap between these two worlds has been the central challenge in battery lifecycle prediction for years. We've had two competing toolkits, and we've been trying to solve the same problem with only half the tools.

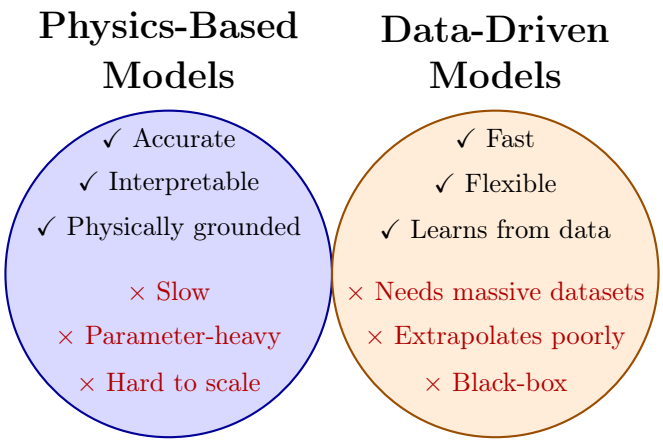


Figure 2: The gap between physics-based and data-driven approaches, with physics-informed machine learning emerging as the synthesis

### The Emerging Synthesis: Physics-Informed Machine Learning

This is where things get interesting. Over the past few years, a new approach has emerged — one that tries to combine the best of both worlds.

It's called **physics-informed machine learning** [9]. The idea is simple: take a neural network, but embed physical equations (like degradation kinetics or ion transport) directly into the loss function. The model still learns from data, but it also has to satisfy the laws of physics.

$$\mathcal{L} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}} \tag{2}$$

where $\mathcal{L}_{\text{data}}$ is the standard data-driven loss (e.g., mean squared error) and $\mathcal{L}_{\text{physics}}$ enforces the governing equations:

$$\mathcal{L}_{\text{physics}} = \frac{1}{N} \sum_{i=1}^{N} \left\| \frac{\partial \hat{u}}{\partial t} - \mathcal{N}[\hat{u}] \right\|^2 \tag{3}$$

where $\hat{u}$ is the neural network approximation and $\mathcal{N}$ is the differential operator.

This changes everything.

Suddenly, you have a model that's fast like ML, but consistent like physics. It can extrapolate to new conditions because it's grounded in physical equations. It can handle sparse data because the physics constraints act as a regularizer. And it remains interpretable because the physics isn't hidden — it's baked into the model architecture.

| **Approach** | **Strengths** | **Limitations** |
|---|---|---|
| Pure Physics | Accurate, interpretable, physically consistent | Slow, requires detailed parameters, hard to scale |
| Pure Data | Fast, flexible, learns from data | Black-box, extrapolates poorly, needs massive datasets |
| Physics-Informed ML | Combines speed with physical consistency, data-efficient | Requires physics knowledge to set up, still emerging |

## What This Means for the Future

This is the approach we're building at Grunuss with PaaS4Bat. We're not trying to replace physics with ML. We're trying to bring them together.

The implications are profound. If we can predict battery lifetime from only the first 10–20% of cycling data, we can compress development cycles from years to months. If we can quantify uncertainty alongside prediction, we can make better decisions about warranty, second-life valuation, and recycling. If we can build models that generalize across chemistries, we can accelerate the discovery of new battery materials.

It's not quite there yet. But we're close.

Battery lifecycle prediction has come a long way — from the drawing board of electrochemistry, through the data deluge of machine learning, to the emerging synthesis of physics-informed AI. The evolution tells a story: that the tools alone aren't enough. It's the integration of physics and data, of understanding and prediction, that finally starts to unlock the answer to that simple question:

*When will this battery stop working?*

### Key Takeaways

- Battery lifecycle prediction has evolved through three distinct phases: physics-based models, data-driven ML, and now physics-informed ML
- Pure physics models are accurate but impractical for real-world applications
- Pure data-driven models are fast but fragile when extrapolating to new conditions
- Physics-informed ML combines the strengths of both approaches by embedding physical equations directly into neural network loss functions
- The future of battery lifetime prediction lies in hybrid approaches that leverage both physics knowledge and data

### References

## References


[1] Newman, J. & Tiedemann, W. *Porous-electrode theory with battery applications*. AIChE Journal, 1975.

[2] Doyle, M., Fuller, T. F. & Newman, J. *Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell*. J. Electrochem. Soc., 1993.

[3] University of Maryland, CALCE Battery Dataset. `https://calce.umd.edu/battery-data`

[4] NASA Prognostics Center, Randomized Battery Usage Dataset. `https://data.nasa.gov/`

[5] Oxford Battery Dataset, University of Oxford.

[6] Zhang, Y. et al. *LSTM-based battery health prediction*. IEEE Access, 2020.

[7] Li, P. et al. *CNN for battery SOH estimation*. J. Power Sources, 2021.

[8] Kim, M. et al. *Transformer-based battery degradation prediction*. Batteries & Supercaps, 2023.

[9] Raissi, M., Perdikaris, P. & Karniadakis, G. E. *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems*. J. Comput. Phys., 2019.

[10] Britannica, "Battery - Rechargeable, Storage, Power."

[11] Reuters, "From curiosity to cornerstone - how batteries went mainstream."

[12] Taylor & Francis, "Fundamentals of Electric Energy Storage Systems."

[13] American Chemical Society, "Columbia Dry Cell Battery - Landmark."

[14] De Gruyter, "From Volta's pile to lithium ion battery: 200 years of energy."

[15] Collège de France, "History and development of battery technologies."

[16] UBS, "Lithium-ion batteries: powering the future."

[17] Taylor & Francis, "Treatment of Battery Manufacturing Wastes."

[18] KAUST, "The history and evolution of batteries."

[19] The Daily Star, "From curiosity to cornerstone - how batteries went mainstream."

[20] BBC, "The evolution of battery tech."

[21] Smithsonian Institution, "Unplugged: Five batteries that gave the world a jolt."

*This is the first article in a series exploring battery physics, AI, and simulation at Grunuss. Follow along as we dive deeper into DFT, DFN, PyBaMM, PINNs, and what it takes to build a platform that connects physics to prediction.*