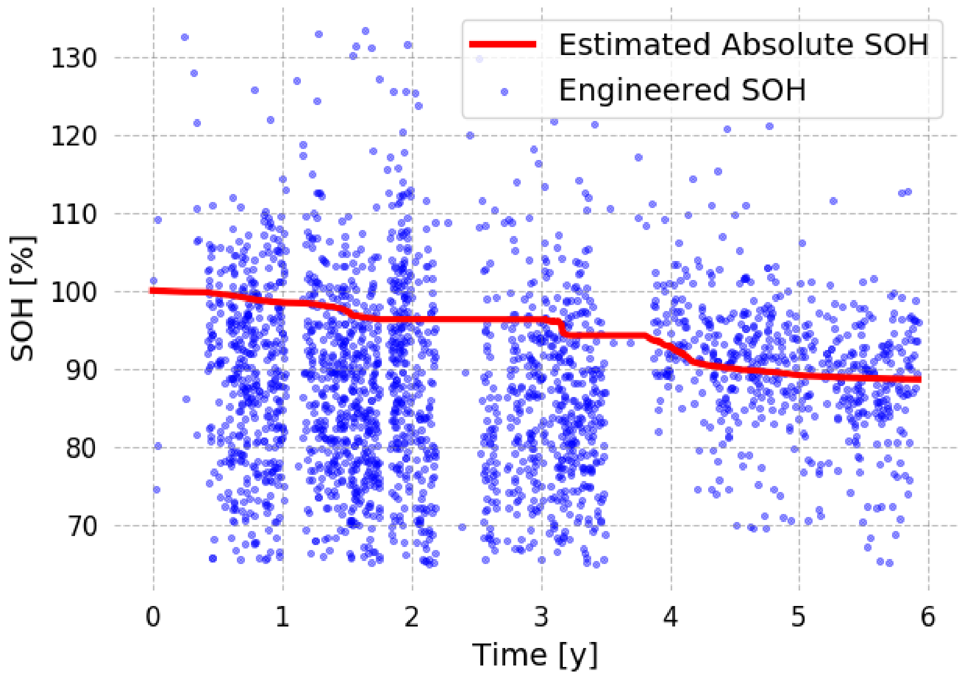

A Neural Network Approach to Absolute State-of-Health Estimation in Electric Vehicles

Battery Degradation Study Based on Fleet Data

HERMAN JOHNSSON

Department of Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2018

Master's thesis 2018

A Neural Network Approach to Absolute State-of-Health Estimation in Electric Vehicles

Battery Degradation Study Based on Fleet Data

HERMAN JOHNSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Physics
Division of Condensed Matter Physics
Chalmers University of Technology
In collaboration with Volvo Cars Corporation
Gothenburg, Sweden 2018

A Neural Network Approach to Absolute State-of-Health Estimation
in Electric Vehicles
Battery Degradation Study Based on Fleet Data
HERMAN JOHNSSON

Supervisor: Christian Fleischer, 97410 Traction Battery Volvo Cars Corporation
Examiner: Patrik Johansson, Department of Physics

Department of Physics
Division of Condensed Matter Physics
Chalmers University of Technology
In collaboration with Volvo Cars Corporation
Gothenburg, Sweden 2018

Cover: The original SOH signal based on capacity together with the network's estimation of an absolute SOH.

A Neural Network Approach to Absolute State-of-Health Estimation in Electric Vehicles
Battery Degradation Study Based on Fleet Data
HERMAN JOHNSSON
Department of Physics
Chalmers University of Technology

Abstract

Electrification is a trend within the automotive industry. Many car manufacturers are launching electric vehicles, which are believed to be more sustainable and environmentally friendly. A major component in these cars is the battery, and its performance is crucial to the success of the electric vehicle. Therefor, the degradation of battery properties is interesting, especially the capacity decline. To understand and counter this degradation it must be measured with high precision in the cars, and be connected to car use. This project approaches this challenge by: using real fleet data, the aggregation of the data into events, and a neural network to estimate the state of the battery. The result is a proof of concept that gives an improved measure of the battery state and how different usage affects the capacity degradation. The result is, however, not validated at this point, shows unexplained properties, and should be further developed.

Acknowledgements

This has been an interesting and challenging project, with a strong purpose. It was initialized by Christian Fleischer, who also supervised the project. A big thank you, Christian, for making the project a possibility and for many good discussions.

Thank you to the many people within Volvo who has been involved with the project. We have met in meetings, discussions, and you have given advises, and criticism. This has been very helpful and interesting. A special thanks to Anders Bern for helping me access the data.
Also, Patrik Johansson, the examiner, has provided good guidance towards a final version of the report and the project, thank you.

Herman Johnsson, Gothenburg, June 2018

Contents

List of Figures

List of Tables

1
Introduction

The automotive industry is experiencing a clear trend towards electrification. Electric vehicles are believed to be more environmentally friendly and sustainable, and more in line with future technologies and regulations, like autonomous cars and emission free cities. The success of the electric car, both from an environmental and commercial perspective, is heavily connected to the sustainability of the battery. If the battery needs to be replaced after a few years, or the performance is severely degraded, the impact on both the environmental footprint and business case will be substantial. Even though batteries have been used for a long time, the effect of highly dynamic and rough use in electric vehicles is not yet fully understood.

Battery degradation, or aging, includes many symptoms, capacity decrease, resistance increase, etc. These are macroscopic effects stemming from multiple processes inside the battery, constituting a very complex system. The degradation is often measured in relation to the initial properties, for instance, the capacity of the battery is 90% of the initial capacity. This measure is in general named the State-of-Health (SOH).

In this report, a data driven approach using Neural Networks is suggested as an alternative to the battery SOH measurement in the car. The approach is chosen to account for the many parameters connecting to aging in the battery, both intrinsic to the battery and associated with car usage. The work is a master's thesis at the Physics Department of Chalmers University of Technology in collaboration with Volvo Car Corporation. The work has been carried out in the spring of 2018 in Gothenburg, Sweden.

1.1 Background

Electric vehicles is a hot topic and the use of both hybrid and fully battery powered cars have increased in recent years. Interest in electric vehicles has sparked a battery cost decline, further research and new mass production methods. Continued improvement of battery properties, such as energy density, is narrowing the pricing gap between electric vehicles and the more traditional combustion engine vehicles. The more environmentally friendly electric car could also be a key asset to better take care of our environment. The number of battery electric cars in the world is now 3 million, and a massive increase is predicted, see figure 1.1.[3] This can also be seen in the number of companies that have decided to launch electric models. Very interesting, and relevant to this project, is Volvo Cars Corporation's goal to have half the fleet electrified by 2025.[4]

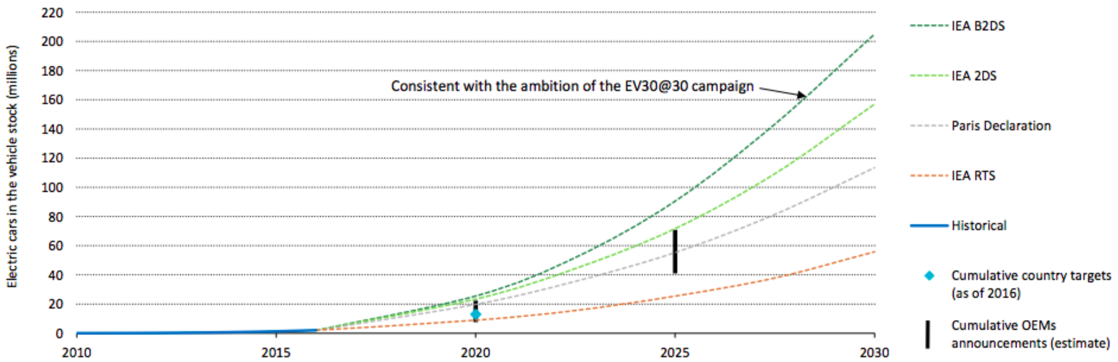

Figure 1.1: According to the *International Energy Agency's* report *Global EV outlook*, the number of electric vehicles used in the world could rise to between 9 and 20 million by 2020.

This is a big change for the car manufacturing companies. Instead of continued construction of in-house combustion engines, well understood and developed under long time, electric engines and batteries are needed, less understood and often bought from suppliers. The nature of the battery, a chemical system with very little transparency, makes it difficult to fully understand. Still, there is a harsh competition in the car market, and companies need to deliver good quality and long warranties. This poses an interesting challenge; to promise quality products without full control over the components.

Batteries, and specifically lithium-ion batteries, degrade with time and use. A new battery has a certain capacity, resistance, etc., that gives certain performance. Over time and with use, these properties changes, often in a way that gives a lower performance. For instance, as the capacity declines, the range of the vehicle decreases. This is generally called aging, and is often measured as State-of-Health (SOH). For instance, a capacity based SOH of 95% means that the current capacity have degraded 5% from the original. An SOH of 70-80% is often defined as End-of-Life (EOL) for a battery. To rival a combustion engine car, a lifetime of tens of years is needed, which limits the annual degradation to only a few percent.[5]

Battery degradation has been studied in laboratories in extent, but this knowledge does not fully cover the highly dynamic and differentiated use in electric vehicles. Some cars are driven carefully, in warm climates, with long drives and long charging events, others are used in a cold climate for short trips and are subject to many short charging events. The number of parameters affecting the battery in this use makes a laboratory testing expensive and tedious. Testing all aspects in the context of aging means artificially age batteries in the laboratory, which can take a very long time.

The capacity degradation is perhaps the most important symptom of aging, but the capacity measurement is not trivial in an electric vehicle. The simplest, and most accurate, form of measurement would be to perform a complete discharge. However, this is not often done in a car, and the capacity might vary with different circumstances. For instance, it is known that the capacity of a battery is dependent on temperature. In fact, it is possible that an older battery in a warmer temperature has a higher capacity than a new battery at a colder temperature.

The high stakes, difficulty to accurately measure SOH, and differentiated use of batteries in electric vehicles, makes the issue of battery degradation in electric vehicles very interesting. It is probably one of the issues that will determine the success of the electric vehicle.

1.2 Purpose

An understanding of how battery usage in electric vehicles accelerates aging could make it possible to avoid certain use, and thus extend the lifetime. From a commercial perspective, knowing the remaining useful lifetime could make it easier to schedule service, and issue warranty. From an environmental standpoint longer lifetimes means a smaller footprint, especially since battery recycling is currently limited.[6]

The project is aimed to evaluate if a data driven approach, founded on vehicle data and machine learning techniques, is usable to determine the SOH and map it to battery usage. The feasibility of the approach is determined by the capability to determine an SOH so that different car usage can be associated with different aging rates. This would allow for in vivo studies of batteries in electric vehicles without having to measure the SOH under specific circumstances.

1.3 Limitations

Due to the exploratory nature of the project, many steps are required from choosing a data set, to estimating an SOH. This poses a limitation since every step could be evaluated, refined, and perfected. Instead, to reach the end of the line, such refinements are omitted, creating a thinner but deeper project. This means that some choices are made and used as long as they work. It leaves it up to future endeavours to examine other approaches to each step.

The SOH can be based on many different properties, but this project limits the search to a capacity based SOH. This means that the SOH is a normalization of current capacity of the battery with respect to the initial capacity. Alternative properties, such as resistance, leakage current, and more, are left for future work.

Machine learning encapsulates an entire field of different techniques. Here, due to time limitations, the focus is on Neural Networks. For optimal estimation accuracy additional preprocessing units, namely data aggregation, scaling, and feature engineering, will be deployed. Other criteria that decides the feasibility of using a machine learning algorithm in a car will not be accounted for. First, the reliability of the system is correlated with the transparency of the algorithm, which means that the workings behind the estimation should be known. Second, to be implemented in a car the software is can not be computationally heavy, needs to be quickly executed, and robust against noise and data loss. These specifications are all left for future work.

Any machine learning process is highly dependent on the data set. Since the project aims to provide information about battery powered cars, the optimal data set comes from a large number of these cars. However, these cars are not yet implemented on a large scale. Therefore, the outcome of the project is very much determined by the data available at Volvo Cars Corporation.

1.4 Problem Statement

The ultimate goal of being able to fully understand battery degradation in electric vehicles depends on the possibility to measure the SOH and compare it to car use. The problem is that measurements of capacity, and thereby the SOH, are noisy and dependent on the measurement circumstances. Even the software used in one of the newest cars can produce values that vary with amounts far exceeding the annual degradation, see figure 1.2. The SOH measurement from this car has a clear seasonal dependence, better SOH in the Summers, and the SOH can also vary up to 20% over shorter periods. It is therefore very difficult so see a general trend in the battery degradation, yet even more difficult to connect certain usage to accelerated aging.

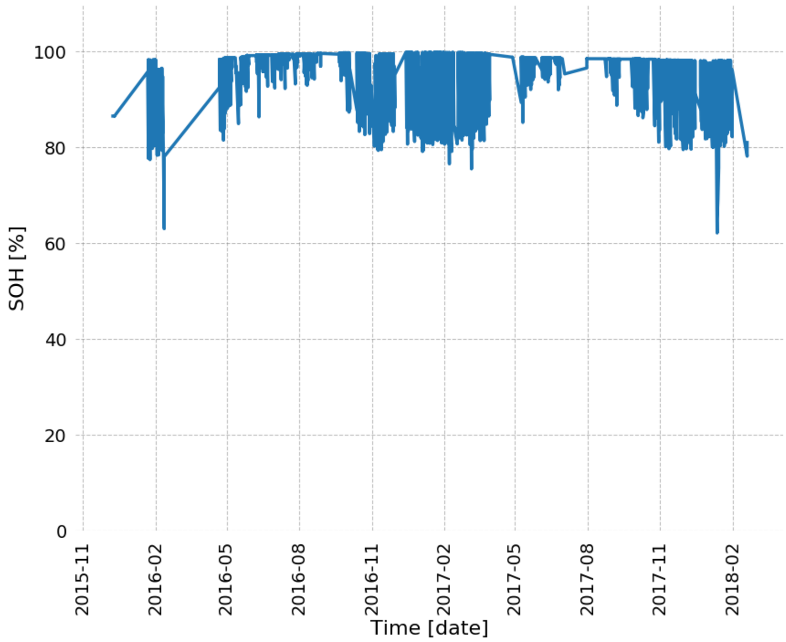

Figure 1.2: The SOH measurement from a 2015/2016 car from Volvo. SOH is based on capacity measurements done while charging. There is a clear seasonal dependence, giving values closer to 100% in the Summer and lower in the Winter. Additionally, there are short interval changes stretching a much as 20%, far exceeding the supposed annual degradation.

Instead, one would like an absolute SOH measurement. A measurement that does not depend on the outer circumstances. An equivalent to perform a standard measurement of capacity, and thereby the SOH, but that can be conducted on-board the car. This measure should be able to be connected to car usage so that different rates of degradation can be associated with different uses. This leads to the problem statement of this project:

> **Could a data driven approach using Neural Networks constitute a feasible option for capacity based State-of-Health estimation in Electric Vehicles and provide information about how car usage age the battery?**

1.5 Guide to understand the report

The project tackles a difficult question, and the approach to finding answers is relaying on multiple steps. The result is a multiple step process from theory and available data, to preprocessing, and the use of a neural network. This is reflected in the report, which has multiple chapters.

After the introductory chapter a brief theory and data description is presented in chapter 2. This is aimed to orient the reader with respect to the starting point of the project.

Chapter 3 presents a prestudy, where an SOH is estimated from lab data. This is done to give a fundamental overview of the different time scales involved.

Chapters 4, 5, and 6 deal with steps done towards the final result. Each chapter includes the method used, results given, and a brief discussion.

The final two chapters, 7, and 8, aim to reconnect the different steps to each other and evaluate their collected success of solving the main problem statement.

2
Theory and Data Set

The topic involves two different fields, lithium-ion batteries and neural networks, and the development of a model based on real fleet data from Volvo Cars Corporation. This means that the project was based on theory about lithium-ion batteries and neural networks, and on the available data set. This chapter aims to orientate the reader according to this starting point.

2.1 Lithium-ion Batteries

The lithium-ion battery is a common choice for electrical energy storage since it has high energy density and slow aging compared to many alternatives.[7] There are many different configurations of cells; different geometries, materials and sizes, but a simple schematics of a lithium-ion battery cell can be seen in figure 2.1.

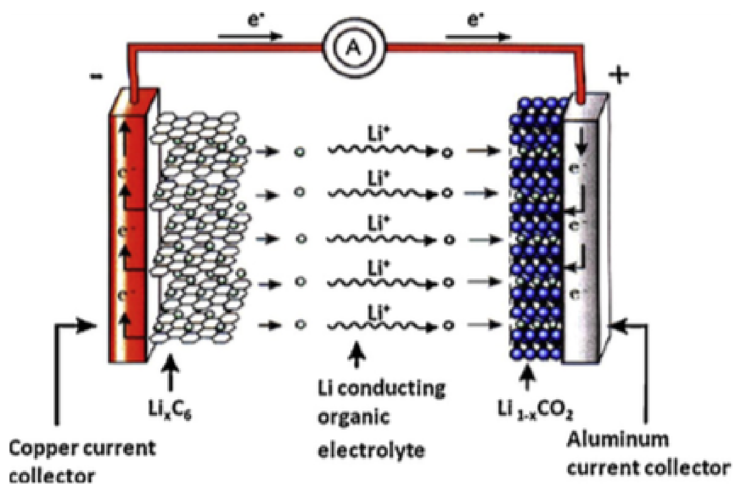

Figure 2.1: A simple sketch of a lithium-ion battery cell. The cell contains a cathode and an anode, and separated by an electrolyte.[1]

All cells have an anode and a cathode side, a negative and a positive pole. Lithium, the active material of the cell, can flow from anode to cathode and from cathode to anode, discharging and charging the battery.[8] Both the anode and cathode materials are layered structures where lithium can be intercalated or extracted. They can be of many different materials, for instance graphite, and various cobalt containing materials.

Between the electrodes is an electrolyte, often a lithium salt. It does not let electrons pass through, but allows a flow of lithium-ions. At the anode, the electrolyte is operated outside the stability window and thus can be reduced. This consumes lithium and creates a solid interface, called the Solid Electrolyte Interface (SEI), that protects from further reduction.[9] Besides fundamental to the function of the battery, this also provides stabilization to the graphite in the anode.[10]

2.1.1 Aging

The properties of battery changes with time and use. Many degrade, causing the abilities of the battery to decline, generally called aging. A battery might age differently depending on how it is used, referred to as cycle aging. However, a battery also degrades without being used, referred to as calendar aging.[5] The extent of these effects differ for different materials, morphology, design etc.[11]

In a broad view, different causes of aging can be put into three categories; Loss of Lithium Inventory (LLI), processes that make the lithium unusable for cycling, Loss of Active Materials (LAM), reduced amount of material enabling the lithium transfer, and structural damage to the components of the battery. All of them include many physical and chemical processes. The LLI is deemed to be the most severe aging mechanism, which is accelerated by high charging rates and very low or very high temperatures.[12]

On a macroscopic level, aging is the effect from many ongoing processes in the battery, see figure 2.2. These constitute the complex inner workings of the battery and makes aging difficult to fully pinpoint. Going through them all is outside of the scope of the project, but this section aims to convey the complexity of the system.

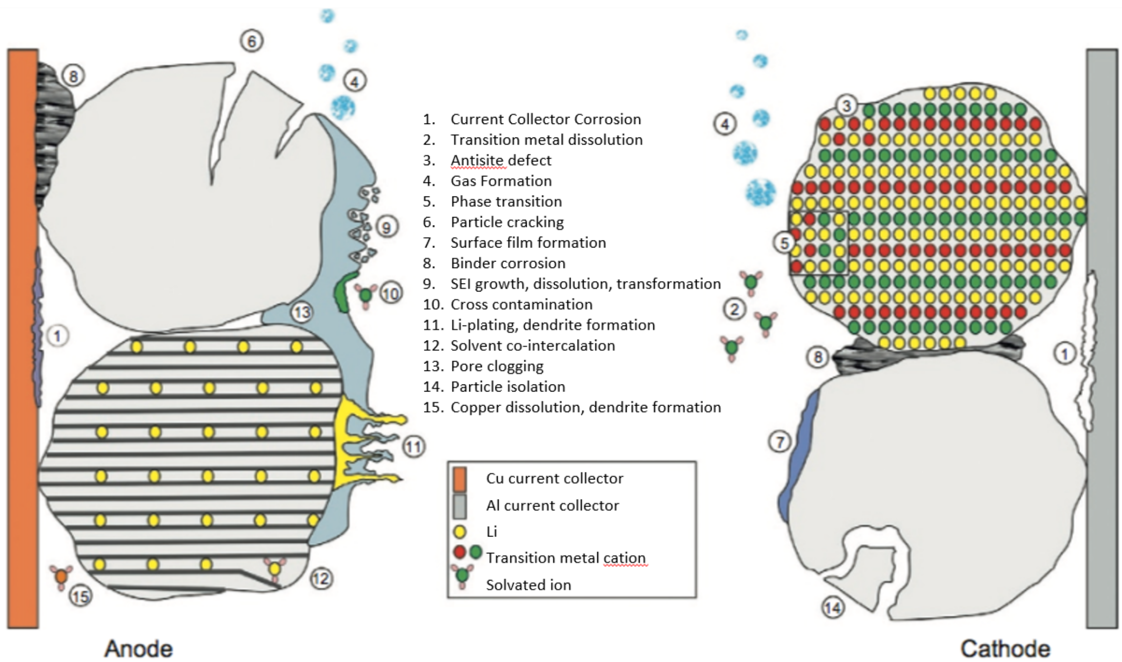

Figure 2.2: Map of different aging phenomenon occurring inside the battery cell.[2]

Just as there are many different phenomena that age the battery, there are multiple parameters that induce or accelerate them. These are often found through empirical studies, or from experiments validating theoretical knowledge. From a literature based research the following parameters were found important to aging in batteries and to the projects.

Parameter	Source
Temperature (High)	[2, 13, 12, 14, 12]
Temperature (Low)	[15, 11, 16, 9, 10, 12]
Relaxation time	[14, 17]
Pressure	[15]
High SOC	[13, 2, 16, 15, 18]
Change in SOC	[19, 18]
Voltage	[20, 21]
Number of Cycles	[17]
Time	[5]
Current density (high)	[11, 9, 16, 10, 12]
Charge rate (high)	[17]
Cell imbalance	[9, 11]

Table 2.1: Parameters affecting battery aging

2.1.2 Battery States

Battery performance is often described in battery states. These are artificial states of the battery that are suitable for describing the operation, but can often not be measured directly. Instead, they are the aggregation of many different measurements. The states are then continuously estimated and used to make sure that the battery is operated within the safe operation window.[19]

State-of-Charge (SOC) is the amount of charge stored in the battery, presented in percent of total capacity, 0-100%. It is the closest analogy to a fuel gauge in a petrol engine car.[19] The SOC can be estimated in many ways. The most reliable and intuitive way is to completely discharge the battery under known conditions and measure the amount of charge drained. The battery will be depleted and the measured SOC, however known, will no longer be relevant.[22]

A related method is Ampere counting, sometimes known as Ampere-hour integral. The idea is to integrate over the current to get the change in SOC from a known starting point. Coulumbic efficiency η is used to scale the integrated current and expresses the success rate of moving lithium ions.[22, 19] The integral is then normalised by C, the capacity. The mathematical expression for coulomb counting can be seen in equation 2.1. Ampere counting suffers if the measurement of the current is imprecise. Also, it is often hard to know the initial SOC and the coulumbic efficiency, but, if these are known, the method could provide an accurate measurement of the SOC.[19]

$$SOC = SOC_0 + \frac{1}{C}\int_{t_0}^{t} \eta I d\tau \tag{2.1}$$

An alternative method is to connect the open circuit voltage (OCV) to the SOC. This is valid if the battery is in equilibrium, which it is after a longer resting period.[22, 19] Because of this, it is not a suitable method for continuous estimations in a car, but could be suitable for parking. The OCV curves are often taken from individual cells and are mapped to the SOC of the entire battery via a lookup table. Novel versions account for how this mapping changes with time, but more often the lookup table is constant throughout the battery life time.

Depth-of-Discharge (DOD) is often used instead of SOC. They are measuring the same change in stored charge related to the capacity of the battery. However, the DOD measures the deviation from a full battery, 100%, and the SOC from a depleted battery. This means that the relationship between DOD and SOC can be described by: $SOC = 100\% - DOD$.

The State-of-Health (SOH) is a general measure of how well the battery is doing. It is often used to compare the current capability of the battery to the initial 100%. To do so, it is connected to a property of the battery, often capacity, resistance, self-discharge rate, or power density.[22, 5, 5, 18] All these will degrade with time and use, and SOH is therefore closely connected to aging. In fact, it might be considered a tool to measure aging in the battery. Unfortunately, these properties will degrade differently, and if the SOH is truly a measure battery aging, it should include, or combine, the degradation of multiple properties, not just one of them.[10]

Perhaps most used is a capacity based SOH, how much charge the battery can store. This parameter will decrease with age, decreasing the amount of energy that can be stored in the battery. Such degradation is important since it is closely connected to the range of the car. It also means that an older battery will charge quicker since there is less available lithium for transport.[19]

Related to SOH are the End-of-Life (EOL) and the Remaining-useful-Lifetime (RUL). EOL is often referred to the point where the battery is no longer useful and is popularly defined as 70 or 80% of the SOH. The RUL is the number of cycles left before reaching the EOL.[19] Being able to estimate these states is very important to operate the batteries in the best way, and for many business related reasons.

2.2 Artificial Neural Networks

Neural networks are in some sense a subcategory of artificial intelligence and a very powerful tool, used in modern search engines and self driving cars, to mention a few. They are often a way to approximate a system without knowing the underlying structure of said system, reminiscent of how human intuition works. Actually, artificial neural networks try to mimic the structure of a human brain, having approximations of neurons and synapses.

The building block of the brain is the neuron cell, see figure 2.3. It is a cell with remarkable properties to receive electric impulses, process them and send an impulse out. The inputs can be received through one or many of the dendrites, attached to the cell body. The pulses are processed and an output is generated and sent through the axon.[23] The actual process is complex and not fully understood, but, we know enough to try to mimic the it.

The artificial neuron, see figure 2.3, is the smallest unit of the artificial neural network. It has, just like its biological counterpart, one or many inputs, which can be continuous or discrete. Each input is multiplied by a specific weight and the result is summed up in the neuron together with a bias term. The sum is often transformed by a nonlinear function to give a nonlinear effect to the system. The nonlinear function is often a logistic function, step function, sigmoid, or tanh. This neuron is often referred to as a Perceptron.[24] Both the weights and the bias term are part of the intrinsic variables of the neuron, and can be changed, which is called learning.

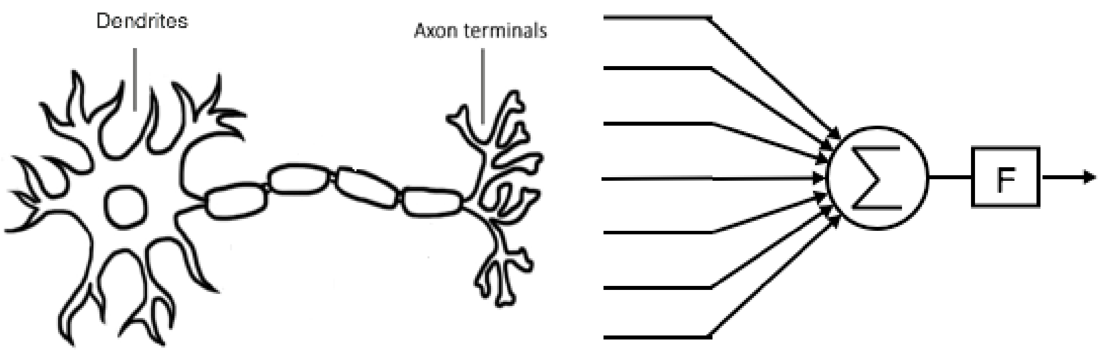

Figure 2.3: Left, a biological neuron. Electrical signals form inputs fed to the dendrites. The inputs are processed and an output signal is sent via the axon. Right, an artificial neuron. It multiplies the inputs with weights, then sums them up, and perform a nonlinear function F.

The artificial neurons can make simple evaluations. As an example, imaging a neuron with two inputs (x,y), representing the coordinates in a Cartesian coordinate system, and a step function as a nonlinear function. A color, blue or red, is assigned to a collection of points so that a point is red if it lies above the line $y = x$. The neuron is fed the coordinates of a point and is asked to predict 1 or 0, red or blue. The output is compared to the actual color of the point and if the prediction is wrong, the weights and bias are adjusted. This is called training, and can be done until a state of minimum error has been found. A training run is called an epoch, and three of them for the posed example can be seen in figure 2.4.

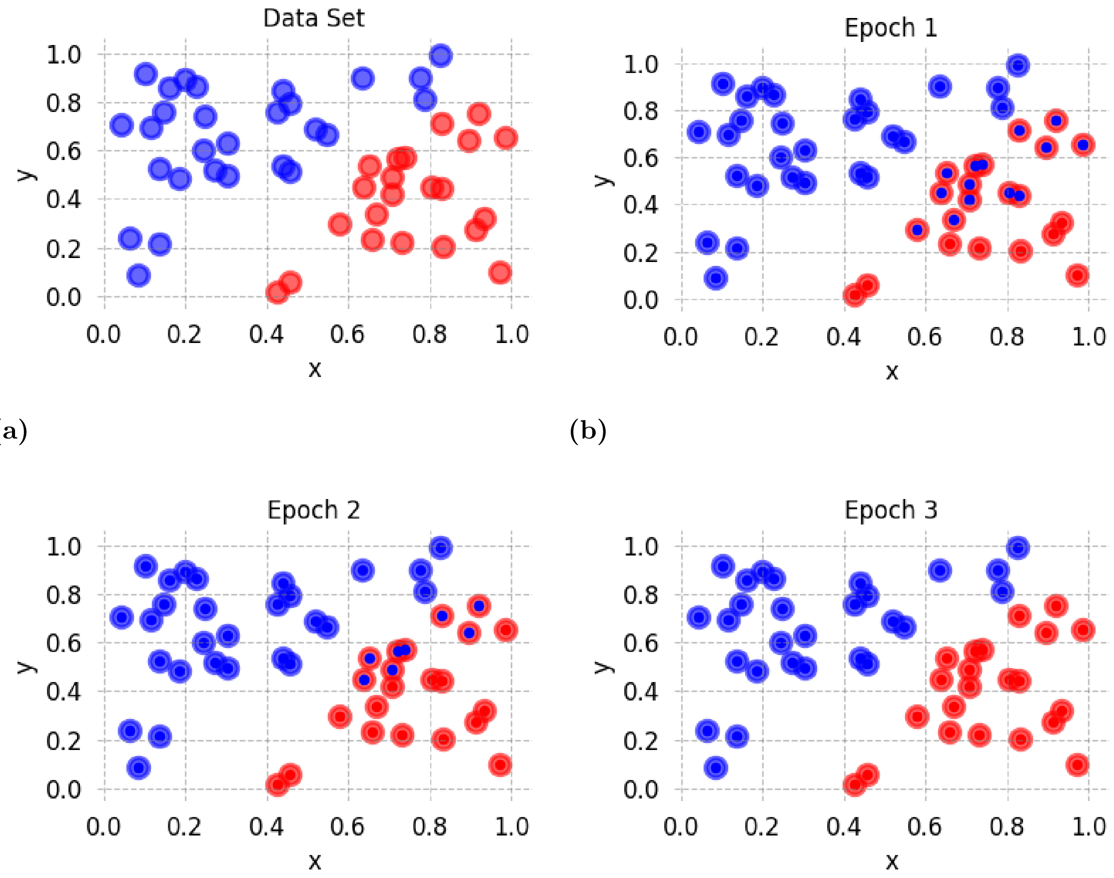

(a) (b) (c) (d)

Figure 2.4: a) the initial data points. b), c) and d) training epoch 1, 2 and 3. The predictions, marked by a smaller point in the predicted color, is improved in every epoch.

In all the graphs of figure 2.4, there seems to exist a clear line, dividing the red or blue predictions. This line is called the decision boundary. This is a property of the artificial neuron, it makes decisions based on a single line, it can thus solve only linearly separable tasks. With every training epoch the predictions become better, and the line changes. By the last epoch, all predicted colors agree with the actual colors for the points.

Not all tasks are linearly separable, and to solve these neurons can be connected to each other. Using multiple neurons provides multiple decision boundaries, this collection of neurons is often referred to as a layer. For real hard problems like image recognition, many layers are joined in a sequence, creating a network. If all neurons in one layer connect to all neurons in the next, the network is called fully connected, see figure 2.5. The decision boundaries of the network is not so intuitive, making the decisions, and thus the system, a black box. A network with many layers are called deep and the use of such is often called Deep Learning.

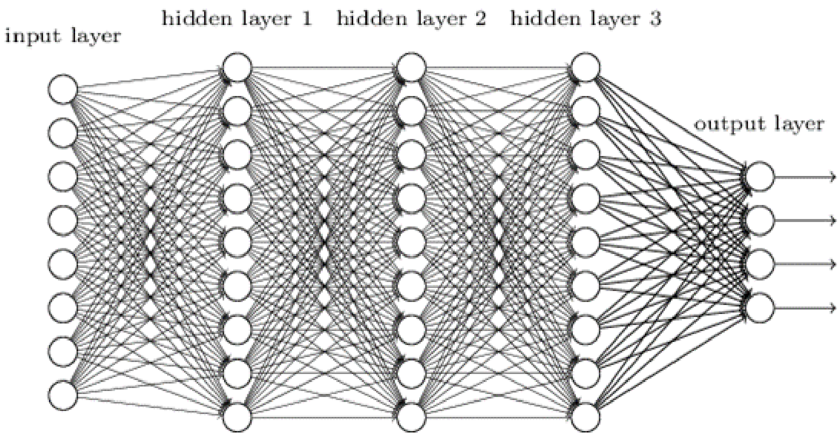

Figure 2.5: Many neurons interconnected form a network. It is fully connected if all nodes in one layer are connected to all nodes in the next. A network with many layers is called deep.

The networks are trained in the same way as a neuron; inputs are fed to the first layer, the information propagates through the network, and the output is given. This output is compared to the correct answer, and the error is often called the cost function. Training means that weights and biases are adjusted to minimize the cost function. The most used tool to make these adjustments is Gradient Descent, meaning that all weights are adjusted along the gradient of the cost function. This provides a descent to a cost function minimum, but it is never guaranteed to be the global minimum.[24]

There are many parameters that specifies the architecture and the training of the network, called hyperparameters. For instance, the number of nodes and layers, and the non-linear function. Further, the adjustments in every training run are scaled by a factor called the learning rate, and the training is often done, not sample by sample, but, by batches. This means that batches of samples are fed to the network and the total cost of the batch is used to adjust the network. This is called batch training, and the size of a batch is another hyperparameter.

The data set used is divided into training set and validation set. The training set is used for training and the validation set is used to evaluate the performance of the network. This is done to make sure that the network learns the structure of the system and not just memories specific samples.

Networks can be trained to well. This is referred to as overfitting and is the process of a network learning the noise of the data, irrelevant to the aim of studying the mechanisms of the system. Often, this can be seen by better performance in the training than validation set. To account for this, drop out can be is used, meaning neurons are reset with a certain probability.

The Universal Approximation Theorem can be applied to a neural network. It states that a network can be built to approximate any continuous function arbitrary well.[24] This means that a network has the capability to mimic any function, it is just a question of having the correct architecture, and training it properly to avoid local minima. This is, however, not an easy task.

The neural networks have an advantage, a model of the system is not needed. Instead, provided a good architecture, the network will learn to map inputs to outputs. This can perhaps be related to humans learning to predict an outcome by experience, regardless of any knowledge of the system's inner workings.

Figure 2.6: A C30 battery electric vehicle from Volvo Cars Corporation.

2.3 Data set from Volvo Cars Corporation: C30 Battery Electric Vehicles

In 2012 Volvo Cars Corporation launched a small fleet of battery electric vehicles of C30 model, see figure 2.6. Around 150 were released in generation 1 and an additional 150 in generation 2, which was launched a few years later. All these cars were equipped with additional measurement capabilities, and the ability to send these measurements to Volvo Cars Corporation via the GSM network.

The data is collected with an event based logging system, meaning that data is logged only when any of the signals changes value. This means that a car can remain idle for weeks if not used, perhaps even longer, and thus no data is logged. On the other hand, intensive use can give a sampling frequency higher than once per second. There are also built-in systems that make the car wake up and send a signal with a set interval, if it is not already in use. Six years of monitoring, therefore, provides a large collection of data points, separated by everything from seconds to weeks.

While all cars in the fleet have been used in the project, only a subset is used to study aging. To obtain a large, homogeneous population of cars, measured for a long time, only data from the first generation car is used. This is done because there might exist differences between the generations, both in the battery and the overall specifications of the car. Further, only cars with all sensors intact for a large part of the time since 2012 are used. This leaves a population of 85 cars.

The 85 cars have been driven around the Gothenburg area. The range is limited to 150km when new and the 25kWh battery takes 7h to charge at the highest allowed charging current. The usage of the cars in the fleet can be seen in figure 2.7, in terms of data points captured, the time between first and last measurement and the driven distance.

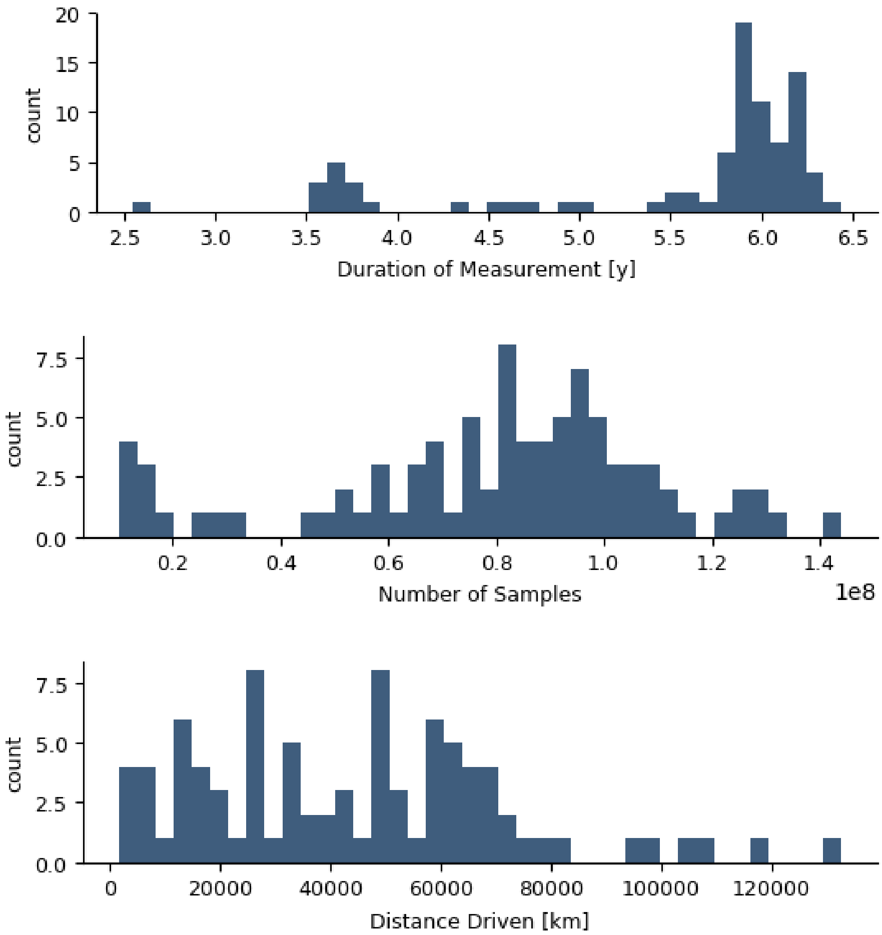

Figure 2.7: Statistics to better understand the C30 data set. The number of measurement points has a large spread, while the number of years monitored has a smaller. This gives a good view over the fact that the cars do not log data while idle.

The cars have been driven in very different ways according to the distance histogram. Some of them have only been driven slightly, around 10000km, while others have been used extensively, around 120000km. This should give a suitable population for studies of aging since both calendar and cycle aging should be prominent. According to the range of the cars, the batteries have been run through hundreds of cycles.

There are about 60 signals recorded in the cars. These are used for many purposes, not only battery monitoring. The signals have been inspected and the ones deemed important for the study can be seen in table 2.2. There are possibly more signals that can be usable but to limit the search, all could not be examined.

Signal	Description	Resolution
Ambient Temperature	Measurements of the ambient temperature	$1^o C$
Battery Current	Current to and from high voltage battery	1 A
Battery Mean Temp	Mean temperature of the cells	$2^o C$
Battery Voltage	Voltage over the battery	1 V
DOD	Depth-of-Discharge	1%
Total Distance	The total distance the car has driven	1 km
Time	Absolute time	1 s

Table 2.2: Selected signals from a total of ca. 60 signals logged in the cars.

The DOD signal is of particular interest since it is estimated on-board the car in two ways: First, coulomb counting is used to change the DOD from an initial value, see section 2.1.2. And then, if DC current through the battery is lower than 2A for more than 45 seconds, the DOD is updated based on the cell OCV and a look-up table.

The battery temperature measurement is an average. There are multiple sensors located inside the battery pack. All of them are measuring the temperature at their respective position, but only the mean of all values is logged.

3

Prestudy: SOH Estimation with Lab Data from NASA

To explore the fundamentals of SOH estimation in batteries the easiest attainable situation is used, lab data. The aim of this chapter, acting like a prestudy, is to explore SOH degradation as a function of time, and get an understanding of the different time scales in play.

The question of different time scales is very interesting when looking at battery aging. Battery operation allows for changes in voltage, current, temperature and other variables in a matter of seconds, maybe less. On the other hand a battery in an electric car is expected to reach EOL after thousands of cycles, which is years with normal usage. With such a difference in time scales, is it possible to measure instantaneously the current, temperature, and voltage of the battery and estimate the battery age?

3.1 Method

For this initial research, lab data from three batteries was used. All of them are cycled until EOL in identical environments, and with the same cycles, except with different DOD. While cycling, the temperature, voltage, current, and time is continuously measured. At discharge, the current drained is measured to get the capacity of every cycle, and the SOH is the normalization of this value. There are roughly 150 cycles per battery. The data set is available at NASA.[25]

To map the features to the measured capacity, stated as an SOH, a simple neural network was used. It had two fully connected layers, and all neurons had a ReLu activation function. It was trained by gradient descent, without any drop out applied or other preprocessing unit. The simple architecture is presented in figure 3.1.

In the first trial, instantaneously measured parameters of the batter: temperature, current, and voltage, were used as input features, and the SOH was used as the label. This was done with hopes that the features combined can provide enough information to estimate the SOH. This can be viewed as a system without any memory, since the history of the battery is not recorded and used.

The second trail aggregated the data into cycles. Instead of instantaneous values, the features were now values per cycle. Specifically, the number of cycles done, maximum DOD, discharge mean temperature and charge mean temperature. The same network was trained in the exact same way.

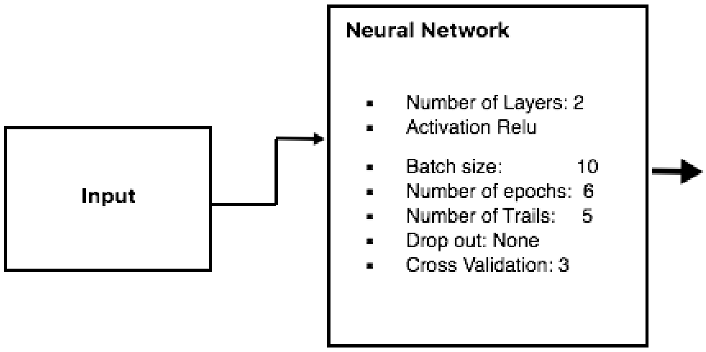

Figure 3.1: The architecture of the neural network used to estimate the lab measured SOH values.

3.2 Results

The results for the first trail, using instantaneously measures values as input features can be seen in figure 3.2. It is very difficult to get the network to estimate the SOH just by the instantaneously measured parameters. This is intuitive since most of the voltage, current, and temperature values will be present in every cycle, regardless of SOH.

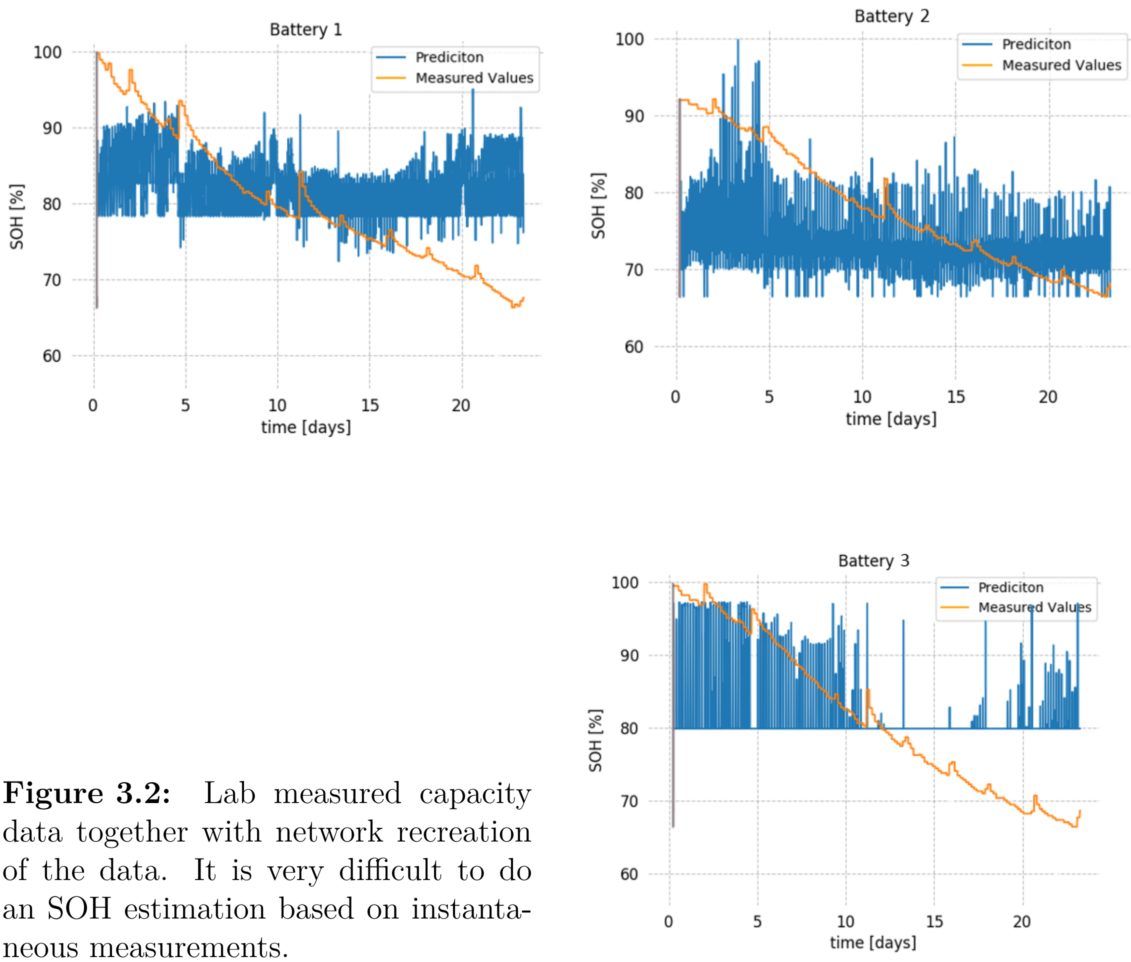

Figure 3.2: Lab measured capacity data together with network recreation of the data. It is very difficult to do an SOH estimation based on instantaneous measurements.

The second trail, using cycle based values as features, did a better job of estimating the SOH, see figure 3.3. The result is still noisy and in some cases deviating from the true value. This can have a number of different explanations, first of them being noisy data. The data is still collected in intervals that are small compared to the aging timescale and the SOH change between cycles is small. Evaluating the graphs, there are also unexplained spikes in the true values. These could be due to unrecorded idle time between cycles, error in the measurements, or errors in the way the SOH is calculated. The model can also be seen to overfit to some of these steps.

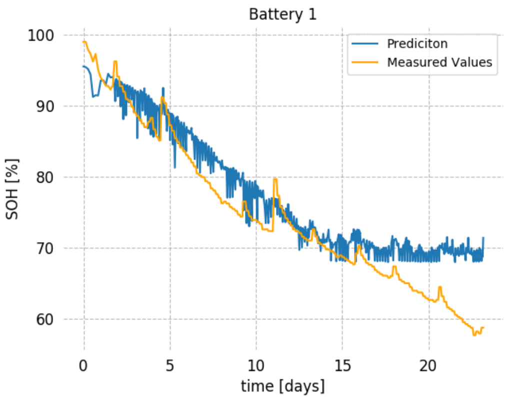

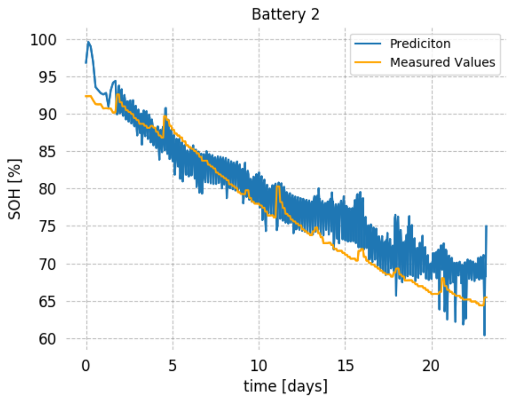

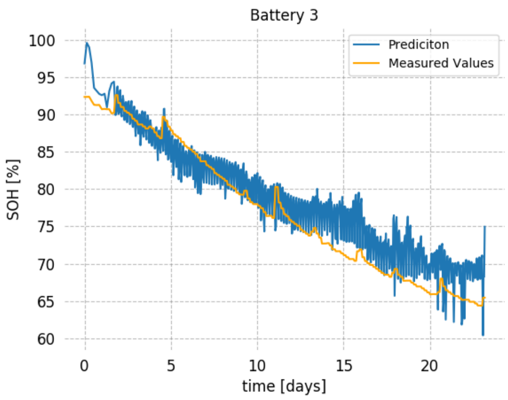

Figure 3.3: Lab measured capacity together with the network created values. The network does a better job recreating the values when the inputs are of the same time scale.

3.3 Discussion

It is evident when evaluating aging of batteries that it is difficult to use instantaneous measurement to detect aging. The neural network does not tell if it could be done, but shows that it is at least difficult. Also, even with a very simple cycling all the way to EOL, the model gives a noisy output that must be improved if it is to be used in a car. In the purpose of progress, it can be concluded that the data from the cars must be aggregated to a time scale that is more compatible with the aging.

The data set for these trails are limited. Three batteries might be too few to correctly capture the mechanics of aging. This could explain the noise results and the poor accuracy.

4

Preprocessing: Washing and Aggregating the Fleet Data

Based on the information in chapter 3, the signals from the cars, being instantaneous measurements, are not easily mapped to aging. Additionally, a car monitored for six years will give billions of data points, too much for normal algorithms to handle. Instead, there are many layers of refinements, categorization, and data aggregation needed before the data is in a suitable format. The following chapter is a walk-through of the steps taken from the time series data to events experienced by the cars.

4.1 Method

The idea is to connect the SOH degradation to the history of the battery. For this, two different types of signals were used; events and accumulated signals. In this project, events are defined as a clearly limited, larger time interval experienced by the car. An event could for instance be a drive or a charge, but also be more specific, for instance a long drive or a charge in cold temperature. The history of the battery is, in this view, described by the number of different events experienced. Accumulated signals are signals like: the total distance driven, total energy throughput, etc. To obtain both sorts of signal, data preprocessing was needed. A schematic view of the preprocessing is presented in figure 4.1.

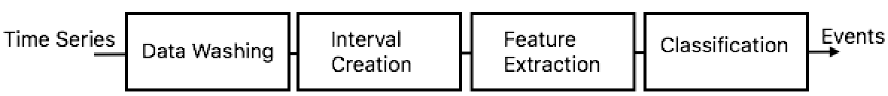

Figure 4.1: Schematics of data preprocessing.

Imperfections in the raw data is a major issue. Signal values are lost or nonphysical. The first step was, therefore, to replace these false values in a suitable ways. Distance, for instance, is an accumulated signal and a lost value was simply replaced by the last known. The same was done for the time stamps. Any lost or nonphysical values of quickly varying signals were replaced by the mean of previous and next true value. Both data recovery techniques can result in data loss, but given the large amount of data, it was believed to be good enough.

The next step was breaking up the time series in shorter intervals. This was based on the DOD signal, which was first smoothed by applying a lowpass filter. It was used on the basic assumption that the DOD can only change accordingly to the current drawn or added to the battery. Very high frequencies were therefore removed. Then only point between which the DOD changes significantly was stored, removing all points collected in idle mode. For the remaining, maximum and minimum were found and the time series was split at these points. Lastly, all created intervals were further split if a long time passed between two consecutive data points.

The created intervals were of one of the following sorts: drive, charge, or idle. From every interval a set of features was extracted, believed to be a good basis for finding the causes of aging in batteries. The list of features can be seen in table 4.1.

Feature	Aggregation
time	Mean, Interval Duration
Distance	Distance Driven in Interval
Current	Mean, Max, Min, Var
Voltage	Mean, Max, Min
DOD	Mean, Max, Min
∂DOD	Mean, Max, Min
$\lvert$DOD$\rvert$	Time Integral
Power	Mean, Max, Min
Ambient Temperature	Mean, Max, Min
Battery Temperature	Mean, Max, Min

Table 4.1: Chosen signal and aggregation to per event features.

The process up to this point had been data washing and aggregation, and could now be followed by classification. The set of features generated above can be a basis for multiple different classifications, but what follows is a simpler version. It was mostly aimed to be a proof of concept.

The categorization was done in a deterministic way with a decision tree, see figure 4.2. It was made on the assumption that a drive will increase the total distance driven, increase the DOD, and have a positive mean current. On the other hand, a charge will decrease the DOD and have a negative mean current. To improve the resolution of the history depiction, instead of counting the number of charges, the number of charges of a particular mean power was used.

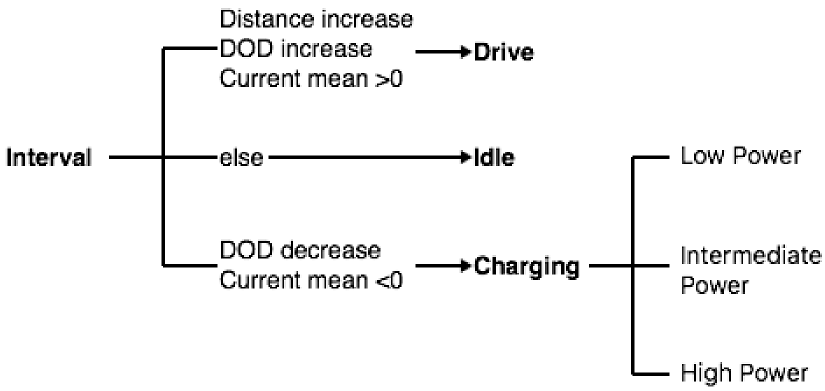

Figure 4.2: Decision tree describing the classification of events. The intervals are inputs and are classified as a type of event.

4.2 Results

The number of charging and driving events, independent of charging power, experienced by each car, over the monitored time, can be seen in figure 4.3. The number of charges and drives are very different among the cars in the population. Of interest is the fact that the number of charges rivals the number of drives. This can, once again, be a result of the limited range of the vehicles.

The measured mean power for all charging events, see figure 4.4, shows three distinct peaks, clearly visible, though not fully separated. These align with the fact that the charging in the cars can be done according to three different current limitations. This is seen as a validation of the charging classes.

Posing as an additional validation of the event creation, the initial DOD for every charging event is inspected, see figure 4.4. The most occurring initial DOD is around 30% and higher initial DOD becomes more and more scarce. This could be a result of the limited range of the car. Users want it to be fully charged all the time, or it is just easy to always connect it to the outlet at the end of the day.

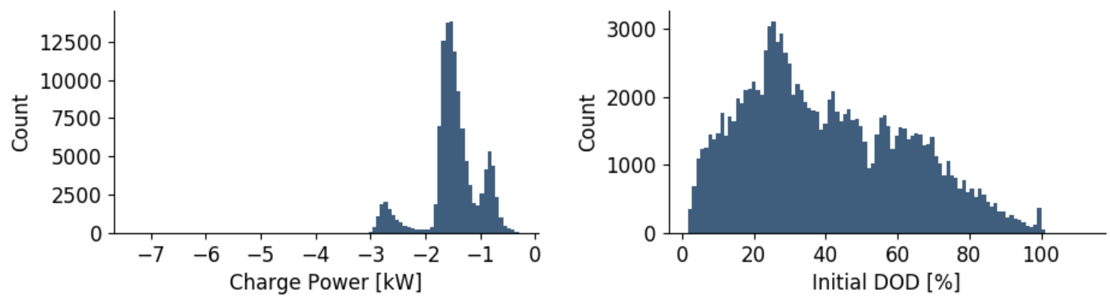

(a) **(b)**

Figure 4.4: a) Mean power in the charging events. There are three distinct peaks. b) The initial DOD at every charging event. Maximum around 30% and more scarce for higher values.

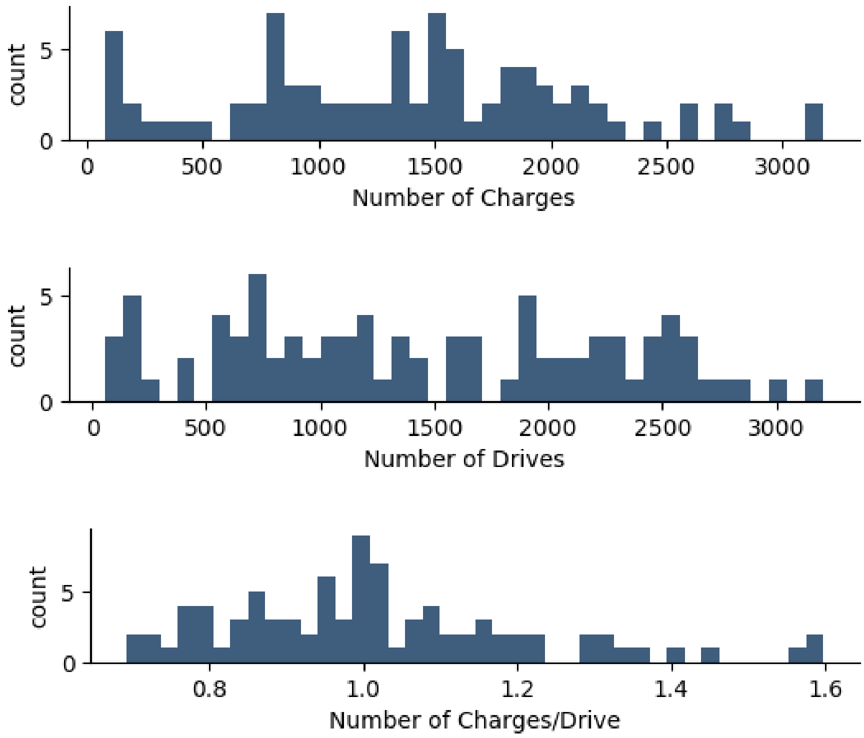

Figure 4.3: The number of different events experienced by each car over the monitored time period. Both the number of charges and drives are well spread out in the population, but the number of charges per drive is distributed around one.

4.3 Discussion

The reality of collecting data from vehicles is challenging. The noise and number missing values are in some cases large, enough to hide any useful information. The signal is first very much determined by the implementation is the cars, sensor etc. This is not easily accessed information and outside the scope of the project. Sending to a server for storage is also an issue, relying on the GSM network. What is left for analysis is often not what first was intended, and needs much work to be usable.

The time series data has an intriguing structure and is difficult to aggregate in an information preserving manner. At a sample level there are many different signals that describe the use and environment of the battery, for instance current, temperature, and voltage. The coincidence of the values of different signals might hold information about how the battery is aged. High current occurring together with low temperate could affect the battery differently from high current at high temperature. This means that every point of the time series holds potential information to the aging.

There is also information in the temporal arrangement, the sequence in which the values occur. The battery might age more if the previous time step included a certain value. For instance, a battery that is standing idle, enabling relaxation in the battery, before charging might be affected differently from a battery that is being charged directly after a driving sequence. On a smaller time scale, a more aggressive driving sequence could affect aging.

So, more than single points are needed to connect to SOH degradation, but using all points is not a valid choice. A car monitored for six years have logged billions of data points, more than can be handled with most algorithms. Logically, there should exist an aggregation of data so that it can be mapped to aging, but with the number of data points small enough to manage. On this basis, the use of events was built.

How these events are created is possibly the biggest challenge. In this project the events are created by domain knowledge and simple decision trees, simple but effective. Given the distribution of the number of charges, drives and the ratio between charges and drives, the event creation was successful. Cars are driven and charged in reasonable amounts over the six years of use. The population is differentiated, some cars used extensively and some limited. Interestingly, and validating for the preprocessing, the number of charges per drive is distributed quite narrow around one. Moreover, the average charging power peaks agree well with the current limitation settings in the cars.

More classes of events could be created, but it should be recognized that there must be sufficient data to study a certain event. For instance, is the seat heater affecting the aging? In fact, taking energy to any operation other than to power the drive train will decrease the range and possibly age the battery like any other use. However, there seems to be three big limitations to asking such specific question. First of all, the contribution might be so small that the above mentioned noise is too large. Driving with or without the seat heater might not make a visible difference. Secondly, there must be enough data to study such effects. If very few people are using the seat heater there will be smaller amounts of data and thus the algorithm will have a lower chance of finding these effects. One can argue, perhaps motivated, that there must be enough data to study the effect of the seat heater, however, say there is an even more specific target like charging with maximum allowed current, directly after a long drive, and at sub-zero temperatures. Such request could be very interesting but might suffer from insufficient data. The third constraint is related to the second in the sense that the data to study an effect must have a sufficiently low correlation to other, more prominent, effects. For instance, the seat heater might be used to high extent when it is cold outside. However, the cold environment itself might damage and age the battery even more than the usage of the seat heater. In this case a study might find that the seat heater has a very large impact, but in reality the cold environment is to blame.

In the purpose of progress, the set of classes and accumulated signals are believed to be good features to map to SOH degradation. They are used in the remaining parts of the report. Additional validation is attainable, one could inspect all the per interval values. This was also done to some extent, but it does not offer any more information and is left out of the report.

5

Engineering an SOH from SOC and Current During Charge

There is no SOH signal logged in the cars. Therefore, this chapter presents a posterior engineered SOH signal that is based on a capacity estimation at every charging event. It uses the DOD, and current signal and is a common way of estimating the capacity without doing a full discharge. The behaviour of this signal is evaluated and found to have dependencies other than the history of the battery.

5.1 Method

The engineered signal was based on coulomb counting, mentioned in section 2.1.2, as a determination of the SOC (DOD). The basic assumption was that the change in SOC is proportional to the amount of current going into the battery, see equation 5.1. η is often called the coulombic efficiency, and the integral is normalized by C, the capacity. By measuring the SOC before and after the charge and the current during the charge, the capacity can be calculated.

$$SOC = SOC_0 + \frac{\eta}{C}\int_0^t I(t)dt \rightarrow C = \frac{\eta \int_0^t I(t)dt}{SOC - SOC_0} \tag{5.1}$$

The SOH signal was defined as the normalized capacity and could be found in every charging event. The normalization should be done to the largest value, which is hopefully the first value measured. The intervals used were the ones created in chapter 4. But, to apply additional filtering, only charging longer than 3% SOC change was used.

The capacity based SOH was assumed to depend on measurement conditions. The signal from other fleets at Volvo Cars Corporation have shown seasonal dependence, and the capacity is known to vary with different parameters. The engineered signal's dependencies to charging conditions was therefore investigated. The following was believed to be parameters that describe a charging event:

- Difference of DOD
- Initial DOD
- Final DOD
- Average charging current
- Rest time before charging
- Average battery temperature
- Average ambient temperature

The dependencies of the engineered SOH signal on each of the above signals were investigated. This was done by observations, graphing the engineered SOH against each parameter. For this, the entire fleet was used, all 296 cars.

5.2 Results

The engineered SOH from the entire fleet, all 296 cars, can be seen in figure 5.1. There is a large spread of values, ranging from well above to well below 100%. Therefore, the normalization to the largest measured value is not suitable since some outliers are very large. Instead, the normalization is to the first standard deviation of the signal for each separate car.

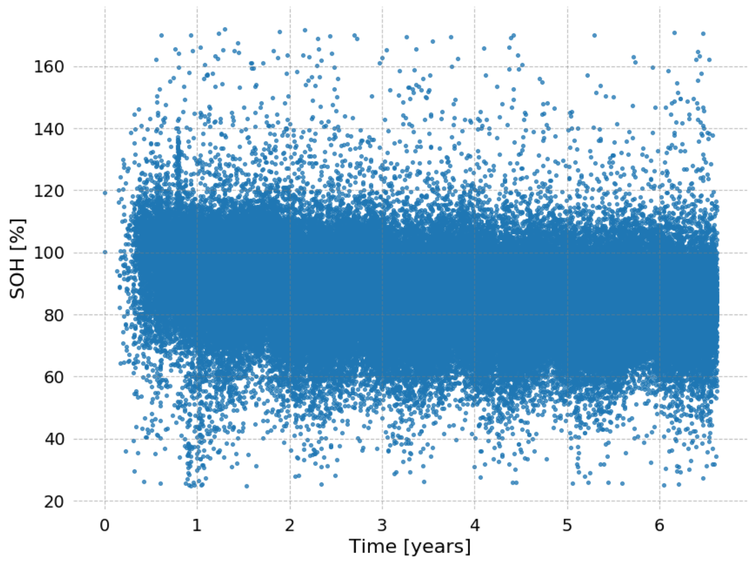

Figure 5.1: SOH signal for all cars in the fleet. The spread of the values makes it difficult to unambiguously state the age in every cycle.

The signal from a single car can be seen in figure 5.2. Points stemming from consecutive charging events can differ from each other with large magnitude. Both large increases and decreases can be observed. This means that usage believed to age the battery has no visible effects for single cycles.

There is a general trend of decline the first five years for this car, based on the most densely populated areas of the graph. This decline appears to reach values below 80%, which is beyond the EOL. However, after five years an increase can be observed. At this point the densely populated parts are around 100%, as good as a new car. This coincides with the usage of higher charging currents, see figure 5.3. There are also differences in other usage parameters, such as SOC levels related to charging events. This supports the premise that charging conditions affect the engineered signal.

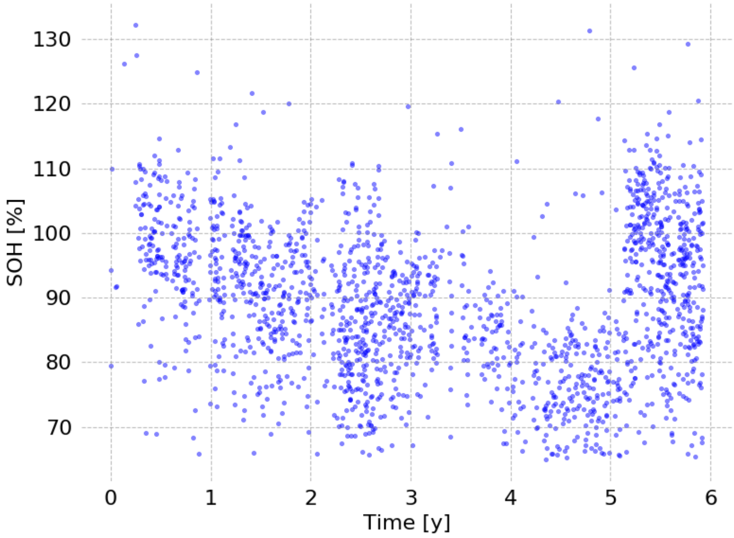

Figure 5.2: Engineered SOH signal from a single car. There is a general decline in the first years, followed by a recovery in the last year. It is very difficult to see degradation between cycles.

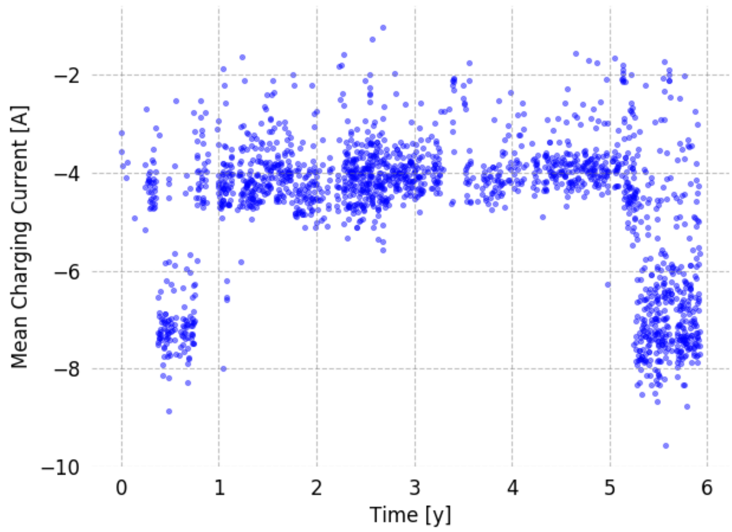

Figure 5.3: Mean charging current for charging events in a single car. There are two periods of considerably higher charging currents, most prominent after five years.

The engineered SOH as a function of charge added in a charging event can be seen in figure 5.4. The signal seems to increase with larger charging events, and the spread of measured values seems to decline. The latter observation is intuitive because the number of samples is much greater for smaller charging events.

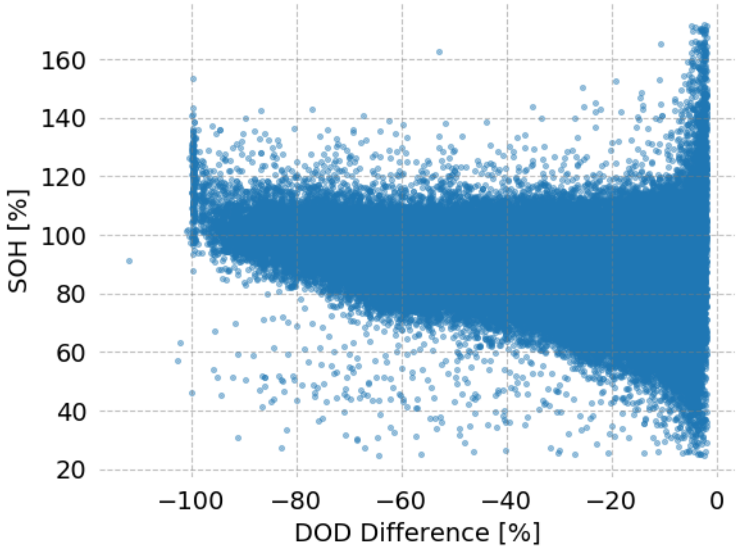

Figure 5.4: The engineered SOH signal as a function of difference in DOD, i.e. the amount of charge added in a charging event. The spread decreases with the amount of charge added, while the engineered SOH seems to increase.

The charging event is also determined by the initial DOD, from what DOD the charge was initialised. The graph for this parameter can be seen in figure 5.5, and the signal's dependency is similar to the above. The spread is once again more prominent for smaller initial DOD, most likely due to more samples, and the engineered signal increases with initial DOD. It should be mentioned that for very small initial DOD, there is a sharp increase in the measured values.

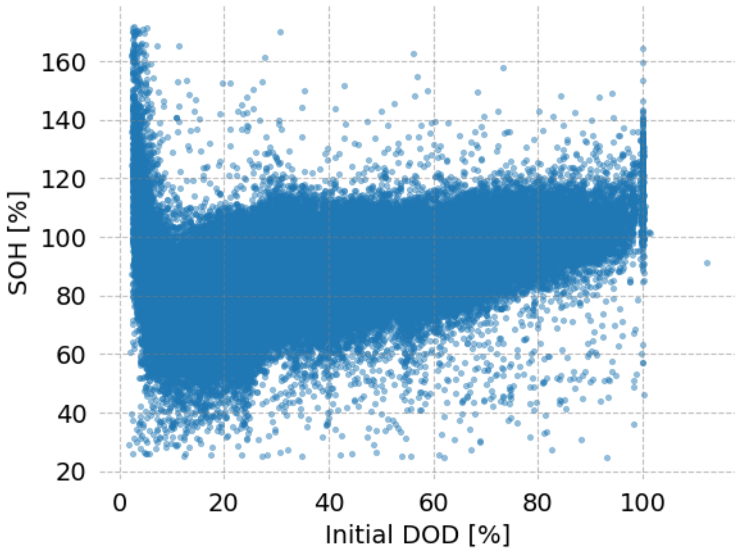

Figure 5.5: The engineered SOH signal plotted against initial DOD. There is a larger spread for small initial DOD and a trend for SOH to be higher for larger initial DOD. Additionally, the SOH increases for very small initial DOD.

Even though the initial DOD and the difference in DOD describes the DOD window of a charging event, also the final DOD is interesting to evaluate. The graph describing the engineered SOH as a function of final DOD can be seen in figure 5.6, and the dependency is very similar to the above, an increase with the final DOD and a larger spread for smaller values. Once again, the signal increases rapidly for very small final DOD values. The ambient temperature is a charging parameter of interest. The available capacity is known to vary with the temperature and the ambient temperature will to some extent affect the temperature in the battery. The

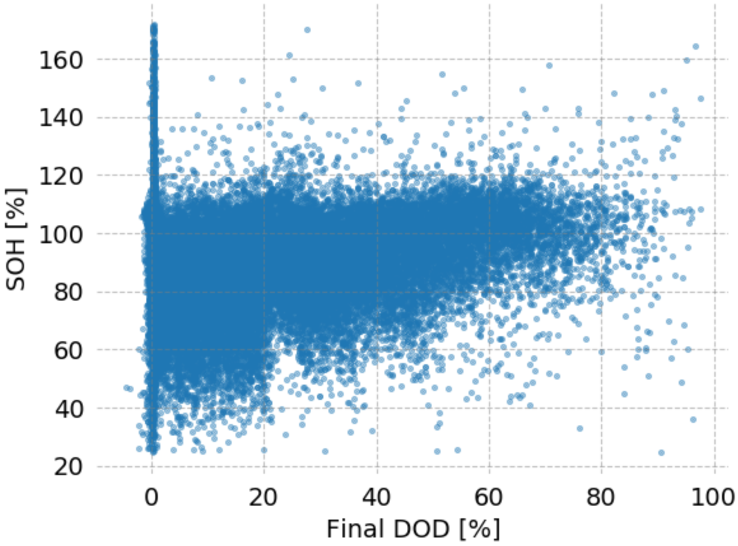

Figure 5.6: The engineered SOH signal plotted against final DOD. The spread is more prominent for smaller values of final DOD. The engineered signal does also tend to be higher for larger final DOD.

engineered signal as a function of ambient temperature can be seen in figure 5.7. There is a trend to measure larger values of the engineered SOH signal at higher temperatures. This trend is, however, harder to evaluate due to the spread of the signal.

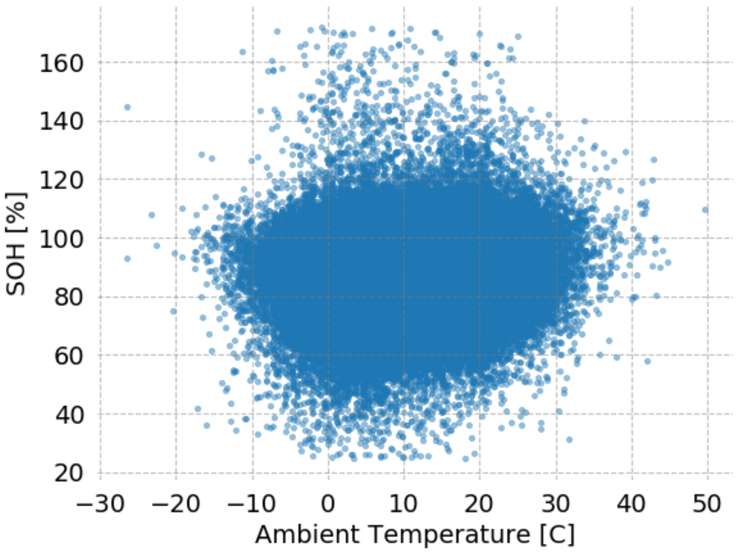

Figure 5.7: The engineered SOH signal plotted against ambient temperature. The structure in the graph is difficult to evaluate, but there is a trend to measure larger values in higher temperatures.

The temperature inside the battery is perhaps the relevant temperature signal. The battery temperature dependence, see figure 5.8, is similar to the ambient temperature dependence and there is probably a correlation between the two parameters. The dependency is even more vague, but visible. The stripes in the graph are due to the rather poor resolution of the battery temperature signal.

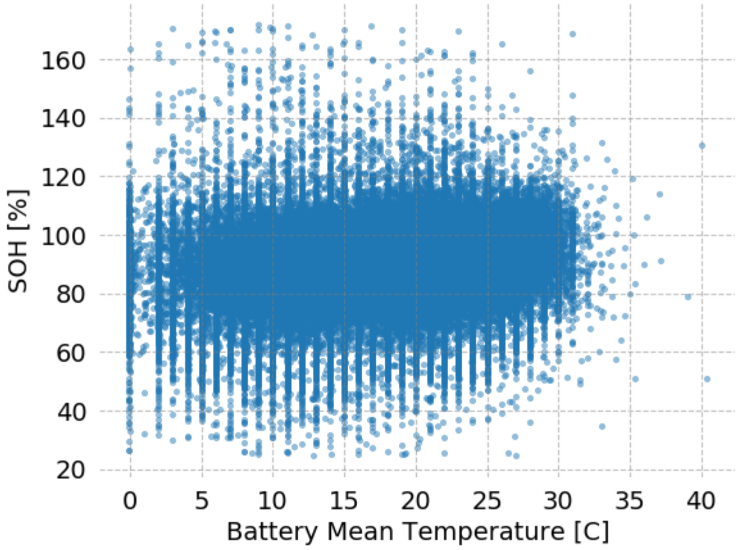

Figure 5.8: The engineered SOH signal plotted against the mean temperature in the battery. The SOH increases slightly with battery temperature, but the trend is difficult to evaluate.

In a world of fast charging it is interesting to evaluate the impact of charging current on the engineered SOH signal. The graph can be seen in figure 5.9. In some aspect, the charging current is embedded in the engineered signal and its effect should by this logic be limited. The charging current seems to have a lower impact on the measured values, as far as can be observed. There is a larger spread for small currents.

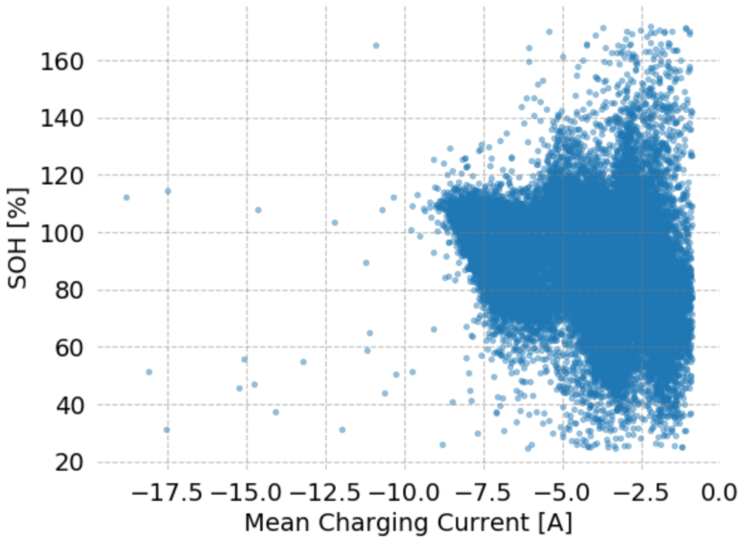

Figure 5.9: The engineered SOH signal plotted against charging current. The signal has a much larger spread for small charging currents. It is, however, very difficult to observe a general difference in measured values.

A battery that is allowed to rest before a charge might respond differently. This is based on the fact that it can take time for a battery to equilibrate. From figure 5.10 a trend can be observed; a car that has been resting before the charge gives larger engineered SOH. The effect seems to be present through all 25 hours plotted in the graph. That there are many more samples of small rest times can be regarded in two ways: the charging events are fragmented, or it is very common to connect the charging cable after the car has been used.

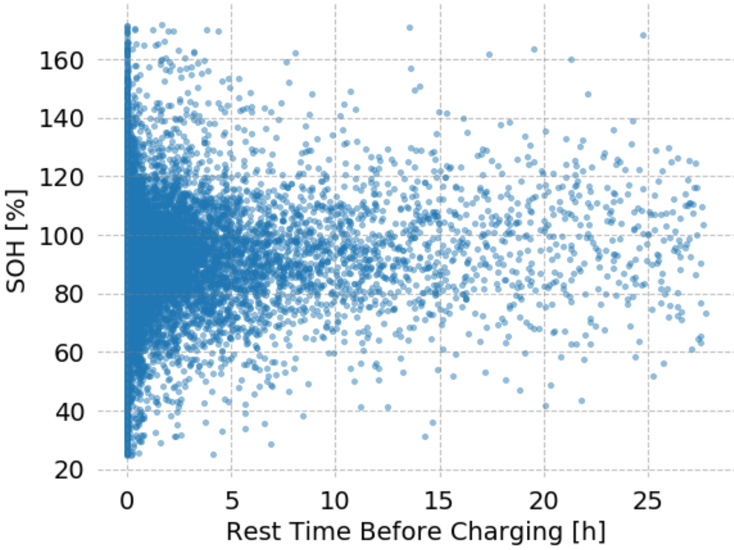

Figure 5.10: The engineered SOH signal plotted against the time spend resting before the charge begins. The spread is much larger for small values. A trend can also be seen linking longer resting times to larger engineered SOH values.

5.3 Discussion

The engineered signal is a common choice of SOH. It is capacity based, a property that is known to decline with time and use, and should be a good measure of aging in the battery. It is also attainable in most charging events, which gives hundreds, sometimes thousands, of samples per car.

The signal is very dependent of the inner working of both the current and SOC signal. These are not fully known and this is an issue. The SOC is calculated onboard with coulomb counting and corrected by measuring the cell OCV. The OCV correction is done as soon as the DC current through the battery has been below 2A for at least 45 seconds. The correction should therefore be part of the engineered signal. The correction does not take aging or charging conditions into account, which is believed to be of great importance to capture aging in the engineered SOH.

The aim to connect certain usage of the car to battery degradation, is not supported by the engineered signal. It is very difficult to conclusively see degradation over the lifetime, and even harder to see a degradation in every cycle. A reasonable assumption is that it is due to one of the following causes:

- Poor quality of the original signals
- Errors in the signal engineering
- Dependencies of other parameters

As mentioned above, an increase in the engineered SOH signal coincided with a change in user behavior. Interesting since charging is a broad term. For instance, charging includes both increasing the SOC 5% at -5°C and going from 10% to 100% SOC in the middle of the summer. Some cars are charged in an indoor garage, making charging temperature fairly constant all year around, and others are exposed to the harsh Swedish winters.

It can be argued that the engineered SOH is a function of: SOC, temperature, rest time before measurement and charging current. Some of the dependencies are vague, and very difficult to evaluate, some are more prominent and easier to argue for. The reality is that the argued existence of dependencies are all interconnected. An evaluation of just one parameter is certainly unfair since the others might obscure the results. It is also possible that some effects on the engineered SOH signal only appear in strict combinations of many parameter values, some of which could yet be undiscovered, and thus would not be visible in the graphs.

A more rigorous analysis would be helpful, but in the purpose of making quicker progress and just inspecting dependence, the approach is deemed more than enough. All the same, the questions remain; Are there more parameters affecting the engineered SOH? Are they logged in the car? And, from the parameters found; Are they properly aggregated?

The parameters affecting the engineered SOH can be concluded to consist of two groups: the history of the battery, and the charging conditions in which the SOH is measured. In some cases the dependencies of the non history parameters exceed the effect of aging. The spread of the signal, and in some cases dependency on charging conditions, can reach more than 50%, while the annual degradation of SOH due to aging is believed to be only a few percent.

6

Neural Network for Estimating an Absolute SOH

The overall aim of the project is to measure the SOH in the cars and connect its degradation to battery history. The project argues that the engineered SOH signal depends on measurement conditions as well as the history of the battery, but that a measure of aging should be independent of the measurement condition.

This chapter presents a neural network that estimates an Absolute SOH, independent on measurement conditions, and connects it to the history of the battery. The inputs to the network are the features developed in chapter 4 and 5. The supervised learning is done with the engineered SOH, from chapter 4, as the label. The network is, in some sense, operated in two different settings; reconstruction of the engineered SOH, and estimating an absolute SOH.

6.1 Method

A neural network model was built and fed inputs based on the affecting charging conditions found in chapter 5, and a simple representation of the history, developed in chapter 4. The history representation was done in a simple manner to provide a proof of concept for the technique. Development of a more complex realisation of the history inputs could be an interesting topic for future work. This view of the causes affecting the engineered SOH is presented in figure 6.1.

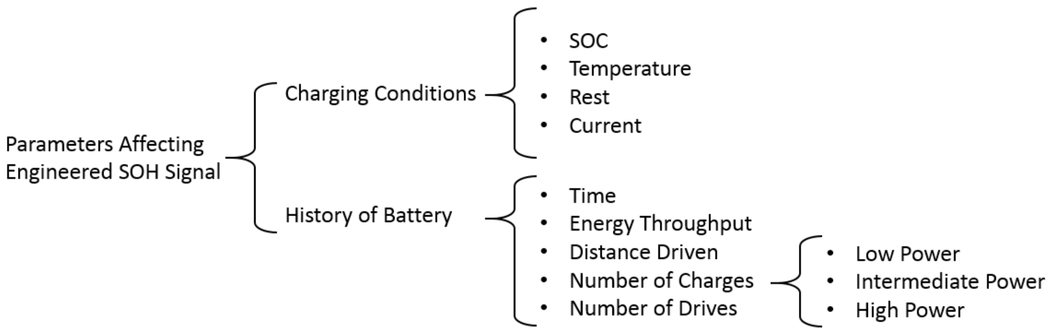

Figure 6.1: The parameters affecting the engineered SOH.

The network was of feedforward design, consisting of 4 layers with 16, 16, 14, and 14, nodes in each respective layer, all having a ReLu activation function. Each node had a 20% dropout rate, and the learning was done with stochastic gradient descent. The inputs were all fed to the same input layer, regardless if it was a charging or a history feature. This must no be the case and future work could include more exotic structures. The architecture was found by trial and error.

The network was used in two different settings (three if training is regarded as a setting), depicted in figure 6.1. The network was first trained with the engineered SOH signal as a label. The aim was to recreate the engineered signal using the features developed in chapter 4 and 5, meaning a specific history and charging conditions. The output was validated by calculating the difference between the recreated SOH and the label.

The training was done in 11-fold cross validation. This meant that the population, 85 cars selected for aging studies, was divided into 11 groups. The network was trained on 10 groups, the training set, and evaluated on the last group, the validation set. Then, the network was reset and another group became the validation set. Every group was, in this way, at some point the validation set. For each specific group, the network was trained for ten trials and the result with the smallest mean square error selected as the winner. Each trial consisted of 20 epochs, used 0.001 learning rate, and had a batch size of 40 samples. The total number of charging events was roughly 120'000.

The use of the network was then altered, by altering the input features. Above, the network was asked to estimate the SOH given a specific history and charging conditions. It could now be asked to estimate the SOH given a specific history and standard charging conditions. This was done by substituting the charging condition features in the validation set to a standard. This standard conditions were from the average charge of the data set, and can be seen in table 6.1.

Feature	Mean Value
Ambient Temperature	9.3 ^oC
Current	-4.6 A
Battery Temperature	15.1 ^oC
Initial DOD	40.8 %
Diff DOD	-36.9 %
Final DOD	3.9 %
Rest Time	846 s

Table 6.1: Features of an average charging event.

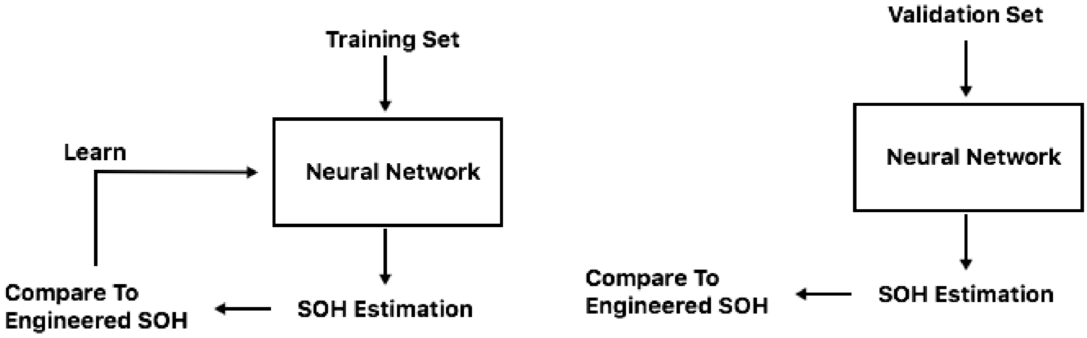

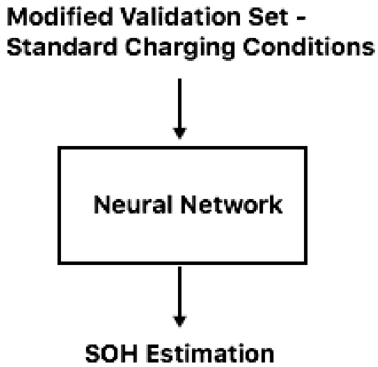

Figure 6.2: Different settings of the neural network. Top left image represents the training to recreate the engineered SOH. The success was then evaluated using the validation set, top-right image. Finally, the charging condition features in the validation set were changed to a standard. In this setting the network produced an estimation of an absolute SOH.

A favorable property of an intelligent model is that the inputs can be altered and the model response studied. For instance, car 7 of group 8 have never been charged with the highest power. The history of this car could now be altered artificially. A certain percentage of the low power charges were replaced by high power charges. The data was fed to the network and the new estimations were recorded.

The same game could be played with the charging conditions, especially the ambient temperature. The history of a single car, at midlife, was used as the basis and different charging conditions were tested. The ambient temperature was ranged between normal values in the Swedish climate, 5^oC to 25^oC. and the network's estimated absolute SOH was recorded.

6.2 Results

The results of using the neural network follows. It is divided into two subsections, depending on which of the two settings was used; recreating the engineered SOH, or estimating an absolute SOH.

6.2.1 Recreating the Engineered SOH

The network was trained according to the above specifications in order to learn how the input features map to the engineered SOH. The error, difference between the network recreated and the engineered SOH, can be seen, collectively for all eleven groups, in figure 6.3. The distribution is Gaussian, with a mean of -1.7% and a standard deviation of 9.9%. These values are to be compared with the engineered SOH, ranging between 20% and 160%.

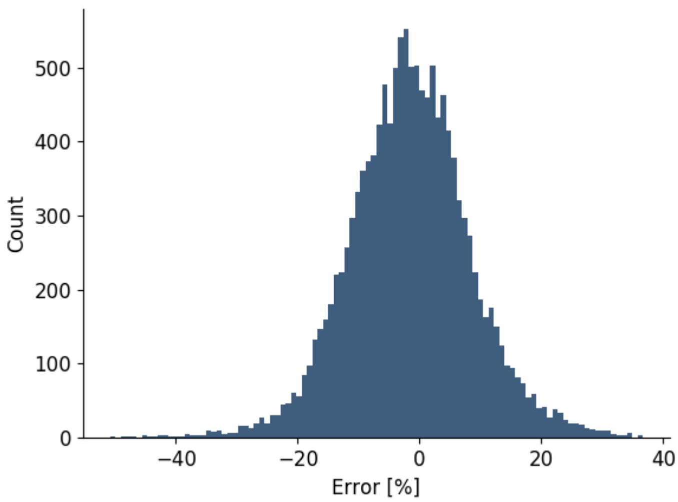

Figure 6.3: The difference between the engineered and the recreated SOH. The distribution is Gaussian with a mean of -1.7% and a standard deviation of 9.9%.

The network recreated SOH can also be compared to the engineered SOH by scatter plotting them against each other, see figure 6.4. In a perfect model, the plot would only contain a diagonal line, but in this model the values are more spread out. However, from the plot it can be seen that some information about the system's inner workings have been picked up by the network since there is a clear diagonal trend. It is also apparent that the network's recreated SOH is within a smaller interval than the engineered SOH, intuitive since a guess from the network would probably ignore the outliers.

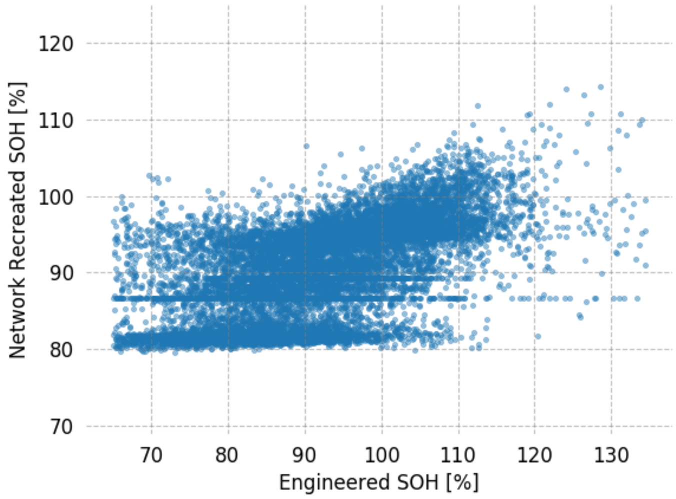

Figure 6.4: The engineered SOH and the network recreated SOH scatter plotted against each other. The diagonal trend is clearly visible.

The network recreated SOH for a single car can be seen, together with the engineered values, in figure 6.5. The network is recreating many of the characteristics of the engineered SOH. Namely, the general trend of degradation, seen in the engineered values from the first five years, seems to be captured. Additionally, and perhaps of more interest, the sudden increase in SOH values seen after five years of measurement seems to be mimicked by the network.

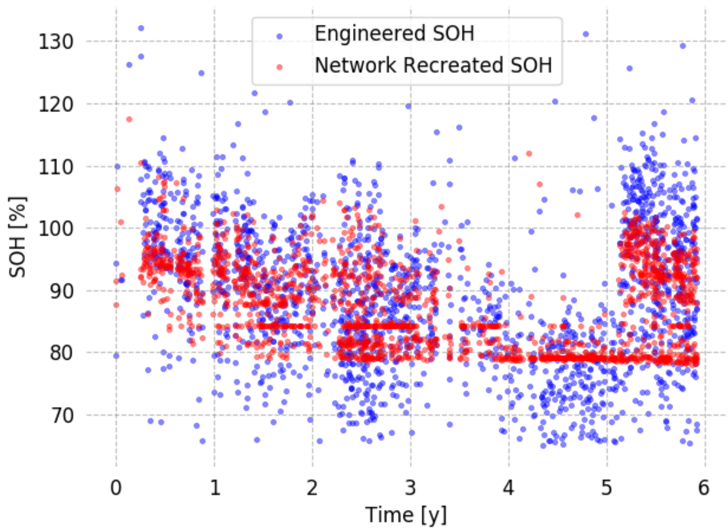

Figure 6.5: Original, engineered SOH of a single car together with the network recreated SOH. Many of the characteristics of the engineered values can be seen also in the network recreated SOH.

Overall, the mean error is close to zero, but this is not true for all individual cars, see figure 6.6. Some cars exhibit a small mean error, while others have a mean error of a few percent, and some outliers have mean errors of up to 10%.

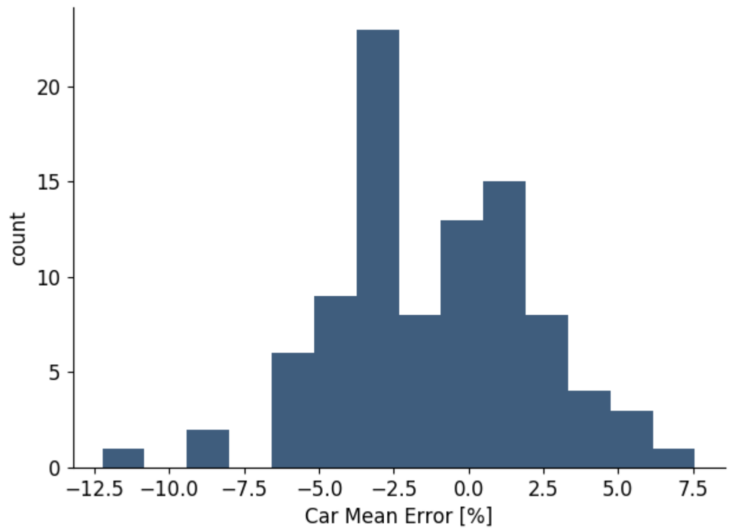

Figure 6.6: Mean errors of individual cars. The errors are distributed around zero, but some deviate substantially.

6.2.2 Estimating an Absolute SOH

For each sample in the validation set, the specific charging conditions were replaced by those of the standard charging, described above. The network was now asked to estimate the SOH given the history of the car and standard charging conditions. The estimation together with the engineered SOH can be seen, for one car per group, in figure 6.2.2. The cars are chosen to give a fair overview of the performance of the network.

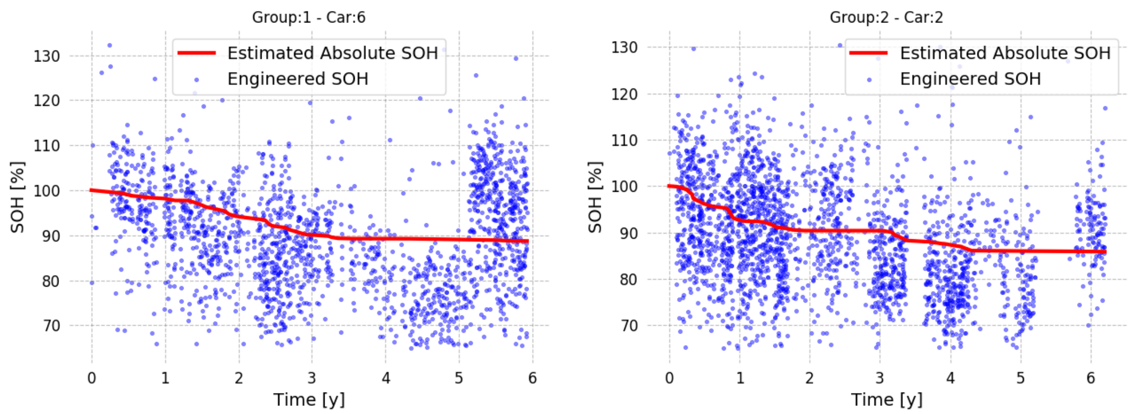

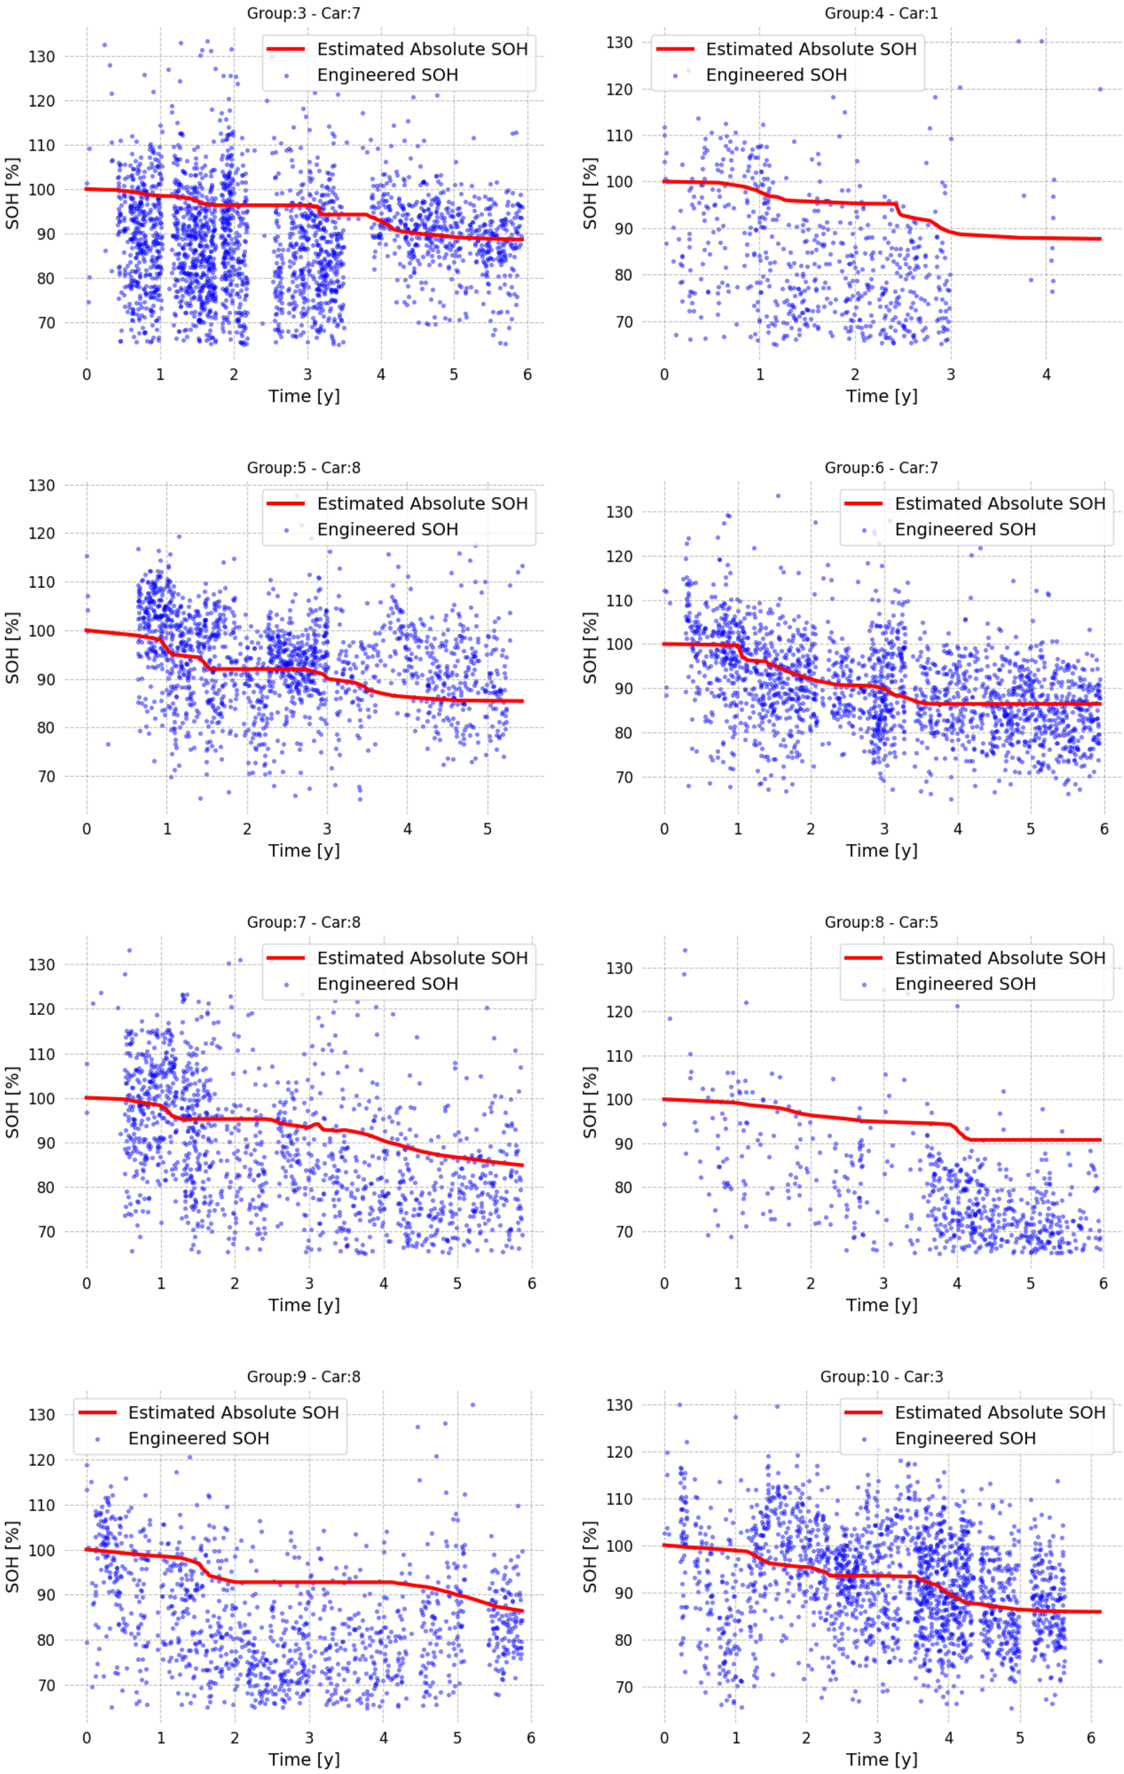
Group:3 - Car:7
Group:4 - Car:1
Group:5 - Car:8
Group:6 - Car:7
Group:7 - Car:8
Group:8 - Car:5
Group:9 - Car:8
Group:10 - Car:3
Estimated Absolute SOH
Engineered SOH
SOH [%]
Time [y]

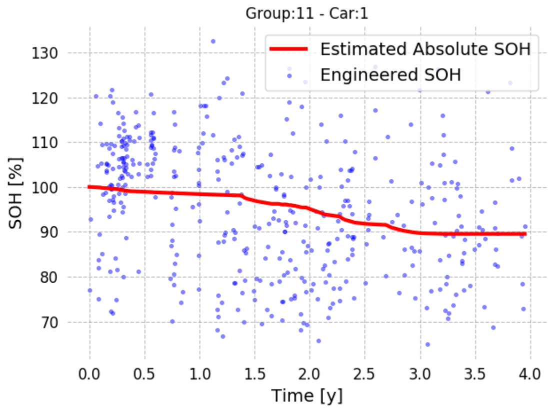

Figure 6.7: Engineered SOH together with the neural network's estimation of an absolute SOH. The curves show a single car from each group in the population.

Estimations in Group 2, Car 6 (G2;C6), (G6;C7), and (G10;C3) do all seem to align with the measured data, and graphs (G4;C1), (G8;C5), and (G9;C8) represent estimations that appear to deviate from the original SOH. The network is not meant to recreate the measured data, but to account for charging condition.

Estimation in (G6;C7) does not change the SOH in the first period of time. This is not physical and is probably a shortcoming of the network. In fact, it is assumed that a threshold exists after which the SOH estimation will be changed. This threshold could be in any of the features. It is possible, for instance, that a certain number of charges are needed to activate the estimation.

The estimations for all cars in a group can be seen, for each group, in figure 6.2.2. Since there are significant differences between the groups, and even between training trials, it is more interesting to compare estimations within a group.

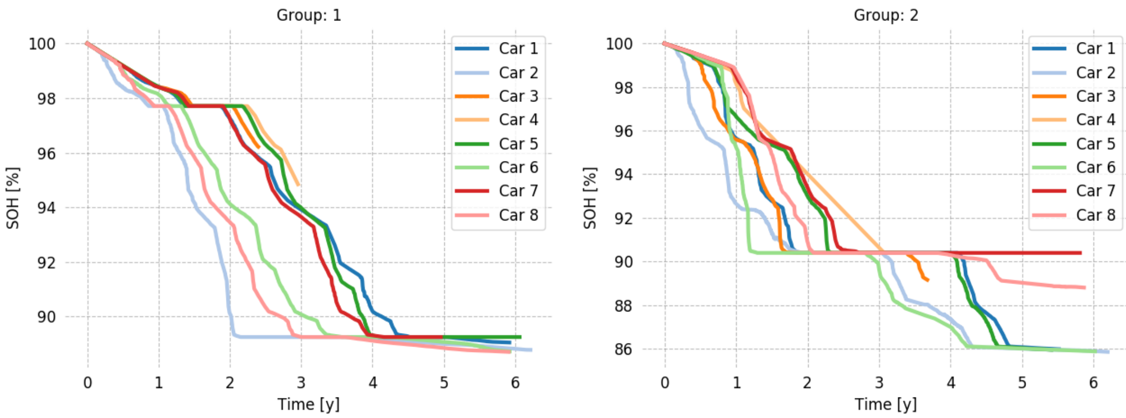

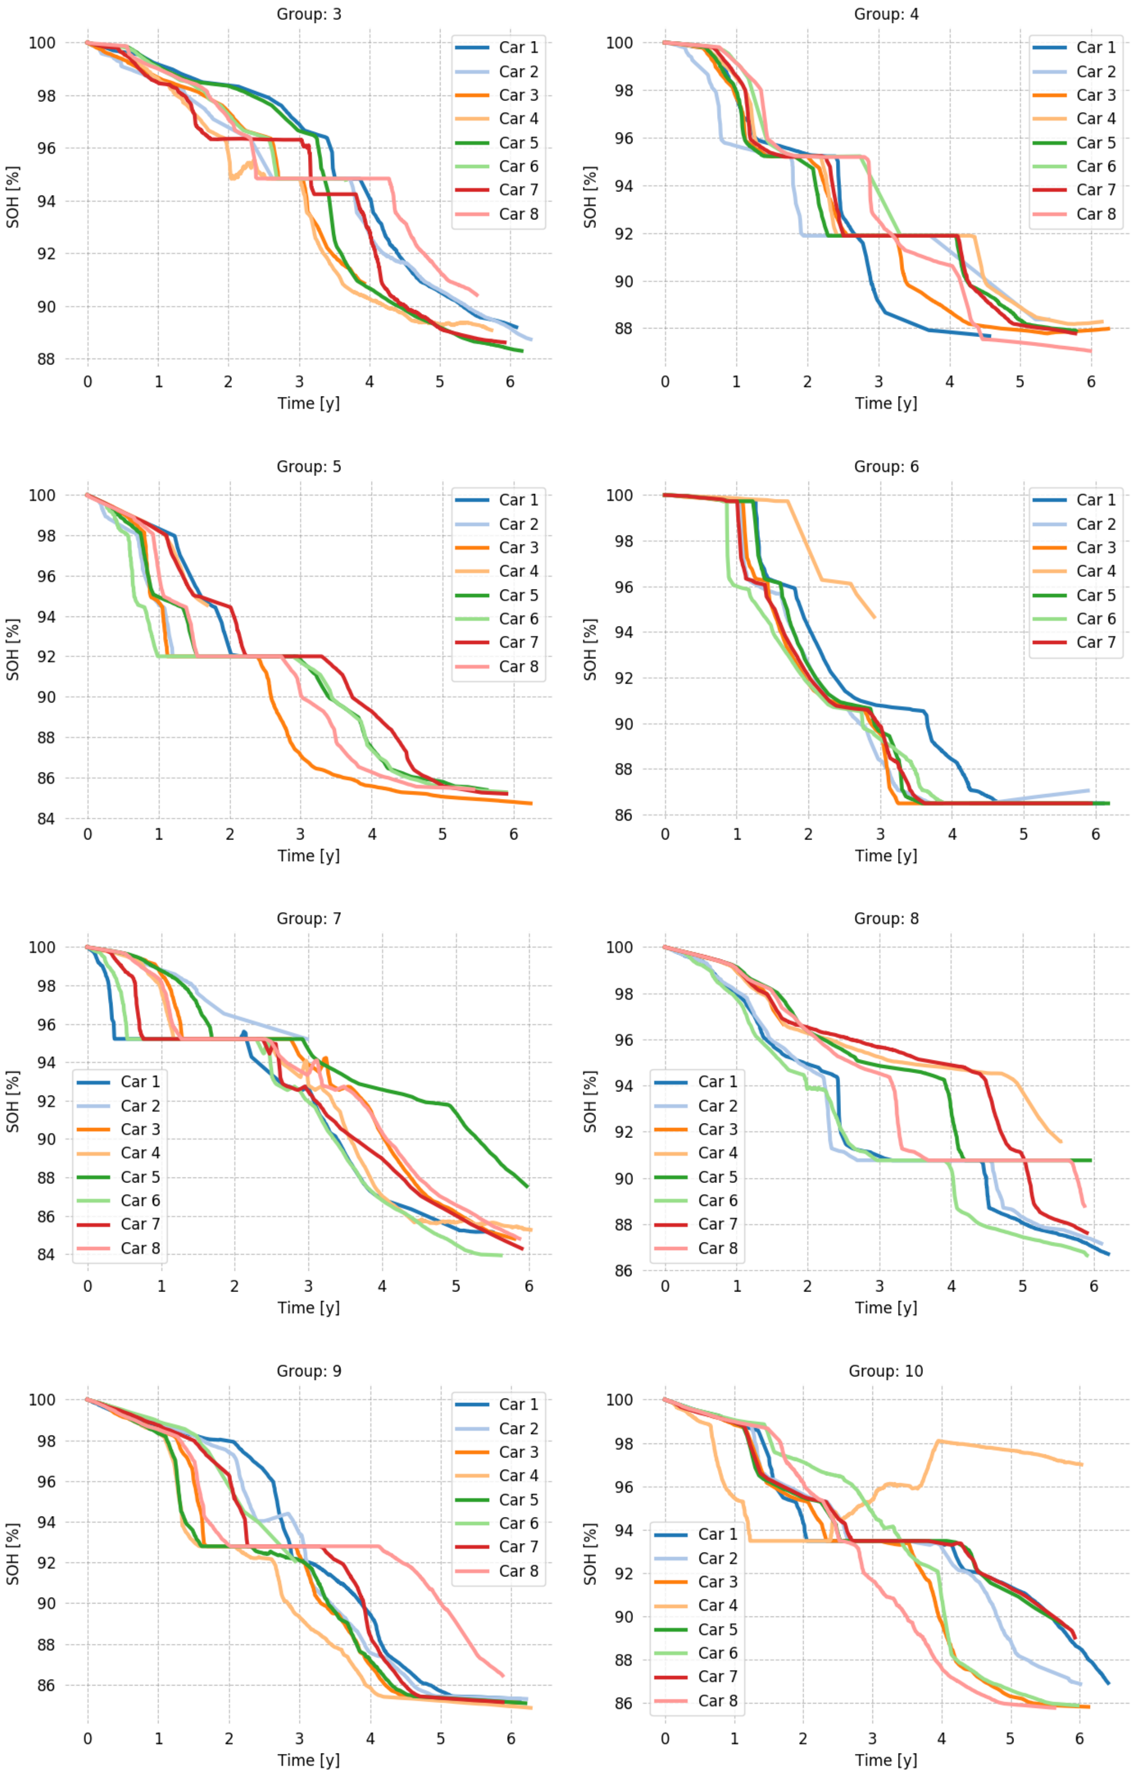
Group: 3
Group: 4
Group: 5
Group: 6
Group: 7
Group: 8
Group: 9
Group: 10
SOH [%]
Time [y]
Car 1
Car 2
Car 3
Car 4
Car 5
Car 6
Car 7
Car 8

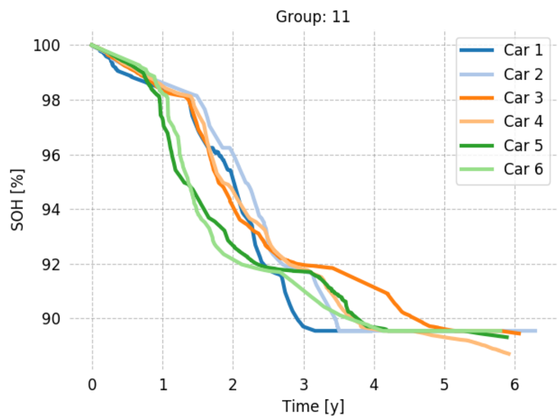

Figure 6.8: The neural network's estimation of the absolute SOH for all cars presented for individual groups. Curves within a group can be seen to share many traits.

The shapes of the curves share some global traits. All curves seem to have a faster degradation in the first years, followed by a slower degradation later in life. The values reached, after six years, range between around 90% to around 82%, but the bulk of the estimations ends up around 86%.

There are features of the curves that could be artifacts. First, group 2 (G2), (G3), (G5), (G7), (G9) and (G10) can all be observed to have midway plateaus. This is, just like the initial plateaus not connected to a physical explanation since even unused batteries will age, earlier described as calendar aging. Second, most of the SOH curves have a tendency to converge for large times. Third, there is one car (G10;C4) that exhibits different behaviour than the rest. It deviates from the common degradation curve and recovers.

The estimations show limited differentiation between cars within a group and their curves seem to be similar. This suggests that the presented aging curves are very much determined by the training. The landscape of the cost function is therefore assumed to have many local minima. However, cars from the same group are still, to some extent, different which means that the network is capable of some degree of differentiation.

Within some groups there are a couple of cars that have a higher SOH than the others. The difference can be above 6% and, if true, should have an impact on the car performance. From the SOH curves, 12 cars have been selected for having higher values of SOH throughout the life time. For these cars the history features are compared to the average of the full population, the 85 cars, and the result can be seen in table 6.2. The values are presented in percent deviated from the population average.

There is a clear trend of the 12 cars being used less than the population average. All but one have been driven a shorter distance, have a lower charge throughput, and fewer number of drives. The average of the winning 12 cars differs significantly from the population average, -40.8% in charge throughput, -53.9% in the distance driven, and -38.8% in the number of drives done.

The model is a more specific for the charging events. The car has three modes of charging. This gives three different charging powers that have been differentiated between in the input features. Since the cars have been charged different number of times, the relevant measure is the distribution of different charging powers found for each car and how it relates to the population average.

Group-Car	Charge Throughput	Distance	Drives	High Power Charge	Intermediate Power Charge	Low Power Charge
1-1	-31.6	-31.1	-40.2	-58.4	15.8	-45.7
1-5	-71.1	-81.5	-75.2	-100.0	-67.3	313.6
1-7	-48.2	-65.5	-49.4	-100.0	8.2	-4.7
2-5	13.2	7.1	81.8	-98.1	10.5	-20.2
2-7	-61.1	-81.6	-59.7	-100.0	3.4	5.8
3-1	-59.0	-72.5	-66.7	-100.0	24.6	-72.5
5-7	4.5	-35.5	-13.1	-89.6	-2.1	35.1
6-1	-8.0	-13.8	-22.4	10.9	13.7	-59.2
7-5	-57.4	-68.2	-58.2	-95.9	-5.5	39.2
8-4	-54.7	-66.1	-48.2	-100.0	-38.6	167.5
8-5	-72.9	-81.6	-74.0	-99.8	14.5	-33.7
8-7	-43.2	-56.6	-39.8	-93.4	-37.4	166.1
Mean	-53.9	-38.8	-85.4	-5.0	40.9	

Table 6.2: The history features of the 12 cars observed to have a higher estimated SOH than their respective group average. All feature values are presented in percent deviated from population average.

The winning 12 exhibits a trend to use the high power charging setting less than the population average. Many of the cars have never been charged with high power and all but one have been using it significantly less than the population mean. This is interesting since charging should have an impact on the SOH. Additionally, also the usage of intermediate power charging is lower on average, but many individuals use it more than the population average. Contrarily, the low power setting is much more represented in the winning 12 cars. On average they use it +40.9% more often than the average of the population.

The impact of high power charging is now studied by using car 7 of group 8. It had never been charged with the highest power setting, but the history was altered so that a certain percentage of the low power charges were replaced. The results for different percentages of high power charges can be seen in figure 6.9.

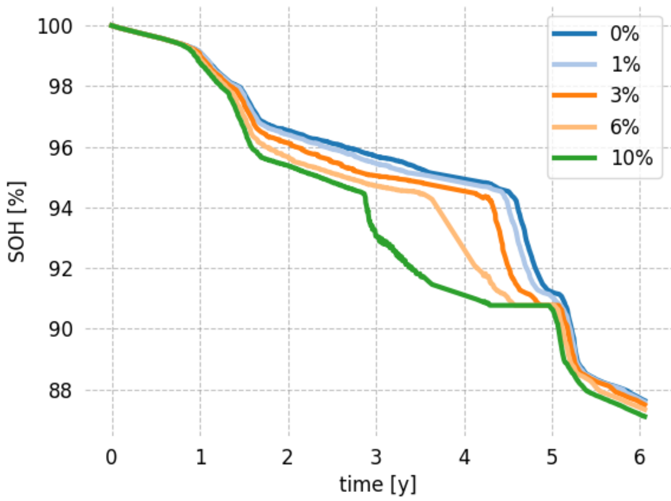

Figure 6.9: Estimated absolute SOH for car 7 of group 8 with different percentage of the low power charges replaced with high power charges. The curve can be seen to drop, treating a high power charge as more damaging than a low power charge.

From the figure it is evident that the model treats a charge with high power as more damaging than the low power charge. The curve drops small amounts with every additional percentage of charges replaced. In the end, 10% of the low power charges are replaced and the curve is now significantly lower for large parts of the time period, with a maximum difference of 4%. The difference becomes smaller for large times.

The same kind of manipulation was done with the ambient temperature. A car midlife was used with its true history. The ambient temperature of a specific charge was altered and the different network responses recorded. The results can be seen in figure 6.10.

There is a clear temperature dependence for the estimated absolute SOH. The difference in estimated values between 10°C and 22°C is 5%. This could very much be the difference between charging in a garage or outside. The effect disappears for lower temperatures.

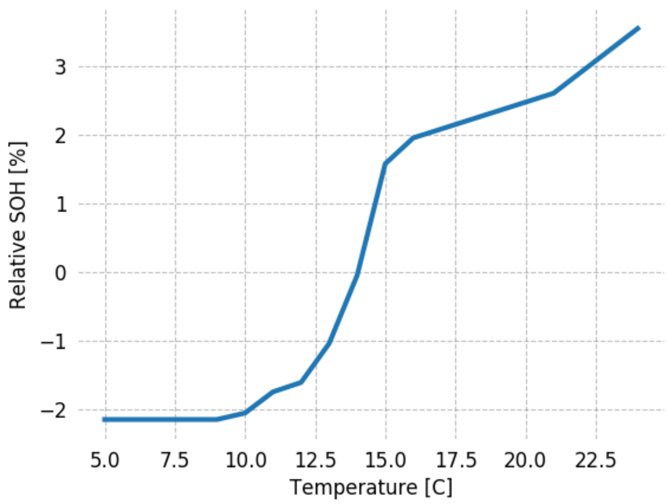

Figure 6.10: Temperature dependence of the estimated absolute SOH. It is clear that higher temperatures give a higher estimation.

6.3 Discussion

The target value, the engineered SOH, is a function of many parameters. This makes the use of a neural network promising. On the other hand, data quality is an issue and the neural network is notorious for bad noise handling. In this aspect, a better choice might be a Random Forrest algorithm, but this would have been outside the scope.

Even if the simple feedforward network is kept, many improvements could be made. For instance, the number of layers and nodes could be altered. This is, at this point, a work of art, and few rules exist to guide the work. Perhaps, the design could be found automatically, for instance by the use of a genetic algorithm. Since there are two kinds of input features, the architecture could probably be even more exotic, treating them differently. Additionally, the SOH could be seen as a continuous function where the next value is very correlated with the previous. This is often used in neural networks by introducing recurrences, a link backwards between the layers. This is creating an artificial memory, an algorithm capable of remembering the previous value.

The recreation of the engineered SOH is interesting, but holds limited value for the aim of the project. It is interesting because it shows that the selected input features can be used to estimate the engineered SOH. However, the recreated values themselves inherit the unwanted property of the engineered SOH; its values alternate quickly and without a general degradation between cycles.

Teaching a neural network the dependencies of the signal poses a limitation. The neural network can be operated with input values seen in the training set. For instance, the voltage is normally around 300V, and there is no way to predict how the algorithm would react to a value of 600V if it is not present in the training set. Luckily, the parameters describing the charging event are purely physical and have defined limits. This means that the network will handle all values from this set of inputs. The history inputs do not have the same limitations. There is no maximum number of cycles.

The neural network is an intelligent system. It learns the dependencies of the different inputs and understands how each individual parameter is affecting the SOH value. Recreating the engineered SOH is equivalent to ask it to estimate the SOH values given a certain battery history and a particular charging condition. It could just as well be asked for an estimation in another charging condition. In fact, it could be a standard charging condition, and be used for all charging events, only the history of the battery changes. This is the main idea behind the estimation of an absolute SOH.

To the sceptic reader, the estimated absolute SOH does not seem to align with the measured values. Neither is this the intention! The network aims to account for all features of the charging condition that affects the measured value. The estimations are therefore only a good fit to the measured data if all charging parameters are evenly distributed around the standard. This must not be the case, and how the car is used can heavily influence this. Graph (G1;C6) exemplifies this since the measured values suffers a significant change after five years. At this point charges are becoming significantly longer, and are being carried out with higher current, which causes the engineered SOH values to increase. This must not be connected to a larger absolute SOH. The network estimation does not make this jump and appears to account for this usage change.

The performance of the network, and the validity of the estimated absolute SOH, is not fully tested. The validation could be done in three ways; evaluating the estimations of the network based on domain knowledge, using the errors of the recreated SOH, or comparing to real world measurements. While all give some credit, the last suggestion is the preferred one.

Evaluating the best performing cars gives a clear picture, the cars are used less. They are driven less, have less energy throughput, and experience less events. This is viewed as validation since it is intuitive that use accelerates the aging. Additionally, the network regards high power charging as more damaging, and the absolute SOH values are larger at higher temperatures. Both these properties aligns with what could be expected.

There are a few properties of the estimations that pose questions. First, curves within a group seem to converge for large times. This is interesting since it is a relevant question whether degradation will eventually approach a steady value. In this case it probably a shortcoming of the system. Second, the plateaus are non physical and the aging should not stop, even if the car is unused. Third, the car that recovers is probably an example of a car exceeding the information in the training set.

Evaluating the model in this manner is very powerful, but also risky. Many of the input features are probably exhibiting a strong correlation in the data set. For instance, the number of drives and the distance travelled. This means that altering one parameter could give the model a set of data points that are outside the data set and which might cause problems. Additionally, what is named the high power charging could be correlated with another usage feature, which is the cause of the accelerated aging.

At this point, an available validation is the accuracy of recreating the engineered SOH signal. By this measure, the algorithm has a global mean error very close to zero, only -1.7%. This is impressing, but the standard deviation does not show the same picture, being 9.9%. This, together with comparing recreated and engineered values, depicts an algorithm that learns some things about the system, but could definitely be improved. The two major areas of improvement are deemed to be: learning from individual cars and a more elaborate history representation.

If the errors from recreating the engineered signal is used to evaluate the network's performance, different cars are mapped with different precision. Some have a smaller mean error and some have an error of 10%. This could be a sign that the model is not detailed enough to enable individuality between the cars. Perhaps, cars with similar history have aged significantly different due to usage, for instance drive aggressiveness, which is not among the input features.

To exactly state the error of the estimated absolute SOH is difficult since there is no true value to compare to and the errors given when the network recreates the measured values are not precisely translated into the estimation of absolute SOH. Better yet, would be to issue standard testing of the cars, a full discharge perhaps, could be a valid approach. Such tests could be compared to the estimations done by the network.

7
Discussion

The battery degradation is one of the biggest challenges in vehicle electrification. It excels the environmental impact, damages the business case, and poses a risk to buyers and sellers. The key to counter the degradation is the ability to measure and connect it to usage of the battery. In this project, this is approached in stages, modules that solve different, consecutive parts of the main problem. The contribution of each module should be evaluated against the final results and to the main problem statement, which is exactly what this chapter aims to do.

7.1 Needed to Measure and Map SOH to Battery Usage

There exists an ambiguity in the term SOH. It refers to the state of the battery and its performance compared to the initial, but is often restricted to a single property, often the capacity or the resistance. Mapping to a single property limits the information described by the state. For instance, SOH based on capacity degradation does not provide all information on increased resistance or leakage current. Perhaps, a single SOH is not enough to accurately measure the age of the battery.

The project has focused on a capacity based SOH and the capacity is believed to vary with ambient conditions, such as temperature. Perfect sensors and minimum noise will not change this underlying insight, and makes it difficult to use this SOH as a reference for aging. Instead, a usable SOH signal should be corrected so that the measurement conditions do not affect the evaluated age of the battery.

Once a pure SOH signal is obtained, one that is independent of the measurement conditions, it can be mapped to the history of the battery. At this stage, it is easy to forget the different time scales involved. Aging will happen in a time scale of years, hopefully limited to a few percent per year, but measurements in the car are done with high frequency sampling. These scales have been found incompatible, meaning that car signals have to be aggregated. The aggregated state can then be mapped to the SOH label using a machine learning approach.

7.2 The Data Set and Data Logging in Cars

Traditionally, batteries are studied in laboratories, and degradation is observed through different cycling patterns, ambient conditions, etc. This gives a reliable source of information, but faces challenges in the number of interesting parameters. In detail, different dynamic uses, in different environments, at different ages, could affect the battery differently. Studying such details takes a massive amount of trails and batteries. A potentially better approach could be to see every car as an experiment. In this view, the dynamic use of the batteries is studied in vivo. This is why we should log battery data.

Data logging in cars is a challenge, and preprocessing is often a vital part in using the data. There are many steps between knowing what should be measured, to actually logging the data in a suitable way. Implementations and sensors will obscure the signals, sending and storing logged data will damage the measured values, and handling big data calls for general techniques that do not treat every sample point in the best possible way. This means that the end results will be dependent on the data available, and a proper preprocessing.

In a larger view, how can we with certainty state the environmental improvement of driving electric cars without additional data. Life cycle analysis includes the battery and how it degrades just as much as where you charge and from what source. It is therefore of great interest to collect more data.

7.3 The Presented Algorithm

Presented in this report is a proof of concept that data aggregation into events, and the use of a neural network can overcome the above challenges. It gives a more usable SOH than the signal engineered from the car data, and maps it to the history of the battery. This gives some interesting results; the degradation is faster in the first years, most cars end up around 86% SOH after six years, and the usage of high power charging is damaging to the battery.

The aggregation of measured values into events, a format more suitable for mapping SOH degradation to battery history, was done by domain knowledge and common sense, which means many assumptions. Based on the results, the used history representation is feasible to map the SOH, despite its simple nature. Future, more elaborate, history representation will give higher precision and better information.

The used label, the engineered SOH based on capacity, poses questions. The SOH is based on the current and SOC during a charge, but the workings, and specific implementations of both signals are not fully known to the writer. Especially, how the SOC is updated with cell OCV measurements has an impact on the project. It is therefore possible that the signal is related, but not equal to the capacity signal. However, according to the results, there is a degrading component in this signal that maps to usage of the battery, very reminiscent of a true aging signal.

The performance of the network itself can to be connected to the recreation of the engineered SOH. In this aspect, a few setbacks are visible; relatively large standard deviation, unexplained plateaus, and convergence of SOH for different cars at large times. There is more to be done in this regard and it is unknown how much the performance can be enhanced by the other stages of the project.

The used label has been found to be affected by measurement condition, and the use of a neural network successfully accounts for this. Both temperatures, SOC levels, rest, and currents are believed to play a role. But, the network learns how measurement conditions affect the signal, and then a standard condition can be used to estimate the SOH in a standard measurement.

The resulting estimation of an absolute SOH is interesting but needs to be validated. It is an estimation of a measurement that never occurred and has no true value for comparison. This is a problem, but one that is present in all attempts to map the SOH degradation to usage. In fact, standard testing batteries in cars would be very useful, measuring the capacity in known conditions. This could be done in workshops and would be excellent for validating the algorithm.

8
Conclusions and Future Work

A definitive answer to the problem statement can, at this point, not be given. The use of data aggregation into events, and a neural network to map and account for non history dependencies of the SOH label, is possible. It is perhaps even useful, but it is not properly validated. Therefore, validation is a priority and a next step.

However, from the findings, these challenging points have been discovered when attempting to measure and map SOH degradation to battery use:

- What property should the SOH be based on? and is a single SOH enough?
- Most SOH signals will be dependent on the measurement circumstances and this has to be accounted for to be able to see aging.
- Mapping the SOH degradation to use of battery calls for an aggregation of measured data to a time scale more compatible with aging.
- A technique for mapping the SOH degradation to the use must be found, a neural network is a candidate.

The project focused on a capacity based SOH, but it would be interesting to develop additional SOH for resistance, leakage current, available power, and more. This could be done in parallel since the above points are believed to be valid for all types of SOH. There is also interest in the degradation differences between different SOHs and the use that accelerates individual degradations. The neural network is a flexible beast. Is it possible to include the estimation of other states as well? Could one software determine them all? If possible, it could provide simplicity, efficiency and be an alternative for online implementation.

The found dependencies of the engineered SOH are not proven to be a complete set. There could be more, or there could be less. Some of the measurement conditions used are probably correlated which means that they could be combined into one feature. At the same time, some of the features exhibits weak effects on the SOH and could be omitted. However, there might also be additional signals, perhaps not even measured in the cars, that could affect the engineered SOH.

The history must be quantified so that measured data is reformed into features compatible to aging. This is a challenge since the collected data, with billions of data points, must be aggregated into a feasible format that preserves all the information about the causes of aging. Is the domain knowledge approach the most suitable? A future topic could be classifying drives, charges, etc., so that they are representing causes of aging in a better way.

The neural network is used as the mapping algorithm and to account for non history dependencies of the SOH label, but it is not the only algorithm with these abilities. Before deploying the software, there should be other machine learning algorithms tested. A random forest, or boosted solutions are interesting, but there are many more.

The presented model is trained on the training set and tested on a few cars, the validation set. But, is this separation necessary? There will be errors due to how specific the history part is and due to differences between the batteries. For instance, batteries could be different due to manufacturing. Could it be possible to let an online model adjust to these differences? This is often referred to as online learners and would be a suitable topic for future projects.

When a pure SOH is found, it can be used as the input for a prediction. For instance, the curve for the first year of a car's life could be an input to predict the SOH the upcoming year. A successful attempt would give interesting applications since it could help to correct driving behaviours, improve service scheduling, and make EOL predictions. However, for such prediction to be of any value, the input must be validated and more data, over longer periods of time, collected.

The project has been of exploratory nature, which is both favourable and challenging. There was no previous work to build this project from, nor a designated data set, or specific method. This meant that much work went into accessing the right data and people. Such work is not contained in the report, but very important. It also argues that solving the challenges connected to battery electric vehicles demands people from many parts of Volvo Cars Corporation, professionals in different parts of the field.

Bibliography

[1] Bruno Scrosati and Juergen Garche. Lithium batteries: Status, prospects and future. *Journal of Power Sources*, 195:2419–2430, 05 2010.

[2] M.A. Danzer, V. Liebau, and F. Maglia. 14 - aging of lithium-ion batteries for electric vehicles. In Bruno Scrosati, Jürgen Garche, and Werner Tillmetz, editors, *Advances in Battery Technologies for Electric Vehicles*, Woodhead Publishing Series in Energy, pages 359 – 387. Woodhead Publishing, 2015.

[3] International Energy Agency. Global ev outlook, 2016.

[4] Volvo Cars Media. Volvo cars mål är att 50 procent av försäljningen ska vara elbilar 2025, 2016.

[5] Anthony Barré, Benjamin Deguilhem, Sébastien Grolleau, Mathias Gerard, Frédéric Suard, and Delphine Riu. A review on lithium-ion battery ageing mechanisms and estimations for automotive applications. *Journal of Power Sources*, 241:680–689, 11 2013.

[6] Lisbeth Dahllöf Mia Romare. The life cycle energy consumption and greenhouse gas emissions from lithium-ion batteries. *IVL*, 2017.

[7] Wikipedia contributors. Energy density — Wikipedia, the free encyclopedia, 2018. [Online; accessed 29-May-2018].

[8] K. Ozawa. *Lithium Ion Rechargeable Batteries: Materials, Technology, and New Applications*. Wiley, 2012.

[9] J. Vetter, P. Novák, M.R. Wagner, C. Veit, K.-C. Möller, J.O. Besenhard, M. Winter, M. Wohlfahrt-Mehrens, C. Vogler, and A. Hammouche. Ageing mechanisms in lithium-ion batteries. *Journal of Power Sources*, 147(1):269 – 281, 2005.

[10] Xuebing Han, Minggao Ouyang, Languang Lu, Jianqiu Li, Yuejiu Zheng, and Zhe Li. A comparative study of commercial lithium ion battery cycle life in electrical vehicle: Aging mechanism identification. *Journal of Power Sources*, 251:38 – 54, 2014.

[11] Victor Agubra and Jeffrey Fergus. Lithium ion battery anode aging mechanisms. *Materials*, 6(4):1310–1325, 2013.

[12] Purewal J. Liu P. Hicks-Garner J. Soukazian S. Sherman E. Sorenson A. Vu L. Tataria H. Wang, J. and M. Verbrugge. Degradation of lithium ion batteries employing graphite negatives and nickel–cobalt–manganese oxide + spinel manganese oxide positives: Part 1, aging mechanisms and life estimation. *Journal of Power Sources*, 269(C):937–948, 2014.

[13] Wen Liu, Pilgun Oh, Xien Liu, Min-Joon Lee, Woongrae Cho, Sujong Chae, Youngsik Kim, and Jaephil Cho. Nickel-rich layered lithium transition-metal

oxide for high-energy lithium-ion batteries. *Angewandte Chemie International Edition*, 54(15):4440–4457, 2015.

[14] Matthieu Dubarry, Bor Yann Liaw, Mao-Sung Chen, Sain-Syan Chyan, Kuo-Chang Han, Wun-Tong Sie, and She-Huang Wu. Identifying battery aging mechanisms in large format li ion cells. *Journal of Power Sources*, 196(7):3420 – 3425, 2011.

[15] Alex Friesen, Fabian Horsthemke, Xaver Mönnighoff, Gunther Brunklaus, Roman Krafft, Markus Börner, Tim Risthaus, Martin Winter, and Falko M. Schappacher. Impact of cycling at low temperatures on the safety behavior of 18650-type lithium ion cells: Combined study of mechanical and thermal abuse testing accompanied by post-mortem analysis. *Journal of Power Sources*, 334:1 – 11, 2016.

[16] Qianqian Liu, Chunyu Du, Bin Shen, Pengjian Zuo, Xinqun Cheng, Yulin Ma, Geping Yin, and Yunzhi Gao. Understanding undesirable anode lithium plating issues in lithium-ion batteries. *RSC Adv.*, 6:88683–88700, 2016.

[17] M. Reichert, D. Andre, A. Rösmann, P. Janssen, H.-G. Bremes, D.U. Sauer, S. Passerini, and M. Winter. Influence of relaxation time on the lifetime of commercial lithium-ion cells. *Journal of Power Sources*, 239:45 – 53, 2013.

[18] I Bloom, B.W Cole, J.J Sohn, S.A Jones, E.G Polzin, Vincent Battaglia, G.L Henriksen, C Motloch, R Richardson, T Unkelhaeuser, D Ingersoll, and Herbert Case. An accelerated calendar and cycle life study of li-ion cells. *Journal of Power Sources*, 101:238–247, 10 2001.

[19] Wladislaw Waag, Christian Fleischer, and Dirk Uwe Sauer. Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles. *Journal of Power Sources*, 258:321 – 339, 2014.

[20] R. Hausbrand, G. Cherkashinin, H. Ehrenberg, M. Gröting, K. Albe, C. Hess, and W. Jaegermann. Fundamental degradation mechanisms of layered oxide li-ion battery cathode materials: Methodology, insights and novel approaches. *Materials Science and Engineering: B*, 192:3 – 25, 2015. Electrical Fatigue.

[21] Barbara Michalak, Heino Sommer, David Mannes, Anders Kaestner, Torsten Brezesinski, and Jürgen Janek. Gas evolution in operating lithium-ion batteries studied in situ by neutron imaging. *Scientific Reports*, 5:15627 EP –, Oct 2015. Article.

[22] Languang Lu, Xuebing Han, Jianqiu Li, Jianfeng Hua, and Minggao Ouyang. A review on the key issues for lithium-ion battery management in electric vehicles. *Journal of Power Sources*, 226:272 – 288, 2013.

[23] R.B. Phillips, J. Kondev, and J. Theriot. *Physical Biology of the Cell*. Garland Science, 2009.

[24] Simon Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2nd edition, 1998.

[25] Mcintosh. Li-ion battery aging datasets. NASA Ames Research Repository, September 2010.