Battery AI

A Self-Learning System
for the Next Stage of Global Electrification

AI and machine learning are all the rage. But for us, they aren't some bright shiny objects. Advanced, sophisticated self-learning (or machine learning) capabilities were baked into the Tanktwo Battery Operating System (TBOS) from the beginning as one of the pillars (along with Battery Security) on which we build our advanced battery management solutions.

Our Battery AI technology goes beyond making batteries work more efficiently. It has a much broader implication as the foundation of our data-driven asset optimization capabilities. It will enable the application of new economic models to make various power and electrification services practical and viable from a financial and operational perspective.

The logic behind the inner workings of our machine-learning algorithm is deceptively simple. Of course, we've spent years refining the capability and fine-tuning the code. But the basics aren't too different from forecasting weather and predicting climate trends. Let's look under the hood.

Battery AI: The inner workings of our self-learning algorithm

Our machine learning algorithm is trained in three stages to provide increasingly accurate predictions on individual cell behaviors over time. Here's an overview:

Stage 1: The scientific method

We start with a textbook implementation of the scientific method. It allows us to understand battery cells' responses to varying parameters such as temperature, supplier, number of discharge cycles, change in workloads, etc.

As the software observes a behavior, it creates a mathematical model to predict how the cell will behave under specific conditions for an application (i.e., forms a hypothesis). The algorithm predicts what will happen in, say, 10 minutes. Then, it observes any discrepancy between the predicted and measured values and draws the appropriate conclusions.

The system retains the hypothesis if the prediction is more accurate than the last one for a given set of conditions. Then, it repeats the process in the next cycle to derive an even deeper understanding of how various parameters impact cell behaviors. However, if the prediction is less accurate, it discards the hypothesis and returns to the old model.

By forming a hypothesis and overseeing subsequent outcomes, this iterative process gives us the fundamentals of a self-learning system that doesn't interfere with reality.

The method doesn't challenge the validity of pre-determined limits. It's like improving the weather forecast: You can't change the weather, but you can improve the forecast's accuracy and how far you can look into the future.

Stage 2: Continuous improvement

We build on the same principles as stage 1 but experimentally vary a set of conditions (e.g., allow a 2% excursion over the typical temperature limit) and observe its impact on performance. This method increases the ability to uncover new patterns emerging by collecting numerous data points.

Over time, we develop a deeper understanding of how a parameter change impacts cell behavior and performance. Some data insights have a near-instant impact (e.g., over seconds to improve vehicle acceleration), while others stretch out over one charging cycle to inform short-term autonomous decision-making.

The process is similar to pharmaceutical drug trials where researchers prove efficacy, dosage, and safety through repeated, random, double-blind experiments.

Stage 3: Long-term behavior modeling

The methodology is similar to stage 2. But instead of measuring input and output in the same charging cycle, we track the impact of actively changing one input over multiple cycles or even the lifetime of a product (e.g., how exposure to deeper-than-usual discharge levels affects cell longevity) — allowing us to refine the accuracy of behavioral modeling.

More importantly, extracting data insights on day-to-day and long-term patterns through this 3-step process allows us to treat battery ecosystems like commodity, mercantile and financial markets, traffic patterns, and climate forecasts through predictive models.

Operators and decision-makers can understand and optimize the lifetime value of an asset (i.e., high-value battery packs), which, in turn, provides transparency for a market to exist in the energy space. Let's delve into the profound implication of this new ability on global electrification.

Battery predictive models: The foundation of global electrification

From car and vacation rental to airline and real estate, a free market in these industries exists because companies can use data insights to inform resource allocation, adjust pricing, and take calculated risks in investing in high-value assets.

Here are two examples to illustrate how the availability of data and the ability to predict behaviors allow businesses to manage physical assets effectively to maintain profitability under shifting circumstances — which reduces risk and gives them the confidence to invest in high-value properties for the market to exist:

Commercial real estate after COVID

The work-from-anywhere trend initiated by the pandemic resulted in excess commercial real estate inventory in many cities. Meanwhile, there is a shortage of residential housing in most markets. The relatively well-understood commercial and residential real estate markets allow owners to make informed decisions about converting some commercial properties to residential ones to minimize loss.

Real estate companies can understand market behaviors and adapt to changes in conditions by analyzing data on many parameters. They can gain the insights and transparency they need to make long-term investment decisions — knowing they can make money if the original assumptions change.

Finnair in the wake of the Ukraine war

> Finnair's business model leverages Finland's geographic location as a gateway between the lucrative European and Asian markets. When the Russian airspace was shut down due to the war in Ukraine, the route from Helsinki to Tokyo increased by almost five hours — making it no longer financially viable.
>
> The company had invested in a large fleet of long-haul planes to serve the Asian routes, and the imploded market made payments on billions of dollars worth of leased aircraft challenging. What can it do in this situation?
>
> The caveat is that the strain on commercial airframes comes from the number of compression cycles (i.e., one takeoff and one landing) and not from the distance flown. A plane essentially wears at the same rate whether it flies from Helsinki to Amsterdam or Tokyo. The question boils down to: Will it make sense to reallocate the assets to shorter and more frequent routes?
>
> Thanks to the availability of data about airline passengers, their price sensitivity, and competing players, Finnair could identify alternatives to reallocate the resources. It pivoted to wet leasing some planes to other airlines (i.e., renting out the equipment and staff) and flying to destinations they typically don't serve (e.g., HEL to LAX).
>
> Additionally, the company could make informed decisions about flying shorter European routes because data analytics suggested that it could generate profits by using long-haul aircraft for short hops, like Amsterdam to Helsinki. Then, it could take some of the passengers to LAX.
>
> These decisions were only possible because the company could leverage vast amounts of data to create financial and pricing models to identify opportunities that previously weren't on the radar.

The mortgage market can exist because a financial model ties the cost of capital, risk, supply, and demand to the physical assets. Mortgage risk is rated using well-understood metrics, which allows the packaging of many individual loans into mortgage-backed securities for trading on an open market.

For players like investment banks to enter the electrification game, we need a similar mechanism to rate the physical assets using a well-defined set of operating conditions to allow a market to define their value.

If you can’t tell the state of health (SoH) of cells in a battery pack or their remaining useful life, no one is willing to buy one that isn’t freshly out of the factory — making it illiquid. Without any rating mechanism, even purchasing new packs can be a toss-up. The lack of visibility has made batteries an expense rather than an investment.

A battery market can’t exist, and no company will back a financial model for leasing without a mechanism like a data-based battery lifecycle index, which did not exist in the energy storage space until the Tanktwo Battery Operating System (TBOS) made it possible.

The index works somewhat similar to a credit rating to guarantee that an asset is of a specific quality so buyers don’t need to know the actual, detailed technical mix of the investment — abstracting transactions from each physical asset to facilitate a free-flowing market and streamline decision-making at scale.

Our Battery AI self-learning system provides the data to inform a theoretical model of economics that will form the foundation for a new range of applications and industries to maximize asset usage and enable electrification at a global scale.

A dynamic system fueled by a continuous learning loop

We can develop new economic models with access to more data, a better understanding of variants (e.g., battery chemistry), and more tightly controlled parameters. The ability to provide an index to indicate a battery investment's risk can expand the market by attracting investors with different risk tolerances.

Less speculative battery investments will create a higher confidence level to attract blue-chip players who rely on having comfortably accurate predictions to enter the market. Meanwhile, high-risk, high-reward battery investments will attract investors with different objectives. Like the stock or real estate market, there's something for everyone.

As with any industry, the data must evolve with the market and each business's circumstances to be relevant and significant. So, how does it work with our technology?

No matter the customization, we provide baseline data to ensure a system is functional and safe in any scenario with a mechanism more sophisticated than any battery management solution on the market.

The system's performance will continue to improve as the algorithm learns the cells' behaviors under specific operating conditions. The improvements are even more pronounced if you have a highly customized implementation (e.g., non-standard chemistry requirements) where off-the-shelf data may be too conservative.

But that's not all. If you enable TBOS's network features, you can leverage insights from aggregated and anonymized training data collected across Tanktwo-supported systems to fine-tune your implementation.

It works like Apple dialing down the peak performance of an iPhone with an aging battery because it has learned from millions of phones in the field that pushing an old cell too hard could cause the phone to shut down unexpectedly.

Meanwhile, organizations with stringent security requirements (e.g., defense) and don't want constant network connectivity can still benefit from the learning process through periodic system updates.

The connectivity, supported by our battery security architecture, enables the use of data to support the development of a global energy storage and power-as-a-service market — made possible by transparent and optimized asset management. Additionally, the transparency into the state of each cell will eliminate the hot potato "second life" problem to maximize asset utilization and minimize the environmental impact of electrification.