June 2020
Ekkono Solutions AB

Edge Machine Learning for Smart Battery Management

In the name of a sustainable future, more and more products are electrified e.g. cars, garden care products, construction and mining equipment, trucks, and buses. The real start of the electrification process was back in 1991 when Sony introduced the first Lithium-Ion battery which enabled the electrification of small handheld devices e.g. the Walkman. An even greater change has happened in the last ten years when we have gone from using batteries in relatively small devices to now use them in cars, boats, trucks etc. What has made the change to larger devices possible is that the energy density in the batteries has increased substantially along with an exponential change in lowered cost measured in cost/kWh. The electrification is now going faster than ever and the International Energy Agency projects that by 2030 a stock of 130 million electric vehicles could be on the world's roads from about 5,2M in 2018. We will also see a large change in power grids moving from a need to balance production and consumption, to enhanced ways of storing energy in more and more powerful batteries in homes and industry.

However, in order to ensure the safe and reliable operation of battery-powered devices, battery management systems (BMS) need to monitor and control several operating parameters, including charge- and discharge current limits, thermal management (high temperature batteries), overcharge protection and over-discharge prevention.

1. Applications

On a higher level the main concerns for a BMS are prediction of range and health. Health is a predictive maintenance case where the remaining life, or remaining useful life (RUL), is calculated based on estimation of health and use. This enables proactive replacement of exhausted batteries/battery cells, but also better use of batteries that can get a prolonged life in one application over another. For any conscious corporation, batteries become a question of sustainability, as life, use and retirement of batteries has a direct impact on the environment. Knowledge, data and insight about battery health are central to the sustainability work of any company going through electrification.

The health of the batteries impacts range. Range is a factor of state of charge (SoC), state of health, ambient temperature, application, and use. The SoC of a battery can normally not be assessed directly, but can relatively accurately be estimated using a number of indicators, by, e.g., measuring the specific gravity of the liquid electrolyte in the battery, measuring the battery voltage, measuring in-and-out-flow of current (so-called Coulomb counting), or using a Kalman filter based on a system model of the battery. However, these traditional techniques all struggle with the settings of IoT devices, i.e. they need a controlled environment, or when the usage pattern and environment changes over time which is the typical for IoT devices. State of health, and thereby discharge capacity, is overtime affected by both how the battery is charged and discharged. More specifically, the current, voltage and cell- & ambient temperature during charge in combination with both the initial and final SoC directly affects the state of health in the long run.

Consequently, we argue that a smart BMS must be able to account for both **usage-**, **charging-** and **environmental** patterns to be truly accurate and smart in its predictions. Furthermore, the nature of a BMS system residing on an IoT device also entails that it is:

- Fast enough to make real-time decisions
- Power efficient enough to be run on battery-powered hardware
- Portable so various battery and device configurations can be supported by the same operating system

2. Edge Machine Learning for BMS

Looking at the setting for IoT devices, big or small, it soon becomes apparent that any BMS system must reside on the device itself. If not, the BMS will be dependent on a stable internet connection which is both costly, energy consuming and hard to achieve in many cases. Machine Learning on the edge, or Edge Machine Learning*, does on the other hand require that the calculations can be done on a microcontroller with limited resources. Furthermore, it is vital that actual learning can occur on the device to account for the action of the user and a changing environment. Learning on the device, i.e. using incremental machine learning techniques, is what Ekkono specialize on by providing a machine learning library for embedded devices.

It's important to note that when we speak about a smart BMS we are not just looking at the internals of the battery in the traditional sense. Instead we extend the concept to the user, the machine connected to the battery and any knowledge of future influencing factors e.g. temperature forecasts or a planned route.

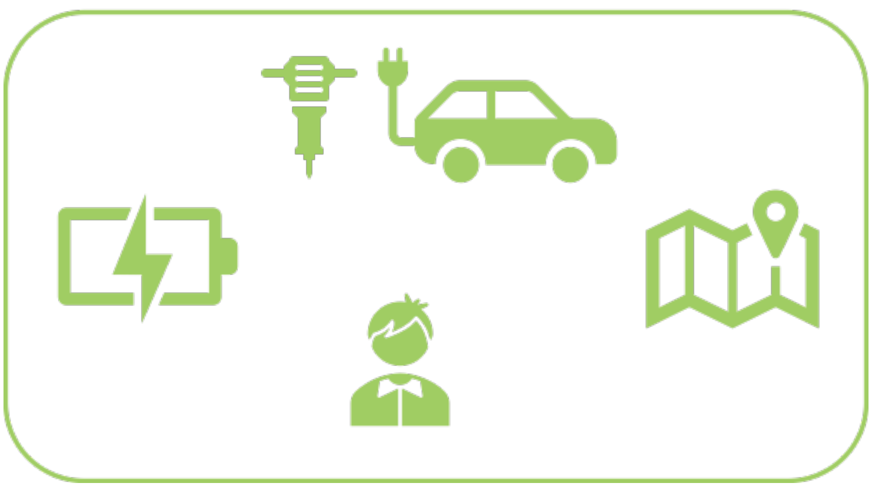

Figure 1 - Smart Battery Management should include the battery, the machine, the user and information of future influencing factors.

* Edge machine learning means running machine learning (ML) at the edge of the network – onboard the connected device. Ekkono develops an edge machine learning software. In Ekkono's case, it is possible to do incremental learning at the edge, which means that the ML model continuously gets better but also that it gets personalized as it is fed with sensor data while in production.

3. Use Cases

Next we will look at a few use cases for smart BMS based on edge machine learning. The main idea is to create a BMS system that can adapt to a changing environment or the user patterns by analyzing sensor data directly on the device using machine learning.

3.1 Remaining Useful Life using Incremental Learning

As mentioned above, one common and important task for a smart BMS is to predict the remaining useful life of a battery, i.e. the number of remaining charge and discharge cycles. The battery is typically deemed to be out of service when it reaches a certain level of degradation compared to its rated capacity. The problem is that the degradation is often very slow, e.g. tens of thousands of hours, which makes collecting data from the whole lifecycle to support traditional machine learning, where a model is trained using thousands of examples i.e. battery life cycles, rather impractical. Furthermore, since the degradations depend on numerous factors such as charging and discharging patterns and which temperatures the battery cells have been exposed too, the collected data may very well be irrelevant for a new device. Finally, since these types of patterns and conditions may change over time, e.g. like for mobile IoT devices, the estimation of the remaining useful life must be adaptive. Figure 2 shows an example of how making and estimation of the remaining life at three different points in time may differ substantially. Obviously, the conditions have changed considerably during these points which explain the big deviations. If the model couldn't adapt, we would have to use a model based on data from another battery and its conditions and apply the resulting model on the new battery. Obviously, this could go very wrong in many cases if the battery data used for training is not representative for the new device, i.e. we would get any of the tree predictions in Figure 2. The only reasonable alternative is to have a model adapt using incremental learning.

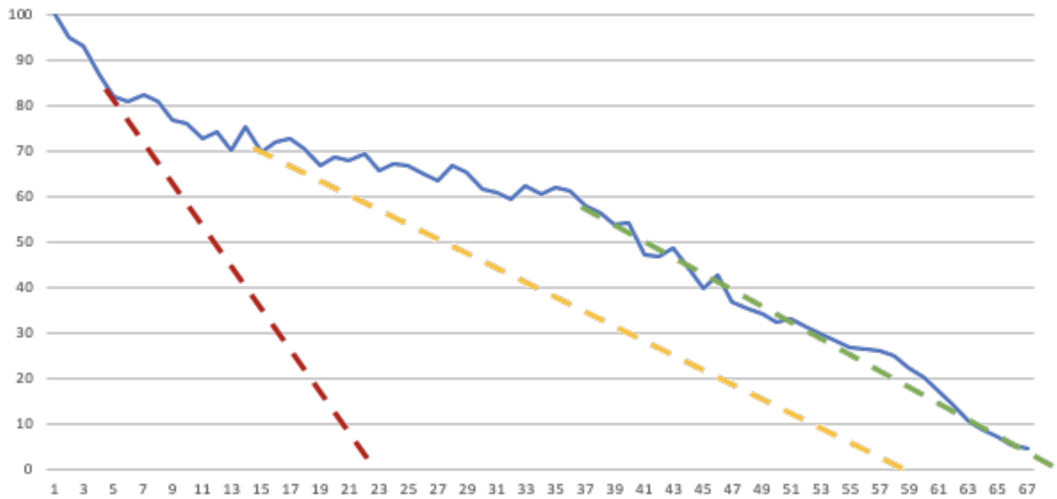

Figure 2 - Estimating Remaining useful Life at different point in time

3.2 Range Prediction by Users Modeling

A common use case for smart BMS is range prediction, i.e. how many kilometers or minutes of use can the battery deliver at the current charge. The use case can seem to be very similar to predicting the remaining useful life but here it is feasible to exploit that several charge and discharge cycles can be observed. In essences this means that we have time to learn the user's behavior between cycles and not only during a single cycle. Figure 3 shows the state of charge over time with an orange line, the power outtake i.e. the user action as a blue line. The two gray dotted lines show the estimated remaining range for two points in time when only considering the current power outtake which obviously would give very different results over time. The yellow line instead shows a prediction based on a machine learning model that takes the users previous action and the characteristics of the battery into account. If the model is accurate it will result in much more reliable and precise estimations.

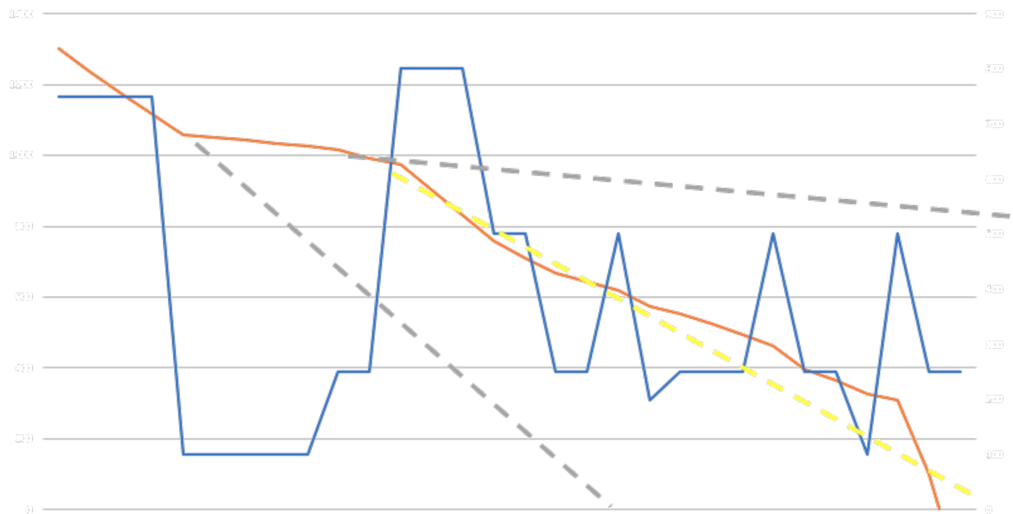

Figure 3 - Range prediction based on usage patterns

3.3 Modeling the Device to Enable Simulation

In some applications the actual power outtake of a device is complex due to several subsystems that consume energy depending on a range of environmental factors. Take an electrical vehicle for example, here the amount of throttle, the speed, inclination of the road, ambient temperature, climate system and the gross weight of the vehicle all affect the amount of energy consumed at any given moment. Furthermore, the battery can often be partially recharged during braking. If these relationships can be learnt using machine learning we can get a better understanding of the device, but more importantly it enables simulations of how the device will behave in the future during different usage patterns.

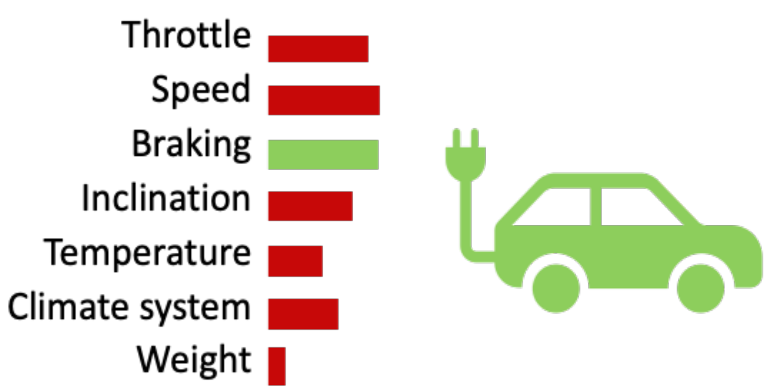

Figure 4 - Modeling power outtake of an electrical vehicle

If we take the example depicted in Figure 3 above, we could for example simulate what would happen if the user switched to "power mode" or "Eco mode" instead of continuing with its normal behavior. This kind of simulation could give valuable information to optimize the operation of the device.

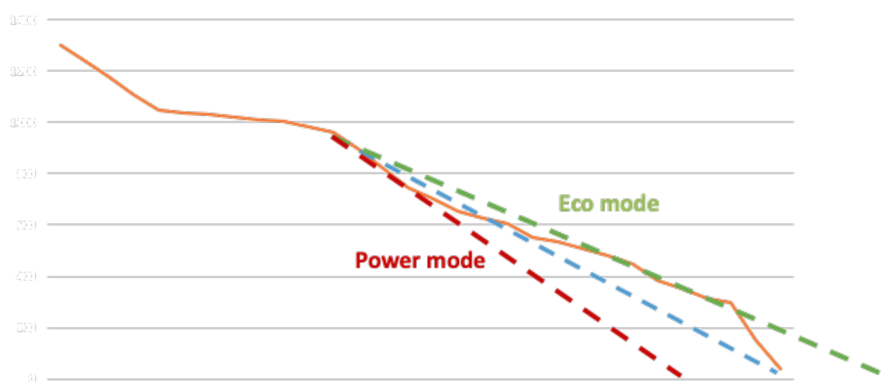

Figure 5 - Simulation of device operation for decision support

It is important to note the simulation part often can include things we actually know, or think that we know, about the future, e.g. a temperature forecast, a planned route including its foundation and predicted traffic. Combining all part's i.e. modelling of the **battery**, the **device**, the **user** and finally including known information about the **future** is what we call a truly smart Battery management system, but each part has a value of its own and can be seen as a good step forward.

Figure 6 - Information sources for a smart BMS predict range for an electric vehicle

4. Conclusions

To conclude, the nature of IoT devices and their environments require adaptive BMS tailored for the specific device and incremental edge machine learning is the enabling technology. To be truly smart a BMS system should include all possible sources of information i.e. data from the battery, the user, the device, the environment and any knowledge about future plans and events. The most reasonable approach for this scenario is to run the machine learning on a microcontroller on the device, but this also entails strict requirements on the software to be feasible. Ekkono's edge machine learning library is designed from scratch to meet these requirements and to enable incremental learning on the edge.

***

Ekkono #openfika is a short open, online fika[†] session, hosted by Ekkono, on hot, contemporary and relevant topics, where a 15 minutes presentation is followed by discussion and Q&A. Keep an eye on www.ekkono.ai and LinkedIn for the next #openfika session.

Ekkono Solutions AB is a software company that develops Edge Machine Learning. Our product is the result of seven years of research at the University of Borås, Sweden, and assists product OEMs in different industries to rapidly develop smart features onboard their products, using machine learning to make them self-learning and predictive. For more information, visit www.ekkono.ai.

[†] fika (wikipedia.org); Swedes have fika (pronounced [ˈfîːka]), meaning "coffee break". The tradition has spread throughout Swedish businesses around the world. Fika is a social institution in Sweden and a common practice at workplaces in Sweden. Fika may also function partially as an informal meeting between co-workers and management people, and it may even be considered impolite not to join in.