



Forecasting Initial Phase Spare Part Demand Using Installed Base Data

A Case Study at Volvo Cars

Master's thesis in Supply Chain Management

AXEL LUNDH
MARTIN MARKLUND

DEPARTMENT OF TECHNOLOGY MANAGEMENT AND ECONOMICS
DIVISION OF SUPPLY AND OPERATIONS MANAGEMENT

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2020

www.chalmers.se
Report No. E2020:080

MASTER'S THESIS NO. E2020:080

Forecasting Initial Phase Spare Part Demand Using Installed Base Data

A Case Study at Volvo Cars

AXEL LUNDH
MARTIN MARKLUND



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Technology Management and Economics
Division of Supply and Operations Management
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2020

Forecasting Initial Phase Spare Part Demand Using Installed Base Data
A Case Study of Volvo Cars
AXEL LUNDH, MARTIN MARKLUND

Examiner: Patrik Jonsson, Chalmers University of Technology
Supervisors: Patrik Jonsson, Chalmers University of Technology
Rickard Holm, Volvo Cars
Advisors: Niklas Österlund, Volvo Cars
Claes Reje, Volvo Cars
Joakim Andersson, Chalmers University of Technology

Master's thesis NO. E2020:080
Department of Technology Management and Economics
Division of Supply and Operations Management
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 (0)31-772 1000

Gothenburg, Sweden
© Axel Lundh, Martin Marklund, 2020.

Abstract

The often inconsistent and irregular demand patterns of spare parts make them hard to forecast using traditional time series methods. This is especially true throughout the initial phase following market introduction, when no sales data exists. The advances in Big Data has given rise to a way of combating these difficulties; installed base data, which is mainly derived from the population and age of products on the market. Together with the development of new machine learning algorithms doors are opening for demand planners to utilize previously untouched data to improve forecast accuracy.

In this thesis, a machine learning model is developed to utilize installed base data for initial phase spare part forecasting to test, analyze and evaluate its performance, compared to a traditional exponential smoothing model. The spare parts involved are categorized according to their demand characteristics, and model performance for each demand patterns. This exploratory case study is performed in collaboration with the Demand Planning & Optimization team at Volvo Cars.

The findings of this thesis clearly show the potential of utilizing installed base data when performing forecasts in the early phase of spare part life-cycles. It is also shown that a machine learning model using installed base data can provide improved results when predicting the first sale of a given part, when no previous sales data exists. The decision to stock or not is crucial for many spare parts, and the classification analysis performed by this model significantly outperforms Volvo's results for a comparable period. Further, the combination of an installed base model with a traditional time series model shows significant promise, outperforming both pure models in all stages of the initial phase. This is especially true for parts with high demand variability.

Although the data used in this study is limited to that collected from Volvo Cars, the challenges faced and practices implemented by the case company are likely similar to those of many other manufacturing companies, and the authors argue for the transferability of its findings. Further, the methodology and analysis performed should be applicable to similar studies and projects elsewhere.

This study presents a way of producing demand forecasts before any sales data is available which outperforms current judgmental practices. The combination of the two models performs impressively, and suggests the combination of time series forecasting with machine learning models holds promise for both Volvo Cars and other similar companies.

The main contributions of this thesis lie in the production of demand forecasts without sales data, as well as the shown potential in combining explicit installed base models with more traditional time-series forecasting.

Keywords: Big Data, Spare Parts, Initial Phase, Spare Part Life-cycle, Demand Forecasting, Installed Base Data

Acknowledgements

The greatest lesson in life is to know
that even fools are right sometimes.

Winston Churchill

As Winston Churchill said, even a fool can be right sometimes. We started this project as fools ourselves, during the spring semester 2020. Luckily we had brave, steadfast and competent advice along the way. In a funny way, this also relates to forecasting, where a lucky guess may be right once or twice, but the name of the game is being right most of time.

First off, we would like to thank our solid supervisor at Chalmers, Professor Patrik Jonsson, for sharing his vast knowledge of the domain of forecasting and supply chain experience. His guidance helped keep us on the right track during some trying months, and to course-correct when we were veering off the path. You helped remind and inform us of key features of the study, and this aided us immensely.

Another massive thank you goes out to Joakim Andersson, Industrial PhD Student at Chalmers, for taking an interest in our project, despite not being under any obligation to do so. The time and effort you spent on us is admirable, and we are very happy for it! Your academic insights, advice on sources and methodology as well as experience was a great support to lean on throughout the project.

Further, a shout out to Rickard Holm, our supervisor at Volvo Cars and in charge of the DP&O team. We have felt and received your support from the very start, even when you've had your plate full with other stuff. Thanks for being a kind, helpful and sympathetic person, always keen to assist. We still hope to test-drive new Volvo's with you in the future!

Lastly, Niklas Österlund and Claes Reje... Where do we even begin? Thanks to the both of you for being our lovely and slightly nerdy advisors at Volvo! Our frequent contact has been vital, and your willingness to go out of your way to be part of the project has been much appreciated. Thank you for lending your expertise, hosting table tennis tournaments during lunch breaks, showing off impressive medieval weaponry and also emphasizing the importance of Sabaton for correct motivation! Hopefully our work can aid you in the future.

All that said, we'd also like to extend our gratitude to friends, family and other dearly beloved for, in one way or another, taking part in our journey.

Axel Lundh & Martin Marklund
Gothenburg, June, 2020

Glossary

Big Data A term describing the increasing volume and variety of available data mainly due to digitalization.

CatBoost An advanced open-source gradient boosting decision tree algorithm. It can be used for Classification and Regression problems. Compared to Random Forrest it provides faster training time, higher accuracy and a number of aids to simplify its implementation.

Causal-based Forecasting Forecasting that uses other data than historic sales to predict future demand.

Classification The process of classifying something based on its features. In machine learning a classification problem is the process of classifying a set of independent variables to a set of pre-defined categories.

Combined Model A model evaluated in this thesis. The model implements the CatBoost algorithm and takes spare part features, installed base data and the outcome of the exponential smoothing model as its inputs and use these to predict future demand.

Confusion Matrix A way to visualize and evaluate the performance of a binary classification algorithm. The matrix shows the four possible outcomes of a binary prediction: 1) The actual outcome is YES and the model prediction is NO, 2) The actual outcome is YES and the model prediction is YES, 3) The actual outcome is NO and the model prediction is YES, 4) The actual outcome is NO and the model prediction is NO. This allows evaluation of the model's performance on both true positives and true negatives.

Cross Validation A robust way to reduce overfitting in an algorithm. Instead of splitting data in a training and test set just once, Cross Validation does this multiple times. The number of iterations is often denoted 'K'-times. The model with the highest performance is deemed to be the one with the highest average performance over the K iterations.

Demand Pattern How the demand for a particular product behaves. Generally defined by the mean time between demand occurrences and the variability of demand size.

Dependent Variable The variable that an algorithm tries to predict with by using independent variables. This is generally the prediction output of an algorithm.

DP&O The Demand Planning & Optimization team at Volvo Cars.

Exponential Smoothing Model A model evaluated in this thesis. The model applies an exponential smoothing Time Series forecasting method to historic sales data.

Feature Another name for independent variable. In this thesis 'Feature' refers to the inherent attributes exhibited by individual spare parts such as what function group it belongs to or its weight.

Grid Search A method to aid in parameter tuning. The method takes a set of values for a various number of parameters, a so-called grid. The model is then trained with each possible combination of these parameter values in order to find the combination with the best performance.

Independent Variable A variable used to describe another. Independent variables are used as input for many Classification and Regression Machine Learning algorithms.

Initial Phase The initial part of a product's life-cycle. In this thesis defined as the first 24 months after market introduction.

Installed Base Data Information about the units of a product currently utilized on the market.

Installed Base Model A model evaluated in this thesis. This model implements the CatBoost algorithm and takes spare part features and installed base data as its independent variables.

Machine Learning Advanced computer algorithms that take data as input in order to learn and improve, while finding patterns and predicting outcomes when presented with new sets of data.

OEM Original Equipment Manufacturer.

Open-Source Computer software openly available and freely distributed, often with several individuals and organizations cooperating and contributing to its development.

Overfitting An issue where a machine learning algorithm becomes too fitted to the data it is trained on and performs worse when presented with previously unseen data.

Parameter Tuning The process of changing and evaluating parameter values in order to increase the predictive power of an algorithm.

Parameters A set of values which modify the behavior of an algorithm. The values impact algorithm performance.

Planning Frequency Determines the time between which forecasts are updated.

Planning Horizon The time covered by the current predictions of future demand.

Python A computer programming language commonly used in data science and many other fields.

Random Forest A kind of machine learning algorithm which can be used for Classification and Regression problems.

Regression A regression problem tries to predict a quantity or other outcome based on a set of independent variables.

SCM Supply Chain Management.

SKU Stock Keeping Unit, a unit for which a stock record is kept.

Test Data Data used to evaluate a machine learning algorithm previously trained on training data. This is done to ensure the algorithm has not only learned to make predictions about the data it has previously been trained on.

Time-Series Forecasting Forecasting based on how demand varies over time. Generally uses historic sales data as its input.

Training Data Data used to train a machine learning algorithm. The algorithm tries to find patterns in the data and in the case of Classification or Regression the algorithm searches for the relation between the independent variables and the dependent one.

Contents

1	Introduction	1
1.1	Background	1
1.1.1	Automotive Spare Part Demand	1
1.1.2	Installed Base Data	2
1.1.3	Demand Forecasting	3
1.1.4	Case Company	3
1.1.5	Problem	5
1.2	Aim	6
1.3	Scope	8
2	Theory	9
2.1	The Automotive Aftermarket	9
2.1.1	Spare Part Demand Drivers	9
2.1.2	Product Life-Cycles and the Installed Base	10
2.1.3	Demand Patterns and Item Categorization	11
2.2	Forecasting Methods	13
2.2.1	Time-Series Forecasting	14
2.2.2	Causal-based Forecasting	15
2.2.3	Forecasting Methods Utilizing Installed Base Data	16
2.2.4	Forecasting Errors and Accuracy Measures	16
2.3	Big Data Analytics	17
2.3.1	Dependent and Independent Variables	18
2.3.2	Training Data, Test Data and Overfitting	18
2.3.3	Random Forest	18
2.4	Conceptual Model	19
3	Methodology	21
3.1	Research design	21
3.2	Workflow and Structure	21
3.2.1	Phase 1: Orientation and Data Collection	21
3.2.2	Phase 2: Data Analysis and Evaluation	21
3.2.3	Phase 3: Results and Findings	22
3.3	Choice of Accuracy Measures	24
3.4	Data Collection	24
3.4.1	Quantitative Data	24
3.4.2	Qualitative Data	25
3.5	Data Pre-processing	26
3.6	Validity and Reliability	27
3.7	Overview of Analysis	28

4	Analysis	29
4.1	Planning Horizons	29
4.2	Forecasting Models	30
4.2.1	The Installed Base Model	30
4.2.2	The Exponential Smoothing Model	32
4.2.3	The Combined Model	33
4.3	Forecast Accuracy	34
4.3.1	Stage 1	34
4.3.2	Stage 2	35
4.3.3	Stage 3	35
4.4	Feature Importance	36
4.5	Demand Pattern Categorization	37
4.6	Model Performance per Demand Pattern	38
4.6.1	Lumpy	39
4.6.2	Smooth	39
4.6.3	Intermittent	40
4.6.4	Erratic	40
5	Discussion	41
5.1	The Implications of Model Performance	41
5.1.1	Stage 1	41
5.1.2	Stage 2	42
5.1.3	Stage 3	43
5.2	The Impact of Model Features	44
5.3	The Imposition of Demand Pattern Categorization	46
6	Conclusion	47
6.1	Recommendation	47
6.2	Limitations	49
6.3	Future Studies	50
	Bibliography	51
A	Appendix	I
A.1	Interview Questions	I
A.2	CatBoost Classifier	III
A.3	GridSearch	IV
A.4	Implementation of Exponential Smoothing	V
A.5	Feature importance	VI

1 | Introduction

This thesis is performed in collaboration with Volvo Cars Group, and their Demand Planning & Optimization (DP&O) team within the department of Supply Chain Management (SCM). This team in particular is tasked with demand planning for Volvo's global aftermarket. In order to carry out their objective they require accurate forecasts of future demand for spare parts. The demand for spare parts has unique characteristics which make inventory management challenging (Bacchetti and Saccani 2012). There are thus financial and customer service interests in evaluating and implementing more accurate forecasts during the initial two years of a spare part's life-cycle. One opportunity is to use installed base data, the effectiveness of which is examined in this thesis.

1.1 Background

This section provides a background to the thesis, as well as presenting an introduction to the topics covered, aspects relevant to forecasting using installed base data and the literature laying the foundation for the thesis in general. This is followed by a presentation of the case company in relation to spare part forecasting.

1.1.1 Automotive Spare Part Demand

The aftermarket for a product is the potential extended customer-buyer relationship after the initial sale of a product. According to Cohen and Agrawal (2006), the aftermarket includes services such as maintenance and repairs as well as the sale and provision of spare parts. Further, Cohen and Agrawal (2006) explain the ongoing shift towards a service centered economy and the aftermarket's role as an important part of this transition for traditional manufacturing companies, such as automotive manufacturers. Koudal (2006) estimates that for producing companies, service offerings can represent as much as 25% of total revenue and have a 75% higher profitability than the sale of finished products.

The life-cycles of physical products can, according to Dekker et al. (2013), be divided into three phases: initial, mature and end-of-life. The demand for spare parts generally follows the demand for their associated product rather closely, albeit with a slight time-delay. This time-lag has been shown by and dubbed the 'Life Cycle Mismatch' by Solomon et al. (2000). The demand patterns throughout product life-cycles obviously vary depending on the type of part, but in general this correlation can be said to hold true (Solomon et al. 2000). However, Dekker et al. (2013) argue that the size and age of the installed base and the distribution of related spare parts provides a more comprehensive explanation of spare part demand than relying on demand derived from the life-cycle phase of the product itself. Installed base information is during the initial phase, also known as 'design/purchase', mainly available from design specifications on expected product and spare part life and the initial purchase data as well as projected sales (Dekker et al. 2013). Combined with the relatively small size of said installed base and the lack of actual sales history,

this makes the initial phase especially challenging in terms of forecasting spare part demand. Effective utilization of the installed base data may still be used to improve demand forecasting in the initial phase (Dekker et al. 2013), and is discussed closer in the following section.

In addition to the issues mentioned above, Van der Auweraer et al. (2019) state that the majority of spare parts exhibit *intermittent demand patterns*. This means the demand for said parts is characterized by sequences of zero demand periods interspersed with non-zero demands. These demand patterns are known to make the approximation of demand difficult using traditional forecasting methods (Boylan et al. 2008). Further, when demand does occur the size of demand may vary widely. Such patterns are called *erratic*, while patterns which are both erratic and intermittent are called *lumpy* (Boylan et al. 2008). All these patterns pose challenges when estimating demand for spare parts compared to the smoother, more stable demands of finished products (Van der Auweraer et al. 2019). Several methods and forecasting techniques have been proposed to forecast intermittent spare part demand, and Boylan et al. (2008) suggest classifying parts according to their individual demand patterns and applying the appropriate methods separately.

1.1.2 Installed Base Data

The digitalization of industries has given way to what is today referred to as Big Data. McAfee and Brynjolfsson (2012) explain Big Data with the help of three V's, Volume, Variety and Velocity. According to Brinch (2018) this general definition of big data also holds true for Supply Chain Management. Today, companies have access to and have started collecting an ever-growing set of data from an increasing number of sources and at an increased speed, often close to real time. A major enabler for this kind of data is the increased connectivity of products which allow real-time and continuous reporting of various data. In the case of the automotive industry, the increased number of internet-enabled vehicles has created the opportunity to collect large amounts of information from and about each individual vehicle. This includes data from various sensors in the car, mileage, fault codes as well as geo-location (Andersson 2019).

An information source emerging along with and as part of Big Data is *installed base data*. The installed base is, in short, the number of sold products that can generate spare part demand (Dekker et al. 2013). Installed base data has potential to be used for spare part demand forecasting, since the demand for said parts is highly dependent on the age and size of the product population and originates from the replacement of existing parts either from preventative or corrective maintenance (Van der Auweraer et al. 2019). Van der Auweraer et al. (2019) summarize previous research in the field, concluding that correct knowledge of the characteristics of these products (such as age and usage) is highly important for correct inventory levels and can lead to stock reductions of up to 25%. Van der Auweraer et al. (2019) define installed base information as data derived from three main sources: i) the size and status of the spare part itself, ii) the chosen maintenance policy, and iii) environmental factors impacting part reliability.

Further, Van der Auweraer et al. (2019) summarize other findings from existing literature. An important aspect for forecasting in general, and the aim of this thesis in particular, is the connection between the size and status of the installed base and the replacement probability of the included parts to the actual demand for parts. This is also related to the life-cycle of the product. As mentioned above, demand for spare parts is related to the life-cycle of a product

but installed base data might be more useful for demand forecasting, especially in the initial phase of said life-cycle (Dekker et al. 2013).

1.1.3 Demand Forecasting

Forecasts are used on various levels of an organization in order to plan production and product availability more efficiently (Sanders and Ritzman 2004). More accurate forecasts allow for higher levels of service and in the case of products a lower need for stock keeping. According to Jonsson and Mattson (2009), forecasting methods can be divided into two major groups: *qualitative* and *quantitative*. Qualitative forecasts are based on human decisions and experience. Quantitative forecasts are instead based on data, generally sales history and past forecasts. Intrinsic quantitative methods, generally used for time series forecasting, look at data of past sales and try to predict future demand. Extrinsic quantitative methods instead utilize additional data besides historical sales to produce forecasts. Sanders and Ritzman (2004) explains that intrinsic methods are a good fit when demand is relatively stable, historical sales data is available and there is a need to generate forecasts for large amount of products. A major drawback of these methods is the need for historical data which make the forecasts hard to apply in situations with little historical data such as during the initial phase of a product or in situations with volatile and intermittent demand (Jonsson and Mattson 2009).

Extrinsic methods on the other hand have little or no need for historical data (Jonsson and Mattson 2009). Traditionally, access to and availability of alternative data has been limited but advances in analytic technology and the trend of Big Data opens up for new possibilities for extrinsic methods in demand forecasting (Andersson 2019). Relevant for this thesis are the possibilities to use extrinsic predictive models on data from the existing and near-future installed base to improve forecast accuracy. This mainly means the size and age of cars and their components sold, combined with numbers of sold spare parts to establish a reliability-based forecasting model.

1.1.4 Case Company

Volvo Car Group is a global automotive manufacturer and traces its roots back to 1927, with the founding of AB Volvo in Gothenburg, Sweden. In 1999, Volvo Cars was split from Volvo Group and sold to Ford Motor Company. Today, the company is owned by the Chinese company Zhejiang Geely Holding Group. Volvo Car Group also includes the electric performance car brand Polestar as well as stakes in the Chinese car brand Lynk & Co and Zenuity, a software company. In 2018, Volvo Cars sold 642,000 cars, had 43,000 employees on average and worked with approximately 2,300 dealers globally, with a net revenue of around 250 Billion SEK (Volvo Cars 2018).

More specifically, the thesis has been written in collaboration with the department for Demand Planning & Sourcing Optimization. This department focuses largely on the optimization and planning of global spare part logistics between suppliers, warehouses and dealers, assuring satisfactory customer service and spare part availability. DP& O have a heavy operational emphasis on the initial phase of the launch of a new automotive model and thereby its associated spare parts. The in-use fleet of Volvo Cars consists of around 6.5 Million cars, where 4 Million were produced and sold within the last 10 years (Volvo Cars 2018). Supporting this fleet, Volvo

provides a range of more than 130,000 spare parts, supplied through 54 warehouses of which 21 are Local Distribution Centers (LDCs). The after market and spare part business has an annual net revenue of 20 Billion SEK, made up of approximately 85% parts and 15% accessories. The service profits are Volvo's largest EBIT factor, and contributes to 60-70% of total dealer profit (Volvo Cars 2018).

Today, Volvo Car Group guarantees the availability of spare parts for up to 15 years after the car model goes out of production. With an average lifespan for most car models of 6-7 years, this means that spare part availability is required for up to 22 years after the car model launches. For Volvo DP& O, the initial phase is mainly considered to be the first two years after the launch of a product. This phase is where they experience most service and availability issues, due to the lack of sales data to base forecasting on. This phase is important for Volvo, since customers are well within their service guarantees and almost exclusively use Volvo's dealers and spare parts. Expectations on fast and effective service are also high¹.

The mature phase concerns the period after the first two years until the car model goes out of production. According to the above, this generally means years 2-7. During this phase, forecast accuracy is generally considered acceptable. The large installed base and existing sales history means that the historic and time-series based forecast are relatively accurate. Lastly, the end-of-life phase starts when the model goes out of production and lasts until the guarantee of spare part availability expires, meaning years 7-22. At this stage in the life-cycle, inventories are usually relatively well-established and the available data is considerable. However, customers generally start shifting from OEM parts to third-party dealers, while also repairing and servicing their car less frequently since the service guarantee offered by Volvo on new cars has expired. A further challenge during this stage is that suppliers of parts to Volvo may wish to discontinue production and/or charge higher prices due to less demand and worsened production economies. Lastly, since spare part demand is difficult to predict at the very tail-end of the cars' life-cycle, there are risks for Volvo and their dealers of being stuck with obsolete stock and associated costs².

The empirical focus of this thesis is Volvo's US market. The US market is an established and large market for Volvo Cars. In 2018, it was the company's third largest market in terms of retail sales, making up 15% of the companies total sales. The automotive market in the US experienced a small growth of 0.3%, while Volvo Cars recorded a much larger growth of 21%. In 2018, Volvo also opened their first car manufacturing plant in the US, located in Charleston, South Carolina (Volvo Cars 2018).

Regarding the aftermarket set-up in the USA, Volvo has five distribution centers throughout the country, see Figure 1.1. Each dealer is assigned to the DC within the closest geographical proximity, the dealer-facing DC, and is supplied from there. When a dealer requests a part, the request is first sent to this DC. However, all DC's do not permanently stock all spare parts. This might be for financial reasons, or physical space limitations. So, if the closest DC does not have the requested part, the request is sent on to one of the other four, creating a so-called 'referral'. Given the size of the US market, the distances to the DC stocking the part may be longer than what can be covered by truck within the required lead time. In these cases, the part

¹Personal Communication, Rickard Holm Volvo Car Group

²Personal Communication, Rickard Holm Volvo Car Group

is sent by express flight to meet the service requirement. This drives costs, and is a challenge Volvo is facing that is likely to increase along with the projected growth of the installed base³.

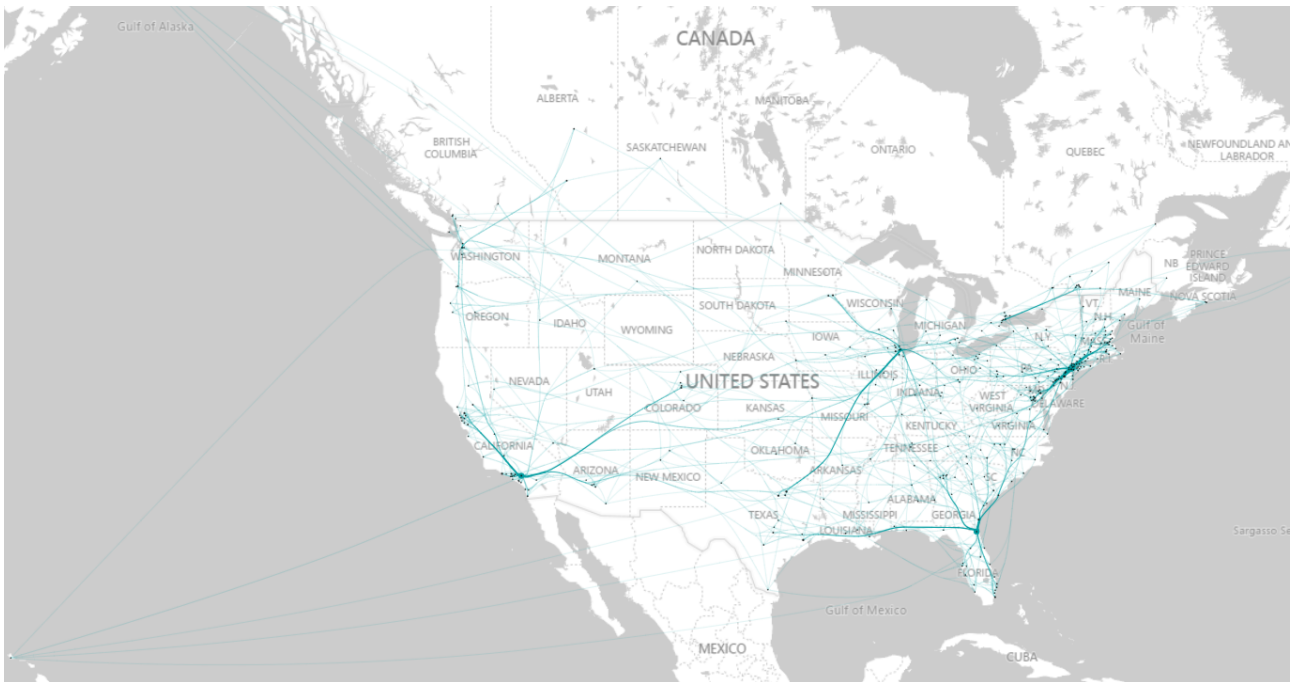


Figure 1.1: Overview of Volvo’s US distribution centers and part shipments.

1.1.5 Problem

Internal metrics at Volvo DP&O show a lack of forecast accuracy during the initial phase of new car models and the associated spare parts, leading to availability issues and expensive express shipments. Studies confirm that this is a common challenge in the spare parts market (Van der Auweraer et al. 2019), since the installed base is generally small and sales data limited. Traditional time series forecasting and intrinsic methods struggle to provide reliable forecasts when demand is slow or intermittent, demand patterns are shifting throughout product life-cycles and short delivery times from stock are important, all of which are common attributes for spare parts (Dekker et al. 2013). As mentioned previously, these issues are exacerbated in the initial phase of a products life-cycle, and utilizing installed base data is a potential way of improving forecast accuracy and mitigating the effects of having little or no historic sales data.

In the specific case of Volvo Cars, the lack of historical sales data renders the company’s current forecasting methods ineffective during this period. The DP&O team is thus investigating ways to improve their forecasting methods, using both new models and new data sources. The company has access to increased amounts of installed base data, in terms of cars sold, planned and spare part usage, especially from the US market. This opens up possibilities to use this installed base data with new forecasting models. As Dekker et al. (2013) show, there are significant savings to be made in inventory and obsolescence costs by collecting and utilizing installed base data. Thus, the potential to use installed base data in spare part forecasting is relevant not only

³Personal Communication, Rickard Holm Volvo Car Group

to Volvo Cars, but to producers and providers of spare parts generally, who often struggle to maintain high part availability while balancing stock-out and obsolescence costs.

1.2 Aim

The aim of this thesis is to examine the possibilities of using installed base data in the forecasting of spare part demand. This means developing a forecasting model using methods from forecasting literature, and evaluating it using actual demand and sales data from Volvo Cars. The project will be performed in collaboration with the DP&O team at the Supply Chain Management division at Volvo Cars, and data from their American market will be used as a case study to develop, test and evaluate findings.

The targeted life-cycle phase is the initial launch phase of new spare parts, often correlating with the phase-in period for new car models. For the purposes of this report, and the organizational structure of the case company, this is considered to be during months 0-24 (year 0-2) after the launch of a new spare part. The initial phase is further divided into three stages; (i) before any sale of the part has occurred, (ii) the part has sold in three demand periods and (iii) after the end of stage (ii) until the part has been on the market for 24 months. These stages have different demand characteristics and challenges, and will all be treated in this thesis. See Table 1.1 for details on each stage.

Stage	Description	Planning Horizon	Problem	What to predict
1	No Sales	6 m	Classification	(Yes or No) Sale in 6 m
2	First sale - 3 m sales	3 m	Regression	Demand next 3 m
3	3 m with sales - 24 m	1 m	Regression	Demand next 1 m

Table 1.1: Overview of the three planning stages during initial phase, the period from market introduction until 24 months (m) after market introduction.

During the analysis, a model will be constructed based on methods and theories from forecasting literature, as well as practical experience from the domain experts at Volvo Cars. This model will be designed with installed base data in mind, and its theoretical suitability to the data provided by Volvo. The model will be iterated upon and developed using more data points, parameters and aspects from other relevant forecasting methods as the analysis progresses. The analysis is performed in order to evaluate the forecast accuracy of the model, and will be discussed and evaluated through collaboration with said domain experts, as well as forecasting accuracy measures.

In addition, the models performance on spare parts with various demand patterns is evaluated. This involves categorizing parts according to their demand patterns, and testing the developed model for each category.

Lastly, the final stage of the thesis is evaluation of the results and potential impact. This results in a discussion of its usefulness and applicability to situations similar to that of Volvo Cars. In addition to the final forecasting model and its application to the selected group of parts, this

includes recommendations to the case company in terms of how to approach and conduct initial phase forecasting in the future, using this or similar models.

In conclusion, the research questions to be analyzed and answered in this thesis are the following:

RQ1: *What forecasting model can be applied to effectively predict initial phase spare part demand for Volvo Cars using installed base data?*

RQ2: *How would the applied model impact forecast accuracy for products with different demand patterns?*

Research Question 1 is answered through a combination of literature review on the nature of spare part demand and relevant forecasting models and the development and evaluation of a model with potential to outperform traditional methods. It also serves to examine the usefulness of installed base data for initial phase forecasting, while showing its theoretical and practical implementation.

Research Question 2 is mainly an extension of the forecasting models used to answer RQ1, with an added emphasis on the demand patterns of the spare parts involved and the impact said patterns have on model performance. This is to provide a framework and context within which the performance of the developed model and similar forecasting methods can be evaluated and practically implemented.

The results of the analysis are discussed in Chapter 5, Discussion, from various perspectives. The paper concludes with Chapter 6.1, Recommendation, and Chapter 6.3, Future Studies, where the implications of the study are summarized, both in regards to Volvo Cars and the field more generally.

1.3 Scope

The scope of this thesis has been set to facilitate the answering of the research questions, considering the current organization and situation of Volvo Cars. The thesis has been written with the DP&O team, and the knowledge and expertise within that team has been the source of qualitative data. Other actors within the SCM department have also been consulted, but to a lesser extent.

To keep the analysis focused on the issue at hand, quantitative data has only been collected through the DP&O team. No data from external sources has been collected for the quantitative analysis. In order to achieve applicability to other contexts with similar challenges, extensive theory and literature on spare parts has been studied, mainly in terms of demand planning and forecasting as well as aspects of spare part demand that complicate the issue. This means that the thesis focuses on the life-cycles of spare parts, mainly the initial phase, and the various demand patterns they exhibit. This has been examined and connected to the development of the installed base over time, in order to both achieve practical use in the case of Volvo Cars and allow for transferability to other situations in similar contexts.

To facilitate the thesis' main focus on the development and analysis of an initial phase forecasting model, both installed base and sales data has been collected from Volvo's US market. This data was both plentiful and easily accessible through the DP&O team, and meant time and resources otherwise spent on the gathering and comparing of multi-national and varied forms of data could be used to further the aims of the thesis.

Lastly, during the writing of this thesis the world suffered through the COVID-19 pandemic, and Volvo and many other companies limited their staff's working hours and disallowed meetings in person. This limited the possibilities of conducting interviews with various parties, which meant that qualitative aspects from other business areas which were initially considered were excluded from the scope of the thesis.

2 | Theory

In the following sections, existing literature concerning the topics of the spare part market and forecasting methods are presented.

2.1 The Automotive Aftermarket

In the automotive industry, spare parts and accessories, make up a vast part of the revenue associated with a certain car model (Souza et al. 2011). Furthermore, after the initial sale, the cars' long term reliability and life-time is dependant on replacement and maintenance of its component parts. The aftermarket and sales of spare parts is thus a crucial part of an automotive manufacturers business and profitability, as well as customer satisfaction and value proposition (Khajavi et al. 2014).

McKinsey&Company (2017) estimate that the aftermarket generated around 20% of the total revenues of the automotive industry in 2015, which meant approximately \$ 760 billion. Service and maintenance makes up 45% of that number, while the remaining 55% is constituted of the retail and wholesale of vehicle parts (McKinsey&Company 2017).

In general, the aftermarket can be divided between the networks surrounding and controlled by the OEM's, such as Volvo, and independent actors. As mentioned above, the market is very lucrative for most actor's, but in recent years several OEM's have strived to increase their influence and market share through various service programs as well as utilizing their greater network power, access to end customers, and data derived from the cars' usage (McKinsey&Company 2017). Customers tend to be especially loyal to the OEM during the early phases of a cars life-cycle, while the warranty is still valid and the value of the car remains high¹. Leveraging these inherent strengths successfully will be necessary in a market with stiff competition and several potential disruptive factors on the horizon. One way of gaining or increasing one's advantage is improved forecasting, making sure that related costs are kept low and service levels high.

2.1.1 Spare Part Demand Drivers

The need to replace a specific part could come from a number of factors. In their work on maintenance scheduling, Yang et al. (2008) identify and explore four different maintenance strategies; Preventive, Predictive, Condition-based and Corrective maintenance. These strategies are also considered the main underlying demand drivers for spare parts, as stated by Andersson (2019). The first three approaches are all different variations of pro-active and preventive maintenance, while corrective maintenance is reactive and applied after part failures (Yang et al. 2008).

¹Personal Communication, Rickard Holm Volvo Car Group

Preventive, or time-based, *maintenance* is a traditional and commonly used maintenance approach usually based on various service intervals, such as mileage or time. The aim of such policies is to reduce the need for corrective maintenance due to breakdowns, by following regular service plans and pre-empting part failures (Kennedy et al. 2002). In terms of forecasting, preventive maintenance theoretically makes the planning of item demand easier since the maintenance is regularly scheduled and also allows for the logging and utilization of part usage data. It does however require high degrees of customer loyalty and adherence to the service plan in order to be effective, and can also result in increased maintenance costs due to over maintenance and pre-mature replacement of parts (Andersson 2019).

An extension of the of the preventive maintenance explained above is called *predictive maintenance*, and utilizes on-board sensor data of monitored parts to predict part failures in advance. Forecasting based on such data will be discussed later in this section, but basically this approach replaces static service intervals with dynamic maintenance schedules based on the monitored status of involved parts (Yang et al. 2008).

Condition-based maintenance, *CBM*, is a similar approach to predictive maintenance, but utilizes real-time data from the on-board sensors. Andersson (2019) mentions methods utilizing CBM to optimize the total cost of maintenance, balancing costs of under- and over-maintenance. Such CBM approaches make use of data such as performance and physical degradation, vibrations and other usage-based factors. This approach is the analytically most advanced of these maintenance methods, and is becoming increasingly feasible in practice with technological advances (Fritzsche et al. 2014).

Lastly, *corrective maintenance* is usually defined as an operation done after the part has broken down or failed. As such, it is obviously a reactive approach compared to the pro-activeness of those above mentioned. This approach focused on repairing malfunctions usually leads to higher costs and longer down-times than preventive methods, and it is generally recommended to move towards said preventive and predictive approaches (Andersson 2019).

2.1.2 Product Life-Cycles and the Installed Base

The classic product life-cycle, as presented by Levitt (1965) consists of four steps and describes how demand and sales volume change over time for a product, from market introduction until it is phased out and no further sales are made. The first stage of a product's life-cycle, according to Levitt (1965), is the *The Development Stage*. This stage is characterised by lack of information and low market knowledge. Since the market has not been exposed to the product before it is hard to estimate the market's reaction and associated sales volumes. The introduction of new products is generally seen as risky, and failed introductions are often costly (Levitt 1965). Should a product introduction be successful, this stage is followed by the *Growth Phase*. During this phase sales ramp up, competitors may take notice and try to copy the new products and market awareness of the product increases (Levitt 1965). Eventually, the growth rate of sales starts to flatten out and the product reaches the *Maturity Stage*. During this stage, the product has generally fulfilled its market potential and substitutes start to catch up and price-based competition increases (Levitt 1965). The fourth and final stage is *Market Decline*, which the product enters when sales start to slow down and decline. During this stage, customers move away from the product, often looking for and finding substitutes or other replacing products.

Lastly, the product reaches its end-of-life and is discontinued (Levitt 1965).

According to Dekker et al. (2013), the installed base of a product closely follows the product's life-cycle, as described above. Further, the authors state that demand for spare parts is connected to sales of the actual product and growth of the installed base. Generally, demand for spare parts will follow product demand, albeit with a slight time delay. This 'Life-cycle Mismatch' as presented by Solomon et al. (2000) is the time and volume gap between product and spare part demand. Dekker et al. (2013) expand on this concept and the product life-cycles explained by Levitt (1965) to present a modified version of the product life-cycle. This was mentioned in Section 1.1.1, and focuses on the size of the installed base over time and corresponding spare part demand. Their visualization is shown, in Figure 2.1. In this case, the *Initial phase* can be said to replace the first two stages of Levitt (1965)'s model, Market Development and Growth. The *Mature Phase* is generally the same in both models, while Market Decline is replaced by the *End-of-life Phase*. Based on this simplified model, Dekker et al. (2013) argue and show that the size and age of the installed base provides a comprehensive explanation of spare part demand, and emphasize the potential to improve demand planning if these factors can be included in spare part forecasting measures.

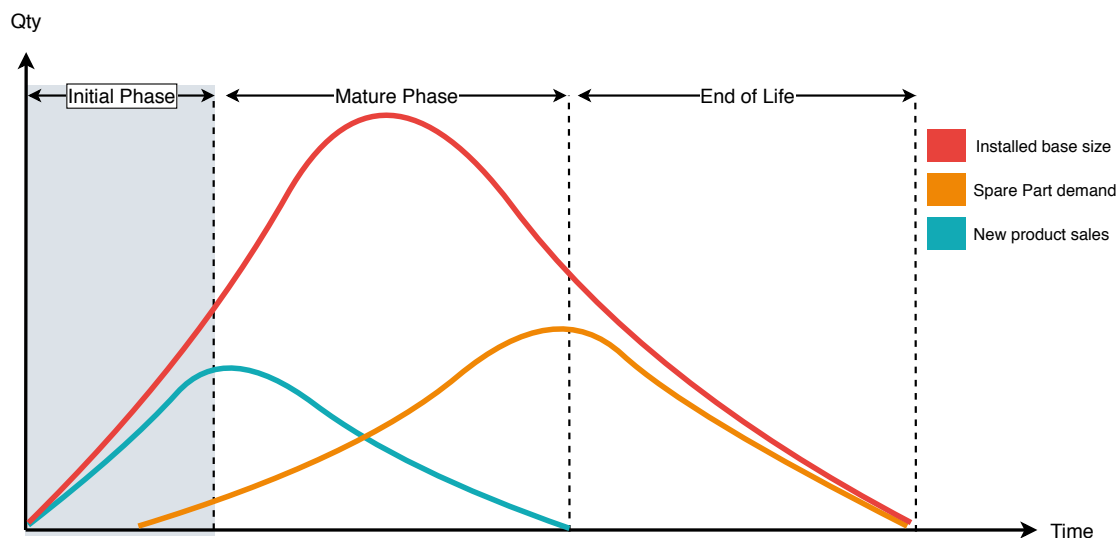


Figure 2.1: Graphical representation of the installed base size, new product sales and spare part demand across the Product Life-cycle, adapted from Dekker et al. (2013). The initial phase is highlighted in grey.

2.1.3 Demand Patterns and Item Categorization

As mentioned in the introduction, there are several factors and characteristics that make demand for spare parts different compared to that of the original products. Bacchetti and Saccani (2012) mention several such aspects that combine to complicate the forecasting of said demand, the presence of so-called lumpy and/or intermittent demand being one of them. The authors go on to discuss and highlight the importance of item categorization, through either quantitative or qualitative approaches, in order to determine appropriate stock-keeping strategies and forecasting methods for various parts, rather than using a uniform approach for all. Commonly used methods include ABC-classification into item categories, often based on

demand volume and part criticality (Bacchetti and Saccani 2012). However, the traditional ABC-approach struggles to take into account the intermittency of demand, the variations in order frequency which is common to many spare parts.

Another approach that includes demand frequency in addition to demand variability is presented by Boylan et al. (2008). The authors suggest classifying the demand patterns of all products in terms of their *mean inter-demand period*, p , and the variability of demand sizes, measured through the *squared coefficient of variation of demand size*, CV^2 . Mean inter-demand period is quite simply the average number of periods between each period with demand. The squared coefficient of variation of demand size is the coefficient of variation, $(\sigma/\mu)^2$, is the standard deviation of demand divided by the mean demand. Important for CV^2 calculations is to note that zero demand periods are ignored, in order to achieve the variation of demand when it actually occurs (Boylan and Syntetos 2007). This framework is conceptual, since high and low values are not generally quantified and thus need to be specified for the specific environment it is to be used in. The authors suggest statistical evaluation of the performance of various forecasting methods, and adjusting the cut-off values according to the superior performance of the evaluated methods. The authors do however suggest specific cut-off values, based on their comparative study of the performance of the (Croston 1972) and (Syntetos and Boylan 2005) methods. The values recommended are $CV^2 = 0.49$ and $p = 1.32$. Regardless of chosen cut-off values, they divide the classification matrix between high and low along both axes, as is shown clearly in Figure 2.2. For the purposes of this thesis, these cut-off values are used for demand pattern categorization. The adaptation of this framework and its application to the examined data sets will be discussed closer in Chapter 4, Analysis.

Since infrequent demand and irregular demand sizes make normal distribution unsuitable for representing the demand pattern, Syntetos et al. (2005) present several alternative patterns which are used in the categorization mentioned above:

An *intermittent demand* pattern means that the item shows a demand pattern with infrequent demand occurrences but with rather stable demand volumes. Intermittent demand, according to the framework, thus exhibits high mean inter-demand intervals with low variability of demand (Boylan et al. 2008).

An *erratic demand* item is an item whose demand size is highly variable, while the inter-demand periods are rather short. Erratic is considered the opposite of intermittent, since the variability lies in the size of the demand rather than varying inter-demand periods (Boylan et al. 2008).

A *lumpy demand* item is an intermittent item for which demand, when it occurs, is highly variable. Lumpy demand pattern exhibit the high inter-demand periods of intermittent demand, while the demand itself varies highly as in erratic demand patterns (Boylan et al. 2008).

Lastly, a *smooth demand* pattern means a more stable demand, with low variability as well as short inter-demand periods. For items with smooth demand patterns, traditional methods such as Croston (1972) are more likely to be effective (Boylan et al. 2008).

A two-dimensional framework with cut-off values for low and high then theoretically determine the appropriate forecasting model to be applied to items with each demand pattern (see Figure 2.2). As mentioned above, the framework provides a useful way of classifying spare parts

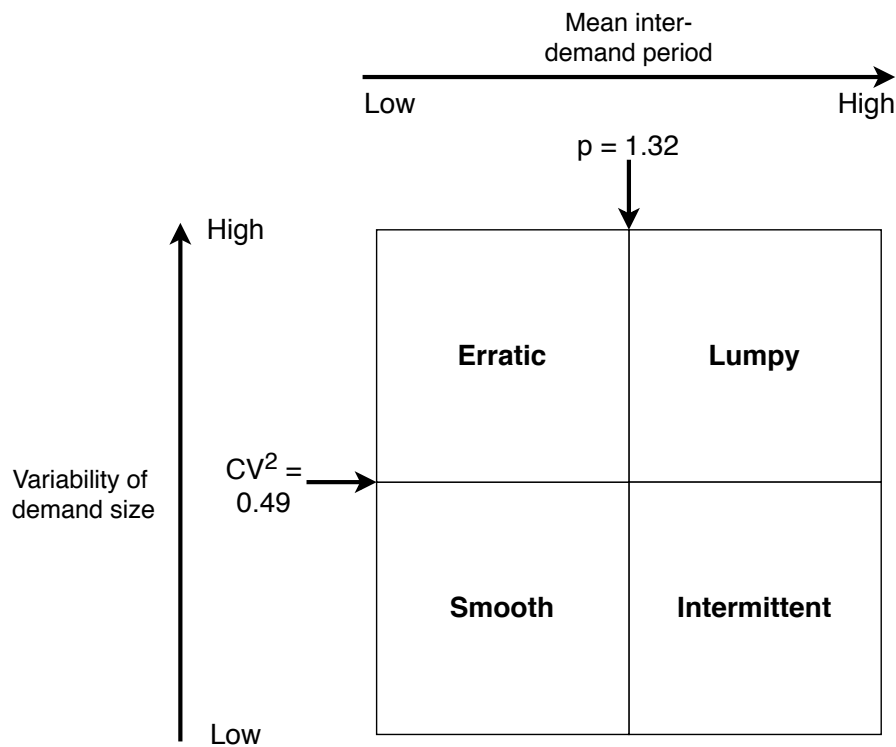


Figure 2.2: Demand-based categorization for forecasting, adapted from Boylan et al. (2008)

according to their underlying demand patterns, allowing selection and analysis of the applied forecasting model on an aggregated level.

2.2 Forecasting Methods

For the purposes of this report and its theoretical framework, forecasting and demand planning will mainly be described in the context of the aftermarket and spare parts. This bears obvious similarities with demand planning for manufacturing, but has several characteristics that differ. According to Cohen and Agrawal (2006), these characteristics make the aftermarket planning more complex and include more unpredictable demand, an increased number of SKUs and shorter expected response times. In the following sections, forecasting methods are presented as they relate to spare part demand with an emphasis on time-series and causal-based forecasting.

According to Jonsson and Mattson (2009), forecasting is a way for companies to make assessments of future external factors not fully under their own control, such as future market conditions. The authors divide traditional forecasting methods into quantitative and qualitative methods. Quantitative methods rely on information such as time series of sales and other historical data to calculate estimated demand, while qualitative methods instead make use of the expertise, experience and subjective judgement of individuals or groups. This includes methods such as market research, focus groups and Delphi methods, which might be especially useful in combination with more quantitative approaches (Andersson 2019). Quantitative methods are the main focus of this report and are therefor elaborated upon below.

2.2.1 Time-Series Forecasting

Quantitative methods use various calculations based on available data to estimate future demand. Boylan and Syntetos (2007) state that quantitative methods can be further divided into two different approaches; time-series based and causal-based. Time-based methods are entirely dependent on the history of demand to calculate future demand, while causal-based methods depend on certain explanatory variables to predict future outcomes. Such causal methods are especially useful in the initial phase of a parts life-cycle since there is no (or very little) demand history, making time-series difficult to apply successfully. Time-series methods are, on the other hand, commonly used for fast moving parts with varying demand characteristics (Boylan and Syntetos 2007). Demand characteristics will be discussed later in this chapter.

Similarly, Jonsson and Mattson (2009) discuss the differences between intrinsic and extrinsic forecasting methods. Intrinsic methods exclusively analyze data inherent to the variable being forecast, while extrinsic methods attempt to create a model connecting the forecasted variable to one or more explanatory variables. Hence, intrinsic as a forecasting concept is very similar to time-series while extrinsic closely resembles causal-based methods.

A commonly used quantitative and time-based method is the *moving average method* (Jonsson and Mattson 2009). This method calculates the demand for a given period by including the average demand from several previous periods. This evens out fluctuations due to random demand variations, smoothing out the demand curve. The formula for moving average is presented below, in Equation 2.1, as described by Jonsson and Mattson (2009).

$$\mathbf{F}(t + 1) = \frac{(D(t) + D(t - 1) + \dots + D(t - n + 1))}{n} \quad (2.1)$$

Where:

$\mathbf{F}(t + 1)$: Forecast demand for period $t + 1$

$D(t)$: Actual demand during period t

n : number of periods in the moving average forecast

It is clear from the formula above that the moving average method puts equal importance on every period included in the calculations, which isn't always desirable. The *exponential smoothing method*, as mentioned by Jonsson and Mattson (2009), is a way of using a simple time-series based method but introduce a weighting of the most recent period compared to previous forecasts. This is done by introducing a smoothing factor, called α , with a value between 0 and 1. A smoothing factor closer to 1 will make the calculation more responsive to systematic changes of the demand, while also becoming more exposed to random variations (Jonsson and Mattson 2009). The formula, as the authors state it, is presented below in Equation 2.2.

$$\mathbf{F}(t + 1) = \alpha * D(t) + (1 - \alpha) * F(t) \quad (2.2)$$

Where:

$\mathbf{F}(t + 1)$: Forecast demand for period $t + 1$

$D(t)$: Actual demand during period t

α : Smoothing factor

The main benefits of the two traditional methods mentioned above are their ease of use and applicability to fast-moving parts. However, Andersson (2019) states that research in both theory and practice shows their insufficiency when forecasting intermittent and/or slow-moving demand. The author brings up several methods that attempt to bridge the performance gap, such as the most commonly used *Croston method* (Croston 1972) and a variant thereof developed by Syntetos and Boylan (2005). Other methods include ones based on the ARIMA methodology (Hyndman and Athanasopoulos 2018) and the use of demand aggregation (Bartezzaghi and Kalchschmidt 2011). Even with such improved and expanded methods, the issue of time-series methods not considering the actual underlying demand remains (Andersson 2019). This fact combined with the increasing availability of various product-in-use and installed base data, as well as increased analytical capabilities, means the case for developing and using causal-based methods is gaining academic traction and practical feasibility (Andersson 2019).

2.2.2 Causal-based Forecasting

Various and plentiful alternative methods are suggested for predicting spare part demand, in cases where more traditional time-series methods are unsatisfactory. Andersson (2019) mentions *causal-based forecasting*, *judgemental forecasting* and *mixed method forecasting*. Causal-based methods will be elaborated upon below.

Causal-based methods differ from time series methods in the sense that they do not depend on historical demand data to predict future demand. Instead these methods aim at finding relations between extrinsic variables and demand and using that connection to predict future demand. Andersson (2019) mentions three categories of causal-based forecasting methods applicable to the automotive aftermarket, which are presented below.

Firstly, *reliability-based* forecasting looks at the expected failure rate or life-time of each spare part and correlates this number with the installed-base of said spare part and can in that way predict future demand. In their work on managing maintenance spare parts, Cavalieri et al. (2008) present a forecasting framework that uses the failure rate of components together with the existing number of products in the installed-base to replace the need for demand history. The authors present and argue for the usefulness of such a reliability-based approach in situations where historical demand doesn't exist as well as if the demand patterns are lumpy and/or intermittent (Cavalieri et al. 2008). According to Andersson (2019), there are two main methods for retrieving the necessary data when including failure rate in forecasting; collecting the data from databases where the various product failure rates are stored or using dynamic data from monitoring equipment to compile Life Data Analysis.

Methods utilizing data from on-board sensors installed in each product, such as vehicles and vehicle parts, are generally known as *on-board sensor based* (Andersson 2019). The general idea is to find causal connections between sensor values and the need for part replacement and hence a demand for spare parts. This category of methods is commonly referred to as *condition-based maintenance*, and the main purpose is to improve the planning of maintenance on single vehicles (Andersson 2019). However, the data collected from individual vehicles could be aggregated to produce a forecast for the total population, according to Andersson (2019).

The third and final category mentioned by Andersson (2019) consist of *regression-analysis based*

methods. As the name implies, these methods use regression analysis including multivariate linear regression, ARIMAX and ARX (Andersson 2019). The main concept here is the attempt to find a causal connection between demand and multiple extrinsic variables. This is sometimes referred to as *leading indicator forecasting*, and can be combined with machine learning methods to further improve the accuracy of demand forecasts (Wheelwright et al. 1998).

2.2.3 Forecasting Methods Utilizing Installed Base Data

From theory and the literature review presented above, a number of promising forecasting methods for utilizing installed base data have been identified. The method selected for further study are presented in the conceptual model section below.

Andersson (2019) and Van der Auweraer et al. (2019) suggest *reliability based methods* have the potential to successfully predict the lumpy and intermittent demand of many spare parts. Further, Liu and Tang (2016) effectively apply a reliability based method on installed base data in order to predict spare part demand. The authors look at the time from the installation of a specific system until it fails and a need for a replacement part arises. Hence Liu and Tang (2016) highlights a way to apply a Reliability-based forecasting method with the help of installed base data and without access to measured or calculated reliability information for individual parts.

Steuer et al. (2018) suggest an alternative causal based method for spare part demand. A *Demand-similarity Method*. Steuer et al. attempts to predict overall life-cycle demand by analysing the demand pattern for spare parts previously sold. This is done in three steps; 1) Cluster or sort previous spare parts based on the shape of their normalized (failure rate) demand pattern. A representative demand pattern is then determined for each of these clusters, 2) Use specific features of products to classify new products which have yet to see a full demand cycle into one of the clusters, 3) Use the normalized demand pattern of the clusters as a predictor by multiplying it with the planned installed base at the particular time of the forecast.

2.2.4 Forecasting Errors and Accuracy Measures

Jonsson and Mattson (2009) mention, forecasts are always estimates of the real outcome and are thus more or less accurate in relation to said outcome. A forecast error is the difference between the actual observed value and its forecast. There are several ways of aggregating forecast errors, and thus measuring their accuracy (Papalambros and Wilde 2018). Forecast errors are normally measured on a period-by-period basis, and are in their purest form the difference between the forecast of one period and the actual demand for that period. Over time, and especially for automatic methods of forecasting, the monitoring of these errors is crucial. The aim is to identify both individual random errors as well as systematic errors creating either too low or too high forecast values (Jonsson and Mattson 2009).

A common method for continuous monitoring of forecast errors is calculating the *mean error*, *ME*, and the average value of the forecast error in absolute terms. This is usually called the *mean absolute deviation*, *MAD*, and does not consider whether the forecast error was higher or lower than the actual demand (Jonsson and Mattson 2009). ME and MAD are scale dependent measures, meaning that the value of the errors are on the same scale as the object of the forecast.

An alternative, but still scale dependent, measure to MAD is the *mean square error*, *MSE*, which estimates the variance of demand when the mean demand is 0. In practice, using MSE as your measure penalizes large forecast errors harder than ME, but is naturally suitable for items with specific demand characteristics (Jonsson and Mattson 2009).

The *mean absolute percentage error*, MAPE, is the most commonly used measure of percentage errors. MAPE; and other percentage errors, are scale independent and are thus not on the same scale as the objects being forecasted. The advantage of such measures is that they are unit free, and can thus be used to compare forecast performances between data sets (Hyndman and Athanasopoulos 2018). A percentage error is calculated as the forecast error relative to the observed value. See the standard formula for MAPE calculations in Equation 2.3, as presented by (Jonsson and Mattson 2009).

$$\mathbf{MAPE} = \frac{1}{n} \sum_{t=1}^n \left| \frac{D(t) - F(t)}{D(t)} \right| \times 100 \quad (2.3)$$

Where:

$\mathbf{F(t)}$: Forecast at time t

$\mathbf{D(t)}$: Actual demand at time t

MAPE calculates the absolute mean of these errors and is thus susceptible to being infinite or undefined for observed zero values, and having extreme values if the observed value is close to zero (Hyndman and Athanasopoulos 2018). According to Hyndman and Athanasopoulos (2018), these factors, in addition to MAPE penalizing over forecasting more heavily than under forecasting, led to the introduction of the so-called *symmetric mean absolute percentage error*, sMAPE, which is widely used by actors throughout industry today (Hyndman and Athanasopoulos 2018), Volvo DP&O being among them. The risk for divisions by zero still exists though, especially for products with intermittent demand characteristics. If both the observed value and the forecast is 0, the denominator will be 0. This means other quantitative measures, such as mean square error, or even qualitative judgements may be more suitable for such products. The equation for sMAPE, according to Hyndman and Athanasopoulos (2018), is presented below in Equation 2.4.

$$\mathbf{sMAPE} = \frac{1}{n} \sum_{t=1}^n \frac{|F(t) - D(t)|}{(|D(t)| + |F(t)|)/2} \times 100 \quad (2.4)$$

Where:

$\mathbf{F(t)}$: Forecast at time t

$\mathbf{D(t)}$: Actual demand at time t

2.3 Big Data Analytics

The below section introduce and describes a number of concepts within the topics of Big Data analytics and machine learning. In addition specific algorithms applied in this study are described in detail. Hastie et al. (2009) and Rebala et al. (2019) provide an introduction to many of the central topics.

2.3.1 Dependent and Independent Variables

During regression and classification analysis, the goal of a model is to understand how a set of variables affect the value of another. The variables that are used for the prediction are called independent variables and the value that is to be predicted dependent variable. The idea is that the value of the dependent variable in some way depends on the values of the independent variables and hence a prediction about the dependent variable is possible.

2.3.2 Training Data, Test Data and Overfitting

Rebala et al. (2019) explain the construction of training and test data to be an important step in data analytic studies. The general idea is to split your data set in two. One set (training set) is shown to the model and from the training set the model construct correlations between the independent variables and the dependant one. The second set (test set) is excluded from the training set in order to allow fresh data to be used to evaluate the model. This is done to reduce the risk of constructing a model which is too tuned to the specific data it has been trained on and hence would perform poorly when shown new sets of independent variables to make predictions on. This issue is called *Overfitting*.

2.3.3 Random Forest

Random Forest algorithms are a commonly used set of machine learning algorithms applied to regressions and classification problems. The fundamentals and implementation is described by Liaw and Wiener (2002). Rebala et al. (2019) list intuitively, ease of understanding and ability to backtrack what made a model make a particular prediction as benefits of a random forest algorithm. A random forest is based on the concepts of decisions trees and ensemble learning, both are discussed below.

A decision is explained by Rebala et al. (2019) as a structure of nodes and edges. The nodes are decisions based on the value of specific variables and the edges links nodes together. This constructs different paths and branches in a tree like structure and the last node in a branch is refereed to as a leaf node and represent a specific value which will be the output of the decision tree. Figure 2.3 shows a simple decision tree.

The example tree shown in Figure 2.3 is simple and can without any difficulty be constructed by hand. However once the number of variables and the amount of data increases computers and algorithms are the only viable option. Rebala et al. (2019) explain the general strategy of such an algorithm. The idea is to construct a tree where each node's decisions adds as much information to the overall model as possible. This can be done in multiple ways and are outlined in detail by Rebala et al. (2019).

A random forest is, as the name suggests, a number of random decision trees. A random forest constructs a number of decision trees based on subsets of the input data. The prediction by the model is then performed by averaging the outcome of each individual decision tree. This reduces *overfitting* by not relying on a single decision tree. The details or the algorithm can be found in Liaw and Wiener (2002) which also explain how a random forest algorithm can be used both for classification and regression problems.

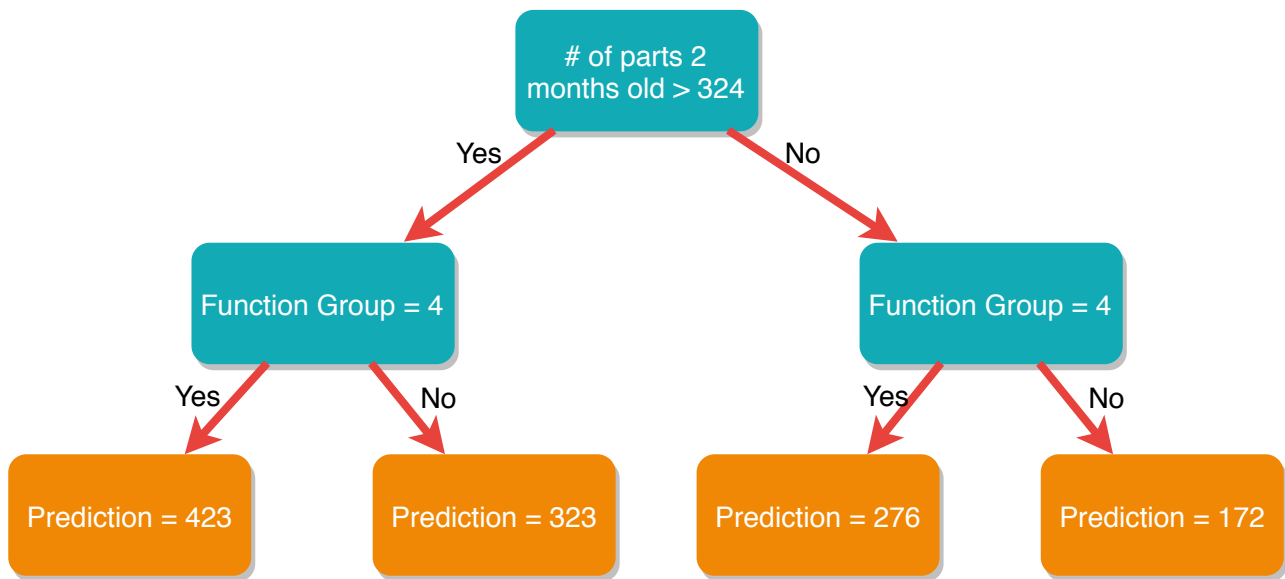


Figure 2.3: Visualization of how a decision tree can be used to solve regression problems. In a random forest algorithm multiple trees are created and the outcome of the model is the average of the chosen trees.

2.4 Conceptual Model

This section describes which of the topics covered in the literature review that are applied in the analysis section of this thesis. Three main topics, Spare Parts, Forecasting Methods and Big Data, are identified and together they form the conceptual model for this study. Finally the derived forecasting method to be evaluated in this study is presented.

The unique characteristics of spare parts are considered when forming the foundation for the analysis. The relation between products and spare parts life-cycle as well as the demand pattern for various types of parts are considered.

Time series and Big Data analytics forecasting are deemed necessary for the analysis of this thesis. Installed base data is a kind of Big Data can be evaluated using Big Data analytics. The performance of a Big Data method is to be compared against a time series method and as such time series forecasting is to be applied in the analysis.

Big Data poses different challenges compared to traditional intrinsic data and have to be evaluated with different tools. In this study installed base data is evaluated as a source Big Data. The understating of how Big Data differs from traditional data is important for further analysis.

Inspiration is drawn from Liu and Tang (2016) when designing the following reliability based method which is to be tested in this study. Neither of the data sets described in Section 3.4 contains reliability data for the various parts. However by combining the data sets in a similar way to Liu and Tang (2016) reliability data for case company spare part can be derived. The derived data is not data over historical sales but instead a kind of extrinsic data for which time series forecasting methods do not work. Instead the suggested method apply a machine learning, more specifically a random forest algorithm to predict demand based on installed base reliability

information. In order to evaluate this proposed method a Time Series forecasting method is applied to historic sales data. Finally the two methods are combined in a method which utilize both historical sales and installed base information. The forecasting accuracy of each of the three method is evaluated on the different stages of the spare part life cycle and in relation to the demand pattern of the specific parts.

3 | Methodology

This sections describes the study's research design as well as how the study was conducted. For an overview of the methodology, see Figure 3.1

3.1 Research design

Bryman and Bell (2003) describe the mixed methods research as the combination of quantitative and qualitative methods. This study was built on that principle and both quantitative data as well as qualitative interview data has been utilized. One concept highlighted by Bryman and Bell (2003) is *triangulation* in which for instance quantitative data was used to verify qualitative findings. During this study qualitative data has been used to determine what quantitative experiments were to be run and quantitative results were analysed with qualitative knowledge accessed through domain experts.

3.2 Workflow and Structure

This section describes the approach and how the work during the thesis was structured. The goal of the structure was to allow for ongoing testing and evaluation of forecasting methods while still ensuring a high level of validity and repeatability of findings. The study and thesis work was split into three main phases; 1) *Orientation and Data Collection*, 2) *Data Analysis and Evaluation* and 3) *Results and Findings*.

3.2.1 Phase 1: Orientation and Data Collection

The purpose of Phase 1 was to give the authors a broad understanding of the field of study as well as the case company situation and to perform data collection. This was achieved through literature studies as well as the conducting of interviews with domain experts. Interviewees were primarily Volvo Swedish and US managers, domain experts and data scientists. The data gathered was thus both quantitative and qualitative. Finally quantitative data was collected from internal Volvo sales and product databases. This process is outlined in detail in Section 3.4.

3.2.2 Phase 2: Data Analysis and Evaluation

The data analysis started with a literature study in order to find forecasting models applicable to installed base data. The methods were to be applied to the data and the outcome evaluated and validated. The purpose of this was to determine if and how well these methods might contribute to a greater understanding and approximation of demand patterns and increased forecast accuracy. The forecasting method selected for study is presented in Section 3.5.

The data analysis was conducted in the Python programming language and the models evaluated in this study were implemented with the help of readily and openly available data science Python packages. For instance *Scikit Learn* (2020), *Pandas* (2020) and CatBoost, Yandex (2020) were utilized. These packages include tools for data processing, implementations of commonly used machine learning algorithms and tools to evaluate the performance of models. The data analysis was supported by two data scientist experts employed by the case company. Their technical knowledge as well as their understanding of the data to be studied were key for the success of the study.

This phase of the study was iterative and structured around the process of plan, develop, test and evaluation of the next step to improve the proposed forecasting model. A phase began with that a potential improvement was identified. The model was modified to reflect the proposed improvement and the performance evaluated. The insight to perform a specific test came both from studying literature and in correspondence with case company domain experts. The performance of each model was discussed in relation to the specific case company situation and the data at hand and was cross-referenced with literature to provide explanations to why each model performed as it did.

There is a continuous dialog with Volvo Data Scientists in a number of informal meetings. This allowed the authors to ask questions about the data and discuss and validate findings with what they knew to be true from their day to day job at Volvo.

3.2.3 Phase 3: Results and Findings

In the last stage, the results and findings the analysis was further evaluated, and benchmarked against Volvo Cars existing forecasting methods. A summary of the result was produced, based on the findings from the analysis and literature. This was presented as a recommendation of how installed base information can and should be used to increase spare part demand forecast accuracy at the case company and in general.

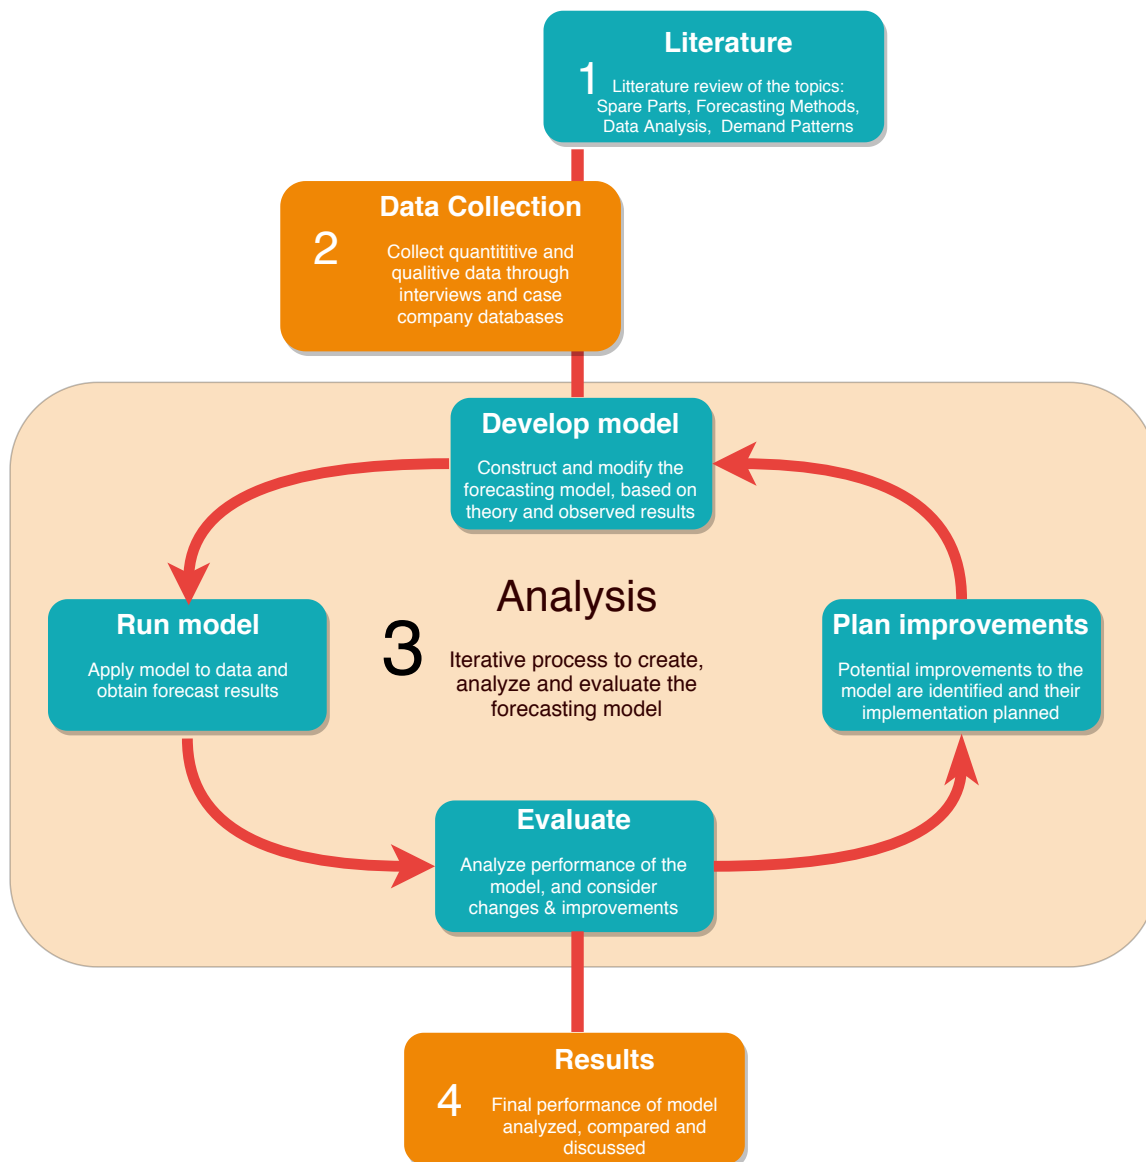


Figure 3.1: Conceptual overview of methodology and workflow

3.3 Choice of Accuracy Measures

Firstly, the three accuracy measures chosen for the analysis were *mean absolute deviation* (MAD), *mean square error* (MSE) and *standard mean absolute percentage error* (sMAPE). This was partly because MAD and MSE are easy to use and common in practice, while sMAPE is the KPI used by DP&O at Volvo. Further, MAD and MSE are both scale-dependent, meaning that errors measured by both are on the same scale as the data. Accordingly, these errors are the mean number of units with which the forecasts deviate from actual demand in the case of MAD and the square of said deviation in the case of MSE. sMAPE on the other hand is a scale independent percentage error, potentially allowing for comparisons external to Volvo.

For comparisons within the same data set, as performed in this thesis, the essential task of accuracy measures is to identify the model with the best performance. As such, the relative and ranked performance of each model is the key result. However, when comparing to other data sets and situations, the scale-dependent measures of MAD and MSE can be misleading. Since the quantity of actual and predicted units may vary greatly from case to case, using them for comparative purposes can lead to wrong conclusions being drawn. The percentage errors of sMAPE do however allow for comparisons external to both the data and specific situation of this thesis, although the exact parameters and causalities may be hard to replicate.

That being said, the differences between accuracy measures are significant and it is important to choose correct and relevant measures for each situation, especially in terms of the magnitude of the errors, to allow for fair and useful comparisons. For the purposes of this thesis and the analysis performed, MAD, MSE and sMAPE were considered to provide a sufficient basis for comparison between the evaluated models.

3.4 Data Collection

The data collection was conducted together with representatives from the case company. Installed base data, as well as product and sales data, was collected from a number of the case company's databases.

3.4.1 Quantitative Data

The first data set contains spare parts master data with descriptions of the parts name, volume, weights, product group, function group and a text description of the part. The product group refers to what general type of product the part was, for instance whether it was a spare part or an accessory. The function group was more detailed and describes where in the car the part was installed and what general function the part has.

The second set of data was historical invoice data. This data contains all historical sales of spare parts going back to 2007. This means at the time of this study, year 2020 the data contains 13 years of sales history. This data was determined to be accurate as it was based on actual sales of products and has been used for accounting purposes. An overview of the data in Data Set 2 is presented in Table 3.1.

Label	District	Part	Year Month	Quantity Sold	Country Code
Description	Market	Article number	Year and month	Number of parts replaced	Country
Example	5413	7846841530	200810	41	SE

Table 3.1: Describes the data in Data set 2. The example row shows that 41 units of part 7846841530 were sold in October 2008 on the Swedish market and that Sweden is in District 5413.

The third data set contains historical car sales and information of which market it was sold on. This set also include the number of units of each potential spare part that was installed when the car was produced. This data was determined to have high accuracy as it the basis for the Bill of Material for each car that is produced. An overview of the data set is shown in Table 3.2

Label	Country Code	Part	Year Month	Quantity car	Quantity part
Description	Country	Article number	Year and month	Number of cars sold	Number of parts installed
Example	SE	78468530	200810	1200	4800

Table 3.2: Describes the data of Data set 3. The example shows that in October 2008, 2000 cars containing article 7846841530 were sold and four were used in each car.

In order to curate the data and reduce irregularities data cleaning was conducted. First Data Set 3 was filtered in order to remove parts which have been replaced with another part. Only parts which had never been replaced were kept. This was done in order to keep only spare parts with a full life-cycle in the final data set and avoid parts which have been introduced as a replacement for another in the middle of the life-cycle. Secondly, Data Set 2 and Data Set 3 were combined to filter out parts which are not present in Data Set 1. This means only parts for which historical demand existed were included in the final data set. An overview of the three data sets is presented in Table 3.3.

Data Set	Source	Kind of Data	Time range
1	Volvo Internal	Spare Part Master Data	-
2	Volvo Internal	Historical Sales	2007 - 2020
3	Volvo Internal	Installed-base	2013 - 2020

Table 3.3: Overview of Data Sets

3.4.2 Qualitative Data

In order to gather information and knowledge from experts interviews were planned. However, this part of the study was cut short due to the COVID-19 pandemic and changed priorities for both the authors and Volvo Cars. Instead, only one interview was conducted, following the

semi-structured approach described by Bryman and Bell (2003). The questions asked to the interviewee can be seen in Appendix A.1.

The interview provided insight into how the logistics of Volvo US is organized. One of the key take-aways was that in the US spare part demand is most often not pre-planned. Instead the customer drives to the repair shop and expects the part to be readily available or at most a couple of days away. This differs compared to Europe, where most repairs are planned and the customer only goes to the shop once the shop was ready to perform the repair. A second point was the longer lead time to the US from CDC in Gothenburg, around two months by boat. The lead time combined with specific tax rules have lead to some vehicles getting their final configuration once they reach the US. This is achieved by ordering more stripped down cars and install customer specific accessories on cars already in US storage.

3.5 Data Pre-processing

The three data sets included in this study lacked direct information of spare part reliability and expected life time. This thesis evaluated the use of installed base data and a reliability-based method. Therefore reliability data was derived by combining the three data sets. Data Set 2, containing historical sales of spare parts is combined with Data Set 3 which contained the installed base and expected future installed base. The combinations were done by deriving the age distribution of installed spare parts from the installed base data at specific dates and considering what demand the particular age distribution produced in the following planning period. Demand for a particular spare part was deemed to represent a failure and hence the life time or reliability of that particular spare part. For each part the first date of interest was the first time the part is installed in a car. Beginning with this date and for each of the next 24 months the part number, date and age distribution of the specific part was recorded. Such a record is referred to as an *Observation*. The date recorded in an observation represented the date for which the age distribution of the installed base was recorded. This is henceforth referred to as *Check Date*. Thus, for each part the Check Date always lies between the first date the part was installed in a car and 24 months forward.

The first observation of a part only included parts with an age of 0 months. The second observation only parts of 0 and 1 months of age, and so forth. For each part 24 observations are generated. Together these 24 observations describe the age distribution of the part for each of the 24 months since the parts first introduction. The population distribution of the installed base for a given date and part are represented by $x_{311} - x_{3...n}$ in Table 3.4. The demand for the following planning period after the specific date was the dependant variable and was represented by y in Table 3.4. This means that the same part is represented multiple times in the initial data set but with different *Check Dates* (the date from which the forecast was to be performed). The age distribution of the installed base as well as the spare part features were the independent variables and the dependent variable was the demand.

Obs	Check Date	Part	qty month old				Sales next month
			1	2	...	n	
1	x_{11}	x_{21}	x_{311}	x_{321}	$x_{3...1}$	x_{3n1}	y_1
2	x_{12}	x_{22}	x_{312}	x_{322}	$x_{3...2}$	x_{3n2}	y_2
...	$x_{1...}$	$x_{2...}$	$x_{31...}$	$x_{32...}$	$x_{3.....}$	$x_{3n...}$	$y...$
n	x_{1n}	x_{2n}	x_{31n}	x_{32n}	$x_{3...n}$	x_{3nn}	y_n
Example	Feb 1 2020	1681531	432	153	543	287

Table 3.4: Explanation of the initial data set, x represents independent variables and y is the dependent variable, demand.

3.6 Validity and Reliability

The overall validity of the data collected was deemed high. The data was collected from internal systems and the data has been used for record keeping and forecasting. Volvo Data Scientists confirm the validity of the data and that it corresponded well to reality. When analyzing the data no prominent inconsistencies were identified. For instance the cleaned data do not contained any negative values for demand and data was present for the whole range that the data sets covered. In addition, findings and results were cross-referenced with previous research and with domain experts at Chalmers and at Volvo Cars. The consensus between this study's results and literature indicates the data and methods holds high validity.

3.7 Overview of Analysis

The analysis was structured around the development of a reliability-based forecasting model, utilizing installed base data. This proposed model was evaluated against a time series forecasting method. The methods were compared both on their overall accuracy and on how well each method performed on spare parts with different demand patterns. An overview of the data included, the methods implemented and their evaluation is shown in Figure 3.2. Each model is explained in detail in the following Analysis Chapter.

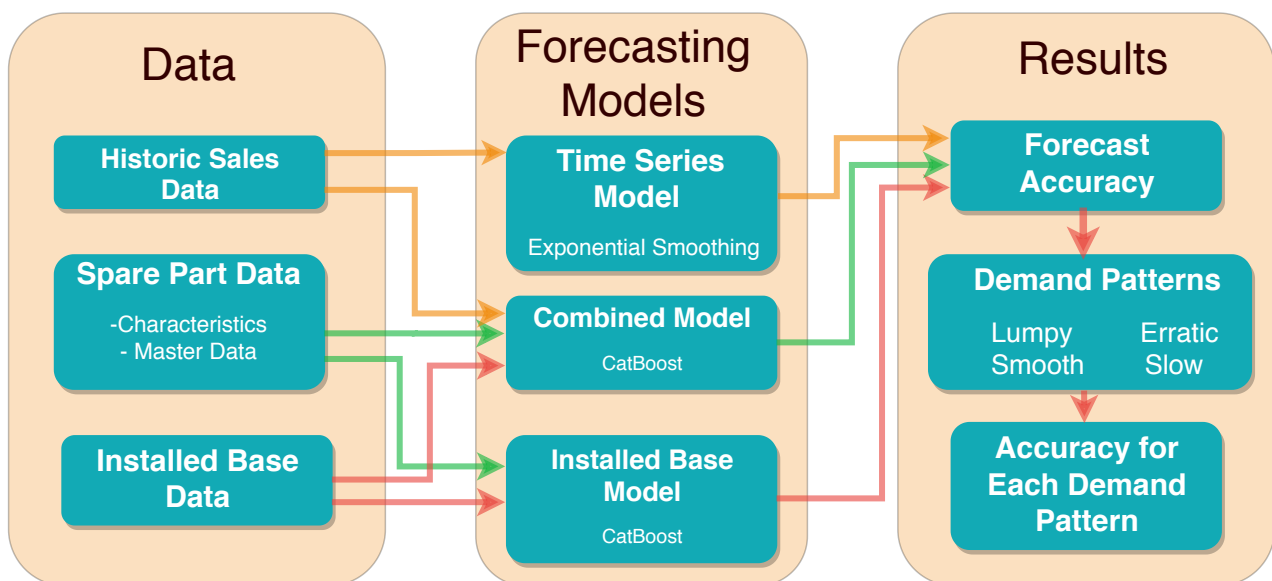


Figure 3.2: Overview of data sources, models and results involved in the Analysis

4 | Analysis

In this chapter, the analysis is presented from the planning of the predictive data study through construction and implementation of the model. The different stages of forecasting during the launch phase of a new spare part are explained. Further, forecasting is performed using an exponential smoothing model representing traditional time series forecasting and a new extrinsic machine learning algorithm. These two models are also combined in a third model, and the results of these three forecasts are compared. The included features are analysed and compared on their impact on model outcome. Lastly, the parts included in the data are split according to their demand patterns, and the results of each model evaluated for each pattern.

4.1 Planning Horizons

The aim of this study is to focus on predicting spare part demand in the launch phase of new spare parts, 0-24 months from its market introduction. This period is divided into three sub-phases. The forecast result varies in these phases, since the decisions made based on the forecasts vary and pose different challenges for Volvo. In discussion with domain experts at Volvo the following horizons and corresponding target outputs (summarized in Table 4.1) were established:

- During Stage 1, which is before any sale of the part has occurred, the decision Volvo Cars has to make is whether or not the part should be stocked at all. Volvo wants to stock the part if it is likely to sell in any of the next 6 months, meaning if any demand exists at all. What is to be predicted is thus whether or not the part will sell in the next 6 months.
- In Stage 2, once the part has seen initial sales and until the part have seen sales in three periods, the total installed base is still small and total sales low for most parts. The planning horizon is set to three months and what is to be predicted is the total demand for the next three months. This stage lasts until a part has seen sales in a total of three periods, meaning demand has existed in three not necessarily connected months.
- The last stage, Stage 3, consists of the months remaining after the end of stage two until 24 months after product launch. During this stage, the demand for the next month is to be forecast. Since the product is slightly more established sales can likely be more accurately predicted.

Stage	Planning Frequency	Planning Horizon	Problem	What to predict
1	1 month	6 months	Classification	Sale in next 6 months
2	1 month	3 months	Regression	Demand next 3 months
3	1 month	1 month	Regression	Demand next 1 month

Table 4.1: Overview of the three planning stages introduced in section 4.1

4.2 Forecasting Models

In the following section, the three models used during the analysis are presented. The first, a machine learning model utilizing extrinsic installed base data (referred to as IB from here on). The second is a simpler time series-based model using exponential smoothing (called ES), representing the current method used by Volvo DP & O. This is to facilitate a comparison between the more complex algorithm and the simpler one, evaluating if forecasting accuracy is improved by the added complexity. For Stage 1, only the IB model is run since no historical sales data exists for the ES model to use. The IB model is instead compared with an estimation of Volvo’s actual results. As mentioned above, Stage 1 contains a classification problem, while Stage 2 and 3 consist of regression analysis. Both the IB and ES models are run in the later stages, as well as a model combining both these methods, called the Combined Model (C).

4.2.1 The Installed Base Model

The Installed Base model is based on a so called random forest algorithm. Liaw and Wiener (2002) explain how a such an algorithm can be applied as both a classification and a regression method. The details of the random forest algorithm are presented in section 2.3.3. In this study an optimized kind of random forest algorithm is applied to the data set described in table 3.4. The specific implementation is described below.

4.2.1.1 Implementation

In this study the random forest algorithm was implemented using the CatBoost, Yandex (2020) python package. CatBoost is based on the concept of a random forest, but is more advanced and applies various optimizations to improve model accuracy and reduce the time needed to train the model. The details of the CatBoost algorithm can be found in Dorogush et al. (2018). Furthermore, the algorithm has built-in support for categorical data such as product groups which further simplifies the implementation. Appendix A.2 shows the Python implementation of the CatBoost algorithm as a classifier. The implementation for Stage 2 and 3 are similar but instead of a classifier a regression model is used. The independent variable in this case is demand for the next month n corresponding to the planning horizon, see Section 4.1.

4.2.1.2 Parameters

The CatBoost algorithms takes a number of parameters and a data set as it inputs. It is possible to omit any parameter and if so CatBoost falls back to a default value. The data set should contain the features or independent variables the algorithm should use to predict the dependent

variable. The independent variables can be both a continuous numeric value or categorical discrete values. The dependent variable is, in the case of a classification problem, a discrete category and in the case of a regression problem a numeric value. The CatBoost documentation (Yandex 2020) explains the parameters likely to impact the accuracy of the model according to below:

- Number of Trees, this refers to the number of trees or estimators that are generated while the algorithm is learning from the data. A high number of trees often results in higher accuracy but increases the risk of overfitting as well as increasing the computational power needed and hence also increases the time it takes to train the algorithm.
- Tree Depth is the maximum depth a decision tree can have in the algorithm. For illustration, a tree of depth 2 is shown in Figure 2.3.
- Learning Rate defines how aggressively the algorithm will learn from each iteration of the algorithm. A lower value increase training time but tend to yield more accurate results. A high value increases the risk of overfitting. This parameter needs to be tuned in order to reach good results within a reasonable time with regards to available computing resources.

4.2.1.3 Cross Validation

The model data is split into a training and a test data set. Though the split is done randomly, there is a chance that the specific split used when training the algorithm is not representative of a new independent data set (Hastie et al. 2009). A common way to make the model more accurately predict new data is to perform cross validation. K-fold Cross validation is one method described in Hastie et al. (2009). It works by splitting data in K parts and uses K-1 parts to train a model. The remaining set is used to evaluate the model. This is repeated K times with each of the sets functioning as the evaluation set. The performance for the model is deemed to be the average of the performance from the K iterations.

In Figure 4.1 below, the results of the 5-fold Cross Validation of the Stage 1 classifier. The 5 lines represent the performance of each of 5 folds. The y-axis represents the objective to be minimized during training, while the x-axis represents the number of iterations. The number of iterations corresponds to the number of trees fitted in the model. It is desirable to minimize the variation between the lines, since low variation indicates that the model should perform similar on data it has not previously been trained on.




Figure 4.1: Visualisation of 5-fold Cross Validation. The x-axis shows the number of iterations and the y-axis the model error (lower is better). All five lines move downwards and the variation between them is low. This indicates that the model is not too fitted to the test data and that the number of iterations is not too high.

4.2.1.4 Parameter Tuning

The performance of forecasting models is dependant on the values of the various parameters defined for the model. In order to optimize the model, parameter values have to be considered and tuned. One way to help find the optimal parameters for a given model and data set is to perform a so called Grid Search. In a Grid Search a set of parameters and related values are defined and a forecast model is trained with each possible combination of the parameters. The idea is that the best of these models should be close to representing the optimal model. The implementation of GridSearch for the Stage 1 classifier can be seen in Appendix A.3. The implementation utilizes 5-fold Cross Validation to determine the performance of each model.

4.2.2 The Exponential Smoothing Model

An exponential smoothing model (ES) is chosen to represent general time series forecasting methods since it is a commonly used approach in practice when dealing with intermittent demand (Syntetos and Boylan 2005). It also represents a decent estimation of the methods currently used by Volvo. The main difference between the model implemented here and the one used by Volvo is that the ES model used here applied a dynamic smoothing value, α , whereas Volvo's is static. Volvo's model is tailored to their own needs, and the dynamic smoothing

value is an attempt to compensate for the general nature of the ES model. The implementation and results of this has been verified with domain experts at Volvo, and is considered a good representative for the purposes of this study.

The ES model was implemented in Python, see Appendix A.4 for details. The basis of the algorithm is described in Equation 2.2, with the addition that the algorithm searches for and finds the optimal α or smoothing value for the forecast. Figure 4.2 shows the forecast and actual sales for a specific part.

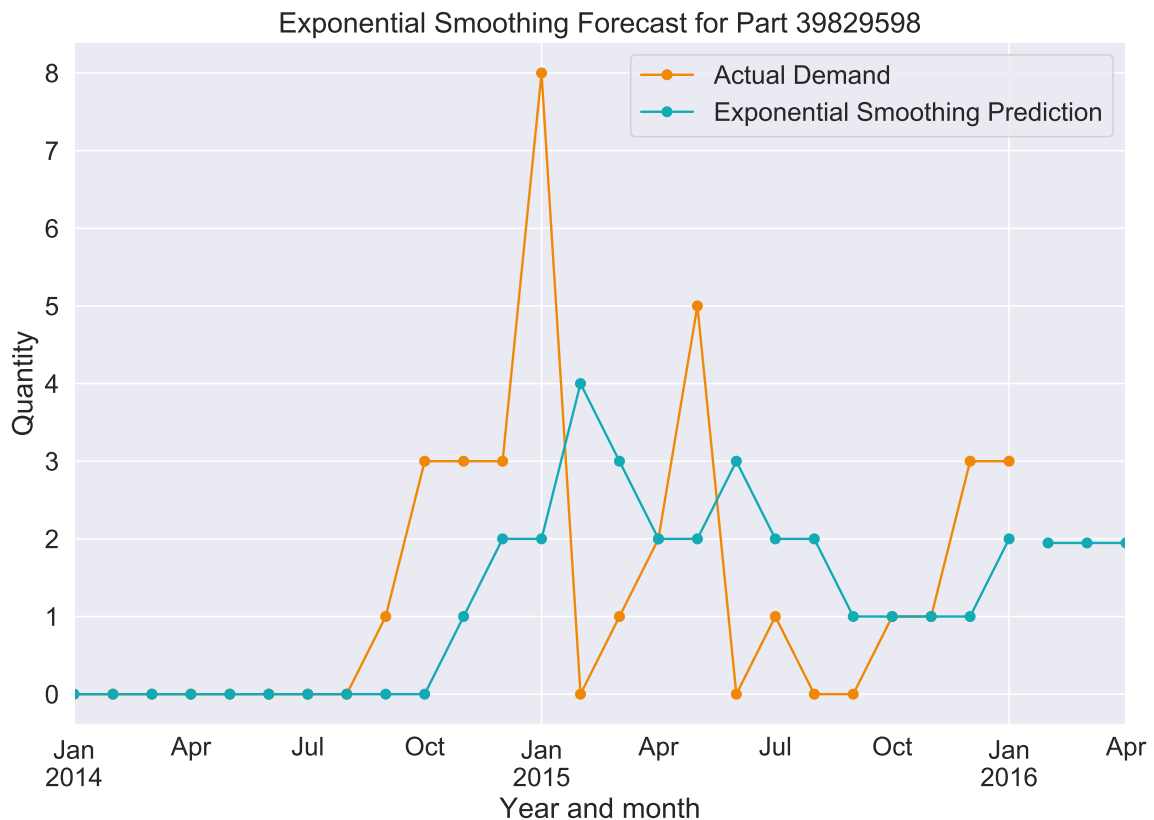


Figure 4.2: Exponential smoothing forecast for part 39829598 during the initial phase (24 months after market introduction). The graph shows how an exponential smoothing model struggles to predict sudden spikes and drops in demand.

4.2.3 The Combined Model

The above described Installed Base model is trained solely on extrinsic data, and is not aware of any previous sales of a product, only the age distribution of the installed base and the spare part features. On the other hand the exponential smoothing model is only fitted to intrinsic previous sales data for each product. In the interest of reaching a higher forecast accuracy these two models are combined to construct a model which take both extrinsic and intrinsic at into account. This is done by to Data Set 1, shown in table 3.4 add an additional column that contains the exponential smoothing forecast for the particular part and month. A new CatBoost regression model is then trained with this new data set.

4.3 Forecast Accuracy

In this section overall forecast accuracy of the three models is presented and analyzed. See table 4.2 for an overview of the performance of the three models in the three different stages. The Installed Base model (IB) is the machine learning model considering only installed base data. The Exponential Smoothing model (ES) is the time series method and considers only historical sales data. The Combined model (C) is a combination of the two, considering both installed base data and historical sales data. A more detailed analysis and breakdown of the results is presented below.

4.3.1 Stage 1

The aim of the Stage 1 analysis is a monthly prediction of whether the part will sell within the six month planning horizon. Due to this, it is a binary classification problem. For each observation there are two possible outcomes: 'YES' or 'NO'. Each observation can be correctly predicted or not. This results in four possible outcomes and can be summarized in a so called confusion matrix (Rebala et al. 2019). The confusion matrix of the Stage 1 classification is presented in Figure 4.3. The y-axis represents the actual value of the observations and the x-axis represents the predicted label. The four boxes show the number and percentage of observations classified correctly and how many were not. The rows of the matrix together represent 100% of actual 'YES' and actual 'NO'. Stage 1 covers the time before the part has seen any sales. Exponential smoothing is a time series forecasting method and uses historical, intrinsic data for its predictions. Hence it is not possible to apply exponential smoothing on Stage 1 data. Further, Volvo does not currently have a systematic way of replacing their time-series method for Stage 1 forecasts, and only decide to stock in the US in specific cases. This relies on the expertise and know-how of the demand planners involved, but in most cases the risk of an item not being in stock at the time of first sale is consciously taken. Due to this, there are no forecasts or predictions to compare the results of the classification with. Instead, Volvo's actual results for the period 2018/04/01 - 2020/02/01 have been used. This limitation in time frame is due to the lack of available data from earlier periods, but is still deemed to serve as a reasonable estimation of Volvo's current performance. The results are presented in Figure 4.4, where 'In Stock' represents Volvo's choice to stock a part or not within six months before the first sale.

According to the results in Figure 4.3 and Figure 4.4, the Installed Base model seems to provide a way of predicting if and when first demand for a part will occur. The model outperforms Volvo's comparable results clearly, in both the desired outcomes of 'YES/YES' and 'NO/NO'. The model manages to correctly predict whether a part will sell or not in the vast majority of cases, utilizing only the available installed base data. It should however be reiterated that the Installed Base model does not take any qualitative or strategic factors into account, and Volvo's figures are not the result of a systematic qualitative approach. A direct comparison is thus difficult. Still, the results seem to indicate clear potential in using installed base data and a machine learning model.

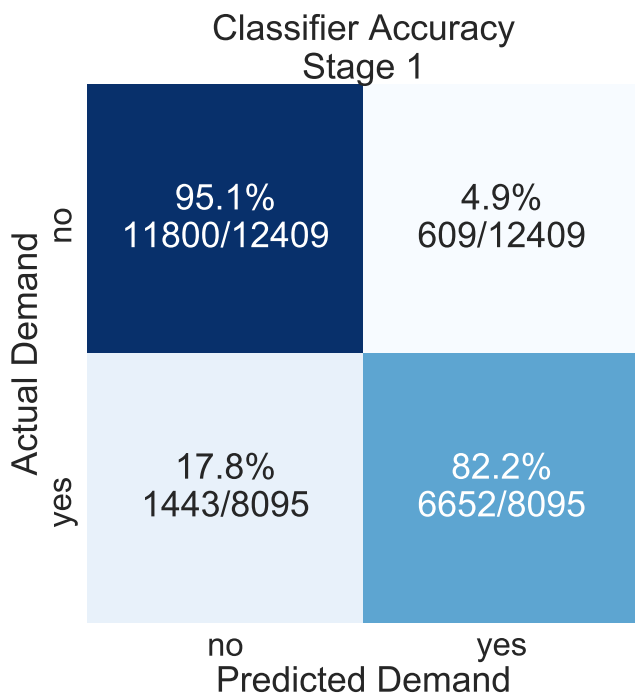


Figure 4.3: Confusion Matrix of the classifier performance for Stage 1.

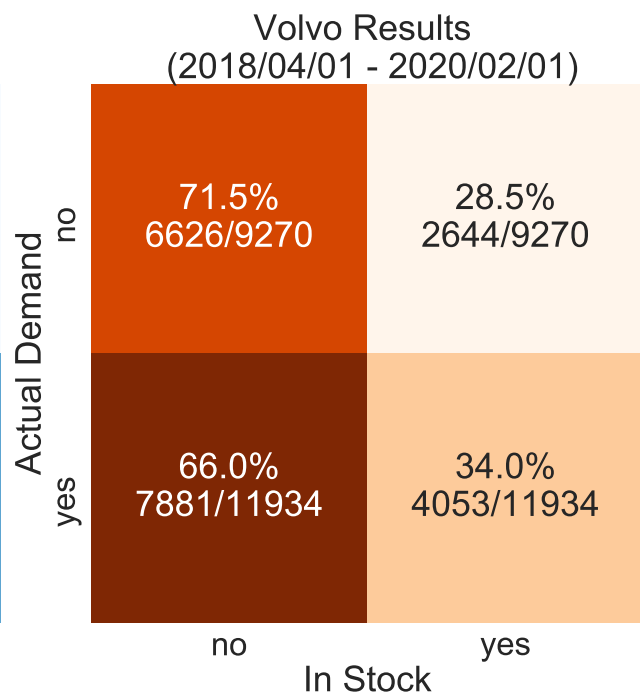


Figure 4.4: Confusion Matrix of Volvo's actual Stage 1 results.

4.3.2 Stage 2

Stage 2 include the the time between the first sale and the third period which have sales of a particular product. This means historical sales exists and hence both the Exponential Smoothing and Combined model can be applied. The goal is to predict the sales of the next three months, with forecasts performed every month. As shown in table 4.2, the Installed Base (IB) model has the worst performance according to all three accuracy measures (MAD, MSE and sMAPE), followed by the Exponential Smoothing (ES). The Combined (C) model performs the best. It is to be expected that the Combined model performs at least as good as the Exponential Smoothing model since one of its features is the outcome of the Exponential Smoothing model. While the inclusion of actual sales data is necessary for any time series method, the performance of the Combined model is a good indication of part features and installed base information contributing further to improved forecast accuracy. A more detailed overview of each feature's impact on the model output is presented in Section 4.4.

4.3.3 Stage 3

Similar to Stage 2, the Stage 3 analysis includes only observations for products with actual sales data, and all three models can thus be applied even here. The goal is to predict the sales for the following month. Table 4.2 shows that the Installed Base (IB) method is the worst performing model for Stage 3 and Combined the best. Again, this implies that the inclusion of parts features and installed base data into a time series method increases forecasting accuracy. It should be noted that while IB improved according to all accuracy measures compared to Stage 2, ES performed worse according to both MAD and sMAPE. C on the other hand improved according to MSE, while retaining similar MAD performance and performing worse according to sMAPE.

	MAD			MSE			sMAPE(%)		
	IB	ES	C	IB	ES	C	IB	ES	C
Stage 1	See 4.3	-	-	See 4.3	-	-	See 4.3	-	-
Stage 2	8	2	1	26097	5310	1666	116	45	44
Stage 3	7	4	1	1242	4550	87	110	58	54

Table 4.2: Overview of forecast accuracy for the three models; Installed Base (IB), Exponential Smoothing (ES) and Combined (C)

4.4 Feature Importance

From a trained CatBoost model it is possible to collect the importance of each feature that was used as an input for the model. This increases the understanding of the model and aids in visualising which variables have the most impact on the predicted model output. The importance of each feature for Stage 1 is summarized in Table 4.3 and for Stage 2 and 3 in table 4.4. In the tables the top 10 most importance features are presented. For a complete list of feature importance, see Appendix A.5. The features are ranked and the feature with the highest importance is ranked as 1. The number next to the feature name represents a normalized value of the importance of the specific feature. The importance of all features adds to 100.

In Stage 1 the only applicable model is the Installed Base model. From Table 4.3 it can be seen that what kind of product and its features are the most importance for the outcome of the forecasting result. The features *function_group*, *description*, *weight* and *volume_cm3* are all inherent features of the part for which future demand is forecasted. However, taken together the existing installed base age features contribute 40/100 of the normalized impact and hence has a greater combined importance than any other single feature.

Stage 1	
Top N feature	IB
1	group (18)
2	description (15)
3	weight_g (13)
4	volume_cm3 (12)
5	parts_0m_old (7)
6	parts_1m_old (4)
7	parts_2m_old (3)
8	parts_3m_old (3)
9	parts_4m_old (3)
10	parts_5m_old (3)

Table 4.3: Ranking of most influential features for Stage 1 and the installed base model. The number next to each feature is a normalized value of the influence of the feature. The importance of all features adds to 100.

Table 4.4 explains why the addition of the result of the Exponential Smoothing forecast has such an impact on the forecast accuracy in the Combined model. In both Stage 2 and 3 for

the Installed Base model the part features are the greatest contributing factors to the forecast outcome. This changes when the Exponential Smoothing forecast result is added in the Combined model. In the Combined model the Exponential Smoothing forecast contributes 67/100 in Stage 2 and 74/100 of the feature importance. In the Combined model all the following features are installed base data. The forecast based on historical sales overtakes the parts features as the most important forecasting factor. This indicates that what group a part belongs to is an indicator of future sales but once actual historical sales are available to the model it becomes a much better indicator of demand and hence the part features become less influential. In the Combined model, installed base information contributes around 35/100 and 25/100. This contribution is enough to influence the model and likely explains why the Combined model outshines the Exponential Smoothing model.

Importance	Stage 2		Stage 3	
	IB	C	IB	C
1	group (28)	ES_forecast (67)	description (18)	ES_forecast (74)
2	1m_old (26)	7m_old (7)	group (15)	6m_old (6)
3	0m_old (18)	1m_old (4)	0m_old (12)	7m_old (3)
4	2m_old (10)	6m_old (3)	2m_old (11)	8m_old (3)
5	4m_old (9)	8m_old (3)	1m_old (10)	1m_old (3)
6	description (4)	2m_old (3)	7m_old (7)	0m_old (2)
7	3m_old (2)	5m_old (2)	4m_old (5)	3m_old (2)
8	5m_old (1)	0m_old (2)	3m_old (4)	2m_old (2)
9	6m_old (1)	4m_old (2)	8m_old (3)	4m_old (2)
10	7m_old (0)	3m_old (1)	5m_old (3)	weight_g (1)

Table 4.4: Ranking of most influential features for Stage 2 and Stage 3 and the installed base and combined model. The importance of all features adds to 100.

4.5 Demand Pattern Categorization

To determine the demand patterns of the various parts, the categorization framework presented Boylan et al. (2008) and discussed in Section 2.1.3 is used. This is done based on the mean inter-demand periods and variability of demand throughout the two year period. The sales data in Data Set 2, see Table 3.1, is used for the categorization of parts according to their demand patterns. Based on this data, the p and CV^2 values are calculated for each part in order to categorize their demand patterns. The values chosen as cut-offs between 'High' and 'Low' are $CV^2 = 0.49$ and $p = 1.32$. The corresponding values for each part is calculated using Python, and each part is then assigned one of following demand patterns, based on Boylan et al. (2008); smooth, intermittent, lumpy or erratic. The period for which the demand pattern is calculated are the two years after the part is initially installed in a car. The parts used for the evaluation is limited to parts which have the full initial phase of their life-cycle within the range of the data available. This means that only parts whose first install happens after the first date of available sales data and whose two years since first install are within the range of available data are included. A breakdown of the demand pattern categorization of included parts are presented in Table 4.5.

	CV^2	p	# Parts	% Parts
Lumpy	≥ 0.49	≥ 1.32	544	13.65
Smooth	< 0.49	< 1.32	291	7.30
Intermittent	< 0.49	≥ 1.32	2,996	75.16
Erratic	≥ 0.49	< 1.32	155	3.89
Total			3,986	100%

Table 4.5: Distribution of parts according to their Demand Pattern Categorization.

Boylan et al. (2008) suggest statistical evaluation of the performance of various forecasting methods, and adjusting the cut-off values according to the superior performance of the evaluated methods. For the purposes of this thesis and analysis, the cut-off values from their comparative work on (Croston 1972) and the (Syntetos and Boylan 2005) methods are used to categorize parts in the data according to their demand patterns. As can be seen in the Table 4.5 above, the vast majority of parts sold show intermittent demand patterns, using the cut-off values provided. This aligns with the expressed experience from domain experts at Volvo, who have repeatedly stated that most parts have long sequences of zero demand periods, interspersed with some non-zero ones. Lumpy and erratic are versions of intermittent demand, and serve to further illustrate the various distributions of spare part demand. Smooth parts are expected to be fewer in number than the three versions of intermittent demand, for spare parts in general and during the initial phase especially. Thus, the cut-off values chosen provide a solid foundation for the evaluation and comparison of the models' performance even though they are not specifically calculated for the data at hand.

The classifications are used to evaluate the performance of the various models in Stage 2 and 3, to determine if the developed model is more or less suitable for parts with certain demand patterns. This aims to answer Research Question 2, as well as provide the thesis with a degree of practical usefulness and transferability.

4.6 Model Performance per Demand Pattern

The performance of the three models with regards to the demand pattern of the forecasted parts are presented below. This analysis allows for further evaluation of the forecast accuracy and applicability of the different forecasting methods. Further it ties back to existing literature presented in the theoretical background of this thesis which states that time series methods struggle to predict demand for parts with lumpy, erratic and shifting demand patterns. The forecast accuracy measures used are explained in section 2.2.4.

The forecast accuracy of the methods for the different demand patterns are similar between Stage 2 and 3. The comparison shows that the CatBoost model using purely installed base data (IB) performs the worst on most demand patterns according to all three accuracy measures. The exponential smoothing method (ES) performs better, while the method with best performance in the majority of cases is the combined Installed Base/Exponential Smoothing method (C). These results will be discussed for each of the four demand patterns below.

	MAD			MSE			sMAPE (%)		
	IB	ES	C	IB	ES	C	IB	ES	C
Lumpy	12	3	1	5174	303	11	123	51	42
Smooth	72	2	2	300,050	25	45	110	21	17
Intermittent	3	1	1	304	2	1	120	53	54
Erratic	123	67	24	1,132,626	332,529	48,401	111	22	15

Table 4.6: Forecast accuracy in Stage 2 for each method and part with different demand patterns. Installed Base (IB), Exponential Smoothing (ES), Combined (C).

	MAD			MSE			sMAPE(%)		
	IB	ES	C	IB	ES	C	IB	ES	C
Lumpy	8	9	2	1,027	11,817	119	120	82	74
Smooth	13	6	3	1,090	500	29	87	45	38
Intermittent	3	1	1	132	4	2	118	68	67
Erratic	18	18	4	6,294	32,819	1,146	89	52	39

Table 4.7: Forecast accuracy in Stage 3 for each method and part with different demand patterns. Installed Base (IB), Exponential Smoothing (ES), Combined (C).

4.6.1 Lumpy

Parts with lumpy demand have high variability of demand and a high mean inter-demand period. Van der Auweraer et al. (2019) list this demand pattern as especially challenging for time series methods. This is confirmed by the result of this study. For lumpy demand the Combined (C) model performs better according to all accuracy measures when compared to the Exponential Smoothing (ES) model. This clearly shows the possibility of improving forecast accuracy by including installed base data when forecasting spare parts with lumpy demand patterns.

4.6.2 Smooth

Smooth demand means that the part has both low variability of demand and a low mean inter-demand period. According to the literature, this kind of demand should be a relatively good fit for time series forecasting methods, especially when compared to the other intermittent types. This is confirmed by this study's results as the Exponential Smoothing model (ES) has comparatively high accuracy for products with smooth demand. The results show that depending on the situation and the accuracy measure applied it can be beneficial to include installed base data. For instance, the combined method shows better results for all measures in Stage 3 but mixed results in Stage 2. The improvement is not as significant as it was for parts with lumpy demand patterns.

4.6.3 Intermittent

Intermittent demand parts have low variation in demand but a high mean demand interval period. According to Syntetos and Boylan (2005) time series methods often perform reasonably well in this kind of data but more suitable methods exist. This is confirmed by these results. The ES model performs well with regards to all accuracy measures and in both Stage 2 and 3. The combined model perform on pair with the ES model and hence there is no real benefit of including installed base data for parts what show an intermittent demand pattern.

4.6.4 Erratic

Erratic demand is characterized by high variation in demand and a low mean intermittent demand period. Turrini and Meissner (2019) list erratic demand nature of spare part as troubling for inventory management. The issue is exemplified by the results in this study. The erratic demand pattern show a high MAD and MSE for all methods. The results show that there are improvements to be gained by using a Combined model of historical sales and installed base data. The Combined model outperforms the ES model with regards to all accuracy measures.

5 | Discussion

In this chapter, the analysis and results thereof will be discussed. It is initially divided into a discussion of the performance results in each stage, based mainly on the models and demand patterns. The discussion then goes into the impact of factors included and excluded in the forecasting models and the importance of chosen accuracy measures for performance evaluation. The discussion concludes with demand pattern categorization and its imposition on the analysed data, before concluding and summarizing the thesis's findings and their relevance for Volvo Cars specifically and the field of initial phase forecasting generally.

5.1 The Implications of Model Performance

Below the results found in Stage 1-3 and their implications are discussed and put into perspective by an analysis similar to Boylan et al. (2008). For each demand pattern the relative performance of the combined model compared to the Exponential Smoothing model is visualized in Figure 5.1 and 5.2. This section focuses on the forecast results before the Demand Pattern Categorization was performed, and attempts to highlight the reasons and potential impact of the findings.

5.1.1 Stage 1

According to the results of the Stage 1 analysis presented in Figure 4.3 and 4.4, where the data includes no actual sales, there are clear benefits to using installed base data, although this does come with some caveats. Firstly, in this stage the Installed Base model is used to classify Sale or No Sale within the coming six months. As can be seen in the two confusion matrices, the model outperforms Volvo's current manual performance both in correctly predicting instances of no sales and instances with sales. The practical implementation of a Predicted Yes is that the part will be in stock, while a Predicted No means the item will likely not be stocked. This means that these two outcomes are of significant importance, since the aim of demand planning is to have demanded parts in stock while keeping unnecessary inventory at a minimum. The Installed Base model's 95,1% correct prediction of No Sale means potential for massive savings in reduced inventories, while the 82,2% hit rate on Sale potentially avoids significant costs for stock-out and lowered customer service due to poor availability of parts. Combined with the estimates of Volvo's own performance, the analysis shows vast potential in these areas. This conclusion is again strengthened by Dekker et al. (2013), who show that the incorporation of installed base data into forecasting can help reduce inventory costs, while balancing the often contradictory risks of stock-out and obsolescence.

In addition, the results of Volvo's current approach show a much higher number of 'Actual No/Predicted Yes' and 'Actual Yes/Predicted No' which are the undesirable outcomes in this scenario. If these estimates correspond well with reality, this means that Volvo's current approach leads to a lot of stock outs while simultaneously holding unnecessary inventory. A worst case scenario, of sorts. However, in practice, Volvo likely anticipates several of these cases and accepts

the risk of back orders or overstocking. Parameters such as the volume, value and criticality of the parts are considered, and some of the undesired results in the estimation may in reality have been expected or at least accepted. It should be noted that the results for Volvo's approach are estimates, and may not completely represent reality. However, it does give a clear indication that the room for improvement by utilizing installed base data in Stage 1 is substantial. It should however be mentioned that the planning horizon used in Stage 1 is 6 months, both for the classifier and the estimation of Volvo's results. A shorter planning horizon would have made the predictions more sensitive to outliers and may have impacted the accuracy of both negatively.

Lastly, one might add that since the demand pattern for most spare parts in the data set contains many zero demand periods, simply predicting No Sales for every part in every period would yield a decent result on its own. However, as the results show, such a simplistic approach can clearly be outperformed by the Installed Base classifier.

5.1.2 Stage 2

The Installed Base (IB) model suffers heavily in comparison to the other two models. In contrast to Stage 1, where the model answered a binary classification problem on a 6 month planning horizon, Stage 2 requires prediction of the demand for the coming three months based on regression. This penalizes IB since it solely utilizes installed base data, meaning cars and parts already on the market. In combination with the findings from Stage 1, this might imply that the model is useful for predicting if a part will sell within a certain time frame, but not the actual number of said parts that demanded.

Historical sales data is necessary for the Exponential Smoothing model. Since the ES model used in this analysis has been confirmed as a decent estimation of Volvo's actual model, it is not surprising that it performs decently in this stage and outclasses the IB model. However, the most interesting results from the Stage 2 analysis is that the Combined model (C) performs significantly better than the pure ES model, especially when measured by MSE. The impact on accuracy measures are discussed below and visualized in Figure 5.1, but in short these results clearly signify the potential value of combining installed base data with more traditional time series based forecasting. It also shows that time series methods provide insight even in situations with more extreme demand patterns, where they normally struggle to accurately predict demand especially if combined with Installed Base data. The improved performance achieved by combining the ES and IB model shows that this weakness can be mitigated by including installed base data into the forecasting model. These insights, together with the work from authors such as Dekker et al. (2013) and Van der Auweraer et al. (2019) highlight the importance of effective collection and utilization of installed base data.




Figure 5.1: The relative performance of the Combined model compared to the Exponential Smoothing model for each demand pattern in Stage 2.

5.1.3 Stage 3

During the Stage 3 analysis, the Installed Base model improved its performance compared to Stage 2. This is especially clear according to MSE, where a significant improvement from 26,097 to 1,242 is achieved. MAD and sMAPE also show improvements, although they are minor ones. This is likely due to Stage 3 comprising the end of the initial phase, meaning the period after the parts have already had three non-zero demand periods. This corresponds with a larger installed base for the IB model to utilize. According to MSE, the IB model even outperforms the ES model, which is remarkable since the increased amount of available sales data should benefit the time series model. MAD and sMAPE complicate the picture, since ES clearly outperforms IB when they are considered. The likely conclusion here is that a time series model is still preferable to a purely installed base one, but the nature of MSE as an accuracy measure penalizes certain errors more heavily than others. Given how the ES method is an estimate of Volvo's current practices, it is especially relevant to compare its performance with the Combined model. The relative improvements of C compared to ES are thus presented in Figure 5.2.

Best in class for Stage 3 is the Combined model. It outperforms the other two with significant margin according to all chosen accuracy measures, and only sMAPE shows a proximity in performance with ES. This reinforces the conclusions drawn from the literature review and Stage 2, where the combination of a time series model and installed base model proved fruitful with similar comparisons drawn between the two. The further increases in available historical sales data and the larger size of the installed base provides the Combined model with more data than the other two, and it delivers accordingly.




Figure 5.2: The relative performance of the Combined model compared to the Exponential Smoothing model for each demand pattern in Stage 3.

It is however interesting to note that the sMAPE results are worse in this stage than in Stage 2, despite the increased amount of data points. This can likely be partly traced to the shorter planning horizon, down to only one month from three. This tendency is only visible in the sMAPE results of ES and C however, which slightly contradicts the idea that shorter planning horizons have a negative impact on forecast accuracy. Since the main purpose of this thesis was to evaluate the potential of utilizing installed base data to improve forecast accuracy, this is not evaluated further but is definitely something that should be taken into account both for forecasting in general and when performing similar studies specifically.

5.2 The Impact of Model Features

The output of any forecasting algorithm is strongly affected by the data used. It is therefore of great importance that the data utilized accurately represents the reality one is attempting to predict. The data cleaning was necessary to facilitate the analysis, but may have made the data less descriptive of Volvo's current situation. For instance, parts which have been replaced or had no sales during the studied period were excluded. For the purpose of this thesis, the cleaned data set is deemed to be sufficiently representative. It is however a slightly simplified version of Volvo's spare part reality. If Volvo or any other organization are to incorporate installed base data or perform other Big Data projects it is of utmost importance to fully understand that a model can only be as good as the data. An example of this is how this study only regards the US market. Including additional countries and a complete set of spare parts, for example, would make the analysis more complex but also increase the installed base data set and potentially further help improve forecast accuracy.

Yang et al. (2008) present various maintenance strategies, while Andersson (2019) discusses these strategies as the key demand drivers for spare part demand. These could be used to add information to the features of a forecasting model, in addition to the age and size of the installed base. Different spare parts can also be categorized based on their overall use. For instance, a part that is replaced at regular intervals could be categorized differently from a part requiring replacement only if the car is involved in an accident. The mentioned categorizations, strategies and similar information could be included in models such as IB and C presented in this thesis and could likely improve forecast accuracy further, if used correctly. They have not been included in this study, mainly due to time and resource constraints and lack of alignment with data available.

In addition to more and bigger sets of quantitative data and features, qualitative data provides further possibilities. Wheelwright et al. (1998) argue for the usefulness of qualitative forecasting, although that is mainly in the medium to long-term and thus outside the scope of this thesis. However, the authors also point out the potential of mixed method forecasting. In this study, the mixed method is C, the Combined model, which combines time series with causal-based forecasting. There could thus be further potential in adding judgemental forecasts and other qualitative factors to the model. The initial planning of this thesis included an additional qualitative data collection phase, which would have included interviews with US car dealers and other parties with first hand knowledge of working with spare parts. The hope was that findings from these interviews could help identify a number of potential factors which affect spare part demand that are not usually visible in raw data. These factors could have been utilized as further data input for the models during the analysis. The effect of including these factors would then have been evaluated, and provided a way of quantifying qualitative and judgmental knowledge. Due to the COVID-19 pandemic however, interactions with other parties in both Sweden and the US became severely limited and hence this part of the study was excluded. The potential impact this would have had is hard to speculate about, but if the interviews had been conducted more factors could potentially have been included and tested besides the data already included in this study.

Given that the input data for each model differs, it is possible to compare how the importance of features varies depending on other data available to the model. For this analysis, it is mainly the Installed Base and Combined model that are of interest since the Exponential Smoothing model only uses historical sales as its input. In Stage 1 and Table 4.3 it can be seen that the top contributing features all are features of the individual parts, and together make up 58/100 of the impact on the model. The remaining impact is from the installed base information. This indicates that even if the most prominent features are not installed base data it still makes up a considerable 42/100 of the impact on model outcome.

In Stage 2 and 3 the Installed Base and Combined model can be compared. Table 4.4 shows just how important the inclusion of historical sales is to the forecast which in turn helps explain why the Combined model performs so much better than the Installed Base model. This is in line with the findings of Andersson (2019), who found that the inclusion of historical sales on a fault code regression analysis improved performance. In both Stage 2 and 3 the function group and description of the individual parts are the most influential features. For the Combined model there is a substantial shift with the Exponential Smoothing forecast on top and the group and description not even in among the ten most influential features. The nature of the CatBoost

algorithm, which continually seeks to prioritize the most influential features, tells us that as soon as historical sales are available the specific information internal to each spare part loses importance. This strongly indicates that the information gained from historical sales is a stronger predictor of future sales than the various spare part features. The installed base information still influences the forecast outcome however, sitting at around 25/100 during both Stage 2 and 3. To conclude, the most efficient way of further increasing the total forecast accuracy would be to fine-tune and configure the Exponential Smoothing model. Another approach, which is not exclusive to the previous one, is to increase the scope of the model and include additional features such as maintenance strategies and technical specifications and evaluate their impact on model performance.

5.3 The Imposition of Demand Pattern Categorization

In the Demand Pattern Categorization-part of the analysis, Section 4.5, the cut-off values of $CV^2 = 0.49$ and $p = 1.32$ were chosen according to the recommendations by Syntetos et al. (2005). This was done despite the authors arguing against the imposition of arbitrary cut-off values for categorization. This is mainly to statistically identify the areas for which each model performs optimally, and adjust the cut-off values accordingly. The alternative for this thesis would have been to evaluate the demand patterns of parts involved in the studied data set separately, and performing a statistical analysis on the forecasting models for each pattern. However, for the purposes of this thesis the categorization was done to evaluate the performance of the models on the various demand patterns, not to assign the models to certain patterns. The use of pre-determined cut-off values was thus determined to better suit the scope of the thesis, while still facilitating the desired analysis.

It should also be noted that in the data sets used throughout the thesis, shown in Table 3.1, only includes parts which have sales in the studied period. Including parts without sales in the Demand Pattern Categorization would however not impact the current distribution, since no CV^2 or p values are calculable without at least two data points. Thus, including such parts would have necessitated a category without any demand pattern, which would have been excluded from the analysis either way. Thus, the impact of this exclusion on the categorization is considered minimal.

The Demand Pattern Categorization, shown in Table 4.5, categorizes 544 parts as Lumpy, 291 as Smooth, 2996 as Intermittent and 155 as Erratic. According to Van der Auweraer et al. (2019), this distribution is similar to other spare part demand patterns. Table 4.6 and 4.7 shows the Exponential Smoothing and Combined models having similar performance on predicting demand of Intermittent parts. This likely means that the high percentage of Intermittent parts somewhat mitigates the shortcomings of the time series model and improves its overall accuracy, compared to the Combined model. Given how the Combined model performs on par or better than the ES model for all demand patterns, it should generate the best forecast regardless of demand pattern distribution. In contexts with increased shares of lumpy or erratic demand the difference between a time series model and a combined model should thus prove even more significant. This is in line with the arguments by Dekker et al. (2013) on the benefits of including installed base data when forecasting intermittent demand items in general.

6 | Conclusion

The answer to Research Question 1, *What forecasting model can be applied to effectively predict initial phase spare part demand for Volvo Cars using installed base data?*, is two-fold. The Installed Base model developed in this thesis provides a way of predicting the first sale of parts when neither sales data or demand patterns are available. This can be clearly seen by comparing the results in Table 4.3 with Table 4.4. During Stage 2 and 3, a combination of the Exponential Smoothing model and Installed Base model is clearly the best choice regardless of demand pattern, as shown in Figure 5.1 and 5.2.

Further, during Stage 1, the results surpass those of Volvo's current approach clearly, and the potential improvements in first-hit availability and inventory costs are significant. The model does however not take factors such as maintenance strategy, calculated risks and the costs of inventory versus back-orders into account. Adding a layer of item value or criticality to the model is probably necessary in any practical implementation. Incorporating further features into the installed base model could likely improve first sales accuracy further.

In Stages 2 and 3, the benefit of using solely installed base data decreases as more sales data becomes available. The time series model performs better than the exclusively installed base one, but the combination of the two outperforms them both. This shows that a combination of a time series model adapted to the data and a model using installed base data has clear potential. The data used in this study was however cleaned before analysis, and any similar studies and practical experiments should thus be planned and examined closely to ensure suitability with the data and products at hand.

Research Question 2, *How would the applied model impact forecast accuracy for products with different demand patterns?* is answered in Section 4.5, as well as in Table 4.6 and 4.7. To conclude, the results show that for parts with lumpy or erratic demand the inclusion of installed base data shows substantial improvements to forecast accuracy. For parts with intermittent or smooth demand patterns the effect of including installed base information is situational.

Lastly, the Demand Pattern categorization and following analysis also shows that the more extreme the demand behavior is, the more helpful the use of installed base data becomes. So, even though the potential is greatest for lumpy and erratic demand, it is likely so that most challenges associated with the forecasting of intermittent demand patterns can be mitigated by similar approaches.

6.1 Recommendation

Dekker et al. (2013) show the benefits of utilizing installed base data, in terms of potential cost savings derived from reduced inventory and obsolescence costs. These benefits are strengthened by the findings in this study. Combined with the overview presented by Andersson (2019) of

when to apply certain forecasting models and data sources, provide ample foundation for the potential of installed base data. A strong recommendation is thus for organizations with access to installed base data to investigate how installed base forecasting in the early phase of spare parts life cycle could be used to improve forecast performance.

When moving forward with the suggestions above, or other Big Data projects, one should be aware of the associated challenges. Big Data provides vast potential, but successful utilization comes with many challenges. Organizational change is often required in order to become more data-driven and to actually trust the data and make decisions based on it. New competences may need to be acquired, and current employees educated and trained. The competences needed for a successful Big Data or Data Science project is, for example, described in the Drew Conway Data Science Venn Diagram, see Figure 6.1.

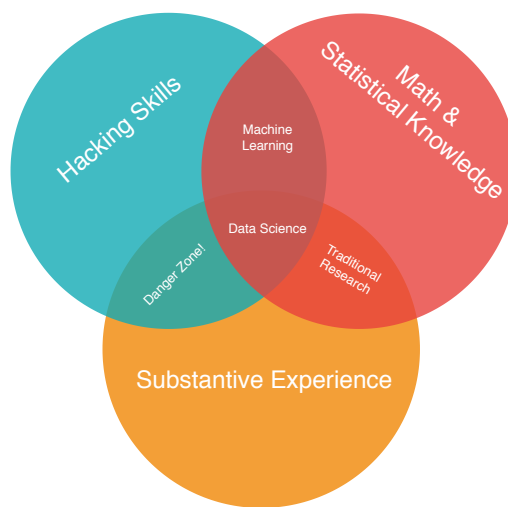


Figure 6.1: The 'Drew Conway Data Science Venn diagram', visualizing the interactions between the three parts of Data Science (Connway 2013)

The diagram shows the interactions of three key knowledge areas in regards to Big Data, and the intersections between them. It is the combination and overlap of 'Hacking Skills' (programming), 'Math and Statistical Knowledge' and 'Substantive Expertise' (domain knowledge) that is the field Data Science. The lack of either of these three skills means Data Science is not truly achieved, and any project performed will yield sub-optimal results. The lack of programming and computer skills tend to lean a project towards research with little real implementation. The lack of domain knowledge may lead to a project with a poor fit to the reality of the business. Finally, the lack of statistical knowledge can lead to a project with low theoretical understanding of the subject and lead to problems of guaranteeing long-term accuracy and ability to explain the outcome of any model. In short, all of these three competences are deemed necessary, and organizational processes for them to coexist and align should be put in place in order to enable a successful data-driven business (Connway 2013). Specific recommendations are outside the scope of this thesis, but it is generally recommended that organizations curate these three competences and work towards collaboration between the three.

The model combining time series forecasting and installed base data provides a simple way of

including installed base information in Volvo Cars forecasting. The current forecast produced by Volvo Cars systems can be used as the base for the historical sales extrinsic variable in the model where it is combined with installed base information. Volvo should move forward and start investigating how a practical implementation of installed base forecasting in the early phase of spare parts life cycle could be done. In order to do so Volvo needs formalized ways to evaluate the forecasting accuracy of new models. The new model should be continually tested over time with real-life data and its accuracy evaluated. Only once the impact of the model and installed base information is fully understood it should be considered for day-to-day forecasting tasks. Lastly Volvo should be aware of the challenges associated with Big Data and work to increase their competence of the field both through new hires and ongoing training of current employees.

6.2 Limitations

The COVID-19 pandemic forced the authors to omit a qualitative interview-based data collection. The qualitative information could have resulted in better forecast accuracy but does not affect the validity of this study since the scope was adjusted to cover only quantitative data.

In Stage 2 and 3 an estimation of the current forecasting method is included as an input to the Combined forecast. In Stage 1 no such inclusion of the current forecasting process is performed. This is done as a result of Volvo's own perception of the low performance of their current manual approach as well as with regards to time constraints and available resources. Still the results of Stage 1 speaks for themselves and show just how powerful installed base data can be for initial phase demand forecasting.

In with the above, another limitation with this study is that no comparison with an actual qualitative method has been done. This is especially relevant in Stage 1, where the comparison performed is to Volvo's results and not those of any established method. Using a structured qualitative and judgmental method and comparing it to the performance of an installed base model could be interesting, and provide further depth to the field.

When developing the Installed Base model, only a simple comparison of various options is done. The choice of the decision-tree based CatBoost algorithm is based on the type of data available and a basic understanding of the theory surrounding it. This thesis argues that machine learning algorithms in general can be utilized with installed base data to improve initial phase forecasting, but a study of which exact algorithms to use has not been done.

The Demand Pattern Categorization was performed by using the values suggested by Boylan et al. (2008). As mentioned in the thesis, this is a somewhat backward order of doing things. It is generally suggested to decide the cut-off values according to the performance of the available forecasting methods, but this was not quite in the scope of this study since a thorough evaluation and comparison of various machine learning algorithms was not performed.

6.3 Future Studies

This thesis only evaluates one source of alternative data, namely installed base data. However the forecasting methodology presented is flexible in terms of the data applied. Therefore, great potential exists in evaluating additional data sources to be used as extrinsic variables for forecasting. One alternative source presented by Andersson (2019) is in-use data, such as real time sensor data collected from vehicles.

A study that collects information from various experts such as car dealerships, mechanics, and R&D personnel with first hand knowledge about the spare parts provides an interesting opportunity. Such expertise of spare part demand drivers and indicators of spare part demand could, if incorporated into similar models, provide vast qualitative data for a deepened quantitative analysis. This could improve forecast accuracy further, and provide systematized ways of incorporating domain expertise into the forecasting process.

On that note, a comparative study of the performance of qualitative methods during the initial phase could be relevant. This is especially true for Stage 1, since they are one of few realistic methods when historical data is unavailable. A study similar to this one, but focusing on the first stage and using various structured and qualitative methods could improve the understanding of the potential of installed base forecasting as well as develop the qualitative methods themselves.

Such a study could be combined with, or separate from, a more thorough evaluation of various machine learning algorithms. Again, using a similar analysis approach of developing, comparing and evaluating but focusing on a set of different machine learning models could help narrow the scope of available models. These models could also include qualitative assessments as features, and potentially improve performance further.

Lastly, a study based on this or the future studies mentioned above could evaluate the performance of various models and then perform the Demand Pattern Categorization based on contextual and adapted cut-off values. This might improve the usefulness of the models further. Utilizing different ones for the various demand patterns is potentially effective.

This is all relevant in order to increase the usefulness of installed base data. As a way of leveraging current competence it is thus of interest to further evaluate how installed base data and judgemental forecasting may be combined in the early stages of the spare part life cycle.

Bibliography

- Andersson, Joakim (2019). *THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING Enhancing aftermarket demand planning with product-in-use data*. Tech. rep. Chalmers University of Technology.
- Bacchetti, Andrea and Nicola Sacconi (2012). “Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice”. In: *Omega* 40.6, pp. 722–737. ISSN: 03050483. DOI: 10.1016/j.omega.2011.06.008.
- Bartezzaghi, Emilio and Matteo Kalchschmidt (2011). “The Impact of Aggregation Level on Lumpy Demand Management”. In: *Service Parts Management*. Springer London, pp. 89–104. DOI: 10.1007/978-0-85729-039-7{_}4.
- Boylan, J E, A A Syntetos, and G C Karakostas (2008). “Classification for forecasting and stock control: A case study”. In: *Journal of the Operational Research Society*. Vol. 59. 4, pp. 473–481. DOI: 10.1057/palgrave.jors.2602312. URL: <https://orsociety.tandfonline.com/action/journalInformation?journalCode=tjor20www.palgrave-journals.com/jors>.
- Boylan, John E and Aris A Syntetos (2007). “Forecasting for Inventory Management of Service Parts”. In: *Complex System Maintenance Handbook*. Springer-Verlag. Chap. 20. ISBN: 9781848000117. DOI: 10.1007/978-1-84800-011-7.
- Brinch, Morten (2018). “Understanding the value of big data in supply chain management and its business processes: Towards a conceptual framework”. In: *International Journal of Operations and Production Management* 38.7, pp. 1589–1614. ISSN: 17586593. DOI: 10.1108/IJOPM-05-2017-0268.
- Bryman, Alan and Emma Bell (2003). *Business Research Methods*. 3rd. Oxford University Press. ISBN: 964-7445-88-1.
- Cavalieri, S., M. Garetti, M. MacChi, and R. Pinto (2008). “A decision-making framework for managing maintenance spare parts”. In: *Production Planning and Control* 19.4, pp. 379–396. ISSN: 09537287. DOI: 10.1080/09537280802034471.
- Cohen, Morris A and Narendra Agrawal (2006). “Winning in the Aftermarket”. In: *Harvard Business Review* 2006.may, pp. 11–11. DOI: 10.1007/0-387-26336-5{_}55.
- Connway, Drew (2013). *The Data Science Venn Diagram*. URL: <http://drewconway.com/zia/2013/3/26/the-data-science-venn-diagram>.
- Croston, J. D. (1972). “Forecasting and Stock Control for Intermittent Demands.” In: *Operational Research Quarterly* 23.3, pp. 289–303. ISSN: 0160-5682. DOI: 10.1057/jors.1972.50.
- Dekker, Rommert, Çerağ Pınar, Rob Zuidwijk, and Muhammad Naiman Jalil (2013). “On the use of installed base information for spare parts logistics: A review of ideas and industry practice”.

- In: *International Journal of Production Economics* 143.2, pp. 536–545. ISSN: 09255273. DOI: 10.1016/j.ijpe.2011.11.025.
- Dorogush, Anna Veronika, Vasily Ershov, and Andrey Gulin (2018). “CatBoost: gradient boosting with categorical features support”. In: pp. 1–7. URL: <http://arxiv.org/abs/1810.11363>.
- Fritzsche, R., J. N.D. Gupta, and R. Lasch (2014). “Optimal prognostic distance to minimize total maintenance cost: The case of the airline industry”. In: *International Journal of Production Economics* 151, pp. 76–88. ISSN: 09255273. DOI: 10.1016/j.ijpe.2014.02.001. URL: <http://dx.doi.org/10.1016/j.ijpe.2014.02.001>.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The Elements of Statistical Learning*. Second Edi. Springer. ISBN: 9780387848570. DOI: 10.1007/b94608{_}4.
- Hyndman, R.J. and G. Athanasopoulos (2018). *Forecasting: principles and practice*.
- Jonsson, Patrik and Stig-Arne Mattson (2009). *Manufacturing, Planning and Control*. 1st ed. McGraw-Hill Higher Education, p. 468.
- Kennedy, W. J., J. Wayne Patterson, and Lawrence D. Fredendall (2002). “An overview of recent literature on spare parts inventories”. In: *International Journal of Production Economics* 76.2, pp. 201–215. ISSN: 09255273. DOI: 10.1016/S0925-5273(01)00174-8.
- Khajavi, Siavash H., Jouni Partanen, and Jan Holmström (2014). “Additive manufacturing in the spare parts supply chain”. In: *Computers in Industry* 65.1, pp. 50–63. ISSN: 01663615. DOI: 10.1016/j.compind.2013.07.008.
- Koudal, P. (2006). “The service revolution in global manufacturing industries”. In: *Deloitte Research* 2006.October, pp. 1–21. DOI: 10.13140/2.1.2631.9042.
- Levitt, T (1965). “Exploit the product life cycle”. In: 43.6.
- Liaw, Andy and Matthew Wiener (2002). “Classification and Regression by randomForest Andy”. In: *Merck Research Laboratories* 2.3. ISSN: 1609-3631.
- Liu, Xiao and Loon Ching Tang (2016). “Reliability analysis and spares provisioning for repairable systems with dependent failure processes and a time-varying installed base”. In: *IIE Transactions (Institute of Industrial Engineers)* 48.1, pp. 43–56. ISSN: 15458830. DOI: 10.1080/0740817X.2015.1055391.
- McAfee, Andrew and Erik Brynjolfsson (2012). “Big Data: The management review”. In: *Harvard Business Review* 2012.October, pp. 60–69. ISSN: 0017-8012. URL: <http://tarjomefa.com/wp-content/uploads/2017/04/6539-English-TarjomeFa-1.pdf>.
- McKinsey&Company (2017). “The Changing Aftermarket Game”. In: June, pp. 1–36. URL: <https://www.mckinsey.com/~media/McKinsey/Industries/Automotive%20and%20Assembly/Our%20Insights/The%20changing%20aftermarket%20game%20and%20how%20automotive%20suppliers%20can%20benefit%20from%20arising%20opportunities/The-changing-aftermarket-game.ashx>.
- Pandas* (2020). URL: <https://pandas.pydata.org>.
- Papalambros, Panos Y. and Douglass J. Wilde (2018). *Forecasting: Principles and Practice*. DOI: 10.1017/9781316451038.010. URL: <https://otexts.com/fpp2/>.

- Rebala, Gopinath, Ajay Ravi, and Sanjay Churiwala (2019). *Introduction to machine learning*. Springer. ISBN: 9783030157289. DOI: 10.4018/978-1-7998-0414-7.ch003.
- Sanders, Nada R and Larry P Ritzman (2004). “Integrating judgmental and quantitative forecasts: methodologies for pooling marketing and operations information”. In: *International Journal of Operations & Production Management* 24.5, pp. 514–529. DOI: 10.1108/01443570410532560. URL: www.emeraldinsight.com/researchregister.
- Scikit Learn* (2020). URL: <https://scikit-learn.org/>.
- Solomon, Rajeev, Peter A Sandborn, and Michael G Pecht (2000). “Electronic part life cycle concepts and obsolescence forecasting”. In: 23.4, pp. 707–717.
- Souza, Robert de, Albert Wee Kwan Tan, Hafidzaturrafeah Othman, and Miti Garg (2011). “A proposed framework for managing service parts in automotive and aerospace industries”. In: *Benchmarking* 18.6, pp. 769–782. ISSN: 14635771. DOI: 10.1108/14635771111180699.
- Steuer, Daniel, Verena Hutterer, Peter Korevaar, and Hansjoerg Fromm (2018). “A Similarity-Based Approach for the All-Time Demand Prediction of New Automotive Spare Parts”. In: *Proceedings of the 51st Hawaii International Conference on System Sciences*, pp. 1525–1532. DOI: 10.24251/hicss.2018.191.
- Syntetos, A. A., J. E. Boylan, and J. D. Croston (2005). “On the categorization of demand patterns”. In: *Journal of the Operational Research Society* 56.5, pp. 495–503. ISSN: 01605682. DOI: 10.1057/palgrave.jors.2601841.
- Syntetos, Aris A. and John E. Boylan (2005). “The accuracy of intermittent demand estimates”. In: *International Journal of Forecasting* 21.2, pp. 303–314. ISSN: 01692070. DOI: 10.1016/j.ijforecast.2004.10.001.
- Turrini, Laura and Joern Meissner (2019). “Spare parts inventory management: New evidence from distribution fitting”. In: *European Journal of Operational Research* 273.1, pp. 118–130. ISSN: 03772217. DOI: 10.1016/j.ejor.2017.09.039. URL: <https://doi.org/10.1016/j.ejor.2017.09.039>.
- Van der Auweraer, Sarah, Robert N. Boute, and Aris A. Syntetos (2019). “Forecasting spare part demand with installed base information: A review”. In: *International Journal of Forecasting* 35.1, pp. 181–196. ISSN: 01692070. DOI: 10.1016/j.ijforecast.2018.09.002. URL: <https://doi.org/10.1016/j.ijforecast.2018.09.002>.
- Volvo Cars (2018). “VOLVO ANNUAL REPORT 2018”. In: *Volvo annual report*.
- Wheelwright, S, S Makridakis, and RJ Hyndman (1998). *Forecasting: methods and applications*. URL: <http://hephaestus.nup.ac.cy/handle/11728/7113>.
- Yandex (2020). *CatBoost*. URL: <https://catboost.ai/docs/concepts/parameter-tuning.html>.
- Yang, Zimin, Dragan Djurdjanovic, and Jun Ni (2008). “Maintenance scheduling in manufacturing systems based on predicted machine degradation”. In: *Journal of Intelligent Manufacturing* 19.1, pp. 87–98. ISSN: 09565515. DOI: 10.1007/s10845-007-0047-3.

A | Appendix

A.1 Interview Questions

Start of Interview

The purpose of the interview and the scope of the project is explained for the interviewee

The Interviewee

1. What is your current role at Volvo?
2. How long have you been at Volvo, and what is your previous experience (internal/external)?
3. What are your current strategic responsibilities?
4. What your current operational responsibilities?

Volvo on the US market

1. What does Volvo's organisation look like?
2. How is your organizational contact and communication with Volvo Sweden, in terms of demand planning etc?
3. How do dealer's work with forecasting demand and communicating that with Demand Planning in Sweden?
4. Can you describe the process from when an error occurs/is discovered until the part is replaced?
5. How large is Volvo's market share?

Do you have any data sheets etc you could provide us with?

6. Could you describe Volvo's aftermarket set-up in the US?
7. How do you measure service level? Is it the same company-wide?
8. What are the specific challenges (and opportunities) associated with that?
9. Do US dealers experience any especially problematic issues?

10. Are there any products that are regularly problematic, for which increased availability would have a significant impact on service levels/business results?

Recommendations

1. Are there any specific data points and/or indicators you think might be of interest for us?

A.2 CatBoost Classifier

```

1 import numpy as np
2 import catboost
3 from catboost import CatBoostClassifier, Pool
4 from sklearn.model_selection import train_test_split
5
6 # Data set for the model, the name corresponds to columns in the data set.
7 modeldata = modeldata[['parts_0_months_old', 'parts_1_months_old',
8     'parts_10_months_old', 'parts_11_months_old', 'parts_12_months_old',
9     'parts_13_months_old', 'parts_14_months_old', 'parts_15_months_old',
10    'parts_16_months_old', 'parts_17_months_old', 'parts_18_months_old',
11    'parts_19_months_old', 'parts_2_months_old', 'parts_20_months_old',
12    'parts_21_months_old', 'parts_22_months_old', 'parts_23_months_old',
13    'parts_24_months_old', 'parts_3_months_old', 'parts_4_months_old',
14    'parts_5_months_old', 'parts_6_months_old', 'parts_7_months_old',
15    'parts_8_months_old', 'parts_9_months_old', 'weight_g', 'volume_cm3',
16    'description', 'function_group', 'will_sell_in_next_6_months']]
17
18 # Split data in test and training set, last column in the set is the dependent variable
19 X, y = modeldata.iloc[:, :-1], modeldata.iloc[:, -1]
20 X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2,
21     random_state=12864)
22
23 # Handle categorical features, all non float(number) features are categorical
24 is_cat = (X.dtypes != float)
25 for feature, feat_is_cat in is_cat.to_dict().items():
26     if feat_is_cat:
27         X[feature].fillna("NAN", inplace=True)
28 cat_features_index = np.where(is_cat)[0]
29
30 #Pool of data for the model
31 pool = Pool(X, y, cat_features=cat_features_index, feature_names=list(X.columns))
32
33 # Fit the classifier with parameters to the data defined in the pool
34 classifier = CatBoostClassifier(max_depth=6, iterations=10000).fit(pool)
35
36 # preds contains the predictions for the test set X_test
37 preds = classifier.predict(X_test)

```

Listing 1: Implementation of Catboost Classifier in Python code, modeldata is the data for Stage 1 predictions

A.3 GridSearch

```

1  # load libraries
2      from sklearn import datasets
3      from sklearn.model_selection import train_test_split, GridSearchCV
4      from catboost import CatBoostClassifier
5  # Split data in test and training set
6      X, y = modeldata.iloc[:, :-1], modeldata.iloc[:, -1]
7      X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
8  #Define a new classifier model
9      is_cat = (X.dtypes != float)
10     for feature, feat_is_cat in is_cat.to_dict().items():
11         if feat_is_cat:
12             X[feature].fillna("NaN", inplace=True)
13     cat_features_index = np.where(is_cat)[0]
14     classifier = CatBoostClassifier(cat_features=cat_features_index)
15 #Parameter to use for GridSearch
16     grid = {'learning_rate': [0.03, 0.1, 0.183766019, 0.3, 0.7],
17            'depth': [4, 6, 10],
18            'l2_leaf_reg': [1, 3, 5, 7, 9],
19            'iterations': [500, 1000, 3000, 5000, 10000, 15000]}
20 # Run GridSearch with deined parameters
21     gridm = GridSearchCV(estimator=classifier, param_grid = grid,
22                          cv = 5, n_jobs=-1, verbose=10)
23     gridm.fit(X_train, y_train)
24 # The best estimator found during GridSearch
25     best_classifier = gridm.best_estimator_
26 # The parameters for the best estimator
27     gridm.best_params_

```

Listing 2: GridSearch for optimal parameter for Stage 1 classifier, 5-fold cross validation is used to determine model performance

A.4 Implementation of Exponential Smoothing

```
1  #Implementaion of Exponential Smoothing in Python
2  import os
3  import numpy as np
4  import pandas as pd
5  import matplotlib.pyplot as plt
6  from statsmodels.tsa.api import ExponentialSmoothing, SimpleExpSmoothing, Holt
7
8  # Historic sales data for single part
9  historic_sales_data_part = traditional_part_data
10 historic_sales_data_part = historic_sales_data_part[['yearmonth', 'qty']]
11 historic_sales_data_part = historic_sales_data_part.set_index('yearmonth')
12
13 #Tells program the frequency of sales data is monthly.
14 historic_sales_data_part.index = pd.DatetimeIndex(historic_sales_data_part.index,
15                                                    freq='MS')
16
17 # Fit Exponential smoothing model to data
18 ExpModel = SimpleExpSmoothing(historic_sales_data_part).fit()
19
20 #The optimal alpha-value found by optimization.
21 smoothing_level = ExpModel.model.params['smoothing_level']
22
23 #The predicted sales based on Exponential smoothing
24 preds = ExpModel.fittedvalues
```

Listing 3: Exponential Smoothing model, implemented in Python

A.5 Feature importance

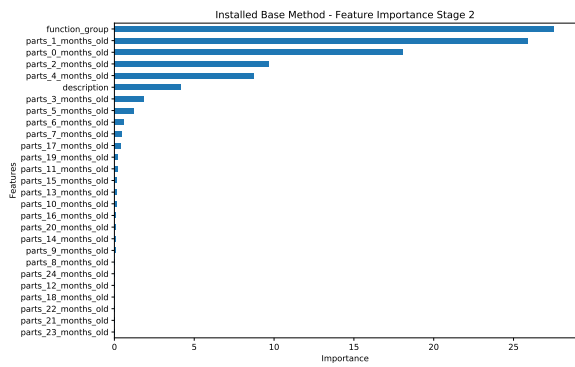


Figure A.1: Installed Base Method Feature Importance Stage 2

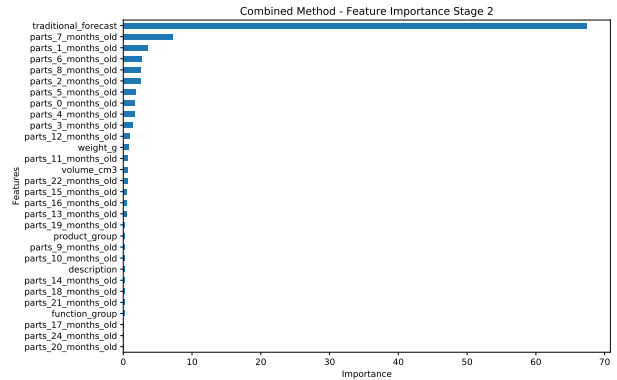


Figure A.2: Combined Method Feature Importance Stage 2




Figure A.3: Installed Base Method Feature Importance Stage 3




Figure A.4: Combined Method Feature Importance Stage 3

DEPARTMENT OF TECHNOLOGY MANAGEMENT AND ECONOMICS
DIVISION OF SUPPLY AND OPERATIONS MANAGEMENT
CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden
www.chalmers.se



CHALMERS
UNIVERSITY OF TECHNOLOGY