

IoT and Predictive Maintenance

Can we extend battery lifespan with it?



Introduction

This White Paper aims to discuss whether we can connect the Internet of things (IoT) with Machine Learning and Predictive Analysis and how. This would allow us to perform Predictive Maintenance and use it to improve the way we handle our infrastructure and other important assets. For easier understanding, we will use a concrete example of an important asset - batteries.

Driven by technology advances and market demands, batteries are becoming increasingly important. Power autonomy is one of the biggest problems limiting us at the moment, starting with smartphones (where we seek longer powered-up time) to electrical cars (where we wish to reach farther destinations). If we look at automotive companies, the ones that have better battery technology and thus can reach longer distances, are more valued by the stock market.

The need for better batteries is just starting to roll-out, and where a strong business need is, investors are more interested to participate. There sure can be expected innovative ideas towards batteries/accumulators, we might be facing a new era in this field, and thus how to manage batteries will be even more important in the future.

Here, we will take a look at rechargeable batteries with emphasis to multi-cell batteries, also known as accumulators; but some of solution's functionality presented here, would benefit for any kind of rechargeable battery.

Target audience for this document are the

people interested in:

- Optimizing current infrastructure/assets (e.g. battery pools)
- Extending the lifetime of infrastructure/assets (e.g. battery lifespan)
- Ensuring smooth operation with predictive maintenance
- Planning maintenance budgets
- Safety

The problem

By rethinking battery management, we are tackling two major issues, and two different perspectives:

- Environmental, and
- Economic

With the help of IoT and Predictive Maintenance, we can work on the economic side by finding out how to prolong the lifespan of batteries. At the same time, longer lifespan will positively impact the environmental side.

So what is it that makes batteries die fast?

Major impacts on battery lifespan

We will not address the chemical and mechanical construction of batteries. We will only try to understand the user's perspective

and will ask ourselves how we can utilize the batteries in a better way.

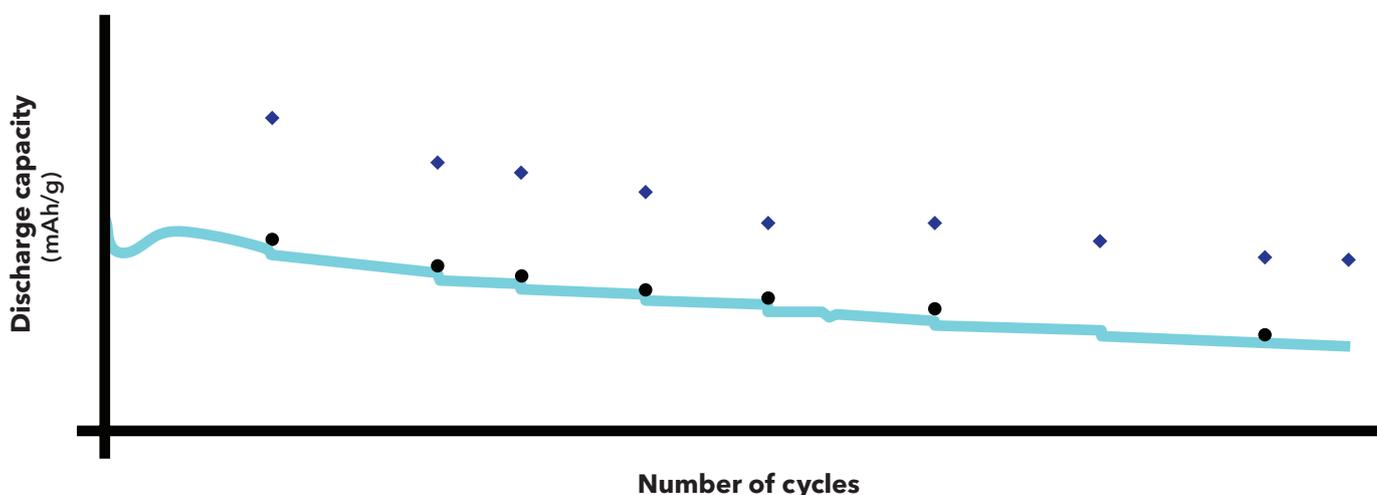


Figure 1: With every new cycle, capacity gets lower

1. How many cycles can a battery withstand?

Every battery has an expected lifespan defined by the manufacturer, meaning a certain number of charge-discharge cycles. One cycle means that the battery is fully charged and then discharged; when we start charging it again, a new cycle begins. Figure 1 shows an example of battery cycle capacity.

In reality, this number depends on its usage (how and where it is used) and its technology. For example, thin film lithium rechargeable batteries may have 40K cycles, while, on the other hand, the Lithium sulfur battery may have 1.4K cycles.

2. DOD or Depth-of-Discharge effect

Depth of Discharge tells us how much capacity can be withdrawn, e.g. 50Ah. If a battery has 100Ah and it is discharged for 20 min, while the withdrawing current is 50Ah, the DOD can be calculated with the following equation:

$$\text{DOD} = \frac{\text{Withdrawing current} \times (\text{dischargingTime}/\text{hour})}{\text{Battery Capacity}}$$

So, if we use different values, we can see that **the higher the load on the battery, the sooner the battery will die** - which means a lower number of charging cycles.

3. Multi-cell batteries

The DOD effect is even more important when it comes to multi-cell batteries, where overall battery voltage is taken into account. **Some of the cells are usually "lazier" than the rest**, which results in the more performing ones being worn out even before the lazier ones are able to reach full capacity - resulting in a dead battery with some cells that are still good. In Figure 2 we can observe the discrepancies between the cells.

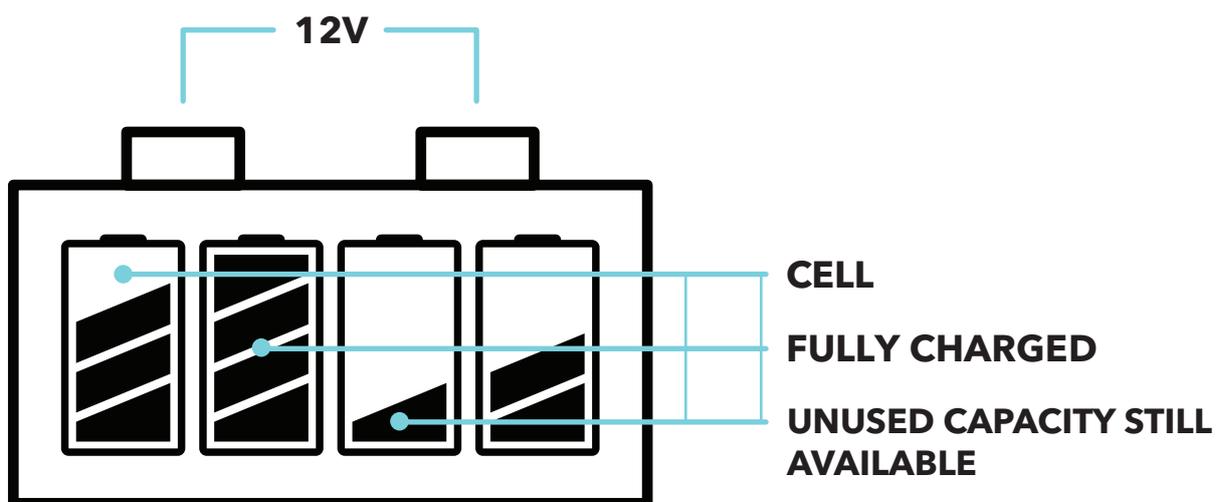


Figure 2: Even though the battery shows an average of 12V, some of the cells are overcharged, while some are undercharged

4. Temperature

Temperature level is very important when charging. If a battery is working outside its best temperature range, e.g. in a colder environment, the capacity of the **battery will be lower. On the other hand, increasing the temperature can help charge the battery**

faster, but here control is essential in order to charge the battery without significant harm (e.g. a cell explosion). So the temperature impact should be taken into account when planning the physical place where the battery will be used (see Figure 3).

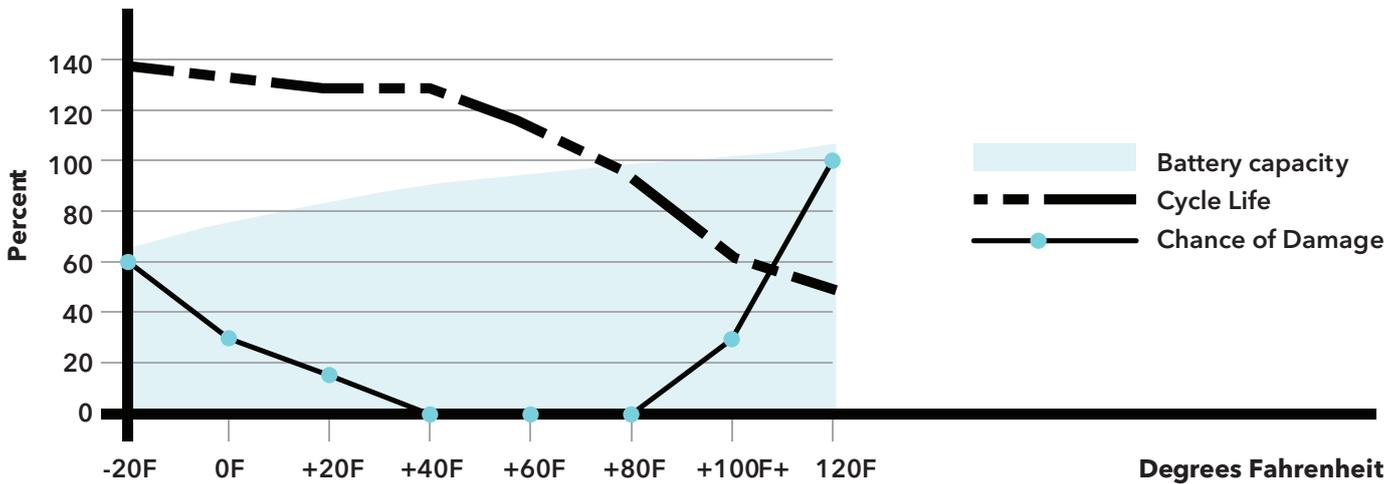


Figure 3: Capacity is higher in higher temperatures, but the chance of damage is also much higher.

5. Charging time

Charging time has an impact on the battery's lifespan. That's why **smart charging is preferable**, to make it safe for batteries and to save time. Batteries are usually charged up to 90% of their capacity in little

more than half of the charging time. Charging the remainder of the capacity takes the remaining time, as it needs to be charged slowly. This is shown in Figure 4.

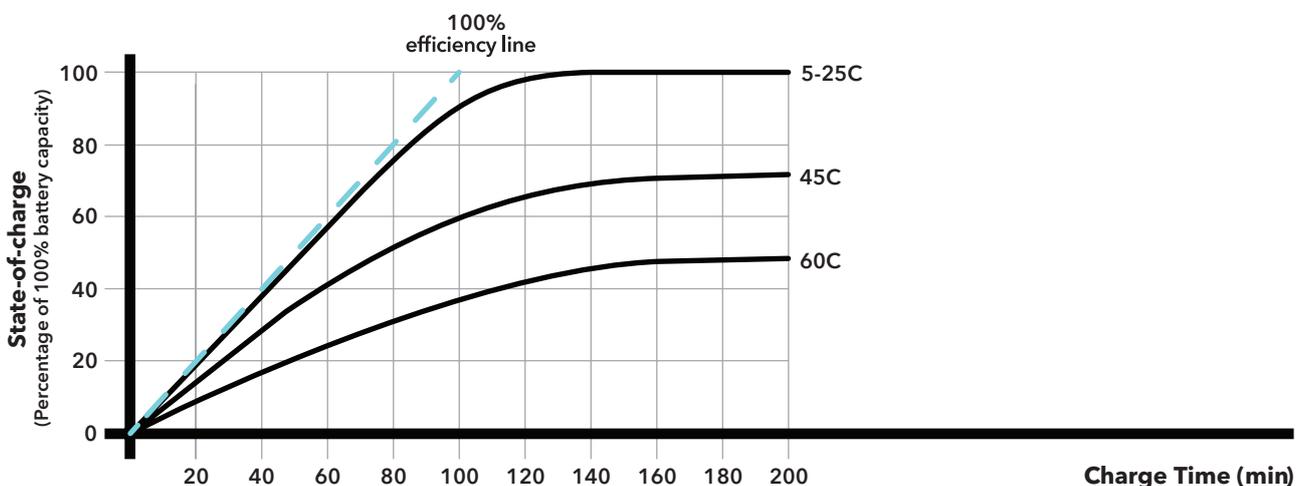


Figure 4: Charging time to reach full capacity.

Brainstorming

How can we use all the facts we identified in our IoT and Predictive Maintenance model? It is clear that the way rechargeable batteries are used impacts their lifespan. For this reason it is important to track the battery usage and to take into account its current state. We can leverage IoT to gather this data. We can later use this data in a smart way in order to:

- extend the expected lifespan of the batteries and
- gain knowledge which ones among them will die sooner.

These insights can help us to:

- ensure smooth operation
- estimate the budget
- plan logistics when needed

1. Battery pool management

As we have learnt, batteries have a certain number of charge-discharge cycles during their lifespan. **If we periodically rotate the**

batteries in the pool, we will impact the quality of usage and can so extend their lifespan tremendously.

In practice, people rather use batteries which are more convenient, while some are just collecting dust in the corner, because they are far away and heavy (e.g. batteries for fork-lifts weigh more than one metric ton).

By introducing a suitable solution that shows the users which battery to use next, based on identifying each battery and its usage data, this problem can be entirely avoided. Furthermore we can also optimize the battery pool based on the following questions:

- **Do we have more batteries than needed?**
- **Are the batteries positioned where they are needed?**

Just by having this data, we can save money on investment by not over purchasing.

Figure 5 shows a sample saving calculation from our partner, presenting a real case:

Project (5 years)	Without Battery MGM	Using Battery MGM
Number of batteries*	108	100
Number of charging stations**	54	50
Cost of purchase	864.000 €	800.000 €
Electricity costs	308.571 €	285.714 €
Investing into Battery MGM	/	28.135 €
Savings	/	58.722 €
Battery lifetime extension	/+	10% or 86.400 €
Total savings	/	145.122 €

Figure 5: Sample of a ROI into such a solution, reference: Ektimo, svetovanje in razvoj, d.o.o.

*Battery price=8.000 EUR
**Electricity per charge station cost = 2.857 EUR

2. Battery charging in a smarter way

When a typical charger is charging the battery, only the general capacity of all cells is taken into account - when the cumulative result is reached, the charging stops.

This results in some cells being overcharged,

while others are not at designated charge level (see Figure 6).

The solution to this would be **to take the data per cell into account** and, while discharging the cells which are over performing, keep charging the ones which need to be stimulated. Implementing this principle would result in a **battery balancer module**.

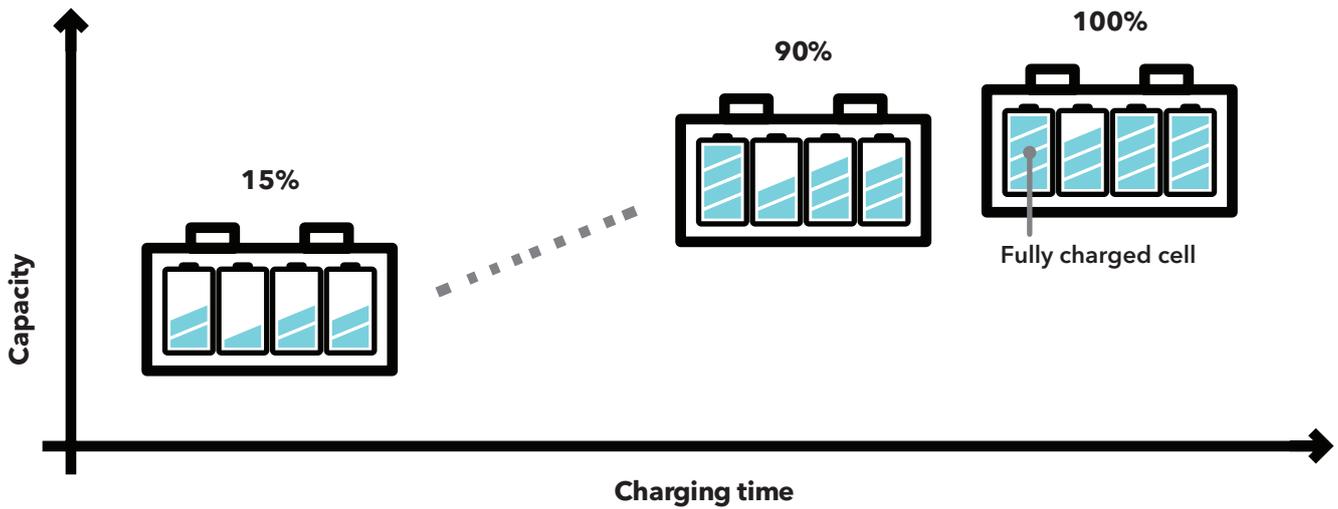


Figure 6: When a multi-cell battery is charged, not all of the cells might be fully charged

3. Battery discharging effect

Another aspect that affects battery lifespan is how we use them. As we have seen in the previous section, not all of the cells in a battery are equally charged. Because of that,

some of the cells (the lazy ones) will deplete faster, after which the burden of the load will fall onto the better performing cells. This **worn-out effect accelerates over time** and results in a shorter battery lifespan. The effect is presented in Figure 7.

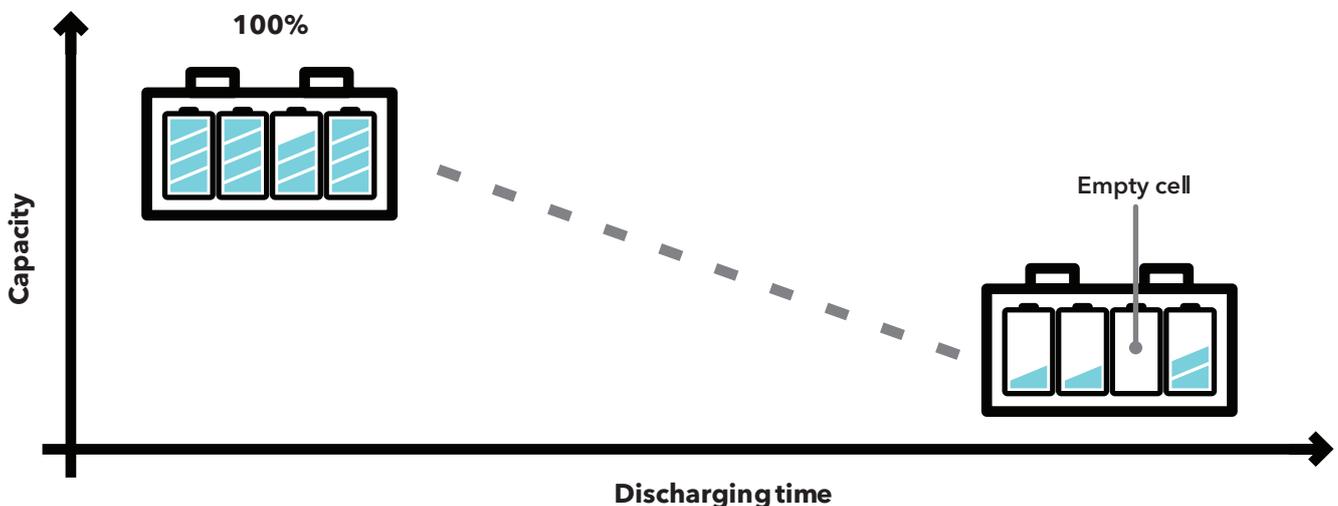


Figure 7: When a battery is being used, some of the cells might be depleted before others

Another aspect that the data can highlight is the load on a battery. As we have seen earlier, the power load on a battery impacts its lifespan. The idea is, if the operating process allows, its **lifespan can be prolonged with**

less stress on the battery. Data gathered via IoT can show these power spikes, which we can take into account and prolong the lifespan of the battery in this way. Figure 8 represents the **spike detection graph.**

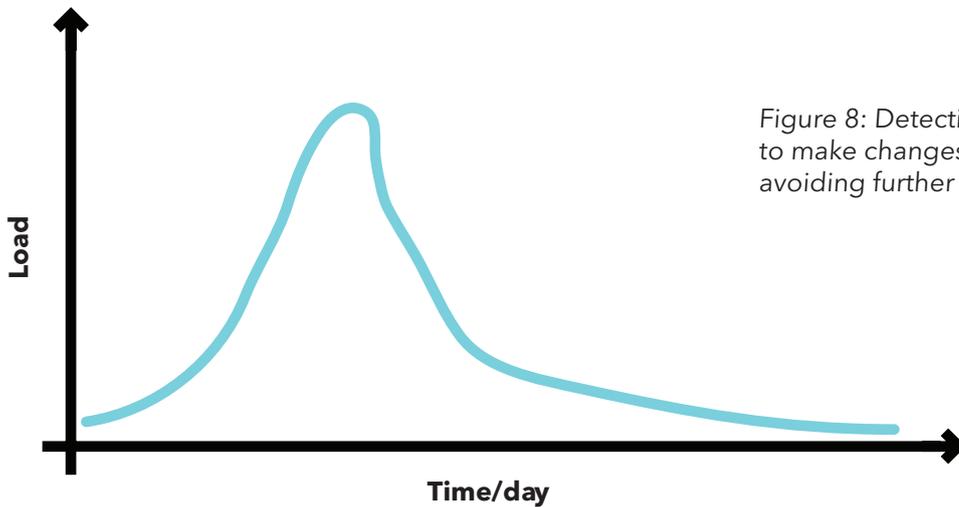


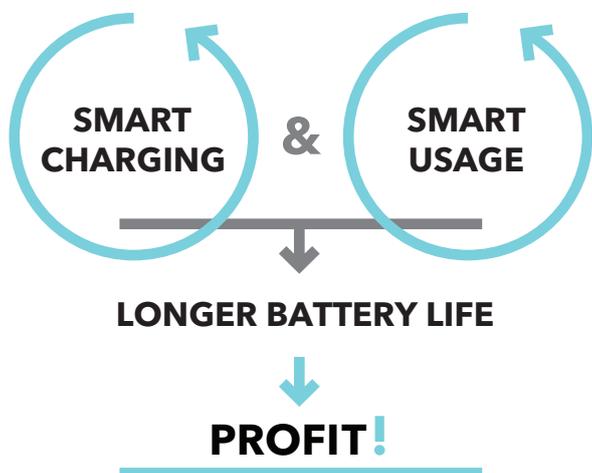
Figure 8: Detecting spikes might give us insights to make changes in the process to avoid them, avoiding further spikes

4. Temperature impact

As we have seen, the temperature affects the lifespan of a battery. Taking this fact into account, combining it with the data from fleet management, and telemetry of the data gathered via IoT, we can **optimize the charging temperature and battery use.**

By introducing logic within the battery, we could:

- **read the temperature per cell and charge it accordingly** (this way we could speed up the charging time of the whole battery in a safe way, without damaging it)
- **use the battery in a different environment**, which is more suitable for it (e.g. in a colder environment, we might use the older batteries indoors)



Faster charge time and smarter usage have an impact on the economic side, which is shown in Figure 9.

Figure 9: Faster charge time and smarter usage will result in profit

Coming up with the solution concept

So far, we gathered some interesting facts and thoughts about prolonging battery lifespan. In order to define the complete

solution and understand it better, we will be using Figure 10. It represents a high-level 10-step solution overview.

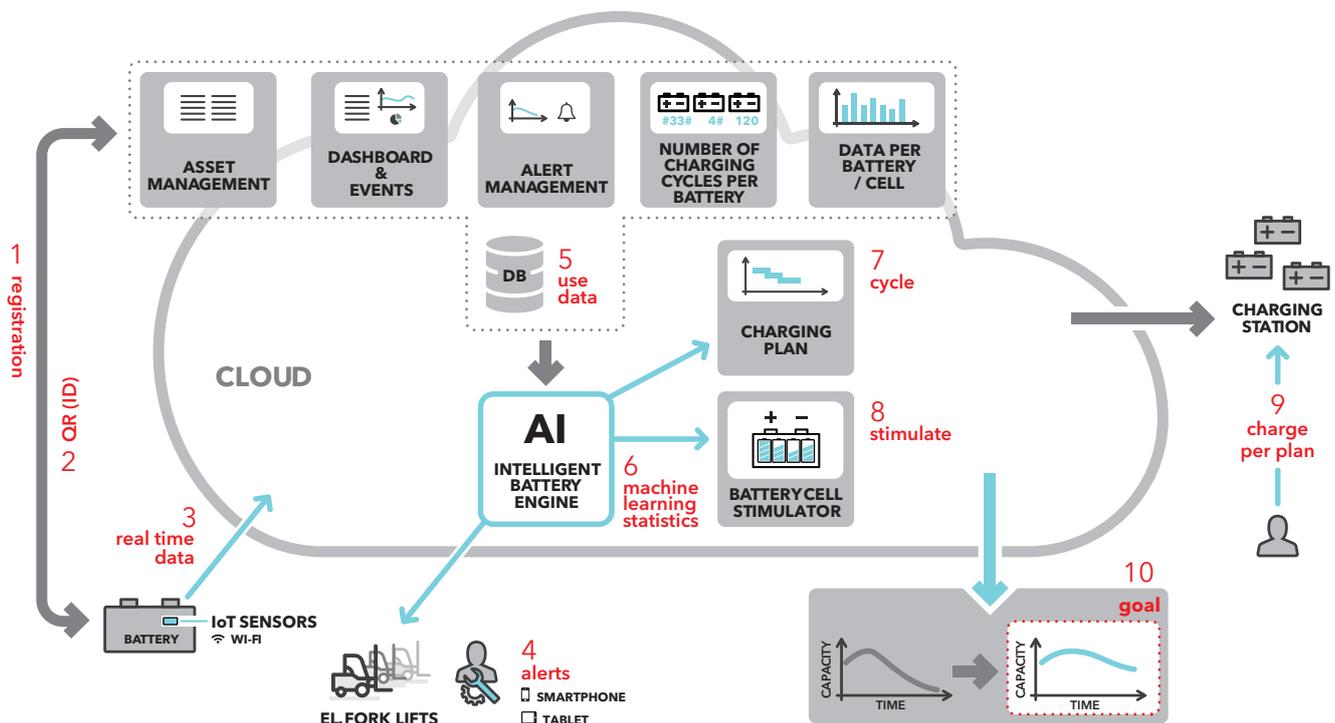


Figure 10: High-level 10-steps solution overview

The high-level 10-step solution overview is based on facts and ideas discovered earlier. In the following paragraphs, we are describing each step:

- Step 1:
Each battery needs to be registered in a management system. Basically, this is one of the modules within Fleet Management. In this module, we can store data about the battery, what type it is, its age, what maintenance was performed on it, etc.

When we register the battery into our Fleet Management system, we receive a unique ID, which is necessary for battery identification during all subsequent steps in this solution.

- Step 2:
QR codes are used, which can be printed on a sticker, or using appropriate printing techniques, like laser-printing, on the chassis of the battery. This QR code represents the unique ID of the asset; it also represents the URL where we can easily see all the related documentation and events concerning this particular asset.

-Step 3:

Batteries are **equipped with IoT hardware**, which enables their real-time tracking. This is **the heart of the solution**. This gives us the option to collect:

- o Real-time data
- o Big data

Real-time data allows us to react to certain events when they occur. This data is connected with an **Alert Management** module, which holds configurations of different kinds of predefined alerts. The module logs the frequency and acknowledgments of the alerts; also, specifics like thresholds, special conditions about repetition and similar. In this way, we can easily define which events require a response, e.g.:

- o When the voltage drops below the set threshold
- o When one of the cells is below designated value and we want to protect it.

We can configure the alerts to be sent via SMS, e-mail or both - or even using push notifications.

Big data is used to collect information on how the battery is being used through time, **down to every single cell level. This is a major difference compared to typical solutions**. This data gives us insights we discussed at the beginning of this paper. We are now able to implement statistical, mathematical and **Machine Learning** algorithms to utilize the knowledge hidden in the data.

- Step 4:

Alerting the team when they need to respond to the events is the key action in the whole solution. When an alert is sent, an action needs to be performed. All alerts and responses are well documented and can be examined retrospectively. Without this, maintenance alerts could potentially

be ignored. By reviewing the alert log, if we discover they had been ignored for too long, we can remotely disable the cell or the battery to prevent further damage.

- Step 5:

The solution consists of several modules, which contribute toward data collection and have key operational features:

- **Fleet management:** besides the data mentioned, also provides the following battery-related data: number of cycles, average charge/discharge time, average current, etc.

- **Dashboard & event module:** the dashboard is very important for monitoring the health of the overall operation. Different kinds of key performance indexes (KPIs) are presented and we are able to assess the state of the current subset, and drill down to the problem that affects the state of the selected fleet; upon opening an asset's page, we can see all the details relevant to that particular asset - its telemetry data, events and alerts.

- **Alert management:** we already explained

- **Asset details module:** provides us with an insight into details of the selected asset that we gathered using IoT. Data is collected about the telemetry and associated information, in our case, the number of cycles and telemetry on specific cells inside the battery.

- Step 6:

The **Intelligent Battery Management** module is based on **AI**, mainly on the **Machine Learning** part, but also includes statistical and mathematical algorithms. In order to be able to define what is best for a particular battery, this module takes into account the following data:

- o battery type and age
- o cycles performed/remaining cycles
- o battery telemetry
- o cell history
- o events

All this data combined is used as an input to the AI module which is able, based on the data and Machine Learning algorithms, able to come outup with a better maintenance plan for each battery.

- **Step 7:**

A charging plan is created by the AI module, based on available data for each battery: cycles, charging details, etc. At this stage, the recommended battery usage is presented in the Fleet management/ dashboard (optimum environment temperature, whether the battery is suitable for spike-heavy load, etc.).

- **Step 8:**

The AI module also recommends **which battery cells to stimulate**. Based on cell data inside the battery, the best treatment for balancing the cells is proposed. The charging station can automatically utilize this, if the interface is introduced; otherwise the stimulation can be determined manually.

- **Step 9:**

A **maintenance team** can leverage the data using the solution in the Cloud or on premise. Using a tablet or smart-phone, they can simply scan the QR code on the battery and be redirected to the asset's page in the Cloud, where they can see how best to maintain the battery.

- **Step 10:**

Here we have reached the goal of the solution: to **prolong the lifespan of the batteries by their smart maintenance and usage, and achieve savings**. This solution also **provides empirical insights helping us to optimize the size of the battery pool** by buying just enough batteries to support the production - and saving us money. It also **helps with budget planning, by enabling us to anticipate when they will reach the end of their lifespan**. **Safety** as well must not be forgotten. We are able to measure the temperature of each cell during charging/ discharging, as a result of this, we can **react in real-time** (e.g. preventing battery explosions).

Prototyping / PoC

Let us explore what the prototype solution looks like, using some cheap equipment available on the market.

EQUIPPING BATTERIES WITH IOT

We need some electronics to gather the data from battery cells. In this case, NodeMCU is used as the base component. It comes in a variety of models. The model can be selected based on how many input/output ports are needed. It already has WI-FI support, which we need. Figure 11 shows us high-level schematics:

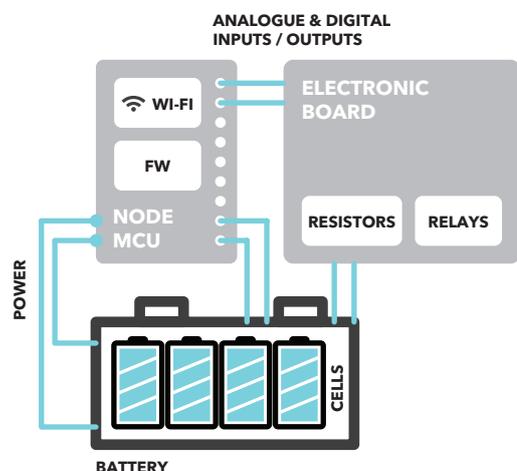


Figure 11: High-level design

POWERING THE CIRCUITRY

The circuitry is powered by the battery itself. It uses 3V to operate when it is not in power-save mode. The telemetry data and current status is pushed to the Cloud via Wi-Fi when the circuitry is in active mode.

WI-FI INTERFACE

Using this approach with a Wi-Fi interface, every battery is able to run its own Web server. In this way, we are able to connect directly to the battery using the QR code on the battery's chassis and see the relevant data gathered by the circuitry.

To be able to configure the Wi-Fi network on the battery, NodeMCU would have to support two modes:

- AP mode
- Web server mode

The firmware is able to determine if it is configured to access the Wi-Fi. If not, it will automatically switch to AP (access point) mode, where we could connect to it and configure its IP and credentials to be able to connect it to the Wi-Fi network.

FW (FIRMWARE)

FW takes into account the current state of the battery, using the analog input ports on the NodeMCU, to track the real time status of each cell. If the value falls out of range, the event is sent to the Cloud, where the alerting module can take over and alert the maintenance team, according to the predefined rules.

The following data is measured for each cell:

- Cell temperature
- Voltage/capacity
- Load

Data is stored on an onboard memory; this enables tracking data over time (telemetry data), so that we know exactly what was going on with the battery at a given time. This data is available for pushing to the Cloud from the onboard memory when needed, using REST API.

As already mentioned, FW takes into account the current state of the battery. If it is out of accepted range, FW disables the usage of the battery, preventing any damage. The state of the battery is monitored to the cell level, so if one cell drops below designated value, an alert is sent to change or charge the battery. If the alert is ignored, the final step is to block the battery usage, in order to prevent battery damage.

During the charging time, FW takes into account the temperature, using an analog temperature sensor connected to the NodeMCU input. It is able to define the amount of current to optimally charge the cells. Cells which are over the designated charge level can be discharged using resistors (in case the charger can only charge the whole battery, not individual cells themselves). The rest are stimulated to achieve the designated capacity level - which results in a longer battery lifespan.

Figure 12 shows what the HW prototype looks like:

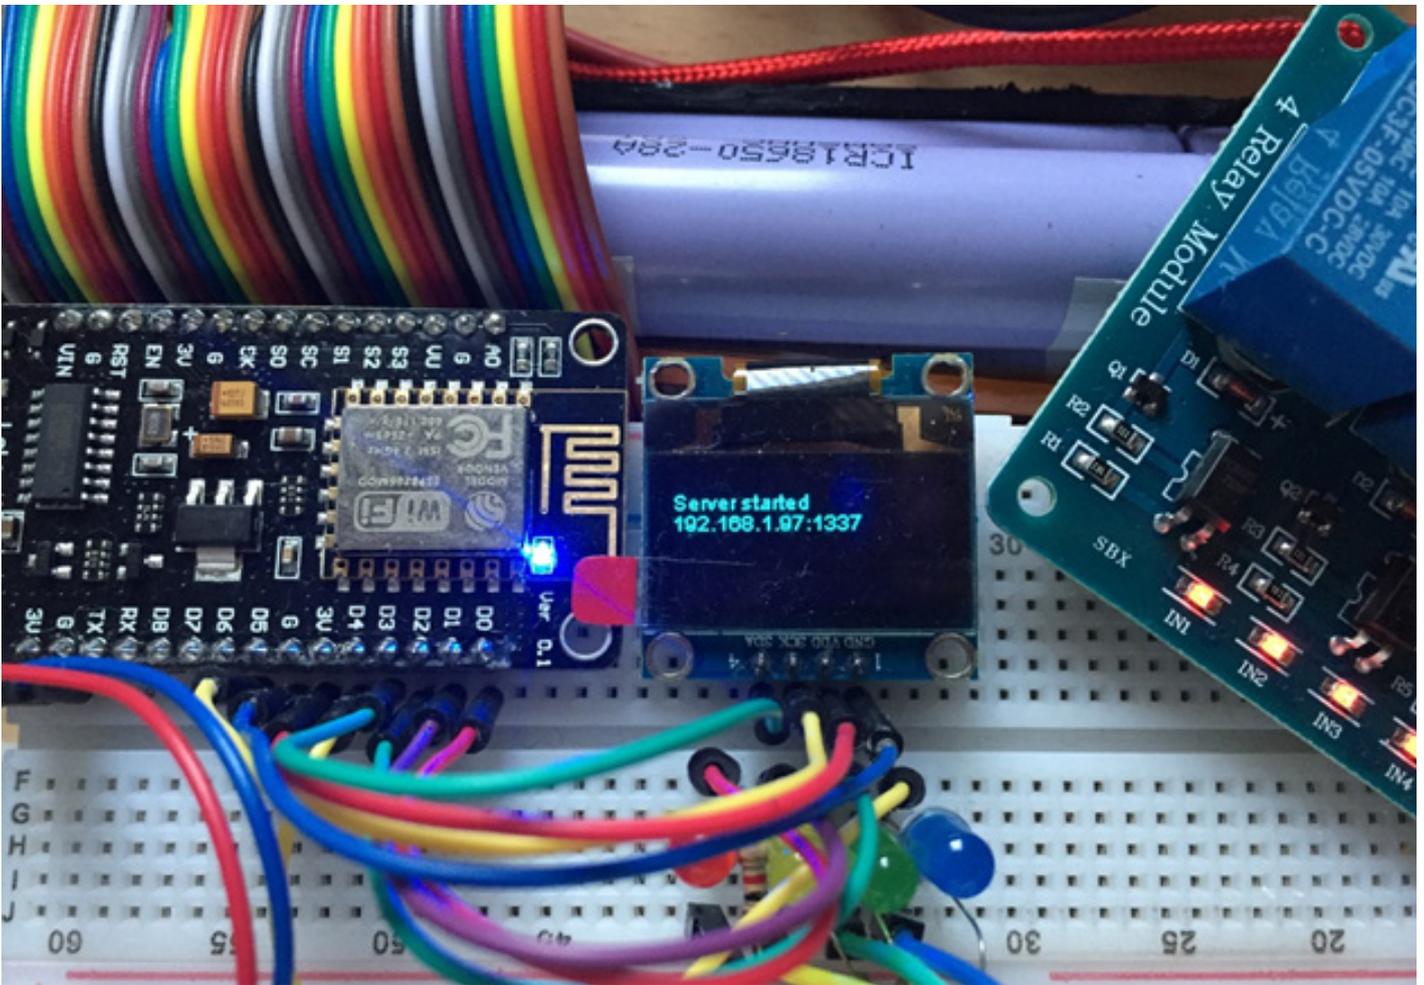


Figure 12: NodeMCU covering analog/digital interfaces, providing a Web Server and other functionality via FW; LCD for getting some direct data from the battery and relay module for opt coupled connections

Figure 13 shows a web page with the battery data, seen on a smartphone, which is running on the NodeMCU installed inside the battery:

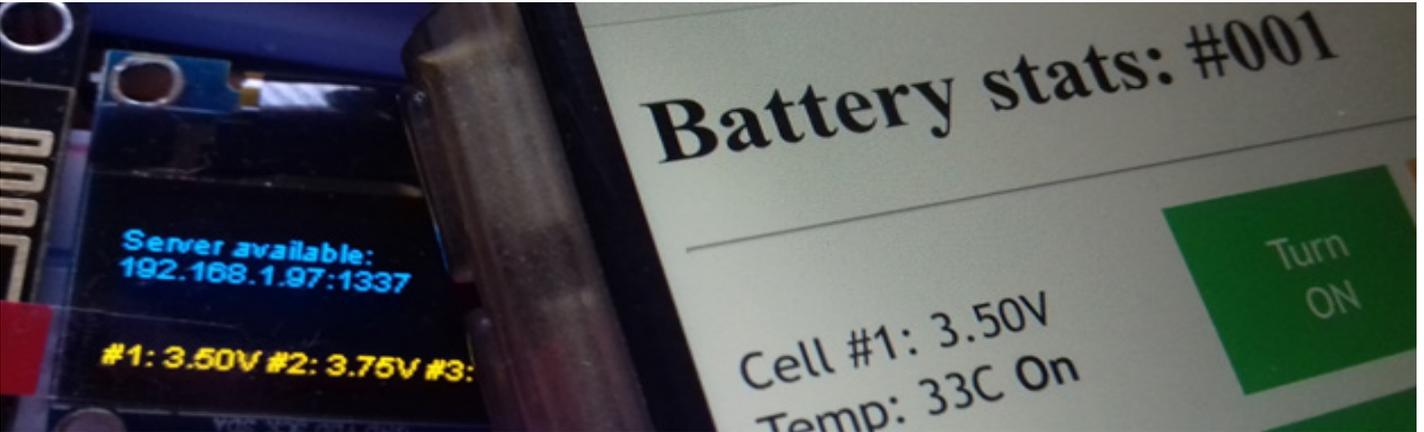


Figure 13: LCD showing IP address and battery insights, Web page showing available data about the battery

Now we are able to gather the telemetry data for a particular battery. The cells in the battery pass the data to the Cloud, where the SW side of the solution can take it over and utilize it for a smarter maintenance approach.

Using AI to enhance the solution

In the first stage, until there is no data available, statistics and mathematical functions are used. Over time, data comes in and we can apply the Machine Learning approach to it.

WHAT WE ARE TRYING TO SOLVE WITH AI

The major goal for our AI module is the **prediction of battery failure**, i.e. generating the following reports:

- **How many** batteries will reach the end of their **lifespan** in the following months (1m, 6m, 12m)
- **Which** batteries (via Fleet Management) **will fail** in the next calendar month.

After further analysis of the specific domain of battery management, the following data is taken into account:

- **IoT side:**
 - o Temperature of each battery cell
 - o Voltage of each battery cell
 - o Capacity of each battery cell
 - o Overall load on the battery over time
 - o Charging time
 - o Discharging time
- **Fleet management side:**
 - o Battery type
 - o Battery age
 - o Location
 - o Battery idle time

AI SOLUTION CONCEPT

Figure 14 pictures our AI solution concept. Battery telemetry data is collected over time. When a battery death event occurs, telemetry data from the past will be labelled and prepared for machine learning algorithms.

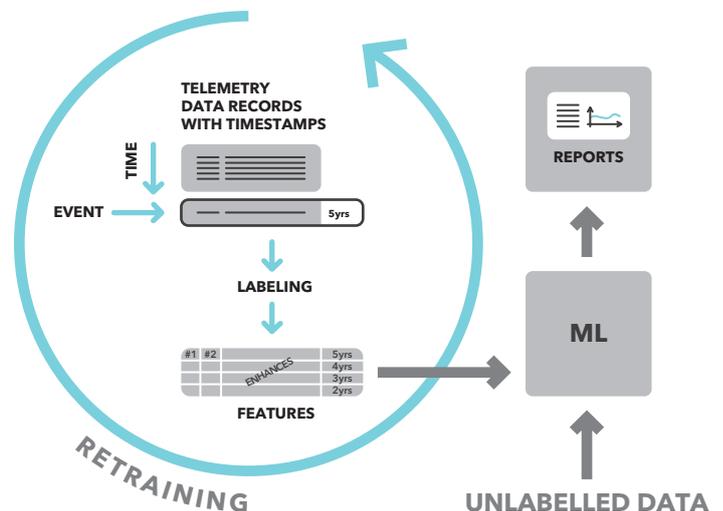


Figure 14: AI concept

Since we collect and store various information about batteries over time, **we can read the historic data about dead batteries** over different time periods and deduce known information about the left-over lifespan of the remaining batteries, which is then **used as the output variable (label) for supervised Machine Learning**. The process is illustrated in Figure 15.

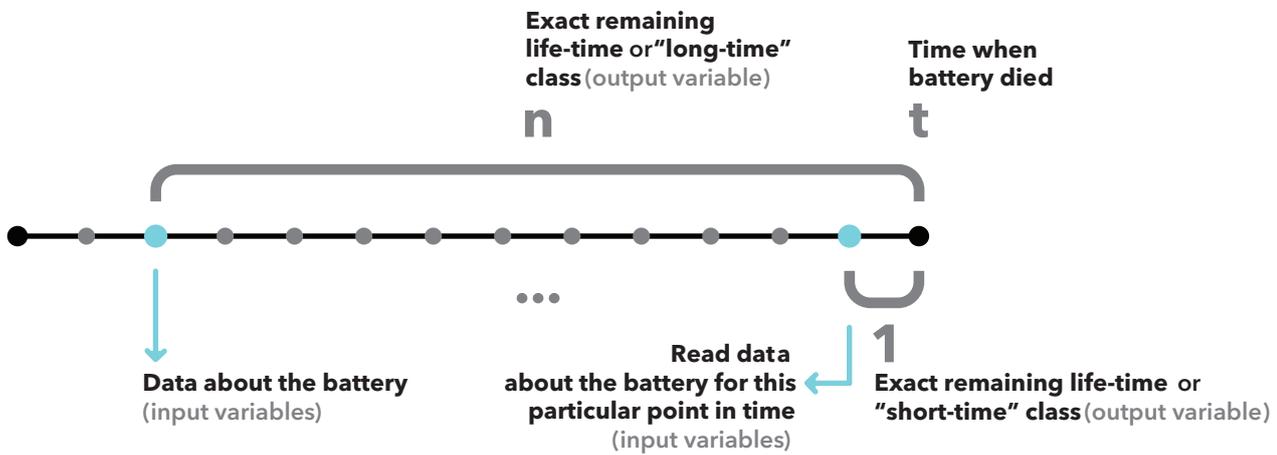


Figure 15: Gathering the Data for Supervised Machine Learning Process

EVALUATING THE DATA

Usually the data has some discrepancies which need to be addressed before we can use it. In Machine Learning, multiple sources are used, thus some fields of data might be missing, have wrong units, etc. All these data issues need to be addressed before the data can be used in further steps.

ENHANCING THE DATA WITH FEATURES

Now that the concept is defined, here are the additional features which are beneficial for Machine Learning:

- Average/Min/Max temperature/voltage/capacity of the battery to date
- Current temperature/voltage/capacity of the battery
- Expected life-time (in cycles) of the battery
- Number of cycles performed
- Average depth of discharge

- Number of time units/percentage of time working in various temperature environments
- Average charging time
- Average discharging time
- Device type/model/age/etc. which uses/used the battery
- Number of cells

Combined data is the starting point for creating a model; over time, the data and predictive models are **periodically updated and revised** by introducing or removing features.

In our case, we are using a supervised Machine Learning approach, where models are trained using a collection of labelled examples. The dataset would consist of battery instances characterized using the data features described above (input variables) and their labels (output variables), i.e. known remaining lifespan. Figure 16 shows an example of how we can represent the data in a tabular way.

Input variables					Output variable
Avg. temp	Avg. charging time	Location	No. of cycles performed	...	Remaining lifetime
...
...
...
...
...

Figure 16: Example of a table structure

PREPARING THE DATA FOR MACHINE LEARNING ALGORITHMS

The acquired data is first analyzed in order to understand it and to gain new knowledge, which is useful in the process of model training and evaluation. Furthermore, various filtering, cleaning, transforming and feature engineering techniques are applied to the data.

MODEL TRAINING AND EVALUATION

This phase includes a number of sub-steps, such as candidate algorithm identification, evaluation method(s) and measure(s) selection, model training and validation, model selection, model testing, etc. The goal is to find the best Machine Learning configuration for the problem at hand and to estimate its future performance. If the outcome of the Machine Learning model belongs to a finite number of categories, the process is called classification. If the outcome is a continuous quantity, the process is referred to as regression. In the case of predicting life-time of the battery, both approaches can be applied:

1) Classifying batteries into a fixed number of classes, such as short, medium and long lifespan (classification task) or

2) Estimate the exact number of remaining time units (hours, days, months, etc.) for the batteries (regression task).

The first option is beneficial, if one is not interested so much in detailed prediction, but in high-level categories. The second provides more specific predictions, but it is challenging to predict the exact lifespan of a battery, especially as there are not many labelled training instances available.

Various algorithms could be applied for estimating remaining lifespan of batteries, which depend on real business cases, for example: K-Nearest Neighbors, Support Vector Machines, Neural Networks, etc. Several candidate algorithms are chosen and applied on a given dataset in order to find the best performing one for the particular business case.

We are splitting the dataset into a training set and a test set. Performing cross-validation on the training part by changing various parameters would help choose the best model. Testing the chosen model on the test set would help estimate its general unbiased performance.

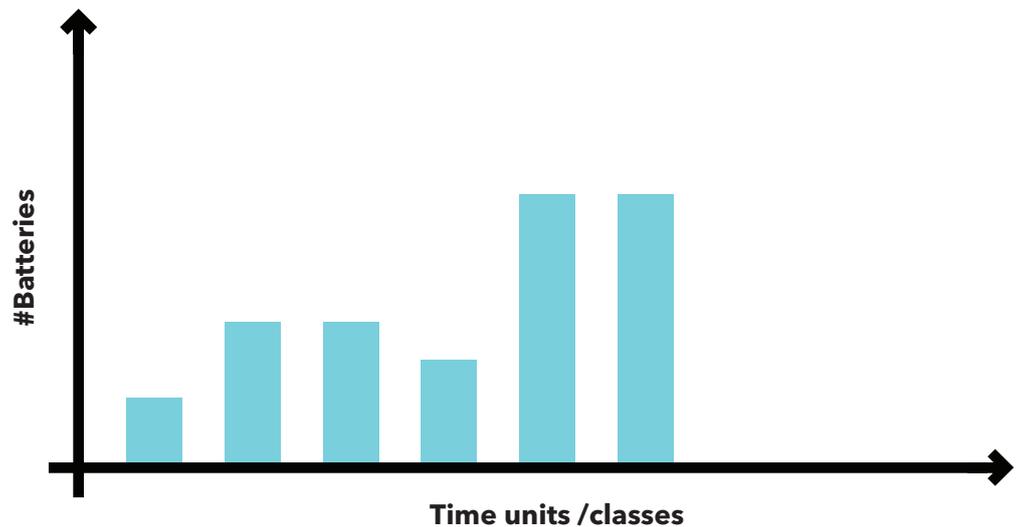
The performance is assessed by calculating evaluation measures, such as accuracy, precision and recall for classification and mean absolute error (MAE), root means square error (RMSE), etc. for regression. Lastly, using the best performing configuration, the final model is trained on all the available labelled data.

APPLYING THE MACHINE LEARNING MODEL ON NEW DATA

After training the Machine Learning model, the model is used to estimate the remaining

lifespan of the batteries which are still functional. The individual results can be aggregated in order to produce a summarized graph (see Figure 17), which shows the estimated remaining lifespan of all batteries per class or time unit.

Figure 17: The summarized graph for battery lifespan estimation



USING THE PREDICTION MODEL

Models created for specific business cases are propagated as Web Services, which makes the solution architecture loosely coupled, so that it can be integrated into any kind of solution.

The principle is that the assets taken from the Asset Module are periodically checked against the prediction model, for reporting/planning needs.

IMPROVING THE MACHINE LEARNING MODEL OVER TIME

The solution doesn't end here, the model needs to be periodically retrained with new data, additional optimization of the data and algorithm tuning is also required.

The model can be improved over time by adding new examples to the training dataset and periodically retraining the model. By applying such an approach, the model constantly gains new knowledge and is updated with the latest trends and changes. The model retraining can be implemented in different ways; for example:

- by applying a strategy which enlarges the training dataset by continuously adding new labelled examples, or
- by implementing a mechanism which forgets the oldest instances, while the new ones are added to the training dataset, etc.

Also, we need to deliberately let some percentage of the batteries die, to check if the prediction model really works.

WRAP-UP

In our concept, we used a step-by-step approach in order to solve the business case:

- we checked what issues contribute to the problem
- we found out how to deal with these issues
- we identified what is needed to find an adequate solution
- we created a Proof-of-Concept solution with minimal effort and cost
- we proved that the solution has true benefits towards attaining our goal.

In light of the information above, we consider the question we posed in the beginning

- can IoT and Machine Learning be used for Predictive Maintenance - adequately addressed.



Comtrade d.o.o.
Letalska cesta 29b
1000 Ljubljana, Slovenia
T: +49 711 22254 233
E: info.de@comtrade.com
www.comtradedigital.com