Sensing the Future

Designing Sensor-based Predictive Information Systems for Forecasting Spare Part Demand for Diesel Engines

Furtak, Simon

Document Version
Final published version

Publication date:
2018

License
CC BY-NC-ND

SENSING THE FUTURE DESIGNING SENSOR-BASED PREDICTIVE INFORMATION SYSTEMS FOR FORECASTING SPARE PART DEMAND FOR DIESEL ENGINES

PhD Series 07-2018

Szymon Furtak

SENSING THE FUTURE

DESIGNING SENSOR-BASED PREDICTIVE INFORMATION SYSTEMS FOR FORECASTING SPARE PART DEMAND FOR DIESEL ENGINES

Doctoral School of Business and Management

PhD Series 07.2018

CBS COPENHAGEN BUSINESS SCHOOL
HANDELSHØJSKOLEN

SENSING THE FUTURE:
DESIGNING SENSOR-BASED PREDICTIVE INFORMATION SYSTEMS FOR FORECASTING SPARE PART DEMAND FOR DIESEL ENGINES

Szymon Furtak

Supervisors:
Rasmus Pedersen
Michel Avital
Samir Chatterjee

Doctoral School of Business and Management
Copenhagen Business School

Szymon Furtak

Sensing the Future: Designing sensor-based predictive information systems for forecasting spare part demand for diesel engines

Doctoral School of Business and Management is a cross disciplinary PhD School connected to research communities within the areas of Languages, Law, Informatics, Operations Management, Accounting, Communication and Cultural Studies.

To Being Relevant!

1 Abstract

As digital technologies become prevalent and embedded in the environment, "smart" everyday objects like smart phones and smart homes have become part and parcel of the human enterprise. The ubiquity of smart objects, that produce ever-growing streams of data, presents both challenges and opportunities. In this dissertation, I argue that information systems extending these data streams, referred to as "predictive information systems with sensors", can generate added value and will be gaining momentum in academia and in the industry. Subsequently, seeing apparent complexity in designing IS artifacts with such functionality, I introduce a framework for Designing Information Systems with Predictive Analytics (DISPA), extending Design Science Research specifically towards rigorous design of predictive analytics.

The framework is evaluated based on a case study of MAN Diesel and Turbo, a lead designer of marine diesel engines generating multiple applicable artifacts in the process. Additionally, the framework exemplification in the case context led to supplementing the framework with a set of Design Principles for Designing Predictive Information Systems as well as a matrix for pre-assessing financial feasibility of predictive information systems with sensor technologies.

This work provides a contribution to information systems research, and in particular to design science research, by introducing a model for Designing Information Systems with Predictive Analytics (DISPA) that can serve as a method for developing IS artifacts. The framework constitutes an Information System Design Theory consistent with the established definitions from the literature (Gregor & Jones, 2007; Kuechler & Vaishnavi, 2012; Walls, Widmeyer, & El Sawy, 1992). In addition, the paper introduces and systematically evaluates a number of spare-part forecasting methods, which can be considered a contribution to operations research literature.

2 Oversigt

I takt med at digitale teknologier udbredes og indlejres i omgivelserne, bliver "smarte" dagligdags genstande som smartphones og intelligente hjem en del af den menneskelige verden. Allestedsnærværende smarte objekter, der producerer en stadigt voksende strøm af data, præsenterer både udfordringer og muligheder. I denne afhandling, argumenterer jeg for, at informationssystemer der bruger disse nye datastrømme, kaldet "forudsigende informationssystemer med sensorer", kan generere merværdi og vil vinde momentum både i den akademiske verden og samt i industrien. Efterfølgende ses på kompleksitet i design af artefakter med sådan funktionalitet, idet jeg præsenterer en ramme for Designing Information Systems med Predictive Analytics (DISPA), der styrer Design Science Research specifikt imod stringent design af predictive analytics.

Rammerne er evalueret baseret på et casestudie af MAN Diesel og Turbo, en ledende producent af marine dieselmotorer, der genererer flere artefakter i processen. Derudover medførte designet at rammen blev suppleret med et sæt retningslinjer for at designe forudsigende informationssystemer, samt et instrument til for vurdering af økonomisk gennemførlighed af prædiktive informationssystemer med sensorteknologi.

Dette arbejde bidrager til forskning indenfor informationssystemer, og især til Design Science, ved at tilføre en model for Designing Information Systems med Predictive Analytics (Dispa), der kan tjene som metode til udvikling af IS artefakter. Rammerne udgør tilsammen en informationssystem design teori i overensstemmelse med de opstillede definitioner udfra litteraturen (Gregor & Jones, 2007; Kuechler & Vaishnavi, 2012; Walls, Widmeyer, & El Sawy, 1992). Hertil kommer, at afhandlingen introducerer, og systematisk evaluerer, en række reservedelsproducenternes prognosemetoder, der kan betragtes som et bidrag til operationer forskningslitteratur.

3 Keywords

Information Systems, Predictive analytics, Design Science Research, Forecasting, Sensors, Information System Design Theory

4 Acknowledgements

5 Contents

6 List of figures

7 List of tables

8 Glossary

Here is the list of project specific terms and their meaning in the context of the project:

- **Activity sensor** – a sensor designed within the period of this project measuring the activity of the engine it is installed in, measured in an equivalent of engine running hours with maximum load. Detailed description can be found in Iteration 3 on Instantiation 1, under Empirical investigations.

- **Croston method** – statistical forecasting method optimized towards infrequent demand, common among spare parts. It is a two-step approach calculating separately expected intervals between demand points and the magnitude of demand, when it occurs. Detailed description can be found in Iteration 1 on Instantiation 1, under Empirical investigations.

- **Hit rate** – percentage of sales offers (quotations) converted to actual sales orders

- **Installed base** – product units currently in use, relevant for aftersales services.

- **Online services** – a name used internally for the Continuous multi-sensor monitoring in the case company. The project was fully implemented in the case company, as a asynchronous satellite communication system able to poll data from the engine installation all over the world. Detailed description can be found in Iteration 1 on Instantiation 2, under Empirical investigations.

- **Phase-out sensor** -– a sensor designed within the period of this project detecting when engine is permanently set not to be in used, i.e. when it is scrapped. Detailed description can be found in Iteration 2 on Instantiation 1, under Empirical investigations.

- **Sensor** - a device that responds to a physical stimulus (as heat, light, sound, pressure, magnetism, or a particular motion) and transmits a resulting impulse (as for measurement or operating a control) (Merriam-Webster, 2015a). In the case context there was often an intermediary between the physical sensor and the case company (e.g. when the data from AIS movement sensor was processed to activity indicator by a third parity provider) leading for simplicity to using notions of "sensor" and "sensor data" interchangeably

- **Slow steaming** – a strategy applied by ship owners in order to conserve fuel and engine wear by limiting the vessels speed.

- **The Red Helmet** – a name used internally for the Flexible Remote Monitoring Interface in the case company. The prototype of the device was an actual helmet with a GO-PRO camera. Detailed description can be found in Iteration 2 on Instantiation 2, under Empirical investigations.

1 Introduction

This chapter provides a general introduction to the research project. It describes how the project was initiated and the initial issue that the project was intended to address. The problem statement leads to the phrasing of the research questions, which leads to defining a research strategy for the project and the structure of the thesis, which is expected to fulfil this strategy by discussing the paradigmatic assumptions. The chapter concludes the limitations of the project scope and draws the conclusions.

1.1 Background and Context

This thesis presents research on using sensors to improve the predictions for demand and, more broadly, predictive analytics in general. The main difficulty is the assessment of the contextual meaning of observed sensor data output to the phenomenon that is being predicted.

As explained by Nate Silver (2012), the author of a popular book on prediction in the age of big data called *The Signal and the Noise: The Art and Science of Prediction*, people are used to mistaking noise in data as a meaningful signal. He noted that our brains have been trained to find meanings and patterns, but the world has considerably evolved since our brain chemistry and common sense were developed. We are bombarded now by substantially more data than we were used to. As a result, there are things we might not understand or that are caused by random chance that we can mistake for a signal (CFA Institute, 2015).

Information overload

The evidence from the industry (including the case described below) shows that increasing data volumes are hindering organisational prediction capabilities due to the overwhelming ratios of false positives in trend detection in predictions based on multiple data sources. More insight into the sensor-monitored phenomena allows better predictions, but only if we have a very good understanding of both the phenomenon in question and the signal describing it. To use Silver's (2015) words, we need to understand what is a meaningful signal and what is just noise. The challenge is crucial, as there are many domains where improving predictions could be beneficial and even lifesaving (like healthcare, failure detection, or disaster forecasting). In every

Context understanding as a prerequisite to effective use of data

case, the contextual classification into the signal and noise can be different and even reverse.

> "I think the challenge people don't realize is that when you have more and more data, in some ways that makes it harder. Now you have to choose which data you're looking at, what you took to be the signal versus the noise." (Silver, 2015)

1.2 Motivations and Case Description

This project was initiated by the industrial partner, MAN Diesel & Turbo in 2011 based on a lack of understanding that they recently experienced regarding the effective use of expensive sensor technology they acquired and its integration into daily operational processes (read more below, in the section 1.2.4 titled 'In the age of sensors'). MAN Diesel & Turbo is the world market leader for large diesel engines for use in ships and power stations and is one of the three leading suppliers of turbo machines. The roots of the company date back to 1758. In the years between 1893 and 1897, Rudolf Diesel and MAN engineers developed the first diesel engine. In 1904, the company constructed its first steam turbine. According to the recent shipbuilding outlook report (Maritime-Insight, 2013), MAN Diesel & Turbo has designed approximately 70% of the engines for active goods-carrying vessels, which means that MAN Diesel & Turbo engines propel more than half of the world trade with almost 90% of the seaborne trade share in world trade! Nowadays, MAN Diesel & Turbo do not build engines. The company's strategy concentrates on engineering-intensive engine design processes and creating revenue from selling manufacturing licences to third parties and from the after-sales part of the engine business, namely, offering spare parts and services.

MAN Diesel & Turbo

When we consider the consequences of moving from a manufacturing business to only after-sales on the supply chain setup, there is a clear increase in the operational difficulty level (Figure 1). In the manufacturing environment, the demand is directly linked to the customers' desire and need to purchase new products, while it additionally includes the condition of equipment for after-sales, which depends on its use and can be a surprise even to the customer.

Aftersales logistics

Manufacturers usually work with required response times, matching the production time, as products are usually manufactured to stock. After-sales products and services depend on the product for which they are intended. For markets with a large, heterogeneous installed base (IB; units currently in use) with large variation, products cannot be produced to stock, creating a gap between production time and the required response. Another consequence of a heterogeneous IB is the substantial number of stock-keeping units (SKUs) and lack of complete control over phasing items out (customers might wish to service their products even after servicing the product stops being profitable for the service provider, but the vender might be contractually bound to offer the service). All those differences lead to huge differences in inventory cost, structure, depreciation risk, and turnover.

PARAMETER	MANUFACTURING SUPPLY CHAIN	AFTER-SALES SERVICES SUPPLY CHAIN
Nature of demand	Predictable, can be forecast	Always unpredictable, sporadic
Required response	Standard, can be scheduled	ASAP (same day or next day)
Number of SKUs	Limited	15 to 20 times more
Product portfolio	Largely homogeneous	Always heterogeneous
Delivery network	Depends on nature of product; multiple networks necessary	Single network, capable of delivering different service products
Inventory management aim	Maximize velocity of resources	Pre-position resources
Reverse logistics	Doesn't handle	Handles return, repair, and disposal of failed components
Performance metric	Fill rate	Product availability (uptime)
Inventory turns (The more the better)	Six to 50 a year	One to four a year

Figure 1 – Comparison of supply chain setups for new building manufacturing and after-sales services (Cohen, Agrawal, & Agrawal, 2006).

MAN Diesel & Turbo focus on introduced challenges in after-sales for their supply chain, particularly in forecasting spare parts and service demand. After-sales-based business models usually involve a higher level of heterogeneity and product variation than initial sales environments, leading to higher levels of demand uncertainty, making demand predictions relatively more difficult (Teunter, Syntetos, & Zied Babai, 2011). Moreover, in the marine business, design changes introduced in the manufacturing process are very common, typically due to local material availability or shipyard manufacturing limitations, causing alterations in instantiations of the same design and additional variation of the IB. Finally, the licence-based business model implemented by MAN Diesel & Turbo creates additional obstructions to after-sales activities, as it introduces the engine builder (MAN Diesel & Turbo licensee) as an intermediary between MAN Diesel & Turbo and the end customer (ship owner) who, in the after-sales market, is a competitor. This setup limits information flow between customers and MAN Diesel & Turbo, additionally hindering forecasting.
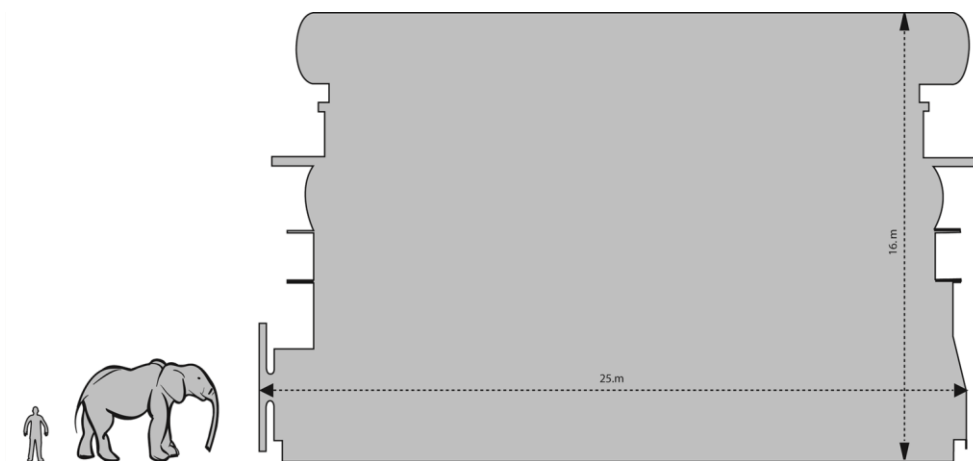
Figure 2 – MAN Diesel & Turbo 2-stroke diesel engine size visualisation (MAN Diesel & Turbo internal).

The emphasis on demand forecasting in the case context is also introduced by the after-sales-oriented business model. In the after-sales environment, the customer purchases spare parts and services based on two main criteria: availability and price. Availability is largely dependent on an accurate demand prediction; if the demand is expected in advance, items or services can be ready at the time the customer requests them, increasing sales probability without the cost of excess inventory (Silver, Pyke, & Peterson, 1998). Moreover, procuring parts in advance (engine spare parts or elements necessary for performing additional engine services, such as retrofit installations) enables stable production pipelines that decrease overall procurement costs by avoiding rush orders and expensive rush transportation, helping to keep the price at a level acceptable to customers. In an environment characterised by a high IB heterogeneity and a high number of offered products and services, additionally underpinned by potentially incomplete information on product build and use, an effective forecasting process can be considered as difficult as it is important. The focus on after-sales, market uncertainty, and sheer size and value of the spare parts (Figure 2) make the forecast error levels that are acceptable in other businesses simply too expensive to bear in the introduced context.

Importance of forecasting

1.2.1 Marine engine market and MAN Diesel & Turbo business model

The market for marine-engine new builds and after-sales is characterised by a high level of complexity. In general, four main actors can be identified: engine designers, original equipment manufacturers (OEM), ship owners, and vessel managers. When a new engine is to be manufactured, a ship owner chooses between offers submitted by a team consisting of a designer and an OEM. After an offer is selected, a team manufactures an engine according to a design provided by the engine designers, using a manufacturing facility and power provided by the OEM. Usually, designers offer their designs to OEMs for a licence fee and provide technical supervision for the entire manufacturing process and the right to offer spare parts for the engine on the after-sales market.

Marine business landscape

After the vessel is ready, the ship owner chooses the vessel manager to control the daily business (in large shipping companies, the ship owner and the vessel manager can be units of the same company), usually for a fixed amount of time and at a fixed rate. Diesel engines require regular spare-part replacement for as long as the engine is in use. As spare-part designs are protected by industrial design rights (a kind of IPR) in the spare-part market, the vessel manager can buy spare parts only from the designer or the OEM (as mentioned above, they buy the rights to offer spare parts under the design licensing agreement). The service to replace the spare part is not restricted and can be performed by any service hub.

In this scenario, MAN Diesel & Turbo is the engine designer. It teams up with manufacturers, such as Hyundai, Mitsui, or

Doosan, and delivers a range of diesel engine models. After the engines are in operation, the partners in the manufacturing process become competitors, offering an alternative for spare-part replacement. MAN Diesel & Turbo's streams of revenue are divided almost equally between the licence fees from OEMs (also referred to as licensees) and the profit from the after-sales services (mostly spare parts but also spare-part replacement services, retrofitting products, or engine management services). This setup introduces tension between MAN Diesel & Turbo and the OEMs because the partners in the new build market become competitors in the after-sales market. This does not create an incentive for OEMs to share all the information regarding the engine manufacturing process (including any eventual changes to the original design), introducing potential data inconsistency in after-sales activities for MAN.

MAN as marine diesel OEM

Recently, this landscape of OEM, engine designer, and engine operator became disrupted by the boom of servitisation and product-service bundles: rather than simply selling products, companies have increasingly started offering product-service bundles. Rather than selling a car, companies are leasing its use on a fix price, or rather than selling JET engines, companies are offering hours of JET engine operation (like GE and Rolls-Royce). Similar concepts started to develop in the diesel engine markets. However, although market players are increasingly offering similar product-service bundles, their popularity on the market is still limited. Nevertheless, the focus on the maintenance market while designing the engine was one of the enablers for conducting this project. More information about servitisation can be found in the literature review.

Servitisation as project enabler

1.2.2 Practical challenges in forecasting diesel spare-part demand in the marine business

The two previous sections explain some of the difficulties introduced by the nature of the after-sales business and the business model implemented by the case company. In addition, there are also challenges introduced by the business environment, which are related to the object of the forecast, the diesel engine spare part, and the marine business itself.

Spare part forecasting challenges
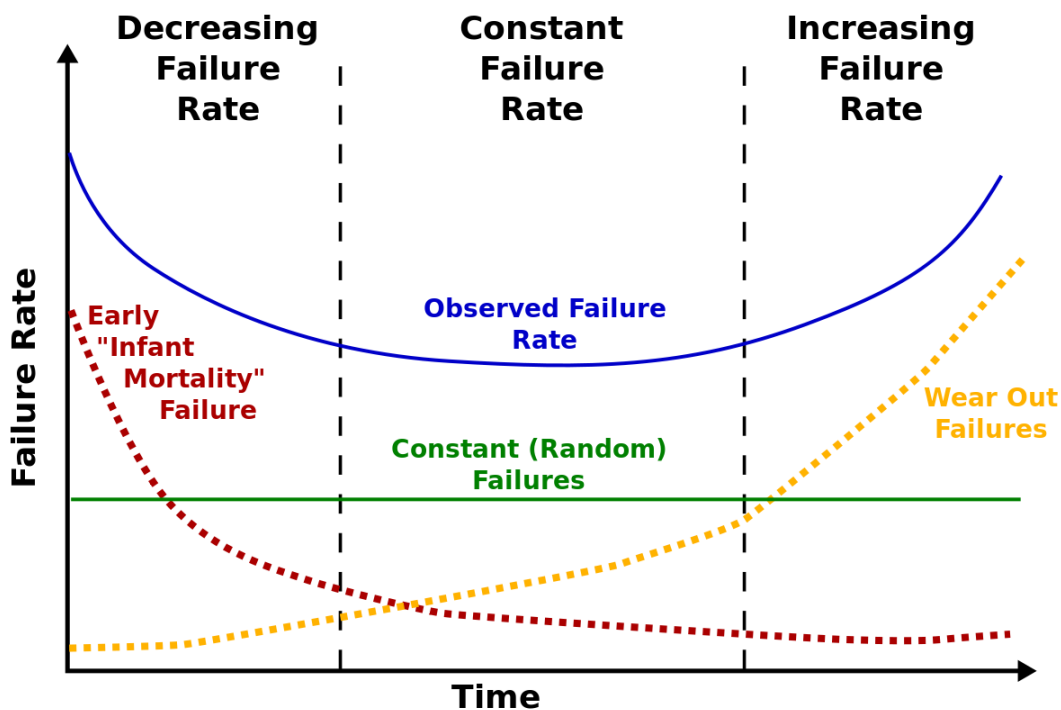
Figure 3 – Diesel spare-part failure rate (Nahim, Younes, Nohra, & Ouladsine, 2015) image source (Wikipedia, 2016a).

According to the engineering literature, the spare-part failure behaviour is driven by a combination of three failure functions: early failure, constant failure, and wear-out failure functions (Nahim et al., 2015; Figure 3). Early failures, often referred to as 'infant mortality' failures, occur in the initial lifecycle phase and are often caused by marginal issues in

Spare-part failure behaviour

spare-part production or installation. The effect of those failures is especially visible for technologically advanced engine components that require complex production or installation and maintenance processes. Constant or random failures can occur throughout the lifecycle with a stable probability. They take place irrespective of the exact state and wear of the component, for example, because of overlooking some maintenance activity or wrongful operation. Wear-out failures are breakdowns that happen when a part is used to the extent that prevents the engine from correct operation. As a result of these three factors, the observed failure rate, presented in blue in Figure 3, has a characteristic bathtub curve (Nahim et al., 2015).

Implications of complexity of failure rates

This complex failure rate distribution has direct implications for forecasting spare-part demand. Predicting future failures based on any historical data is then hindered by the fact that the consumption could have a different nature. Out of the three reasons for failure, only the wear-out scenario could have some repetitiveness in the pattern and, seeing only the patterns of placing orders, it is extremely difficult to assess which orders were driven by wear. Additionally, marine customers are required to maintain a certain quantity of spare parts on board the vessel and often have warehouses where they store parts waiting to be used. As a result, all statistical methods using historical sales data for extracting spare-part use patterns amplify an error present in the input data, due to a gap between sales and replacement (for parts sold to customer stock) and due to early and constant failures.

One more aspect that is almost unique to marine supply chains is that the 'customers' are in constant motion, and they roam around the world. This introduces another dimension

to the forecasting process; it requires not only a prediction of 'when' and 'how much' but also 'where'. Transportation takes time and increases cost. Additionally, if a main engine is not operational, and the vessel cannot be controlled, especially in proximity of land, the lives of hundreds of people can be in danger, and an immense cost for the ship owner can be incurred. It is not uncommon to fly, often tens of kilometres into the ocean, to deliver spare parts to operational engines with a helicopter. Alternatively, a single tug operation of a gigantic uncontrollable containership can cost millions of dollars and bankrupt a significant fraction of shipping companies.

One more source of difficulty when forecasting diesel spare-part demand in the marine business is the changing business environment in which the customers operate. Up to half of the ship operation costs and about one-third of the total shipping company running costs are for the fuel (or, as often referred to in the marine business, the bunker price). With the ever-changing price of crude oil (Figure 4), the most common fuel in the marine business, shipping companies move from the edge of bankruptcy to historically high incomes year-to-year, even without any changes in operation. This leads to changing customer behaviours, introducing another variation to the data used to forecast future demand.
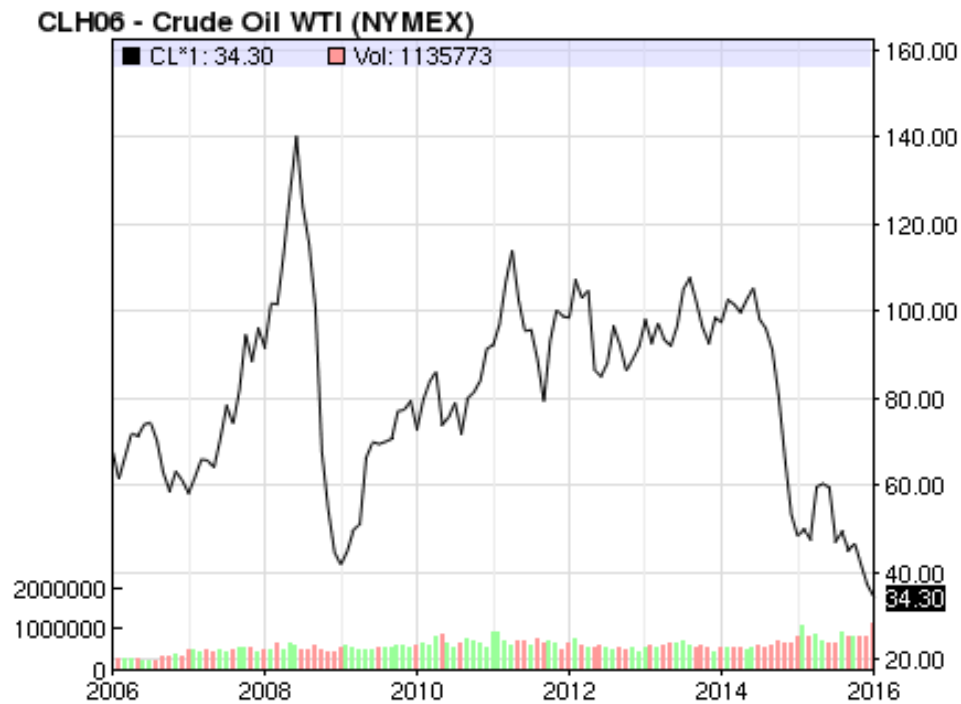
Figure 4 – Crude oil price between 2006 and 2016 (NASDAQ, 2016).

Despite the significance of all those mentioned challenges, probably the most serious source of difficulty is access to only partial information. As explained in the previous section, not building engines in house creates an information gap regarding engine designs and exact spare parts in use. Clearly, the variety of applications and operational scenarios creates *Incomplete* substantial differences in spare-part wear patterns. At the *data* *challenge* same time, the information remains unknown to the spare-part provider, adding another complexity level to the forecasting process. Even historical sales data are incomplete; the customer could replace spare parts from other unauthorised sources, not leaving any trace of replacement to the original spare part provider. Each of these sources of variance creates noise that is input into the forecasting system, creating a bullwhip effect (Lee, Padmanabhan, & Whang, 1997) and amplifying any input noise on the output (Figure 5).
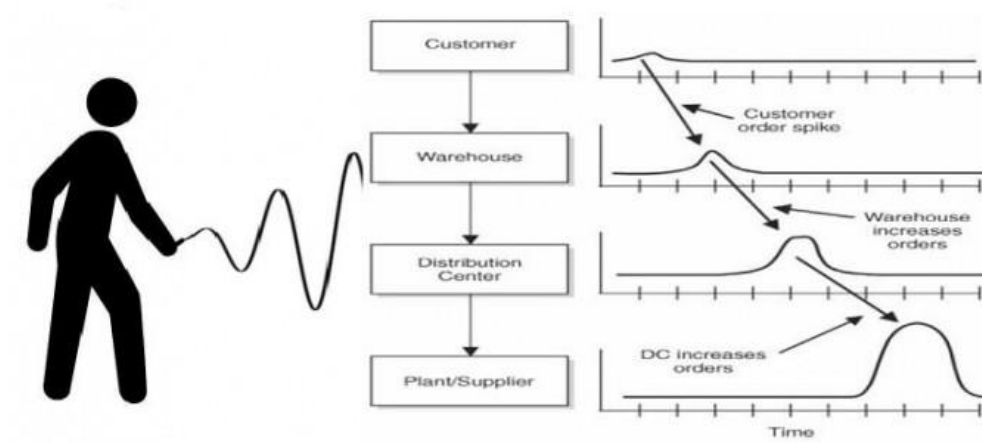
Figure 5 – Illustration of the bullwhip effect (image source: Wikipedia, 2016b).

1.2.3 Engaged scholarship in practice: Practical project setup

This research project is an example of an engaged scholarship approach (Van de Ven, 2007), which comprises research activities conducted by a researcher embedded in a real organisational context, dealing with current problems at hand. Throughout the period of the project, the researcher has been employed at the case company as an industrial PhD student, allowing the researcher to benefit from being embedded in the case environment and allowing the researcher to develop a business understanding of the academic problems discussed in this work. The researcher has been employed in the case company prior to the project period and took active part in defining the scope of the initial problem framing. This understanding and the relations built before and during the initial framing helped to inform the project with perspectives of multiple relevant departments in the case company.

Engaged scholarship in MDT

From the overall scope perspective, the most relevant stakeholders of the project are in the supply chain area. The Supply Chain department has a section dealing with demand forecasting for components and spare parts (called demand planning), running the daily process of forecasting the most common spare parts, and coordinating the monthly S&OP process. The successful execution of the research project is aimed to inform the daily forecasting process of demand planning and could even support it with new support tools. This overlap in objectives led to placing the PhD student with the demand planning team in the supply chain organisation.

MDT Supply Chain department involvement

Another department with interest in the project is the department for after-sales: Spare-Part Sales and Technical

Service. The department is the commercial arm of the case company, selling spare parts for equipment in use and providing advanced maintenance and repair services for those installations. Close to half of the workforce of the department have a very strong technical background, typically complemented by years of experience as a superintendent, a technical engineer responsible for operating an engine on board a vessel. About half of that team and one-fourth of the department workforce comprise the Technical Service team, the most experienced diesel engine engineers who spend about 200 days a year travelling around the world and solving the most complex technical problems on the engines in operations. The setup of the project has been defined in the way that the PhD student was formally affiliated with the department, with two intentions: access to technical experts experiencing daily customer behavioural patterns, while being able to provide qualitative insight into the models under development, and the formal affiliation with the after-sales department, which was established to facilitate knowledge transfer from the project to the sales process, enabling proactive sales activities linked to improved planning and forecasting processes.

MDT aftersales department involvement

The last department identified as key to the success of the project was the department of Business Intelligence (BI). The department is anchored in the business end of the organisation (with strong support by the IT department for infrastructure and technical support) and works with mapping transactional and master data (from ERP, CRM, PLM, and shipping commercial data sources) to enable business-value-creating analytics. Prior to the project period, the PhD student worked as an architect in the BI department. The project was defined with key stakeholders from the department, and their full support has been ensured.

MDT BI department involvement

The main benefit of carrying out an academic project in the industrial setting is daily access to the context details that allow understanding of the details of the problems and quick feedback on any developing ideas. As I sat in the office neighbouring all the mentioned key stakeholder departments, I had a chance to see the challenges myself, allowing me to really understand them. Any idea for improvement could be pre-evaluated quickly, and results could be discussed with stakeholders from multiple functional areas. The three departments discussed in this section were also those evaluating artefact designs proposed in the scope of the project for improved spare-part forecasting.

Involvement of key stakeholders

In addition to the access to people, I also gained almost unlimited access to data. Being an employee of a company removes the legal burden of non-disclosure agreements and trust, speeding up analysis and insight discovery processes. With the support of the BI team, I was able to request the addition of many structured data sources, like tables from the MAN Diesel & Turbo ERP system SAP R/3, which were previously not used for analytics, and semi-structured data sources. I triggered the first-time load of the online service (sensor data collection system) data to the BI environment and coordinated the load of textual service reports.

Daily access to data

Clearly, there are also challenges for an industrial PhD setup that materialised in my setup at MAN Diesel & Turbo. Primarily, there was a significant difference in the scope interests between the company and researcher. The company focus is on the problem at hand, in this case, the challenge in forecasting spare parts. However, the academic search for knowledge requires at least exploring the scope of generalisation and aims at addressing classes of problems, rather than just instances. This mismatch in focus often

Misalignment of focus between academia and industry

18

resurfaced when discussing actions to be taken between the academic and industrial supervisors.

The dual focus, both on the concrete case of 'issue solution' and the more research-relevant investigation of a class of problems introduced an additional challenge: the division of time between the two scopes. Although there is academic value to the rigorous system design of a case-specific challenge, there is clearly more academic value in researching a more general approach, like a framework to design systems for similar problems that, by design, are less interesting for the company. In contrast, the organisational work in ensuring *Dual focus in* the suggestions are embedded in the daily working process is *the project* not that interesting for academia but is crucial for achieving business benefits. There is also an interesting issue in terms of results dissemination. Due to the significant differences in jargon, there is a need for two sets of documentation, as the academic documentation is difficult to understand for the company employees, while the industrial documentation is too unprecise for academia and too focused on operational use. In this case, the industrial PhD student must work as a bridge, 'translating' the documentation and focusing on the relevant elements for both environments.

1.2.4 In the age of sensors

Recently, MAN Diesel & Turbo decided to tap into over 400 sensors installed on every engine for the local monitoring of engine operating conditions by rolling out a project to collect data from those modules and send them to a central data warehouse (DWH), which is a project referred to as the Online Services project. The sensors continuously collect hundreds of various operational measurements, such as temperature and pressure at multiple points, and were only available as a support for an on-board engineer; however, they can now be polled remotely from every engine installation in use all over the world using satellite communication. The project was initially acknowledged as a major success, with all the complex technological elements able to interplay as intended.

Nevertheless, despite terabytes of data collected every day, in time, it became apparent that it is difficult to translate this project into company profit. Although experienced engineers familiar with the context of the data were able to see signs confirming the context, all attempts to infer past maintenance activities based only on the data yielded substantial errors and failed prediction attempts. The only product successfully rolled out based on the online service data is operational reporting, which is elaborate reporting presenting raw collected values, but profit from this product can barely cover the costs of running the project and cannot justify the substantial initial investment.

Initially, the main objective of the project was to improve MAN Diesel & Turbo's chances of winning big maintenance orders. Knowing the demand in advance was supposed to improve availability and allow proactive sales campaigns. In

Sensors in MAN diesel engines

MAN sensor challenge

Value of sensors for MAN Diesel & Turbo

addition, better understanding of demand patterns was expected to stabilise procurement processes and lower internal costs. MAN Diesel & Turbo management saw the project as part of a strategic transformation that could lead to enriching the current business model by additional revenue streams coming from sensor-based, knowledge-intensive products. Unfortunately, due to the inability to draw general insight from the collected data, none of these benefits were realised. Not being able to find a solution on its own, MAN Diesel & Turbo turned to industrial specialists concerning the matter. The company invited leading practitioners offering state-of-the-art IT tools for analytics in the planning context, including SAP APO,[1] SAS Asset Performance Analytics,[2] and others, but none of the implementations seemed to fulfil the planned objectives.

Not being able to find answers in existing industrial solutions, the company turned to academia by initiating an industrial PhD project.[3] The idea behind the project was to identify challenges in designing a satisfactory solution and to subsequently close the gap by designing and developing the required tool. Moreover, as some other internal operational areas could benefit from a similar solution, guidelines on how to proceed in such situations were also expected.

---

[1] SAP APO: business software from SAP, a leading ERP software provider, for advanced planning and optimisation (APO), the core SAP software for supply chain activities.

[2] SAS Asset Performance Analytics: a module of enterprise statistical software SAS to maximize asset performance.

[3] Danish PhD project schema where the student is conducting research activities in a company. More information can be found on the website of the Ministry of Higher Education: http://ufm.dk/en/research-and-innovation/funding-programmes-for-research-and-innovation/find-danish-funding-programmes/programmes-managed-by-innovation-fund-denmark/industrial-phd

1.3 Problem Statement

An early analysis of the industry proved that MAN Diesel & Turbo's difficulties in applying sensor-collected data in predictive analytics are not unique and can represent a common class of problems dealing with encapsulating sensor data in prediction models (Ghodrati, Benjevic, & Jardine, 2012; Ploesser, 2013; Welbourne et al., 2009). As digital technologies become prevalent and are embedded in the environment, an increasing number of everyday objects, such as smart phones, smart cars, smart homes, and even smart clothes have become part of the human enterprise. In this context, the adjective 'smart' denotes that an object can collect, process, and often communicate data regarding its functionality and operating environment (Cook & Das, 2004). Subsequently, all smart objects must be equipped with sensors that can collect various kinds of data.

Predictive analytics with sensors as a class of problems

Although there are many examples of the successful utilisation of the current snapshots of such data, identifying patterns from historical sensor data to make predictions is only now entering everyday applications. For example, GPS data on phones can provide a current location, but they cannot currently indicate where one is going. The main players in the mobile market, Google and Apple, are trying to close this gap by introducing services such as Google Now and Apple Frequent Locations, and collecting data with similar functionality in mind. However, usable applications based on predictions seem significantly more difficult to implement than those using a snapshot picture (Woollaston, 2013).

Forecasting an event upfront is especially important if there is a substantial cost or gain associated with that event. With vast amounts of data collected for snapshot analysis and with

the main players clearly looking to extend it towards the future, it is apparent that more applications designed to benefit from the predictive analysis of historical sensor data will be entering the market.

A preliminary review of literature has shown that sensors and predictive analytics are present in multiple streams of the current research. In this context, predictive analytics refers to empirical methods that are aimed at creating empirical predictions and assessing their quality. Although the application of predictive analytics toolkits is not common in IS research, it is quite prevalent within the body of management science and particularly in operations research (OR), both as a method and as a topic, namely, forecasting.

Predictive analytics in literature

Demand forecasting, one of the key processes in supply chain management, offers generic forecasting algorithms that became a part of the standard ERP system implementations (Bingi, Sharma, & Godla, 1999), making them easy to use in virtually any business. The growing interest in sensor technologies and their ramifications, such as data volume and velocity or information processing capabilities, moves the process of demand forecasting to the forefront of current IS research, where sensors and the Internet of things (IoT) (Atzori, Iera, & Morabito, 2010) are already given significant attention.

1.4 Problem Applicability to a Class of Problems

Although an increasing number of industries embrace the utilisation of sensor devices and IoT, there is no framework that is based on data collected from the sensor architectures in the literature that can guide designers and developers of predictive models. The intended framework solution applies to a class of problems where the objective is to design an IT artefact that relies on collected sensor data and IoT architectures.

The IT artefact is a key contribution of a design science research (DSR) project (Hevner & Chatterjee, 2010). This section describes specific aspects of context-aware predictive systems that make them a class of IT artefacts. A specific definition of the properties of context-aware predictive systems is a prerequisite to specifying the scope of a design framework and to establishing the common features that the design framework can address. The concept of an IT artefact is well established in IS research, and some of the most common definitions are presented in Table 1.

Although the diversity of definitions of IT artefacts in IS research is apparent, there are some commonalities that can be pointed out. An IT artefact has both material and abstract components. They allow processing information to enable or support tasks, and they can take the form of constructs, models, methods, or instantiations. As a sub-class of IT artefacts, context-aware predictive systems possess these characteristics, but their definition could be narrowed further based on their properties. All context-aware systems collect context information using sensor technology to sense specific kinds of changes in the environment in which they are embedded.

Predictive systems have a common modus operandi; they convert input data into a prediction about the future. In congruence with Shmueli and Koppius (2011), we assume that the prediction task is well defined and that the quality of the outcome (i.e., the prediction) can be measured. Thus, we define context-aware predictive systems as systems that process information collected by sensors or other technology using a predictive method to support a well-defined and measurable forecasting task.

Table 1. The IT artefact in IS literature (Zhang et al., 2011).

Source	Description
Orlikowski and Iacono (2001)	'bundles of material and cultural properties packaged in some socially recognizable form such as hardware and/or software' (p. 121)
Benbasat and Zmud (2003)	'the application of IT to enable or support some task(s) embedded within a structure(s) that itself is embedded within a context(s),' whereby its hardware/software design 'encapsulates the structures, routines, norms, and values implicit in the rich contexts within which the artifact is embedded' (p. 186)
King and Lyytinen (2004)	'systematic processing of information in human enterprise' (p. 541)
Hevner et al. (2004)	'constructs (vocabulary and symbols), models (abstractions and representations), methods (algorithms and practices), and instantiations (implemented and prototype systems)' (p. 77)
Agarwal and Lucas (2005)	'the integration of the processing logic found in computers with the massive stores of databases and the connectivity of communication networks', so that it 'includes IT infrastructure, innovations with technology, and especially the Internet' (p. 394)

1.5 Research Questions

To approach the class of problems described above, I will frame my research by aiming to answer the following question:

How can predictive information systems (IS) artefacts using sensor technologies be designed in an organisational context?

Specifically, the problem can be further narrowed to the question:

How can organisations design IS artefacts using sensor technologies to forecast spare-part demand in a marine business?

Although the first, broader question is much more interesting in the context of academic research, the involvement of an industrial partner requires a certain focus on the case environment. As a result, I will combine the broader focus on predictive information systems with the narrower scope regarding MAN Diesel & Turbo throughout the project.

1.6 Paradigmatic Assumptions and Underlying Philosophical Approach

1.6.1 Importance of paradigmatic assumptions

To uncover some of the underling limitations of this work, the paradigmatic assumptions defining the approach to answering the previously stated research questions are discussed. The paradigmatic assumptions are the structuring assumptions used to order the world into fundamental categories (Brookfield, 1995). As they are central to our perception of the world, they define our understanding and guide us in any actions we take, including embarking on research quests. Paradigmatic assumptions are so embedded in us that they are often hard to notice. It takes a great deal of resistance to face and examine one's own paradigmatic assumptions.

Paradigmatic assumptions

The purpose of this exercise is not necessarily to overthrow these beliefs and values but rather to assess their effect on the investigation and findings to 'unbias' them from me and my own beliefs. Brookfield called this act 'hunting assumptions'. He defined it as a prerequisite to any act of critical thinking:

Defining paradigmatic assumptions as an initial research step

> Critical thinking happens first when we try to discover the assumptions that influence the way we think and act. Pretty much every action we take is based on assumptions that we have accepted, sometimes unthinkingly, as accurate. Critical thinking involves us deliberately trying to find out what these assumptions are' (Brookfield, 1987, p. 7).

1.6.2 Understanding the most relevant school of thought in philosophy of science

Paradigmatic assumption as a reflection of our philosophical standpoint

When defining the research method, our paradigmatic assumptions underline a philosophy of science that informs the aspects of our perception of the nature of the examined phenomenon (ontology) and the methods for understanding it (epistemology; Van de Ven, 2007). Other important aspects related to discussing paradigmatic assumptions are what constitute value and valuable (axiology) and the language used to communicate the research (semantics).

This research project employs the engaged scholarship perspective, as it is a collaboration of industry and academia on a real-life problem. This fact alone introduces multiple perspectives (industrial and academic) and the requirement for comparative understanding of different philosophical perspectives, which need to be considered when identifying a suitable approach.

Four most common schools of thought in engaged scholarship

Van de Ven, in his seminal work on engaged scholarship, featured four philosophies of science that reflect current practice and debate among social science involved in the engaged scholarship movement: positivism, relativism, pragmatism, and realism (2007 p. 38). He described the historical review of the development of the four schools and defined them with four discriminating characteristics: ontology, epistemology, knowledge, and language. To organise this complex picture, I present his comparison, simplified to only two dimensions: ontology (the perception of the nature of things) and epistemology (how we gain knowledge about those things) in Table 2, following the logic of Johnson and Duberley (2000, 2003), who distinguished

28

between the four schools based on their ontological and epistemological perspectives.

Logical positivism is ontologically and epistemologically objective. It assumes a world independent of cognition and impartial knowledge generation using inductive reasoning. On the other hand, relativism is ontologically and epistemologically subjective. It perceives reality as socially constructed, dependent on the viewer's perspective, and denies any objective representation of social reality, making impartial knowledge acquisition impossible. The pragmatic school includes streams that adopt either subjective or objective views on ontology, but consistently adopts subjective epistemology, based on the relevance of knowledge to its application and actions that it can trigger. In that context a theory can be 'truthful' (meaning useful) in the case in which its use yields positive results but not necessarily in all other cases. Finally, realism adopts an objective view of ontology (there is a 'real' world that exists independently of our perception), but it does not reject the metaphysics, leading to mixed views on epistemology (although objective knowledge could be acquired, it is often surrounded by a subjective discourse that assesses the accuracy of knowledge in a subjective manner).

Table 2 – Comparison of characteristics of logical positivism, relativism, pragmatism and realism, based on Van de Ven (2007, p. 39).

Dimension	Definition	Ontology	Epistemology
Logical Positivism	Philosophical movement inspired by empiricism, instrumentalism, and positivism (Vienna Circle, Berlin School).	**Objective**: Reality is the empirical world (the world of the senses, i.e., the rejection of the metaphysical).	**Objective**: The correspondence between our statement and reality through inductive verification or deductive falsification.
Relativism	Contemporary intellectual movement characterised by its scepticism about the foundations of Western philosophy (historical relativism, social constructivism, postmodernism, critical theory, and hermeneutics).	**Subjective**: Reality is socially constructed.	**Subjective**: There is no privileged epistemology due to incommensurability (lack of a common measure) of discourses.
Pragmatism	Philosophical movement characterised by the relation of theory and praxis and specifically in the predominant outcomes of an inquiry (relativism: Dewey and Rorty, realism: Peirce, James and Rescher).	**Subjective**: Similar to postmodernism. **Objective**: Reality places limitations and constraints on our actions.	**Subjective:** Dependent on the practical implications (actions that it allows one to take).
Realism	Philosophical movement characterised by the existence of a mind-independent reality and ability of theory to capture partial aspects of reality (conjecture realism, realistic pragmatism, critical realism, etc.)	**Objective**: Reality exists independently of our cognition. Thus, there is no basis to reject the metaphysical (epistemic fallacy).	**Objective or subjective**: The world about which science seeks the truth is objective, but there are limitations in predefined or predetermined methodology or criteria to judge the veracity (accuracy) of knowledge.

Even more clearly, Shollo (2013) used Van de Ven's classification to visualise this grouping of four schools of philosophy in a two-dimensional (2D) coordinate system, presented in Figure 6.

Figure 6 – Visual representation of the four schools of according to their views on ontology and epistemology (Shollo, 2013; Van de Ven, 2007).

1.6.3 Paradigmatic assumptions in a design research project

In addition to using engaged scholarship principals, this research project aims at designing an artefact solution to a class of problems observed in organisations. In this context, appending the previously given overview by examining the research perspective related to design can produce useful insight. In the design perspective, research is a knowledge-using perspective (March & Smith, 1995), aiming at changing existing situations into preferred ones (Simon, 1996).

Specifically, regarding information systems, it 'involves the design of novel or innovative artifacts and the analysis of the use and/or performance of such artifacts' (Vaishnavi & Kuechler, 2004). Vaishnavi and Kuechler (2004) provided a very illustrative comparison of the design research perspective for positivist and interpretative (in Van de Ven's work it is referred to as relativist) perspectives. The comparison is presented in Table 3.

Table 3 – Philosophical assumption of the design research perspective with positivist and interpretivist perspectives (Vaishnavi & Kuechler, 2004).

Basic Belief	Research Perspective		
	Positivist	**Interpretive**	**Design**
Ontology	A single reality, knowable, probabilistic	Multiple realities, socially constructed	Multiple, contextually situated alternative world-states, socio-technologically enabled
Epistemology	Objective, dispassionate, detached observer of truth	Subjective, i.e., values and knowledge emerge from researcher-participant interaction	Knowing through making: objectively constrained construction within a context, iterative circumscription reveals meaning
Methodology	Observation: quantitative, statistical	Participation, qualitative, hermeneutical, dialectical	Developmental, measure artefactual effects on the composite system
Axiology	Truth: universal and beautiful, prediction	Understanding: situated and description	Control, creation, progress (i.e., improvement), understanding

The interesting element of Vaishnavi and Kuechler's (2004) analysis is that it presents the design perspective as an extension to the positivist-interpretivist dichotomy. It takes

advantage of some elements of the positivist perspective but is situated in a context (linked to the interpretative perspective) and aims at progress within the context. As a result, the design perspective can be used in prescriptive research as an instrumental way to create solutions for management problems (Aken, 2004). It is also important to draw a relation of this design paradigm to the behavioural science paradigm. Hevner, March, Park, and Ram argued that behavioural and design paradigms are 'two complementary but distinct paradigms' (2004, p. 10). A combination of the two paradigms can provide a useful and rigorous lens to research-relevant problems.

Design paradigm as neither positivist nor interpretivist perspective

In the context of the project, my crucial paradigmatic assumptions revolve around my problem-solving attitude. It is natural for me to believe that a problem has solutions, set-ups that allow addressing tasks introduced by the problem and that the solutions can be assessed as 'better' or 'worse' in the context of the problem, by the degree to which they address the problem in question. More generally, examining my attitude towards research through the lens of ontological and epistemological assumptions and its relation to most common schools of thought, my approach is closest to pragmatism: I believe that 'truthfulness' (or more generally, the importance or value) of a theory or a proposition is constituted by its utility. If one approach seems to provide more satisfactory results than another, then it is 'better' in the context of the given problem. On the other hand, a theory that does not lead to guiding any actions would not be perceived by me as important. This way of thinking is consistent with the previously introduced design perspective and, more generally, the constructivist approach, as it sees propositions as constructed by the context and effects, but with empiricism, as this implies believing in ways of objective

My own paradigmatic assumptions and most related school of thought

observing and assessing performance. One way of expressing the pragmatic principle (PP), according to Almeder, is:

> (PP) A person will be rationally justified in accepting a proposed proposition P as true if
>
> (a) After exhaustive research, there is at that time no currently available conscious inference, either inductive or deductive, from other antecedently known or justified beliefs that would either confirm or disconfirm the proposition P; and
>
> (b) There is some real possibility that accepting P as true, or very likely to be true, will have a tendency to provide behavioural consequences more productive of cognitive or moral utilities than would be the case if one had accepted instead either the denial of P or nothing at all. (2007 p. 172)

Pragmatism in design research

Pragmatism uses the reductionist approach to isolate a domain (a problem) that can be constructed in a context and empirically evaluated or simulated. It is not explicitly associated with qualitative or quantitative methods; in fact, it is very common to see pragmatists implementing mixed methods. The roots of pragmatism go back to the nineteenth century, but it remains an inspiration to scholars in the twenty-first century. It is consistent with the previously presented design paradigm. It is the ontological foundation of the DSR paradigm (Cole, Purao, Rossi, & Sein, 2005; Hevner et al., 2004; Iivari, 2007; March & Smith, 1995), one of the current IS research paradigms (Iivari, 2007), linking the design paradigm with the behavioural research paradigm.

1.6.4 Consequences of applying a pragmatist lens in a research project

The paradigmatic assumption and the school of thought that is most closely associated determines the feasibility of some

34

parameters of any research project. Assumptions regarding ontology, epistemology, or axiology determine what kind of methodology is feasible for the project and hence what kind of research questions and overall objectives a project should have. There are literature streams defining the right fit between research objectives, research questions, research perspectives, and methodologies. I provide an overview based on Punch and Oancea (2009), which is appended by my own assertion of the feasibility of the pragmatist perspective for a given research project class (Table 4).

My classification of pragmatist feasibility shows that the key for its applicability is the contextualisation of a problem. With its epistemology based on practical implications of research, there needs to be a context to experience the relevance of the research. For the context of this research project, the pragmatist perspective provides a strong fit, especially considering its engaged scholarship perspective and application of the design paradigm. The combination of the three provide a genuine business need that requires a constructed artefact as a 'solution' and a fixed context environment that can be used to evaluate project utility and implications.

The academic challenge in such a research project is two-fold. On one hand, academic rigour needs to be held to ensure internal validity of the results, and on the other, the case must be an example of a meaningful class of problems to ensure the generated knowledge is usable beyond the discussed case (external validity).

Table 4 – A research question: philosophy of science fit according to Punch and Oancea (2009); the last column added by the author.

Research Aims and Claims	Kinds of Research Questions	Examples of Research	Pragmatist Feasibility
Explanatory	What is the relationship between?	Survey, experiment	For relationships in a context
Explanatory Descriptive Prescriptive	What happens if?	Experiment, participatory research, action research	Within experiment, a case or class of cases
Descriptive Explanatory	'What' and 'why'?	Mixed-methods research	Within a case or class of cases
Explanatory Descriptive	What happened in the past? How to make sense of the past?	Historical research	Within a case or class of cases
Understanding Interpretative	How can we understand a situation?	Ethnographic and interpretive/case study	Within a context
Critique Emancipatory	How to disrupt convention and empower participants?	Critical approaches	Limited to a single well-defined context

Further narrowing this project's research perspective, in terms of defining the research paradigm, the methodology will be presented in the third chapter titled 'Research Design'. A thorough summary of IS researchers' discussions on the epistemological tradition of the field (discussing theory in IS) in the context of this project's contributions is presented in the discussion chapter.

1.7 Research Strategy

After stating the research questions and carefully evaluating my paradigmatic assumptions and ontological and

epistemological views on the world, I am ready to establish a process of conducting research activities. Consistent with my pragmatic viewpoint, I see the 'truthfulness' of the knowledge I aim to generate in terms of its utility. Thus, while investigating ways to rigorously design predictive systems with sensors, I will assess their value through the lens of practical utility in the context. Nevertheless, to satisfy the criterion of external validity (understood as the extent to which the case-based results can be generalised to other situations), I will follow a two-fold approach. Based on the properties of sensor technologies and the current state of knowledge (applying state-of-the-art design models), 1) I will work towards researching a general model for designing useful predictive systems using sensors. Accordingly, I will use the case of MAN Diesel & Turbo engines to instantiate the newly constructed model 2) to validate whether the results obtained using the model outperform the state-of-the-art benchmarks.

The first step in approaching my research goal is obtaining a solid understanding of the topic within the scope of the work. According to the work of Lars Mathiassen (2002), understanding of a phenomenon can either be reached by 1) synthesising documented empirical studies and broader literature on the subject or 2) by engaging in interpretation of practice. To ensure the understanding of both academic and industrial perspectives, both suggested paths will be taken. For the interpretation of practice, a very close, action-research-like (Cole et al., 2005) cooperation with the industry will be initiated, allowing me to understand the properties constituting an effective solution to the problem in the context. In parallel with familiarising myself with the practical context, I will conduct a systematic literature review for areas constituting the research context, such as

Developing of general model for designing predictive analytics with sensors

forecasting (especially in the service and after-sales context, including servitisation of manufacturing), sensors, and big data, using data for business analysis and supporting business decisions (BI and decision support systems [DSSs]). Specific search terms and the process of the literature search will be further refined based on the results of the initial approach.

Using theoretical understanding to design artefacts for the case company

After attaining a thorough academic understanding of the concepts, ensuring cultivation of the cumulative research tradition, an artefact or artefacts in response to the state-of-the-art need coming from the case company will be designed. To facilitate this process, I plan to utilise the DSR paradigm. The overall aim of any DSR project is designing IS artefacts, which is consistent with the two-fold goal of this research project: both context-specific designs of predictive sensor systems and the general model for designing useful predictive systems using sensors can be considered valid IS artefacts. Additionally, several proponents of design science have suggested that it is associated with pragmatism as a philosophical ontology in its attempts to bridge science and practice (Cole et al., 2005; Hevner et al., 2004; Iivari, 2007; March & Smith, 1995), ensuring harmony between my personal views and the methods in use.

1.8 Scope Definition and Intended Contributions

The vision behind setting up the project is that it could lead to the introduction of a framework that can be used to facilitate the process of rigorously designing predictive IS using sensor technologies, addressing a class of problems by investigating the prediction of sensor-based variables. More specifically, the research questions stated in the previous section were:

(RQ1) How can predictive IS artefacts using sensor technologies be designed in an organisational context?

For the specific context of the case company:

(RQ2) How can organisations design IS artefacts using sensor technologies to forecast spare-part demand in a marine business?

In this context, sensor technologies can refer to any IT artefact providing data regarding functionality and the operating environment of the monitored object. The general intention of the project is to answer this question by introducing a more general theoretical framework that will facilitate designing predictive IS artefacts that use sensor data (RQ1) that will also be feasible for spare-part demand prediction in a marine business (RQ2). Additional contributions that I can foresee at this stage of the project are the forecasting sensor-based methods that will be developed using the model in the case context and a thorough theoretical evaluation of the research aspects in the state-of-the-art literature.

1.9 Structure of this Dissertation

As this research project deals with designing an IT artefact, the thesis will follow the DSR publication schema suggested by Gregor and Hevner (2013). Table 5 presents the original description of every chapter and the chapter of this dissertation matching it.

A quick scan of the first and last columns in Table 5 shows that the structure of this dissertation follows the schema suggested by Gregor and Hevner (2013) very closely. The only explicit difference is the combination of artefact description and evaluation into one section (Chapter 4). The reason for this modification was foreseeing multiple iterations in the design process that would be easier to describe without repetition if the information was merged into a single chapter. Other changes are only minor phrasing differences.

Table 5 – Thesis structure (Gregor & Hevner, 2013): last column added by the author.

Section	Contents (Gregor & Hevner, 2013)	This Thesis
1. Introduction	*Problem definition, problem significance/motivation, introduction to key concepts, research questions/objectives, scope of study, overview of methods and findings, theoretical and practical significance, structure of remainder of paper.* For DSR, the contents are similar, but the problem definition and research objectives should specify the **goals** that are required of the artefact to be developed.	Chapter 1: Background and Context
2. Literature Review	*Prior work that is relevant to the study, including theories, empirical research studies and findings/reports from practice.* For DSR work, the prior literature surveyed should include any prior design theory/knowledge relating to the class of problems to be addressed, including artefacts that have already been developed to solve similar problems.	Chapter 2: Theoretical Foundations
3. Method	*The research approach that was employed.* For DSR work, the specific DSR approach adopted should be explained with reference to existing authorities.	Chapter 3: Research Design
4. Artefact description	*A concise description of the artefact at the appropriate level of abstraction to make a new contribution to the knowledge base.* This section (or sections) should occupy the major part of the paper. The format is likely to be variable but should include at least the description of the designed artefact and, perhaps, the design search process.	Chapter 4: Empirical Investigations
5. Evaluation	*Evidence that the artefact is useful.* The artefact is evaluated to demonstrate its worth with evidence addressing criteria such as validity, utility, quality, and efficacy.	Chapter 4: Empirical Investigations
6. Discussion	*Interpretation of the results: what the results mean and how they relate back to the objectives stated in the introduction section. Can include: summary of what was learned, comparison with prior work, limitations, theoretical significance, practical significance, and areas requiring further work.* Research contributions are highlighted and the broad implications of the paper's results to research and practice are discussed.	Chapter 5: Discussion & Conclusion
7. Conclusion	*Concluding paragraphs that restate the important findings of the work.* Restates the main ideas in the contribution and why they are important.	Chapter 5: Discussion & Conclusion

1.10 Chapter Summary

In this chapter, I introduced the problem domain and drew a story line concerning how the project came to life. Furthermore, I introduced a class of problems dealing with encapsulating meaningful, sensor-based variables into forecasting models, exemplified by my case environment, for the context of spare-part demand forecasting. I pin-pointed the goals of my research investigation in the research questions, discussed the path to addressing them, including the limits to the scope of the investigation and envisioning the intended contributions. Finally, I referred to the DSR publication schema by Gregor and Hevner (2013) and used it to structure the remainder of this thesis accordingly.

Summary of the chapter

2 Theoretical Foundations

This chapter presents an overview of related academic literature and concludes with outlining a research framework that will guide the following empirical part of the thesis. Given that predictive analytics had rarely been discussed in the IS discourse, either as a method or a subject, I examined other management disciplines for similar problems. I initially reviewed general, well-accepted spare-part forecasting methods and later moved towards more innovative ones, utilising multiple data sources and trying to explore underling mechanisms beyond historical demand data analysis, especially in the context of spare parts and after-sales services.

Chapter overview

The context of after-sales led me to servitisation of manufacturing and new business models for the manufacturing industry. A literature search of IS contributions related to DSSs follows because, through an IS lens, systems supporting organisations in tasks like demand forecasting are normally branded DSSs. Subsequently, the sensor context introduced issues of IoT and big data. In addition, I examined the methodological and paradigmatic foundation of designing predictive IS by revisiting the DSR stream and IS contributions to predictive analytics. Figure 7 presents the graphical presentation of the discussed areas and their relation to the project.
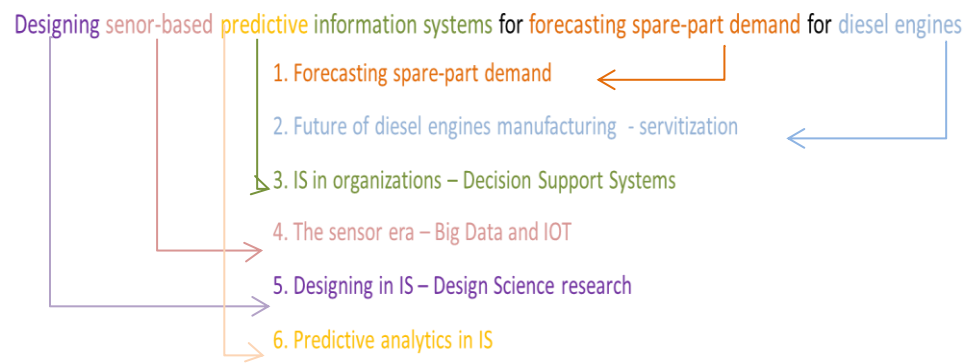
Figure 7 – Literature review topics and their relation to the overall project scope.

2.1 Forecasting Spare-part Demand

2.1.1 Traditional approach to forecasting spare parts

Within OR, forecasting has been thoroughly studied, and multiple literature reviews on spare-part demand are available. Selected methods from Callegaro (2010) are presented in Table 1 below. The purpose of presenting this list is to show that, in the mainstream of current OR predominant forecasting, the methods are limited to the transformation of historical data. Scholars in OR primarily develop clever and complex algorithms to predict the next items in a series based on previous values. Conceptually, this means predicting an output only by analysing its previous outputs and seeing the system as a whole or as a black box. My general idea is that the prediction can be more informed if it is based on an understanding of the activity within the black box. An extended search of the literature shows that a limited number of contributions attempt to conceptually assess this black box (Dolgui & Pashkevich, 2008; Ghodrati & Kumar, 2005; Hellingrath & Cordes, 2014), and I will examine them in the following section.

Forecasting spare-part demand

The simplest forecasting methods are based on a simple historical demand average calculation (arithmetic or weighted). In addition to the average component, trend extraction and extrapolation are common, as in the exponential smoothing method. More specific methods can include multiple trend extraction, such as recent trend extrapolation and seasonal factors (the additive winter method). A slightly more advanced approach that is specifically oriented towards the spare-part business is based on the Croston model. A typical average assumes demand occurring in every period.

Common methods in forecasting spare-part demand

Table 6 – Spare-part forecasting methods as listed by Callegaro (2010); model classification added by the author.

Method	Inputs	Description	Model Classification	Important Features
Weighted moving average	-Historical sales data -Weights (constants)	Mean of past data points with weights (usually the older the sample, the smaller the weight)	Time series – arithmetic average	-Stresses recent trends -Easy to compute
Single exponential smoothing	-Historical sales data -Smoothing constant	Computes moving average of demand with smoothing constant	Time series – average with exponential smoothing	-Works with few samples -Easy to compute
Additive Winter	- Historical data - Smoothing constant - Trend constant - Periodicity constant - Width of periodicity	Variation of single exponential smoothing with additional trend term (for seasonality)	Time series – average with exponential smoothing	-Considers seasonality
Croston's method	-Historical sales data -Smoothing constants	Computes SES for both typical demand magnitude and typical periods between demand points	Croston-based two average value with exponential smoothing	-Intended for materials with intermittent demand (many periods without demand)
Syntetos-Boylan approx.	-Historical sales data -Smoothing constants	Extension of Croston removing the positive bias	Croston-based two average value with exponential smoothing	-Statistically proved bias reduction resulting in lower forecast error
Box-Jenkins method	-Historical data -Constants for regression and average	It chooses between two models, moving average and auto-regression, alternatively selected based on historical error	Time series – average, either moving or weighted average	-Can capture complex trends and seasonality -Requires a lot of history to perform well
Bootstrap method	-Historical sales data -Limit for number for resampling	Randomly chosen subset of historical samples (forecast for next 3 periods is 3 randomly chosen periods from the past)	Stochastic – probabilistic	-Probabilistic approach -Needs few samples
Neural networks (NN)	-Historical sales data -Neural network layout	It infers connection between input and output from the training set and uses it to estimate future values	Stochastic – black box	-Inspired by human brain -Tested in various areas as a predictor
Grey prediction model	- Historical data	Adaptive time-series approach using least square estimate as feedback to correct for the error	Time series – average with least square feedback	-It is designed to work under massive uncertainty and was intended to predict hurricane occurrences

46

The essential idea behind Croston's method is the separation of the calculation of demand occurrence in a period and its magnitude. Croston-based methods use historical data to extract these two values, the interval between demand occurrence and its magnitude if it occurs, typically using simple exponential smoothing. The original Croston method has a systematic positive bias that was removed during its modification using the Syntetos and Boylan approximation, which itself was modified multiple times (Callegaro, 2010; Teunter & Duncan, 2009). Croston-based methods perform significantly better for scenarios in which demand is infrequent or intermittent (Teunter & Duncan, 2009). Due to feasibility in particular scenarios, Croston-based approaches are sometimes combined with time-series or stochastic approaches, based on sale frequency measures. Another method combining multiple approaches is the Box-Jenkins method, which uses average or auto-regression depending on which model delivers a lower historical error value.

Common methods in forecasting spare part demand - Croston

Another approach is to use a stochastic method, one that uses probability distributions and random variables and thus does not guarantee the same output forecast for the same input data. The most common stochastic forecasting method in the spare-part business is bootstrapping, an approach that estimates a probability of demand occurrence in a given period and executes a random event with this very probability for every forecasted period (Silver et al., 1998). If the random event associated with the probability occurs, the period assumes the demand occurred and a demand magnitude from a subset of historical demand magnitude values is selected and assigned to the forecasted period. This approach is similar to the Croston-based methods, as it also models demand dually, separating demand occurrence and demand value. Nevertheless, Croston-based methods are

Stochastic methods in forecasting spare-part demand

much more regular and transparent. Another stochastic black-box approach is neural network forecasting (Silver et al., 1998). In this method, historical values are processed multiple times through a neural network-like structure, computing the forecast by adjusting neuron weights with the objective of minimising prediction error. After an optimal setup is established for the historical data (the test sample), the same neuron setup is used to compute the actual forecast.

To generalise, the selected models can be broadly divided into three classes:

(1) Models that are based on a computing forecast as a single-dimensional aggregation of previous observations were classified as a 'time series' cluster;

(2) Models that are based on computing demand magnitude and interval demand points separately and later combining them in the prediction were clustered as 'Croston-based'; and

(3) Models that are based on calculating a forecast value based on other properties of the previous value set, rather than on the raw values, were grouped into a 'stochastic' class.

My review shows that benchmarks of intermittent demand forecasting are inconclusive regarding the relative performance of any of these models (Ragnarsdóttir et al., 2012). Petropoulos et al. (2013) benchmarked time-series and Croston-based methods and concluded that their relative performance depends heavily on the parameters used in the implementation. On the other hand, Kourentzes (2013) presented a study in which a stochastic solution, namely, neural networks, outperformed both time-series and Croston-based algorithms. Finally, in the study by Teunter and Duncan (2009), time-series methods performed significantly worse than the two other classes, while there was no significant difference between two Croston-based methods and

bootstrapping. In summary, the benchmark of forecasting methods for spare-part demand predictions seem to produce mixed results, possibly showing an advantage for Croston-based and some stochastic methods over simple time-series implementations.

2.1.2 Forecasting spare parts: Beyond time series

Overview of explanatory methods

After an extensive literature search, some forecasting methods taking advantage of data that are different from historical demand data were identified and analysed. The methods are summarised in Table 7.

Table 7 – Forecasting methods using more than historical demand data, based on Bacchetti and Saccani (2012) and Hellingrath and Cordes (2014).

Class	Forecasting Method	References	Data Required, Aside from Demand Data?
Advance demand information	Order over-planning, early sales	Verganti, 1997; Bartezzaghi, Verganti, & Zotteri, 1999	No (extrapolates early orders)
Analysis of reliability	Failure rate analysis, operating condition analysis, installed base (IB) forecasting	Tibben-Lembke & Amato, 2001; Ghodrati & Kumar, 2005; Dekker, Pinçe, Zuidwijk, & Jalil, 2010; Jalil, 2011; Lapide, 2012; Minner, 2011	Equipment failure rates, environmental condition factors, IB
Expert forecast	Three-point estimate, Delphi forecast	Clemen & Winkler, 1993; Rowe & Wright, 1999	All structured and unstructured data
Condition monitoring methods	Integrated forecasting method, intelligent maintenance systems, proportional hazard model	Hellingrath & Cordes, 2014; Hua, Zhang, Yang, & Tan, 2007; Louit, Pascual, Banjevic, & Jardine, 2011	Plant/equipment overhaul arrangements, condition of technical system

Forecasting methods using advanced demand data

One approach that I identified relies on utilising orders submitted with expected delivery far into the future. The authors argued that the proportion of so-called 'early sales' (orders to be delivered a given number of months ahead) is quite stable month to month; therefore, systematic 'over-planning' (scaling the early sales up by a predetermined factor) can be an interesting way to predict total future sales but, as pointed out by Bacchetti et al., it 'seems ill-suited for

forecasting spare part requirement in durable customer goods industries' (2012, p. 726). The idea is very interesting, but its application is limited to stable environments. They pointed out that uncertainty will not be addressed well by early sales, as the demand in the spare-part context is triggered by an unexpected event, such as a product failure, making early planning difficult. In addition, for environments with long lead times, sales submitted sufficiently early to have an influence on availability might be so marginal that any extrapolation based on them may result in unnecessary fluctuation. Although the approach uses only demand data rather than transforming historical demand, as it extrapolates sales to be delivered in the future, it was included in this review.

Another approach relies on extracting mechanical properties of spare-part wear in given operating conditions. Conceptually, the idea is to define, based on historical data, the probability of failure in a given period (usually based on age) and to estimate the aggregated predicted demand value based on an estimate of the number of pieces of equipment in use with a given characteristic, such as age. The exact implementation might be more complex than the description *Forecasting* above, and it varies between methods (for details, see *methods* Ghodrati & Kumar, 2005; Hua et al., 2007). This approach can *reliability* include data related to mechanical wear patterns of spare *data* parts (in this approach, extracted from historical data) as well as equipment condition (such as a reference to a lifetime estimate from the last inspection or the date of the last replacement) and data related to the equipment population currently in use, the IB. This approach opens a black box of demand occurring mechanisms as it observes some variables that are predictors of the demand and produces predictions that can adjust to dynamic market trend changes as they

happen. Unfortunately, as pointed out by Bacchetti et al., 'very specific information is needed' and 'this information is rarely available, except when the supplier is in charge of maintenance activities at the customer's, or when the customer itself forecasts its spare part requirement' (2012, p. 726).

Another extremely powerful set of forecasting techniques is the expert forecasts. The expert methods are structured techniques to generate and combine predictions from subject matter experts. The most common expert methods include the three-point estimate, allowing an estimate of a distribution only based on maximum, minimum, and median observation estimates, or the Delphi forecast method, allowing an efficient method of combining opinions of relevant stakeholders despite possible vast differences in the predictions of individual experts. The biggest challenge in the expert forecasts and the reason they cannot be used in the case context is the lack of scalability that is necessary for forecasting 10,000 or more SKUs.

Forecasting methods using hybrid approach The final group of methods relies on a hybrid approach integrating additional data sources, such as condition monitoring data, into the spare-part forecasting process. It uses current condition information, usually collected from sensors monitoring the equipment and, based on a predefined link of those values to the expected replacement/breakdown date, creates an estimate of demand for monitored equipment. The block diagram of such a solution, an intelligent maintenance system, is presented in Figure 8 (as presented by Hellingrath & Cordes, 2014). This approach shares the disadvantages of solutions based on reliability analysis (data access limitation) and, additionally,

it requires a complex sensor infrastructure for equipment condition monitoring.
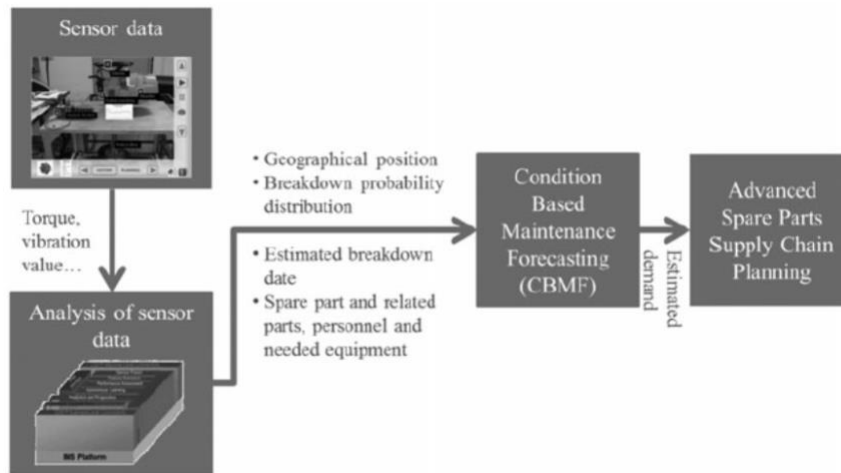
Figure 8 – Block diagram of intelligent maintenance systems, as presented by Hellingrath and Cordes (2014, p. 731).

To summarise, the methods evaluated in this section can be divided into those extracting advanced demand information (methods based on early sales and equipment failure rates, if they are computed based on historical data) and those that require information related to the equipment's condition or overhaul arrangements. The initial approach is generic and easy to implement, but its prediction quality might deteriorate in environments undergoing transformation or that are generally characterised by a high level of uncertainty. On the other hand, the latter approach requires infrastructure to monitor equipment condition and/or additional information related to pieces of equipment in use, and processes related to spare-part replacements. Unfortunately, as pointed out by Bacchetti et al., the OEMs do not usually have access to this information, making this approach more feasible for equipment owners (2012, p. 726).

Summary of explanatory methods

2.1.3 Summary of findings related to spare-part forecasting

Generally, when considering the degree to which the demand for forecasting methods reveals mechanisms causing this demand, this can be divided into two groups. The predominant group consists of generic algorithms, utilising historical demand data (as depicted in). The methods present in this cluster could not determine the demand for a black box, with their predictions based directly on previous observations, such as a simple moving average, neural networks, or exponential smoothing. Alternatively, they could explicate the black box by extracting some characteristics from historical data that could help to explain the underlying mechanisms, such as Croston-based methods in which considering the interval between sales helps to compute a reasonable forecast for parts with low sale frequency. The greatest strength of these methods is that they can easily be adjusted to almost any context; the only requirement is to accumulate some history. On the other hand, as prediction is only based on historical values, any external change, such as a global market trend inversion or changes in customer behaviour, will never be predicted and can only be discovered long after they occur.

The alternative to the generic demand for forecasting algorithms is explanatory methods. They explain the black box regarding the demand-occurring process and observe some variables that are predictors of this demand, such as the condition that equipment should be serviced or have overhaul arrangements preceding the spare-part replacement job and that predictions should adjust to dynamic market trend changes as they happen. Unfortunately, this also makes

Spare part forecasting methods as time series or explanatory

them very context specific and hence very difficult to re-implement in a different environment.

In addition, as pointed out by Bacchetti et al. in their work in 2012, the information is rarely available to the supplier, except when the supplier is responsible for maintenance. Although this literature review was intended to be about explanatory methods, only a handful of examples were found in the literature, as opposed to the extremely popular generic methods. Clearly, the lack of popularity of this kind of solution can, to a substantial extent, be attributed to limited access to explanatory variables by suppliers and to the difficulty and cost of developing application-specific solutions.

	Generic Methods (based only on previous observations)		Explanatory Methods (beyond previous observations)
	Simple Demand Methods	**Advance Demand Methods**	
Explanation	Prediction based directly on previous observations	Prediction based on characteristics extracted from previous observations	Prediction based on variables that are predictors of demand (like equipment condition information or overhaul arrangements)
Example	Exponential smoothing, neural networks	Croston, early sales, failure rate analysis, bootstrap	Integrated forecasting method, intelligent maintenance system, installed base (IB) forecasting

Table 8 - Overview of forecasting methods in OR literature.

The initial barrier seems to be disappearing with the rapid increase in popularity of servitisation strategies. The overall finding from the OR literature review is that the climate

change surrounding explanatory methods caused by servitisation will create a more feasible solution for many spare-part demand-forecasting problems, with the final gap to be closed being the difficulty and cost of developing application-specific solutions.

2.2 Servitisation and New Forecasting Opportunities

In recent years, we could observe a strong trend among manufacturing companies to shift from selling produced goods towards embedding them in services related to their operations. This strategy is often called servitisation, developing product-service systems (PSS) or, more generally, solution-selling. The most well-known example of successful servitisation adoption could be IBM: The company that used to be the biggest computer manufacturer in the world has restructured its turnover in such a way that more than 90% of it came from non-manufacturing activities in the fiscal year 2011 (Ahamed, Inohara, & Kamoshida, 2013).

What is servitisation?

Many market-leading companies, including GE, Rolls-Royce, and Siemens, or in the Danish market Vestas and FL Smith, have been transforming in a similar direction. 'This value chain concept, first introduced by Vandermerwe and Rada in the late 1980s, is now widely recognized as the process of creating value by adding services to products' (Ahamed et al., 2013). In the context of this work, servitisation causes a paradigm shift, as it destroys the distinction between the OEM and the equipment owner, giving the OEM access to data and infrastructures they were not able to explore beforehand.

The main reason that servitisation is gaining momentum in the industry seems to be the expectation that such a move will increase a company's financial performance. This anticipation seems quite feasible; servitisation allows the maintenance of a regular revenue stream for the entire lifetime of the equipment, while the traditional product-selling strategy offers only a one-time certain income.

Motivations to apply servitisation principles

Moreover, for market leaders, there is a strong competitive argument. Services are much more difficult to imitate than products; therefore, they can become a sustainable source of competitive advantage (Heskett, 1997; Oliva & Kallenberg, 2003). Such a radical change, however, has many implications; some of them being somewhat undesirable.

The transformation of company's mind-set is the biggest aspect that can go wrong; the starting point is the manufacturer of a product, the quality of which speaks for itself. They do not put much effort into customer relationship management, which results in customers being unfamiliar with them. The intended outcome is an image of an excellent consultancy bureau, in which the product is merely an addition to expert knowledge, adding value for customers. Failing to accomplish such a transformation will have a significant chance of resulting in neither gaining a service market nor retaining the manufacturing profit.
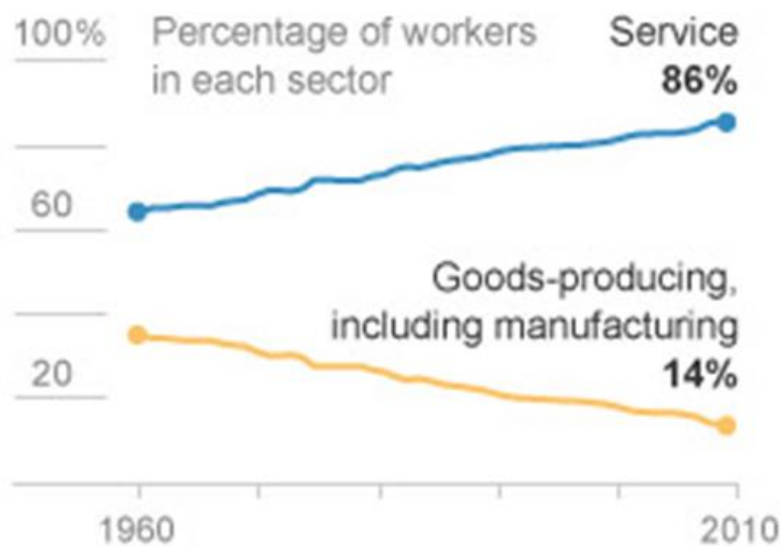
Figure 9 – Trend line of employment in the USA according to the sector. Source: NYT (Duhigg & Bradsher, 2012).

On a global scale, the trend in Western employment structures is clear; while employment in services is constantly growing, fewer people are working in goods production, including manufacturing (for data presenting employment in the USA, see Figure 9). The change can partially be attributed to outsourcing and offshoring trends, but it also clearly shows that servitisation is not just a theoretical construct but is also a unique real-life phenomenon that can have serious implications in the lives of millions of people.

As pointed out in a recent *NYT* article (Duhigg & Bradsher, 2012), following the predictions of the US Bureau of Labor Statistics, 1,000 manufacturing jobs in the US create about 4,710 other jobs for the economy, while the same number of service jobs results in only 700 other jobs. This raises questions about the general, long-term consequences of the massive change that we are witnessing. Embedding manufactured products in value-adding services can give Western manufacturers a chance to compete with their competitors from the Far East, despite their labour cost disadvantage.

Macro-view on the servitisation landscape

Production as we know it is disappearing from Western countries (Duhigg & Bradsher, 2012). Creating product-service bundles is one of the few options for manufacturing companies in Europe, the United States, and many other regions with elaborate production facilities. Developed regions such as Europe and the US are struggling to stay competitive to maintain high living standards for millions of people. The Western world simply cannot compete on the cost of labour; the only argument to justify employing more expensive Western employees is their knowledge and experience. Offering well-designed product-service bundles

Servitisation as a way to manufacture in the developed world

changes the markets in which they are offered (Ahamed et al., 2013; Baines, Lightfoot, Benedettini, & Kay, 2009), offering a unique opportunity to separate labour-intensive manufacturing activities that can be outsourced from knowledge-intensive design and service activities that can be kept.

2.2.1 What are the benefits of servitisation adoption?

In the literature, the key implications of servitisation adoption that were perceived as positive for manufacturers are the drivers of a financial, strategic, and marketing nature. This statement was made by Baines et al. (2009) in their literature review covering work until 2008 and is also confirmed in more recent literature (see the references below). The financial drivers most frequently mentioned in the literature are higher profit margins and stability of income (Baines et al., 2009; Gebauer & Friedli, 2005; Wise & Baumgartner, 1999) as well as the cost reduction achieved through economy of scale utilisation and 'service gain' exploration (Colen & Lambrecht, 2013), as shown in Figure 10.

Positive implications of servitisation adoption

Later studies add examples of such effects, as in the work by Visnjic Kastalli and Van Looy (2013), showing a positive cubic relationship between service sales and profitability. Among strategic drivers, the service component is considered one that can differentiate the offering, leading to gaining a competitive advantage (Barnett, Parry, Saad, Newnes, & Goh, 2013; Gebauer & Fleisch, 2007; Mathieu, 2001b; Oliva & Kallenberg, 2003). Marketing factors introduced when diversifying the portfolio by introducing services are related to using services to increase product sales (Gebauer & Fleisch, 2007; Mathieu, 2001b; Visnjic Kastalli & Van Looy, 2013), creating customer loyalty through services (Baines et al., 2009; Corrêa, Ellram, Scavarda, & Cooper, 2007; Vandermerwe & Rada, 1988), and gaining more insight into customer needs, which allows manufacturers to tailor their solutions and to improve the relationship with customers. From the perspective of this research project, one more gain from applying a servitisation-based strategy is gaining

unlimited access to data describing the use of equipment, enabling more precise explanatory forecast analytics.

When discussing negative consequences of servitisation adoption, the literature seems to be much more scattered. Fang, Palmatier, and Steenkamp (2008) pointed out the loss of strategic focus caused by introducing services in the offering and the organisational conflict introduced by differences in processes, cultures, leaderships, and structures (Deshpandé, Farley, & Webster, 1993; Fang et al., 2008). Neely (2008) established that service manufacturers obtain lower net profits as a percentage of revenue, attributing it to, among others, the higher cost of knowledge-intensive employees in the service industry.

Many other studies examined the challenges faced rather than the negative implications of the servitisation adoption. Baines et al. (2009), in their synthesis of academic contributions, listed the following three aspects: difficulties related to designing value-adding services (Oliva & Kallenberg, 2003; Vandermerwe & Rada, 1988), preparing strategy that consistently supports services (Gebauer & Fleisch, 2007; Mathieu, 2001b; Wise & Baumgartner, 1999), and an organisation's cultural transformation (Mathieu, 2001a, 2001b). Although it is difficult to present a clear agreement among researchers, later studies such as that by Barnett et al. (2013) seem to confirm these three factors as dominant challenges for manufacturers adding services to their offerings.
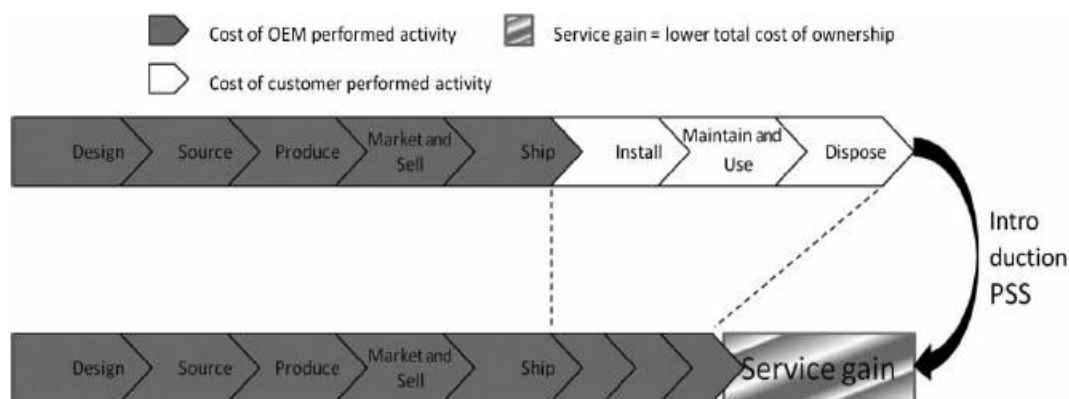
Figure 10 – Service gain or value added for customers by introducing product-service systems (PSS), as depicted by Colen and Lambrecht (2013).

2.2.2 Classification of servitisation

Clearly, the meaning of the adoption of servitisation depends heavily on the kind of services introduced into the portfolio. One of the classic classifications of servitisation efforts is Mathieu's (2001b) split between services supporting products (SSPs), sometimes called 'product-centric services' (Baines et al., 2009), and services supporting customers (SSCs).

Oliva and Kallenberg (2003) introduced another dimension into this picture, namely, the interaction with customers, specified as transaction- or relationship-based. With respect *Classification* to this interaction, the starting point was usually a *dimensions in* transaction-based service, such as individual repairs or *servitisation* upgrades, which could lead to relationship-based services, *activities* such as complete spare-part management or fixed maintenance contracts.

I argue that, when discussing the implications of servitisation adoption for manufacturing companies, the meaningful

63

dimension is the latter dimension, namely, the sustainability of interaction with the customer, as only relationship-based interactions affect the factors listed above. The implications, in contrast to transaction-based interactions, are that they introduce more knowledge about customers that can lead to developing more value-added products, optimising internal resources, and creating customer loyalty. Relationship-based contracts lasting longer than a single transaction are also a great way to stabilise revenue streams.

2.3 Decision Support Systems

Although DSS work draws on multiple management disciplines, including operation management, the foundation of the work and the main intended contributions are anchored in the domain of the management of information systems. To ensure the contributions are in line with the cumulative tradition of IS, a thorough review of the related IS body of literature is provided.

Decision support systems as a predecessor of design research

As a starting point, I identified that, in the IS domain, systems supporting organisations in tasks such as demand forecasting are normally called DSSs (Alter, 1980), and they have been present in the literature for more than 40 years. As a point of departure, I will review this work as a typical IS scope of topics within my domain of interest, which is predictive analytics in IS. While reviewing the relevant literature, I will look for trends reflecting a clear difference in maturity between research areas like servitisation or spare-part forecasting and DSSs.

Methodology of DSS literature review

There are significantly many more sources on DSSs; therefore, conducting a review without a structured framework could lead to presenting a skewed picture. To systematically explore the richness of DSSs, I used a systematic literature review approach (Webster & Watson, 2002). I reviewed articles published in the IS basket of eight, as I assumed that most of the seminal IS work should be present there, and I searched for those having the keywords 'decision making' or 'decision support systems', as presented in the EBSCO Business Source Complete database.

Clearly, limiting the journal set to only eight influential journals leaves out sources such as *Information and*

Management, the *Harvard Business Review*, or all journals dedicated to DSSs, but my intention here is not to discuss all DSS-related contributions, but rather to present an outline of the development of the field and to break it down into certain trends. The initial search produced 511 papers. Overall, DSSs have been very popular in many influential journals, and the trend seems to have declined in the early 2000s but picked up again in 2005. The analysis presented below will concentrate on the summary of concepts, as instructed by Webster and Watson (2002). There is also an analysis concentrated on trends in DSS literature presented in the appendix.

Scope limitation and initial evaluation

I decided not to consider work published before 1990 to limit the number of articles and to not dwell on concepts that were older than 25 years. The articles were then annotated to three relevant clusters based on the main object of theorising: theories about decision making, about generic DSSs, or about a particular decision. Those not explicitly theorising about any of the three were grouped as 'not related' and will not be presented in the analysis below. After excluding 103 papers published before 1990, 211 for not being relevant according to the object of theorising, and nine because they were duplicate entries, I reviewed 188 papers that were assigned to one of three relevant clusters (see Table 9).

Table 9 – Overview of papers according to the object of theorising.

Category	Particular decision	Theorising about Decision Making	Theorising about DSS
Count of papers	**123 (65%)**	**11 (6%)**	**54 (29%)**
Examples	(Keil et al., 2000; Miranda & Yong-Mi Kim, 2006; Xiao & Benbasat, 2007)	(Lerch & Harter, 2000; Todd & Benbasat, 1992; Zhang, 2013)	(Nissen & Sengupta, 2006; Silver, 1991; Turrof, Hiltz, Bahgat, & Rana, 1993)

Clearly, the most common among the three clusters is the one discussing a particular decision context as a topic of

developed theory, corresponding to about two-thirds of the published work. Decision support systems are very strongly anchored in the system development community. Hence, it is very common to see DSS projects that include an element of system development. An interesting observation is that very few authors, only about one in 20, extend their theories to the decision-making process, going beyond the particular decision context or a given DSS.

The three classes presented above are very broad. To present more insight into what the research projects were about, a detailed clustering was performed. Theorising about decision making was broken down to two groups, namely, the influence of DSSs on the decision-making process and the effect of various cognitive factors on the decision-making process. For theorising about DSSs, three meaningful groups were formed: the overall understanding of DSSs, understanding of DSS use, and influence of various factors on DSS performance. In the cluster of particular decisions, the seven most common decisions were extracted (e.g., outsourcing, group decision making, and e-commerce), and the rest were grouped in the 'other' category. The top 10 detailed clusters by popularity are presented in Table 10.

The detailed pictures reflect the popularity of particular DSSs in the last 25 years (e-commerce, outsourcing, and group support systems) as well as three aspects of theorising about DSSs, namely, the influence of various factors on DSS performance, understanding DSS use, and DSS implementation tips. This method of presenting DSS literature confirms the limited popularity of theorising about decision making using DSSs. None of the three detailed clusters of theorising about decision making appeared in the top 10 list.

Table 10 – Top 10 detailed clusters by popularity.

Detailed Cluster	Relation to Decision Making	Count	Example
Group decision making	Particular decision	33 (18%)	(Dickson, Patridge, & Robinson, 1993; Sengupta & Te'eni, 1993)
Influence of various factors on DSS performance	Theorising about DSS	22 (12%)	(Speier & Morris, 2003; Ye & Johnson, 1995)
Other particular decisions	Particular decision	21 (11%)	(Aggarwal & Singh, 2013; George, Carlson, & Valacich, 2013)
Technology adoption	Particular decision	18 (10%)	(Duan, Gu, & Whinston, 2009; Sun, 2013)
Understanding of DSS use	Theorising about DSS	14 (7%)	(Belcher & Watson, 1993; Gill, 1996)
E-commerce	Particular decision	12 (6%)	(Kamis, Koufaris, & Stern, 2008; Xiao & Benbasat, 2007)
Outsourcing	Particular decision	12 (6%)	(Ang & Straub, 1998; Miranda & Yong-Mi Kim, 2006)
DSS implementation tips	Theorising about DSS	10 (5%)	(Kasper, 1996; Sinha & May, 1996)
Investment decisions	Particular decision	10 (5%)	(Abdul-Gader & Kozar, 1995; Taudes, Feurstein, & Mild, 2000)
Information management	Particular decision	10 (5%)	(Balakrishnan & Whinston, 1991; Parssian, Sarkar, & Jacob, 2009)
Other	Other	26 (14%)	(Silver, 1991; Zhang, 2013)

Predictive analytics in DSS literature

In parallel to clustering according to the objective of theorising, the body of DSS literature was also annotated according to its use of predictive methods. All papers having a reference to using a predictive method in the abstract were marked as using predictive methods. Surprisingly, only 6 out of 399 papers were marked as using predictive methods; three theorise about a particular decision and another three did not theorise about DSSs despite having a DSS-relevant keyword. The six relevant papers are presented in Table 11.

Table 11 – DSS papers using predictive analytics.

Relation to Decision Making	Details Clusters	Papers
No relation		(Bansal, Sinha, & Zhao, 2008; Brown, Gatian, & Hicks, 1995; Zhang & Seidmann, 2010)
Particular decision	**Information management**	(Saar-Tsechansky & Provost, 2007)
Particular decision	**Other particular decisions**	(Baars, Gille, & Strüker, 2009; Martens & Provost, 2014)

Although the six papers using predictive analytics are clustered around the most resent years in the period of the study, indicating the increasing interest around that topic, it is challenging to discuss a stream of literature with only six identified publications. Additionally, all the mentioned papers use predictive analytics as the data transformation technique not as an object of study. Combining these observations, despite the increasing popularity of predictive analytics in the DSS community, I found no work using DSSs to theorise about predictive analytics. There also seems not to be a distinctive research stream about predictive analytics within the DSS literature.

Looking for predictive analytics research steam

2.4 Decision Support in the Big Data Era

In the last few years, the volume of data has exploded. As reported by Brown, Chui, and Manyika (2011 p. 2), in 15 of 17 US economy sectors, companies with more than 1,000 employees store 235 terabytes of data on average, which is more than the entire collection of the US Library of Congress. In addition, in scholarly publications, big data analytics have been attracting a lot of attention. As reported by Chen, Chiang, and Storey (2012) in their introduction to the MISQ special issue on big data, the number of publications with big data as a keyword doubled every year between 2006 and 2011 (Figure 11). This research project could be placed at the heart of this wave (Pedersen, Furtak, Häuser, Lauth, & Van Kranenburg, 2013), confirming that the research is timely and that its findings might be useful for the rising stream of IS works.

Keyword	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
Business Intelligence	113	104	146	159	229	330	346	394	352	201	334	338
Business Analytics	0	5	43	4	5	2	9	6	19	16	17	126
Big Data	0	1	0	0	7	4	3	26	11	41	44	95
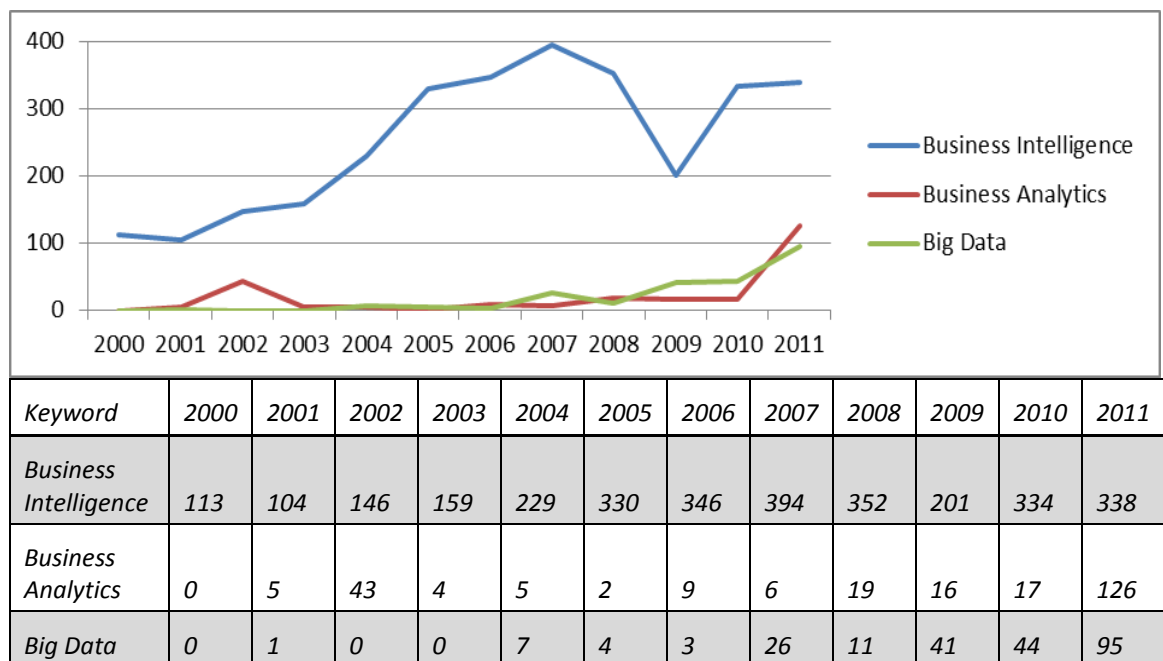
Figure 11 – Popularity of keywords related to business intelligence research (Chen et al., 2012).

2.4.1 How big data brought sensors and context into IS research and artefact design

The terms 'context' and 'context-aware' are widely used in the literature of many fields related to IS, from computer science (Dey, 2001), through human-computer interaction (Nardi, 1996) and sense-making (Narock, Yoon, & March, 2012) to general management literature (Johns, 2006), although the exact characterisations seem to vary. The definition proposed in the introduction (Cook & Das, 2004) is rooted in the framing by Dey, who defined context as 'any information that can be used to characterize the situation of an entity' (2001).

When considering technological artefacts, context awareness implies using sensors (Pedersen & Pedersen, 2014) or digital traces (Karanasios et al., 2013) from other IS to collect information characterising the operational environment. Generally, following an observation of Narock et al., 'information systems are typically confined to a single context and a single perspective of interpretation' (2012 p. 172), which is a direct result of the 'context specificity of IT artifacts' (Orlikowski & Iacono, 2001, p. 131). Nevertheless, for IS artefacts embedded in real organisations, it is increasingly crucial to be able to respond robustly to the constantly changing environment surrounding them (Ploesser, Recker, & Rosemann, 2011). The well-known examples of such operational-level attempts could be using RFIDs for triggering actions based on a product flow (Welbourne et al., 2009) and, more strategically, on the popularity of BI&A, both in the industry and academia.

Context awareness of IS artefacts

The technical development in storage capacity, as well as the increased popularity of digital traces and sensor data, has led

to an explosion of the volume of data, often referred to as the dawn of the big data era (Chen et al., 2012). The fact pointed out in the previous section regarding the amount of data stored by US economy sector companies is a direct result of that.

Chen et al. (2012) pointed out that, historically, BI&A can be presented in three stages. First, BI&A 1.0 refers to the data-centric approach found in data management and warehousing and is concerned with relational databases, popularised by Chen et al. in the 1990s. Second, BI&A 2.0 focuses on text and web analytics for unstructured web content, gaining traction in the early 2000s in relation to the development of the Internet. Third, the recently acquired technological capabilities and increasing business needs can result in emerging opportunities related to BI&A 3.0, understood as mobile and sensor content, also known as the IoT (Pedersen et al., 2013).

Following the emergence of BI&A 3.0 (big data and the IoT) and the increasing interest in the robustness of IS artefacts in dynamic organisational settings, automatic context inclusion in IS artefacts has been of current interest for researchers in the IS community. In their recent work, Ploesser et al. (2011) discussed the challenges of embedding context into business processes, suggesting business process conceptualisation as complex-adaptive systems and introducing a set of prescriptions for their design (Ploesser, 2013).

In contrast, Narock et al. (2012) discussed the role of context and subjectivity on scientific IS, introducing a framework based on multiple perspectives. In addition to the mentioned lenses, Hong, Chan, Thong, Chasalow, and Dhillon (2014)

discussed the effect of context on the process of theory generation in IS and put forward guidelines for developing context-specific models. Altogether, despite the agreement that developing a context-aware IT artefact is difficult (Narock et al., 2012; Ploesser et al., 2011), IS as a community does not offer much assistance to practitioners in terms of how to design such IS artefacts.

2.4.2 Summary: Literature about decision making in the project context

The overall conclusion from reviewing the literature related to DSSs is that, despite the high maturity of the decision support research area, the current stream of scholarly work is rich and is not declining. Within the stream, it is clear that the

Expected development of mobile and sensor data, IoT

focus is shifting towards research linked to BI&A (Figure 11), predicting dynamic growth, especially in the area of BI&A 3.0 (mobile and sensors), sometimes also called context-aware systems. In addition, it seems that the scope of the projects is increasingly focused on a decision (see Figure 44). In the context of the project, these observations lead to the conclusion that a useful framework to rigorously develop a DSS using mobile and sensor content in the context of a particular decision could gain popularity in the DSS research community.

2.5 Design Science Research

Design science research (DSR) 'seeks to create innovations (...) through which the analysis, design, implementation, management, and use of information systems can be effectively and efficiently accomplished' (Hevner et al., 2004 p. 76). It aims at '"utility", i.e., at the construction and evaluation of generic means–ends relations' (Winter, 2008 p. 470), while theory building is still considered a central *What is DSR?* activity within DSR (Venable, 2006). The most commonly accepted reference process model for conducting DSR comes from Hevner et al. (2004) and presents the process as a two-state cycle: develop/build and justify/evaluate, informed by the environment and the knowledge base elements.
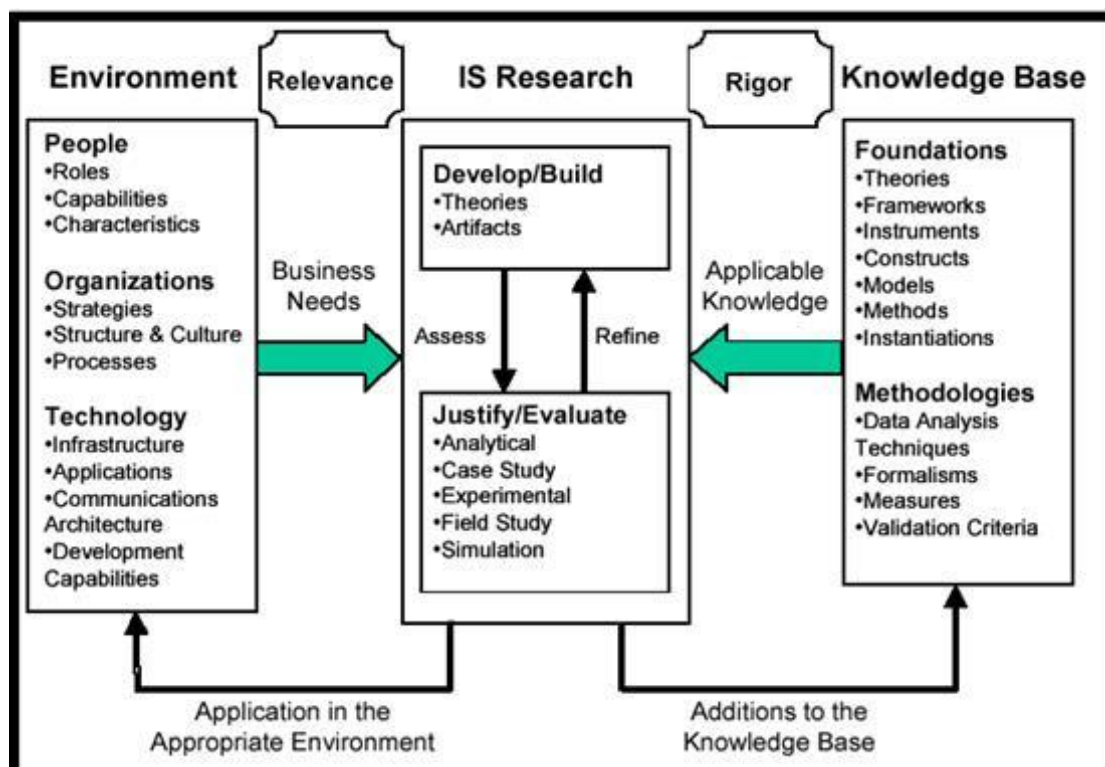
Figure 12 – DSR research framework (Hevner et al., 2004).

Although the DSR process model provides a solid foundation for the design process, it is very generic. To obtain more

specific guidelines for the design process, further exploration of the IS knowledge base needs to be conducted to find a model to combine with DSR.
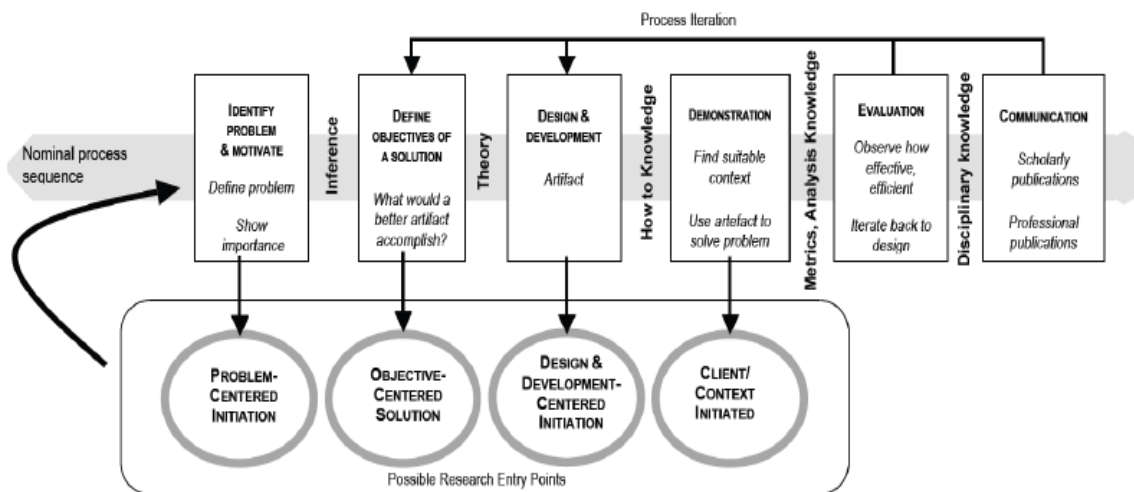
Figure 13 – Design science research process model (Peffers et al., 2007).

Design science has also been successfully applied as a research methodology. In its most common version, the DSR process model by Peffers, Tuunanen, Rothenberger, and Chatterjee (2007) is a transformation of the original guidelines by Hevner et al. (2004) into steps in a process model. In the work by Peffers et al. (2007), the process model is introduced as 'a nominal process model for conducting DS research in IS', depicted in Figure 13. It consists of six steps

DSR as a methodology covering the entire research project activity spectrum from motivations to knowledge dissemination. In addition to the sequential process steps, the model is also equipped with four possible research entry points, allowing researchers to apply the methodology in diverse conditions and with multiple intentions and perspectives. At present, after only about seven years, the original publication of the methodology gained almost 1,300 citations (according to Google Scholar),

making it arguably one of the most common research methodologies in recent IS scholarly work.

The methodologies introduced by Hevner et al. (2004) and Peffers et al. (2007) have attracted a lot of attention that led to extensive feedback on their use. One of the most important observations on the DR methods, made by Sein et al. (2011), is that awareness of the problem precedes the development of the artefact in most DR efforts, which is then followed by evaluation (Vaishnavi & Kuechler, 2006). This prevailing approach of DSR is captured in the 'build and *then* evaluate' cycle proposed by March and Smith (1995) and advocated by Hevner et al. (2004). Thus, DSR places axiological emphasis on utility (Venable, 2006) traced to the problem identified at the beginning of the research project (Sein et al., 2011, p. 3).

Post Hevner and Peffers DSR

Sein et al. (2011) called this 'the problem of sequencing and separation' and it is really an observation that the sequence in which the actions of problem formulation, build and evaluate, are too rigid and do not match reality. To address this shortcoming, they decided to take from the action research paradigm (Baburoglu & Ravn, 1992; Baskerville & Wood-Harper 1998), creating action design research (ADR). The model of the ADR stages and principles is presented in Figure 14. You can see that it is less sequential, allowing for various cycles in the process of designing.

Action Design Research
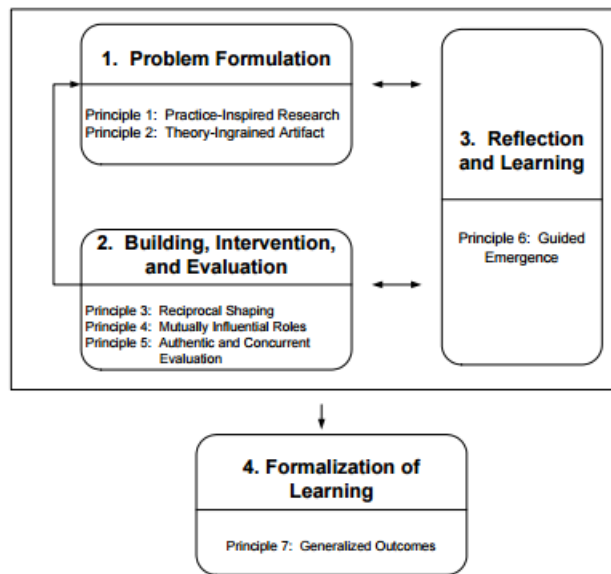
Figure 14 – ADR stages and principles (Sein et al., 2011).

2.5.1 Requirements for high quality design science research

According to Hevner et al. (2004), the fundamental principal of design science is that knowledge and understanding are acquired in the building and application of an artefact. This understanding led them to the development of seven DSR guidelines (G) to 'assist researchers, reviewers, editors and readers to understand the requirements of effective design-science research' (Hevner et al., 2004, p. 82). According to the guidelines, DSR requires an artefact (G1) for a specific problem (G2). The artefact must be evaluated to be proven useful for the problem (G3), providing an improvement or novelty in the approach (G4). The artefact must be rigorously defined (G5) in the process in which a search space for a solution is created and a mechanism to choose the most optimal one is introduced (G6). Finally, the output of the research must be communicated clearly and effectively (G7). The descriptions of the guidelines as presented by the authors are listed in Figure 15.

DSR guidelines

Guideline	Description
Guideline 1: Design as an Artifact	Design-science research must produce a viable artifact in the form of a construct, a model, a method, or an instantiation.
Guideline 2: Problem Relevance	The objective of design-science research is to develop technology-based solutions to important and relevant business problems.
Guideline 3: Design Evaluation	The utility, quality, and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods.
Guideline 4: Research Contributions	Effective design-science research must provide clear and verifiable contributions in the areas of the design artifact, design foundations, and/or design methodologies.
Guideline 5: Research Rigor	Design-science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact.
Guideline 6: Design as a Search Process	The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment.
Guideline 7: Communication of Research	Design-science research must be presented effectively both to technology-oriented as well as management-oriented audiences.

Figure 15 – Design science research guidelines (Hevner et al., 2004).

I present the guidelines because they define what constitutes a good design science project. To ensure that this project, conducted in the spirit of DSR, fulfils all quality requirements for an effective DSR project, I will refer as much as possible to the guidelines in the process of the development of the IS artefact. In addition, after the process is concluded, I will re-evaluate the outcome and the process using the guidelines to show that all the guidelines have been followed.

Another interesting observation following the initial use of the DSR methodologies is that the key to a rigorous and useful DSR project is the evaluation step (i.e., the step in which the relevance, rigour, and relation to all other guidelines is discussed). It is also the most difficult step for most researchers, as the guidance for evaluating design research is insufficient for a complete evaluation. It is this observation that led John Venable, Jan Pries-Heje, and Richard Baskerville to develop a framework for evaluation in DSR (FEDS; see Figure 16). Alongside FEDS, they introduced a four-step evaluation design process, consisting of the following steps: explicate the goal, choose a strategy or strategies for evaluation, determine the properties to evaluate, and design the individual evaluation episode. The work by Venable et al. (2016) can be considered further specification of the evaluation step in all previously mentioned DSR methodologies.
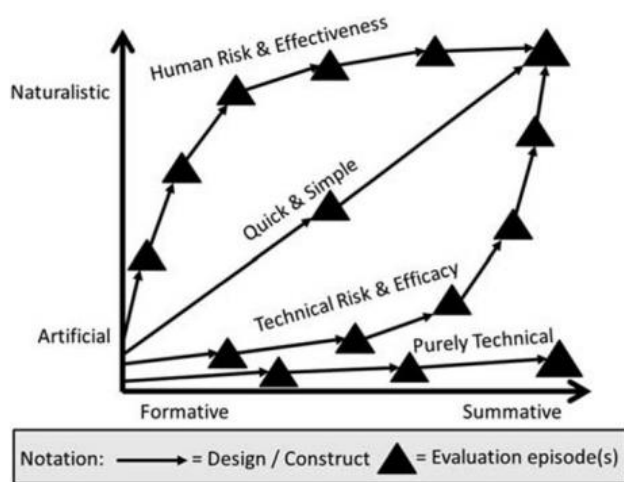
Figure 16 – Framework for evaluation in design science (FEDS) with evaluation strategies (Venable et al., 2016).

2.6 Predictive Analytics in IS

In their seminal work from 2011, Shmueli and Koppius critically evaluated IS research regarding the correct use of the predictive analytics empirical methods and highlighted how rare it is in IS. They defined clear distinctions between empirical models for explanation and for prediction and, acknowledging the strengths of explanatory model, showed the need to enhance the use of predictive methods in IS. The key differences between explanatory and predictive models are summarised in Figure 17.

Step	Explanatory	Predictive
Analysis Goal	Explanatory statistical models are used for testing causal hypotheses.	Predictive models are used for predicting new observations and assessing predictability levels.
Variables of Interest	Operationalized variables are used only as instruments to study the underlying conceptual constructs and the relationships between them.	The observed, measurable variables are the focus.
Model Building Optimized Function	In explanatory modeling the focus is on minimizing model bias. Main risks are type I and II errors.	In predictive modeling the focus is on minimizing the combined bias and variance. The main risk is over-fitting.
Model Building Constraints	Empirical model must be interpretable, must support statistical testing of the hypotheses of interest, must adhere to theoretical model (e.g., in terms of form, variables, specification).	Must use variables that are available at time of model deployment.
Model Evaluation	Explanatory power is measured by strength-of-fit measures and tests (e.g., R^2 and statistical significance of coefficients).	Predictive power is measured by accuracy of out-of-sample predictions.

Figure 17 – Key differences between explanatory and predictive models (Shmueli & Koppius, 2011, p. 557).

Shmueli and Koppius (2011) see six distinctive roles for predictive analytics in scientific research: 1) generating new theories, 2) developing measures, 3) comparing competing theories, 4) improving existing models, 5) assessing relevance, and 6) assessing predictability. They stress different possibilities that predictive analytics give, in comparison to explanatory analytics, presenting how it can use a new dataset for discovering patterns facilitating theory generation 1), improving existing models 4) or allowing new construct operationalisation 2). It exhibits predictive

analytics as a tool to compare the predictive power of competing theories 3), and to assess their relevance 5) and predictability 6). Those six use cases strengthen the importance of predictive analytics in scientific research and exemplify how it can be used to develop strong and informative theories.

To facilitate diffusion of predictive methods in IS, Shmueli and Koppius (2011) provided a model for building an empirical model and explicit guidelines on how to execute it for designing predictive models (Figure 18). It presents an empirical research as comprising eight steps. For predictive models, the initial one defines what it is that needs prediction. Next, some important issues regarding data collection need to be addressed, including using experimental versus observational settings, choosing data collection instruments and the size of the sample, or choosing candidates for observed variables. The next step is defining actions for data quality (DQ) issues, including missing values, and choosing the partitioning strategy. After that, data need to be evaluated to define variables for the analysis. When this is concluded, a data transformation method needs to be selected, along with an evaluation strategy for model selection. Finally, the strategy for research dissemination needs to be chosen and executed.
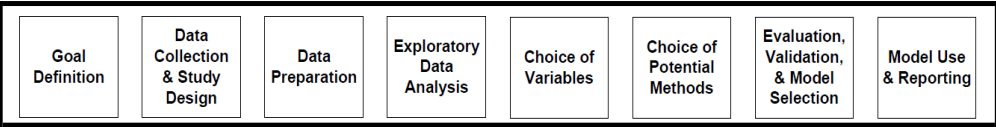
Figure 18 – Steps for building empirical models (Shmueli & Koppius, 2011).

Comparing this framework to the DSR model by Hevner et al. (2004) reveals many similarities: the evaluation and

validation step in the model matches the justify/evaluate step from the DSR almost exactly, and the five preceding steps in the model can be considered a more detailed, application-specific version of the develop/build step of DSR. The original model by Shmueli and Koppius (2011) is not cyclic, as it considers one iteration of the predictive design; however, envisioning multiple iterations, I can imagine cyclic arrows pointing from the data collection step to the evaluation step and back, which is similar to the logic of DSR.

In summary, the model seems to be a more specific guideline for a single develop/build–justify/evaluate cycle of the DSR model by Hevner et al. (2004). This observation underlines the foundation of the merging of the two models to structure a multi-iterator design of predictive IS artefacts. For a more comprehensive review of the usage of predictive techniques in IS, I recommend Shmueli and Koppius' review (2011).

2.7 Chapter Summary

This chapter summarises relevant literature in the spirit of the design science paradigm and follows the recommendations by Gregor and Hevner (2013) for the literature review section. In compliance with their recommendations, I summarise the justificatory knowledge, the kernel theory of my domain, which is spare-part forecasting, stressing the opportunities arising from new business practices (servitisation) and technological changes (sensor popularity and big data).

Chapter summary

The scholarly context of the class of problem (introduced in the previous chapter) is delineated via the introduction of DSSs and their successors, and the emerging importance of the class is shown by establishing the popularity growth of sensor technology and data acquisition techniques. The chapter concludes by presenting frameworks already established in the literature that can be utilised to guide the rigorous process of the development of predictive information systems, which will be the foundation of developing a sound methodological approach in the following chapter.

3 Research Design

3.1 Introduction

Research rigour is a driving goal of research approach selection (Gregor & Hevner, 2013). Gregor and Hevner defined research methodology as 'the research approach that was employed' (2013, p. 350). More broadly, Jonker and Pennink defined methodology as 'the way in which a researcher conducts research' (2009, p. 17). They offered a distinction between a more general term of research methodology and a more specific research method. Gregor and Hevner pointed out that a research approach should be built 'with references to existing authorities' (2013, p. 350). To satisfy this requirement, I decided to employ Jonker and Pennink's (2009) research pyramid framework to develop, explain, and justify the structuring of the research approach. The graphical presentation of the framework is presented in Figure 19. Although there are many alternatives to the pyramid, such as the research onion (Saunders et al., 2007), the overall outline of the model is not conceptually different; they include the same or similar components.

The authors describe the purpose of the pyramid framework as helping researchers to learn and concisely structure their approach to research. The pyramid is made up of four levels, namely, the research paradigm, the research methodology, the research method, and the research techniques. The units are designed to illustrate the abstraction level of the research approach from highest to lowest or from the most general descriptions to the most concrete and specific ones. The top level, the research paradigm, can be understood as 'underpinning values and rules that govern the thinking and behaviour of researcher' (Gummesson, 1999). Self-awareness

86

of the paradigm or paradigms used is important for any researcher, as it might influence his or her actions and thus introduce certain limitations to the overall scope and findings. Following a research methodology is explained as a system of methods and principals for doing something (Jonker & Pennink, 2009).
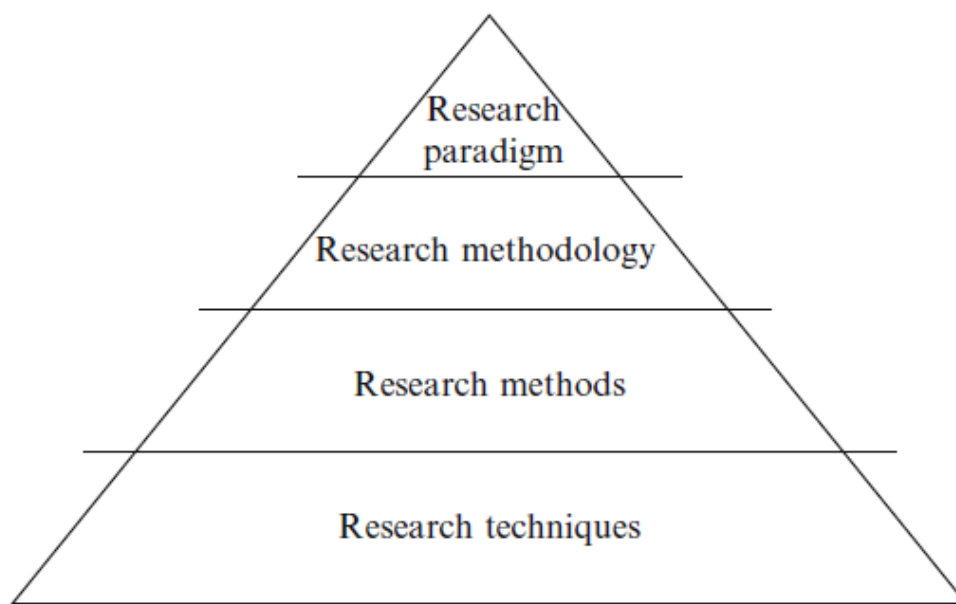
Figure 19 – Research pyramid (Jonker & Pennink, 2009, p. 23).

More visually, the framework's authors refer to methodology as 'not a map, but a domain' (Jonker & Pennink, 2009), reflecting on the fact that methodology is still quite abstract and does not provide exact 'direction', but rather an overall understanding. 'The map' is the research method, the pyramid's next level. The method describes how the research is performed, further specifying assumptions made at higher levels of the pyramid. The framework concludes with the research technique level, which is concerned with the instruments and tools used.

3.2 Research Paradigm

The research questions are aimed at defining the ways to design and the industrial practical settings of this research in the heart of the design research perspective, as described in Table 3 on paradigmatic assumptions. The practical angle requires an applied approach and the need to design enforces using a more 'developmental' methodology (see Table 3), which is provided by DSR. Clearly, this choice of DSR as a research paradigm can be perceived as a rather unorthodox one, knowing that, historically, it has sometimes been treated within IS with a note of scepticism (Galliers & Land, 1987).

Design science as a research paradigm

However, recent research seems to have established sound ontological and epistemological foundations for DSR, and has presented DSR as a paradigm (Iivari, 2007). Judging by the number of outlets currently accepting DSR (including the most traditional ones, such as MISQ and ISR), it seems that it has gained acceptance as a research paradigm. As March and Storey mentioned, 'Design science research is increasingly recognized as an equal companion of behavioral science research in the information system field' (2008, p. 726). With less controversy, several proponents of design science have suggested that it is associated with pragmatism as a philosophical ontology in its attempts to bridge science and practice (Cole et al., 2005; Hevner et al., 2004; Iivari, 2007; March & Smith, 1995).

Design science is an applied discipline (as opposed to 'pure' theoretical disciplines; see Iivari, 2007), solving real-life problems with a strong focus on academic rigour. The description of design science as a framework was presented in the previous chapter, with its original depiction by Hevner et al. (2004; see Figure 12).

3.3 Research Methodology

Research methodology is explained as a system of methods and principals for doing something and is the concretisation of the selected research paradigm (Jonker & Pennink, 2009). Design science research is also considered a research methodology, most commonly with the methodology presented by Peffers et al. (2007), as discussed in the previous chapter. Recently, design research projects carried out in a practical setting started to apply ADR more commonly. Although both of those methodologies would be feasible, they are generic and do not include specific aspects of design predictive systems. Moreover, as the problem at hand is merely a representation of a class of problems, there is an opportunity to adjust the methodology to popularise predictive information system development and facilitate the research process for other instances of the same class of problems. When investigating methodologies for IS research with predictive analytics, I came across a framework for building a predictive or explanatory empirical model artefact in IS by Shmueli and Koppius (2011).

The authors stated that their method is the only IS framework for predictive analytics introduced to address 'near-absence of predictive analytics in mainstream empirical IS research' (Shmueli & Koppius, 2011, p. 554) so there is no obvious alternative to it in IS. The use of this methodology alone could explore specific aspects of the class of problems under investigation, but it would not be consistent with the selected research paradigm, and it would be only loosely connected to the current stream of IS research, not building on the discipline's cumulative tradition. To mitigate the shortcomings of either of the approaches, I intend to combine

the DSR paradigm with Shmueli and Koppius's (2011) predictive analytics methodology.

3.3.1 Combining DSR with PAIS: Designing information systems with predictive analytics

Combining DSR and steps for building empirical models

In this section, I will analyse how DSR (Hevner et al., 2004) and steps for building empirical models by Shmueli and Koppius (2011) can be utilised to devise a framework for structuring a rigorous design process of IS artefacts via predictive models using sensor data. The DSR framework by Hevner et al. (2004) is intended to help to design any IS artefacts, while Shmueli and Koppius's (2011) framework was developed with building predictive empirical models in mind. As my overall goal, designing IS artefacts using predictive models includes the goals of both frameworks. Combining the two should fulfil my intentions, resulting in a more specific version of DSR for predictive analytics.
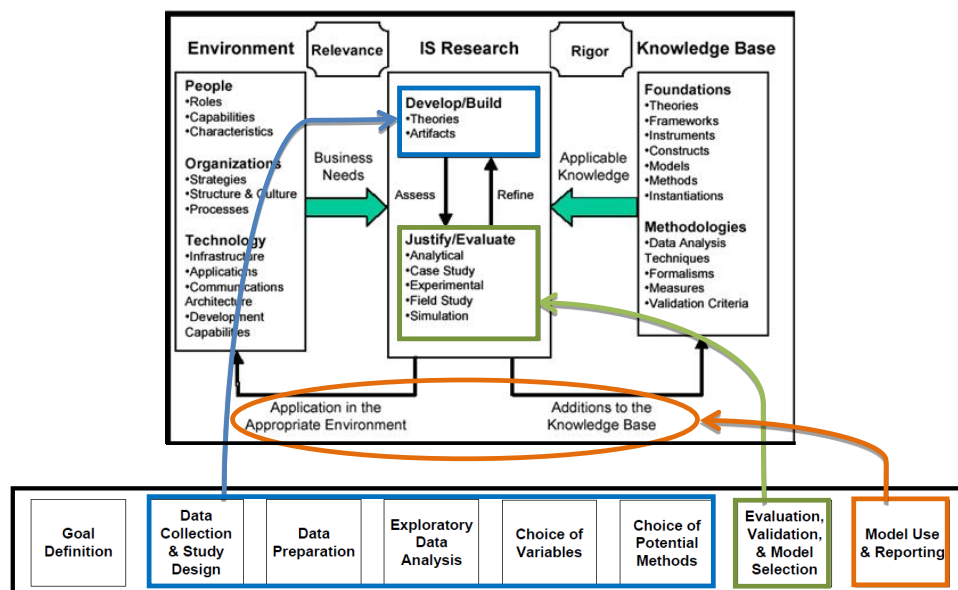
Figure 20 – Combining frameworks by Hevner et al. (2004) and Shmueli and Koppius (2011).

As a first step to combining the frameworks by Hevner et al. (2004) and Shmueli and Koppius (2011), we analyse how closely the respective suggested steps match. The initial step, goal definition, understood as defining the purpose of the design process and properties constituting a good design for that purpose, does not have an explicit match in the framework by Hevner et al. (2004). The following five steps, namely, 1) data collection and study design (initial design choices for the model, such as simulation versus experimental study, data collection strategy, and sample size), 2) data preparation, 3) exploratory data analysis, and 4) variable selection, and 5) choice of a predictive method, seem to be conceptually included in the develop/build step.

Nevertheless, in the context of IS artefacts using predictive models, the output framework could potentially benefit from defining at least some of them more specifically as sub-steps. Evaluation, validation, and model selection from the model by Shmueli and Koppius (2011) seem to correspond to the justify and evaluate step in the model by Hevner et al. (2004). Finally, the model use matches 'application in appropriate environment', and reporting corresponds to 'addition to the knowledge base' by Hevner et al. (2004). The graphical matching of the two frameworks can be seen in Figure 20.

Based on these observations, I started to construct the model for designing IS artefacts using predictive analytics. As a starting point I decided to explicitly include the previously missing goal definition step. In this context, the designer needs to answer questions at this step, such as what needs to be designed (including what needs to be predicted), and what makes a design suitable for the purpose. The next step to follow was develop/build by Hevner et al. (2004), but with

sub-steps inspired by the model by Shmueli and Koppius (2011).

I noticed that four steps from the model by Shmueli and Koppius (2011; data collection and study design, exploratory data analysis, choice of variables, and choice of potential methods) are very tightly coupled, lacking the required flexibility in step ordering. In some scenarios, the nature of available variables and methods heavily affect the study design, and the choice of a method might change the choice of variable. To avoid this ordering dilemma, I suggest structuring the develop/build step as three sub-steps: model definition, data preparation, and model implementation. Model definition is where the abstract model is presented. Data preparation is where data pre-processing for model implementation and/or evaluation is described, and model implementation details the actual implementation process. Specific to sensor data, in the second sub-step (data preparation), an investigation of the match between sensor-measured quantities and predicted values should be discussed.

Although both models specify the validation step as one of the keys to conducting a rigorous study, I further structured the validation process. I defined the intended validation process as an objective (and quantified) comparison of various models, but I also wanted it to extract insight concerning why different methods produce better or worse quantitative results, hoping to identify systematic biases that could be corrected later. The combination of quantitative and qualitative elements points towards the mixed-method approach (Ågerfalk, 2013; Tashakkori & Teddlie, 1998). According to Ågerfalk (following Creswell, 2013), the four

central parts of the design of a mixed-method study are as follows:

1) the sequence,

2) the relative priority and the stage of the project,

3) the stage in which qualitative and quantitative components will be integrated, and

4) the extent to which the components will be embedded in an overarching framework (Creswell, 2013; Ågerfalk, 2013).

With this guideline in mind, I initially structured the validation as a quantitative and qualitative evaluation. I then decided to specify the process further by defining quantitative evaluation in terms of a single meaningful dimension to facilitate comparisons and general understanding, and my choice was the cost of prediction error. The cost, measured as a financial implication of the forecast error, is a meaningful dimension for both business and academia and enables presenting predictive analytics challenges in a way that attracts attention. The purpose of the qualitative step is the analysis of the context of the study to identify and extract any systematic bias and, to emphasise this, the step was renamed accordingly.

In addition, I discovered that, to generalise the contextual findings of the qualitative step, some validation via a general knowledge base might be necessary, which led to the introduction of the third sub-step. The framework is concluded with a process evaluation and conclusion step. Specific to the sensor data, in this step, an investigation clarifies whether there is a more direct way to monitor the predicted value. The final version of the framework developed in this theoretical iteration is presented in Figure 21.
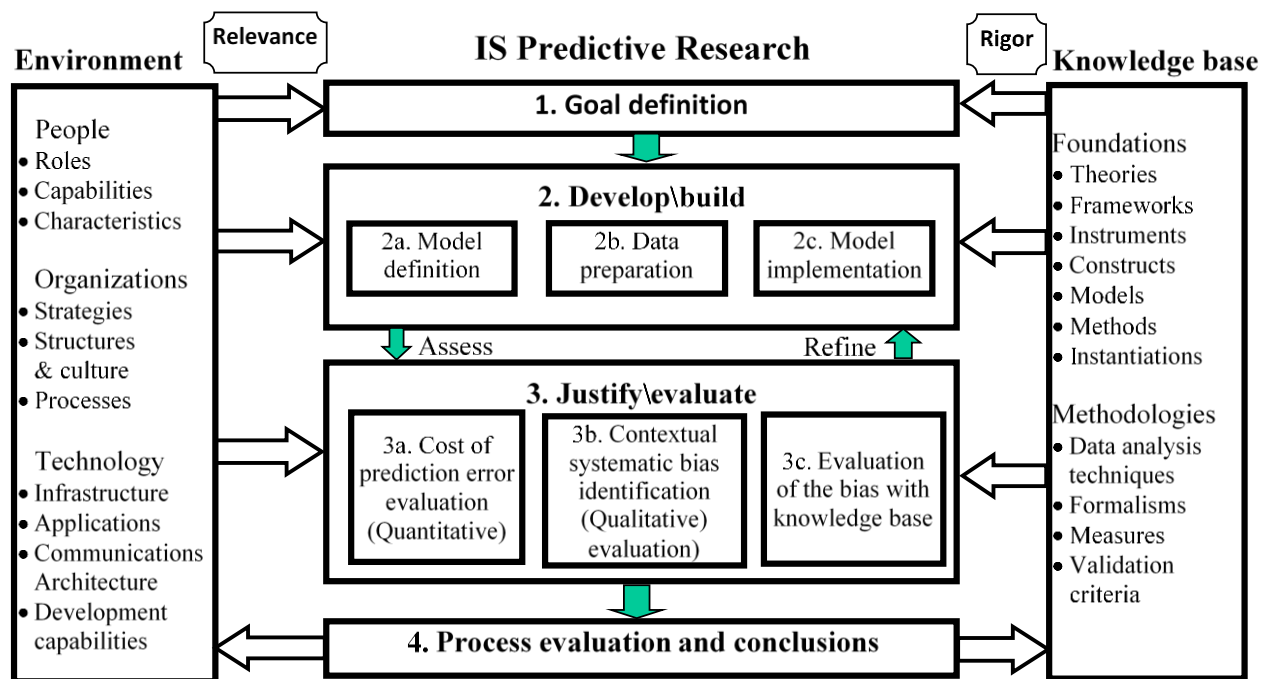
Figure 21 – Model for designing information systems with predictive analytics (DISPA).

3.3.2 Framework evaluation – MAN Diesel & Turbo's expert panel

To evaluate the newly developed framework I decided to collect feedback from MAN Diesel & Turbo's experts on predictive analytics. The group of experts consisted of two demand planners, running predictive models daily for demand predictions and market evaluations; a supply chain data analyst/scientist, working daily with various predictive analytics techniques for inventory modelling and market scenario planning; two master planners, running predictive models for factory utilisation and supplier capacity analysis; and a product portfolio analyst, running daily predictive models for market developments. The justification for choosing them for the expert panel was that even though they have not used the framework, they can reflect on the previous projects they conducted and assess how useful the framework would have been for them. Moreover, as they work with predictive models, they are all potential users of the framework. The most relevant input from the experts is presented in Table 12

Expert panel setup

One of the comments from the product portfolio analyst was that the early goal definition, informed by the environment, not only helps the structure in the design process but is a tool to ensure buy-in from relevant stakeholders throughout the design project. The demand planner pointed out that the graphical presentation of the framework 'crowds' the prescriptive part of the framework with the context elements, making it more difficult to grasp the meaning of the framework for the designer and hindering the diffusion of the framework. The supply chain analyst pointed out that referring to the 'model definition' in Step 2a can be confusing

Feedback from experts

because most predictive data models are well defined. He proposed 'model selection' to ensure reuse of exiting models.

Table 12 – Expert evaluation of the framework (most relevant comments).

Comment	Author	Details	Change Implemented?
1. Goal definition as a first step with environment involvement	Product Portfolio Analyst	Introducing an explicit goal definition as the first step of the project not only facilitates focus and goal orientation but ensures buy-in from business stakeholders on the later stages of the project.	No
2. Central element of the framework visually undermined by the informing context (environment and knowledge base)	Demand Planner	The prescriptive part of the model ('the doing' as the demand planner called it) is visually undermined by the context (environment and knowledge base), which can hinder the understanding of the framework and its utility.	Visually highlight the central part of the framework (blue background added)
Change 'model definition' to 'model selection' in Step 2a	Supply Chain Analyst	Calling Step 2a 'model definition' might suggest custom developments of the data model, while most models are standardised. Calling the step 'model selection' should be more informative and also promote the best practice (using standard data models).	Change 'model definition' to 'model selection' in Step 2a
Positive evaluation of qualitative/ quantitative evaluation steps	Master Planner	The combination of systematic, quantified evaluation with qualitative and descriptive methods has been noted as a good way to foster improvements in data processing and predictive models, as there is a link between the 'score' and system descriptions.	No
Framework works as a project management tool	Product Portfolio Analyst	The steps from goal definitions to process evaluation guide not only the development of the artefact but also a compete project, replacing the need for the project management framework.	No
Arrows between Steps 1, 2, and 3 and the environment and the knowledge base should be bidirectional	Product Portfolio Analyst, Master Planner, Demand Planner	The panel discussed whether the arrows between Steps 1, 2, and 3 and the environment and knowledge base should be bidirectional, as the knowledge developed in the design process should be input into the environment and knowledge base. However, after an intense discussion, the conclusion was reached that the feedback arrow from Step 4 is an optimal solution, as the feedback is only ready to be shared after evaluation.	No

The master planner pointed out on the positive aspect of the dual qualitative and quantitative evaluation as a good way to link systematic evaluation with descriptive insight. Another comment from the product portfolio analyst was that the framework serves as a project management tool. It structures all steps around the design project from requirement specification (the goal) through the evaluation criteria to the process evaluation. This helps to ensure there is structure and rigour in the design project.

The item discussed most on the expert panel was the direction of the arrows in the graphical representation of the framework, connecting Steps 1, 2, and 3 and the environment and knowledge base. The initial comment from the project portfolio analyst was that, during the goal definition, there are some inputs from the process to the environment and knowledge base as well as from the environment to the knowledge base through the goal definition process, which justifies a bidirectional arrow between Step 1 and the environment and knowledge base. Analogically, the knowledge generated in Steps 2 and 3 is also input into the context, suggesting a need for bidirectional arrows.

Discussion: arrows from steps one to three bidirectional

In contrast, the panel discussed the limited utility of the feedback from Steps 1, 2, and 3 before completing the evaluation. Some ideas in the design process could prove to be less useful than expected, and the early expectations could deviate from reality. As a result, the experts agreed that, to ensure clear communication, the findings should only be shared after a thorough evaluation, leading to keeping the current arrow directions.

The framework, after introduction of the changes suggested by the panel of experts, is presented in Figure 22.
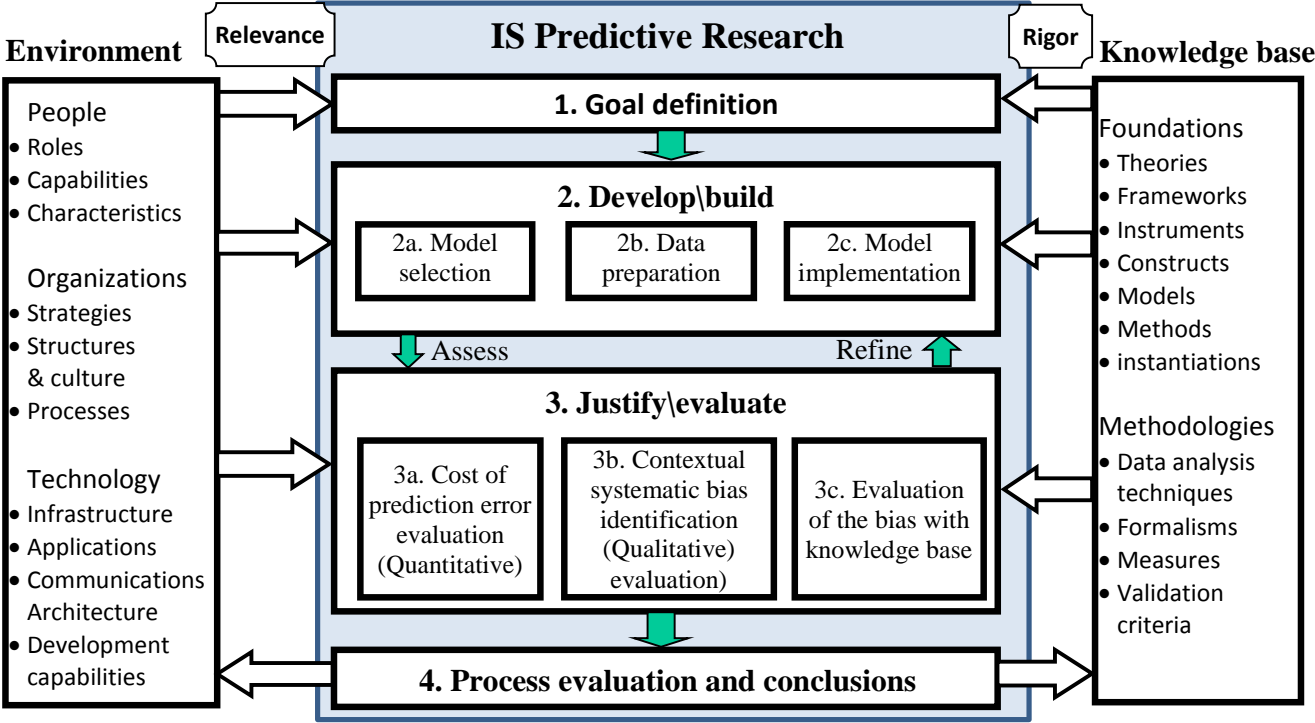
Figure 22 – Framework for designing predictive information systems after expert evaluation.

3.4 Chapter Summary

In this chapter, I used the research pyramid framework (Jonker & Pennink, 2009) to establish the project research paradigm and research methodology. I selected design science as a research paradigm, arguing that it is well accepted for the rigorous investigation of real-life problems. Furthermore, I introduced a combination of the design science model (Hevner et al., 2004) and Shmueli and Koppius's (2011) model for the empirical analysis of predictive IS to further guide this research project. As an observant reader will notice, the pyramid has four levels, while I have only explained two at this point. The missing two steps, the research method and the research technique, will be introduced in the following chapter. The research method will be the instantiation of the methodology introduced in this chapter (including instantiation-specific parameters, such as the evaluation function), while the research techniques will be actual instruments and tools to be used to run instantiations and iterations of the DISPA framework executions.

Chapter summary

4 Empirical Investigations

4.1 Designing Spare-part Forecasting for the Marine Heavy Machinery Industry

After understanding the case environment and relevant knowledge base, the model can be instantiated in this context. In the initial step, goal definition, a thorough evaluation of the environment must be performed to determine what makes the predictive design suitable for the given context and how to quantitatively measure the cost associated with the prediction error. The following two steps, design/build and justify/evaluate, are executed iteratively for the multiple designs under evaluation.

DISPA instantiation project plan

I suggest starting with a state-of-the-art solution from the knowledge base to provide a baseline and ensure the necessary grounding in previous academic work. The evaluation step should use the evaluation function, instantiated as an example in the following section, and identify variables that are not monitored. This adds human cognition to discover a systematic bias that could be removed in the following design iteration. When iteration cycles provide satisfactory output, the environment and IB can be input into the newly designed predictive model with insight acquired during the design process.

Justification for instantiating the model twice

Due to vast differences in the potential solutions under evaluation, multiple model instantiations need to be performed. As model instantiation includes defining a goal and an evaluation function to quantify the progression towards the goal, comparing the solutions that can be developed completely within the outline of the design ideas in a single instantiation will not be feasible. Instead, I will

instantiate the model twice: once for currently implementable solutions and again for ideas for future designs. To help the reader follow regarding potentially confusing instantiations and iterations, I depict the process of artefact development (i.e., the search process as it is called in Guideline 6; Hevner et al., 2004) in Figure 23.

Instantiation 1 Currently Implementable Designs			Instantiation 2 Future Designs	
Croston method	Phase-out sensor	Activity sensor	Continuous multi-sensor monitoring	Remote monitoring interface

Figure 23 – Map of the search process for the artefact, as suggested in Guideline 6 of DSR Guidelines (Hevner et al., 2004).

4.1.1 Instantiation 1: Currently implementable designs:

Goal definition (including evaluation function; Step 1)

Overview of the evaluation process

The goal of the empirical part of this paper is to design, develop, evaluate, and continuously improve a system to predict the frequency of sales of a selected product in the given case context. Initially, a state-of-the-art solution will be selected from the literature as a benchmark. The process of evaluation requires additional explicit structures; a quantitative analysis will be performed in an experimental setting. Data will be partitioned into learning and test periods, and predictions will be made for test periods based on parameters extracted from the learning sample. The prediction will be evaluated by an objective evaluation function. Two reliability tests will be repeated three times for three learning/test samples. A qualitative evaluation will follow, collecting insight concerning systematic advantages and disadvantages of the chosen approach and the possibilities to improve it. Based on these suggestions, verified in the existing literature, refinements leading to new designs will be made, which will be finalised and implemented and will undergo the same systematic evaluation process. Linking qualitative feedback to quantitative results should enable the evaluation not only of holistic solutions but also of their systematic properties.

Developing context specific evaluation function: under-forecasting

The key to a legitimate quantitative evaluation of a design is a meaningful evaluation function. There are many standard measures for a prediction error. However, because of their generic properties, they are not able to capture highly context-specific factors, such as asymmetric error cost. To cater for this diversity, two separate evaluation functions for over- and under-forecasting scenarios are necessary. The

actual cost associated with an under-forecasting situation occurs due to missed sales potential due to product shortage. As goods are not available when demand occurs (I will refer to this event as a 'stock-out'), some customers will decide to relinquish the order rather than wait for the items. The percentage of these customers can be determined by the difference in the conversion ratio of quotes to orders (also referred to as the hit rate) for in-stock quotes versus the stock-outs.

In the case of an example component group, piston ring sales at MAN Diesel & Turbo, in a scenario in which goods are in stock, the average hit rate will oscillate at around 39%, but only 30% of quotes would convert to orders in the case of a stock-out. In the same conditions, I assume that 9% of customers relinquished their purchases due to the lack of availability. To calculate lost profit, the average hit rate difference between in-stock quotes (HR_{is}) and the stock-out hit rate factor (HR_{so}) needs to be multiplied by the under-forecasted volume (U_{vol}) to compute sales volume missed due to stock-outs. To convert sales turnover to EBIT profit, the hit rate difference between in-stock quotes must be multiplied by the average contribution margin (CM).

The cost associated with over-forecasting can be divided into two categories: opportunity cost (OC), also known as the cost of frozen capital, and the cost of potential depreciation and scrap, both proportional to over-forecasted volume (OF_{VOL}). The opportunity cost is experienced because, with an over-forecast, the investment in inventory was unnecessary and the money could have been invested differently, bringing certain profit to the company. Most of the firms have some baseline working capital ratio to be used for such calculations. For MAN Diesel & Turbo, the opportunity cost

Developing context specific evaluation function: over-forecast

103

($OC\%$) was set to 10% (internal) per year in 2014. In this context, sales are expected every month; thus, over-forecasting in one month will lead to a lower replenishment cost in the following month, so that the frozen capital cost will always be calculated for a single month. Regarding the deprecation and scrap factor ($DF\%$), this reflects the possibility that unsold inventory will not move for a period, leading the inventory to be written off by a certain depreciation factor or, if parts are no longer saleable, even to being written off completely, and scrapped. For MAN Diesel & Turbo, the depreciation and scrap factor for 2014 was set to 5%. Putting all the parameters together, the cost of forecast error, COST$_{FE}$, can be described in Formula 1 below:

$$COST_{FE} = (HR_{IS} - HR_{SO}) \cdot UF_{VOL} \cdot CM + OF_{VOL} \cdot (OC\% + DF\%) \qquad (1)$$

Iteration 1 – State-of-the-art solution (Croston forecasting system)

Develop/build: Model definition (Step 2a)
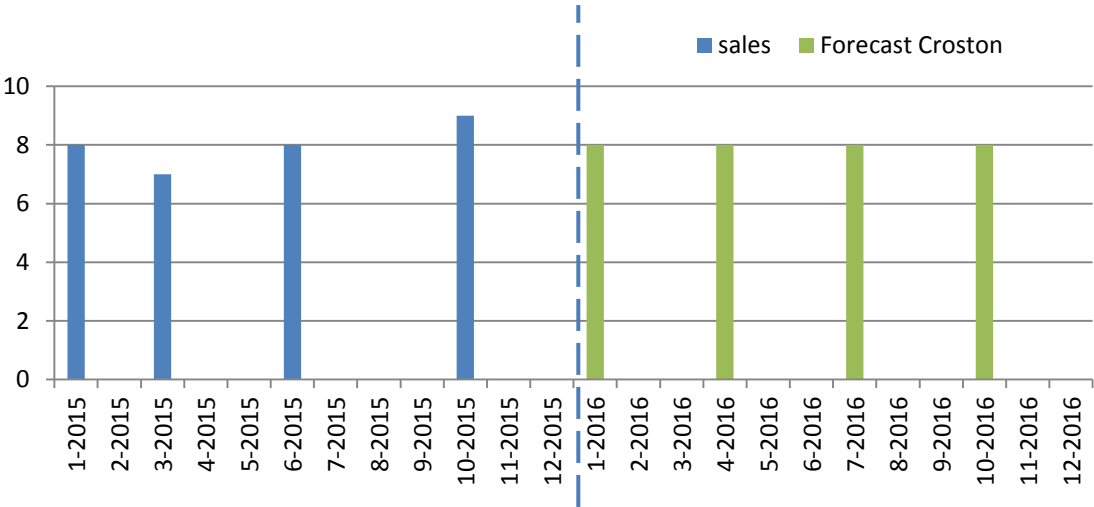
Figure 24 – Croston method example.

The initial method is selected based on the literature review (see Section 2.1), where the Croston method was selected as the state-of-the-art solution. It is a two-step approach, intended for products with infrequent demand, calculating separately expected intervals between demand points and the magnitude of demand, if it occurs. We illustrate this with an example depicted in Figure 24: a product is sold four times in 2015, in January, March, June, and October, in quantities of 8, 7, 8, and 9 pieces, respectively. The forecast for 2016 will be calculated in two steps: by calculating a demand magnitude when demand occurs and an interval between demand points. The forecasted magnitude will be computed based on historical quantities (the original method uses exponential smoothing; here, for clarity, I use an average) resulting in eight pieces (average of 8, 7, 8, and 9). The same is done for intervals between historical sales; the sale in March came two months after the one in January, the one in June was three months after the one from March, and the one

from October was four months after the one from June. The mean of these intervals gives three months between demand points. The forecast is the calculated magnitude spaced by the interval.

Data preparation (Step 2b)

The only input data required by the Croston method is the historical demand for spare parts for a customer and equipment. These data are available from the company ERP system, SAP, and are extracted through the company business DWH. As MAN Diesel & Turbo maintains a complete set of historical orders, no missing data treatment was needed. Furthermore, during the loading of the data from the ERP to the DWH system, data are cleansed and 'dummy orders', used for internal purposes only, are excluded. The time span for available observations is from the beginning of 2008 to the end of 2014. Based on this period, three data-partitioning scenarios are defined to ensure result reliability, with test periods in the years 2014, 2013, and 2012, and the learning periods are 2008–2013, 2008–2012, and 2008–2011.

Method implementation (Step 2c)

The design was implemented in MS Excel. The input data, which are the historical sales data from the industrial partner, were extracted directly from the business DWH to Excel, using dynamic data sources. Based on these data, the expected interval between demand points and the demand quantity were calculated for the training sample and extrapolated to the test sample, resulting in predictions. The procedure was performed for three sets of learning and test samples.

Justify/evaluate: Cost of prediction error evaluation (Step 3a)

Table 13 – Cost of prediction error for the baseline Croston method.

Costs	2012	2013	2014	Sum
Cost of over-forecasting	€ 86.726	€ 140.729	€ 135.505	€ 362.961
Cost of under-forecasting	€ –	€ 13.850	€ 2.620	€ 16.470
Cost total	€ 86.726	€ 154.579	€ 138.125	€ 379.431

Contextual systematic bias identification (Step 3b)

The Croston method is well documented and easy to implement, but it is very generic. It only uses historical data, so it will never predict trend inversion, even if it can be expected; therefore, from the beginning, I started looking for case-specific information that could inform the prediction. The first promising idea coming from the demand planners was to include the lifecycle information, by monitoring the IB (equipment in use) with phase-out sensors. Moreover, in the cases where the equipment did not have a phase-out sensor installed, reports on scrapped vessels and engines were used. I took advantage of the fact that marine diesel engines need to be legally insured and supervised by a third party (a classification society), which allowed me to remove them from the pool of potential after-sales customers. In this way, all data regarding dead installations could be completely removed from both training and test sets.

Explanation to a potential source of bias

Evaluation of the bias in the knowledge base (Step 3c)

Forecasting spare-part demand using IB information and knowledge of the age and status of products and systems in use and customer maintenance and replacement policies (Minner, 2011) has been a subject of recent academic research. The main stream of this research is concerned with optimising inventory policies using detailed geographical information about customers and equipment (Ihde, Merkel, &

Installed base in the literature

Henning, 1999; Jalil, 2011; Song & Zipkin, 1996). The academic efforts to develop a demand-forecasting method exploring IB are limited. As Dekker et al. pointed out, 'scientific research on Installed base forecasting is limited and the term is pretty scarce in the operations literature' (2010, p. 2).

The outline of the idea that demand forecast can be based on the IB was proposed by Lapide (2012), but the approach is so simplistic that it cannot be considered an applicable method. An extremely interesting theoretical forecasting framework was presented by Minner (2011); the framework estimates the probability of spare-part sales for equipment of a certain age and, based on the age of equipment in the field, estimates the total demand for spare parts. All those occurrences of IB use in the context of forecasting make it a promising candidate to include in the forecasting method.

Iteration 2 – Phase-out sensor forecasting system

Develop/build (Step 2): Model definition (Step 2a)

Based on the insight regarding an IB from the case and the knowledge base, an enhancement of the Croston method will be developed, including phase-out sensor output. For all pieces of equipment that are already not in use, no forecast will be calculated. For all others, the same algorithm by Croston will be used. Note that only data related to the phase-out of engines are used. Although phase-in data are available, they are unusable for the Croston algorithm, as no historical sales data are available to compute the forecast.

Detailed model definition

Data preparation (Step 2b)

The IB information, namely, the list of all the pieces of MAN Diesel & Turbo equipment in operation, was also extracted from the business DWH, and it is regularly uploaded there from classification societies, legally supervising the use of engines in marine applications. Data are loaded by an external data provider to SAP every quarter and, from there, they are sourced to the DWH. In this context, there is a perfect match between the sensor data usage and the measured variable. The data describe scrapped installations and are used to exclude them directly from the predictions.

Data preparation details

Method implementation (Step 2c)

The implementation is very similar to a previous Croston method, with the only difference being that, for engines not in use, the forecast will always be set to zero. To ensure consistency, the data related to currently dead installations are removed from the learning data sets as well.

Method implementation details

Justify/evaluate: Cost of prediction error evaluation (Step 3a)

Table 14 – Cost of prediction error for the phase-out sensor forecasting system.

Costs	2012	2013	2014	Sum
Cost of over-forecasting	€ 47.164	€ 125.819	€ 137.708	€ 310.691
Cost of under-forecasting	€ 31.929	€ 20.700	€ 599	€ 53.228
Cost total	€ 79.094	€ 146.519	€ 138.307	€ 363.919

Quantitative evaluation summary

Table 14 shows the cost evaluation of the Croston method enriched with the phase-out sensor. It is important to note that, compared to the baseline method, the cost decreased by more than 4%. The improvement is visible for two data-partitioning scenarios (2012 and 2013), while it remained almost constant for the last partition (2014).

Contextual systematic bias identification (Step 3b)

Identifying potential sources of bias

Service engineers and sales personnel drew my attention to the market diversity during the period of the study. As an effect of the lower demand for transportation services because of the global financial crises, the sailing patterns of most of the customers changed. Rather than travelling at maximum frequency and speed, the vessel management companies were concentrating on cost reduction, maximising the load per vessel, and cutting the duration of a transport. Moreover, to optimise fuel consumption and decrease wear on a vessel, ships would travel at the most efficient rather than the maximum speed. This phenomenon is often referred to as 'slow steaming' in the shipping industry.

As the global economy started to recover from the crisis, the situation started to return to the previous status quo. All these changes could potentially lead to a very significant change in the demand for spare parts. This means that, for example, if a ship owner would ordinarily replace a given

spare part every five months, under the slow steaming scenario, assuming a vessel activity reduction of 20%, the matching period would be six months. In the context of the Croston method, this suggests that data collected in the slow steaming period must be somehow 'normalised' to be comparable to the previous observations. This will be the aim of the next iteration.

Evaluation of the bias in the knowledge base (Step 3c)

Slow steaming has been widely recognised in recent shipping literature (Notteboom & Cariou, 2013; Woo & Moon, 2014; Yin, Fan, Yang, & Li, 2014). Three reasons for the popularity of slow steaming among ship managers are the oversupply of shipping capacity, an increase in the fuel price, and environmental pressure (Yin et al., 2014). According to Notteboom and Cariou (2013), the strategy has been gradually implemented by the main liner shipping companies since 2008, considerably affecting the analysed dataset. These observations demonstrate the potential of including activity-sensor information and directly observing ship engine utilisation patterns rather than inferring them from time intervals between replacements in the predictive model.

Slow steaming in the literature

Iteration 3 – Activity-sensor forecasting system

Develop/build: Model definition (Step 2a)

To compensate for the bias in the data caused by a changing market pattern during the period under study (the application of slow steaming), a method of normalising data periods based on actual engine activity needs to be introduced. The simplistic idea behind this approach is that, if an engine was used 20% less often in a given period than in an ordinary period, the calculation of intervals between the replacement of spare parts would extend the expected lifetime of spare parts in that engine by 20%. To achieve this goal, the unit of interval between replacements will be changed from time (in months) to engine running hours with

Detailed model definition

equivalent maximum revolutions. Intuitively, an engine can accomplish one running hour with maximum equivalent revolutions by either running one hour at full speed or two hours at half of the maximum speed, and so on. To measure running hours with maximum revolutions, engine activity sensors were introduced. The prediction model will predict the magnitude of sales in the same way as the traditional Croston approach, but the interval between sales will now be predicted based on the engine activity rather than on time. Conceptually, this modification can be considered a next step, after implementing a phase-out sensor (phase out monitors engine activity in a binary fashion in terms of 'in use' versus 'not in use'), while the activity sensor creates a more continuous scale of engine activity.

Data preparation (Step 2b)

The running hours at the full-speed resolution equivalent were estimated based on the engine application. Due to their customer privacy protection policy, MAN Diesel & Turbo do not have access to all the activity sensors that are mounted on

the equipment (although they are installed on practically all engines). The values are estimated using the average value expected for an engine application: Typically, according to MAN Diesel & Turbo's service engineers, an engine in a stationary plant will run at full speed almost all the time (about 8,600 hours of full-speed resolution per year), and the main engine on a ship for around 6,000 hours, while an auxiliary engine will only run for about 3,000 hours.

Furthermore, I introduced scaling factors, based on market behaviour according to the literature and MAN Diesel & Turbo experts for a given year (the slow steaming scenario popularity grew from 2008 and peaked in 2011; from then onwards, the average engine activity stabilised at levels slightly below normal from before 2008). Running hour values are also extrapolated to the future, using the same estimation logic. The values are extracted for the same periods as the historical sales data, dating back to 2008. Because of the estimation factor, the match between sensor data usage and the measured variable is not perfect. Clearly, a more optimal method would be to use 'real' running hours measured on every engine.

Method implementation (Step 2c)

The implementation is identical to the baseline Croston method for the calculation of the expected demand magnitude, but the expected interval between replacements is now calculated using the elapsed running hours value. For every month when a replacement occurred, an elapsed running hours value is saved. In this way, the average running hours elapsed between replacements is obtained. The forecast is the extrapolation of that average, with the

113

magnitude set using exponential smoothing, as in the original Croston method.

Justify/evaluate: Cost of prediction error evaluation (Step 3a)

Table 15 – Cost of prediction error for the Croston method with activity-sensor output (using running hours).

Costs	2012	2013	2014	Sum
Cost of over-forecasting	€ 59.999	€ 90.321	€ 124.276	€ 274.596
Cost of under-forecasting	€ 6.438	€ -	€ 20.475	€ 26.914
Cost total	€ 66.437	€ 90.321	€ 144.752	€ 301.510

Quant. evaluation summary

Introducing activity-sensor output has improved the baseline cost derived by the Croston method by as much as 20%, and the previously proposed phase-out sensor implementation by 17%. The improvement is visible for two data-partitioning scenarios (2012 and 2013), while it remained practically constant for the last partition (2014).

Contextual systematic bias identification (Step 3b)

Identifying weak spots and potential biases of the method

The overall impressive improvement is achieved despite the assumption of a somewhat rough estimation, leading to a potential DQ issue. Although engine application seems to be a good estimate of engine utilisation, there must be a variation within segments of the same engine application, so the full potential of this solution could be achieved if estimated values were replaced by real observations from in-situ installations. Nevertheless, the realised forecast quality improvement is significant, although further model development will require significant investment in infrastructure. Thus, at this point, it is not feasible to run another design iteration.

Evaluation of the bias in the knowledge base (Step 3c)

As no new systematic bias is identified for implementation, this step can be omitted.

4.1.2 Process evaluation and conclusion (Step 4)

Cost of Prediction Error per Implementation
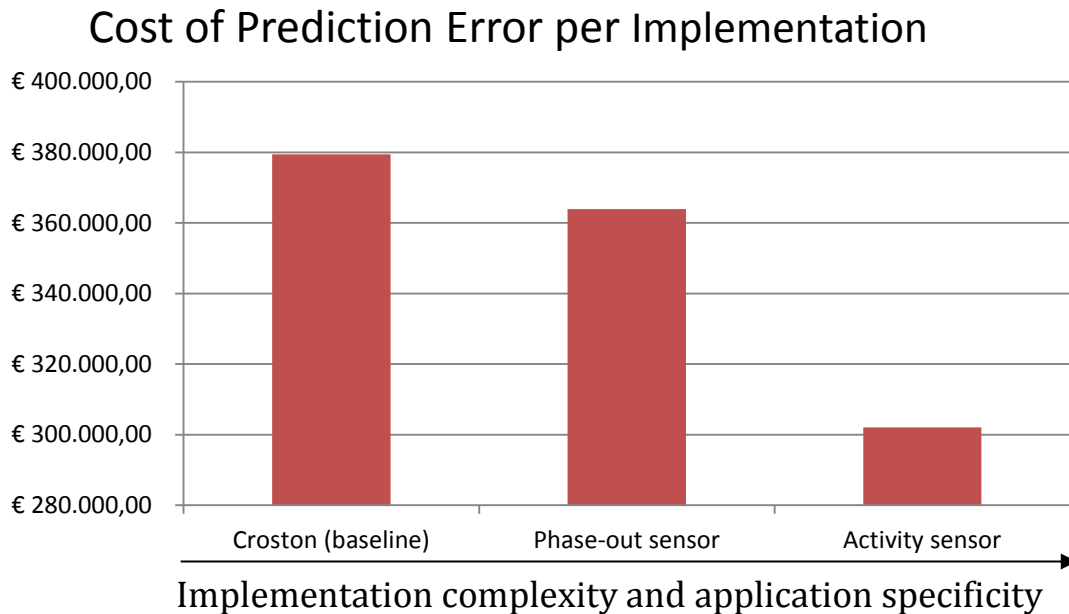
Figure 25 – Qualitative result summary for Instantiation 1.

Summary of the first instantiation

In summary, all the implemented sensor-based designs show prediction quality improvements when compared to the baseline Croston solution. Unfortunately, the quality comes at the price of complexity and specificity to a given environment. An initial Croston solution could easily be implemented for any data series. Sensor-based solutions require extremely specific additional information, and the quality improvement they provide is gradually coupled more tightly with the application. This tight coupling and specificity increase in tandem with the increased prediction quality.

Furthermore, the additional information comes from sensor installation (or its simulation), which needs to be pre-installed, and this complexity introduces costs not present in the Croston scenario. Moreover, those observations suggest that sensor-enabled forecasting solutions would be financially feasible in environments in which the gain in

forecast improvement outweighs the cost of solution implementation. This kind of environment would be characterised by a high level of uncertainty and have a high cost associated with forecasting errors.

The quantitative output of the three implemented designs is presented in Figure 25. The phase-out sensor design improves the baseline prediction quality by 4%, while an activity sensor beats the baseline by 20% and the phase-out solution by 17%. This activity-sensor example shows that, including the additional dataset, the DQ problems can be faced, as the data might have been estimated or generated from a source not intended for that specific data usage. In that case, new data management routines should be implemented, leading to gradual DQ improvement.

4.1.3 Instantiation 2 – Future designs:

Goal definition (including evaluation function; Step 1)

With the insight from the first instantiation, I move towards more complex system designs that are impossible to develop fully or to evaluate within the time frame of the project. This means that the overall goal of this iteration has not changed. It is still designing, developing, and continuously improving a system that can predict sales in the given context. However, there will be no means to quantitatively evaluate the new design performance in the way this was done in the first instantiation. As a result, a new quantitative evaluation function that can provide insight regarding the evaluated designs and that can be executed even without complete system implementation needs to be developed.

Understanding the purpose of the instantiation

The idea behind the proposed evaluation function revolves around the concept that, despite not being able to calculate the cost reduction, due to the improved forecast, there still is a possibility to define what the reduction should be in order to make a given investment in a new system feasible. Thinking along these lines, a requirement for developing a sensor-enabled forecasting IS should be that the cost difference ($\Delta COST_{FE}$) between the cost of the forecast error before system implementation ($COST_{FE\_BEFORE}$) and the cost of the forecast error after implementation ($COST_{FE\_AFTER}$) should outperform the system implementation cost ($COST_{INV}$) within a reasonable time.

New evaluation function definition

My enquiry to the organisation concerning what the exact reasonable period should be produced mixed results with a maximum allowed payback period varying from two to five years. I decided to use a five-year threshold period because

the financial benefit of improved forecasting could be underestimated by my forecast error function (there are multiple potential benefits not included, such as greater stability in opportunities for production or proactive sales campaigns), knowing that innovative, 'smart' solutions regarding better customer intelligence are currently within the strategic focus of the company. As a result, the new evaluation function appears as follows:

$$5 \cdot \Delta COST_{FE} > COST_{INV} \iff$$
$$COST_{FE\_AFTER} < COST_{FE\_BEFORE} - \frac{COST_{INV}}{5} \qquad (2)$$

Although the sensor infrastructure required for models discussed in Instantiation 2 was not developed to the extent of allowing the computation of predictions and measuring the actual error, all presented sensor designs had at least prototype implementations. Therefore, in sections related to the step develop/build, I will discuss the conceptual vision of an abstract model in Step 2a (model definition). In Step 2b (data preparation), I will discuss in detail the data extraction for evaluation function. In Step 2c (model implementation), I will discuss the current prototype implementation, explaining in detail the limitations of the current approach. Step 3 (justify/evaluate) remains structured as it was in the previous instantiation.

Detailed plan for execution of the model in second instant.

Iteration 2.1 – Continuous multi-sensor monitoring

Develop/build: Model definition (Step 2a)

Detailed model definition

Returning to the conclusions from the first instantiation, the issues to be addressed in the activity-sensor model revolved mainly around potential DQ problems and the lack of clarity concerning what the exact value observed from a remote installation would be. To prevent these issues in the new model, a customised system of monitoring sensors will be designed, while observing measures that could be used to predict wear to spare parts. Reliable information about maintenance activities provided in advance would not only limit breakdowns, which are extremely expensive for equipment owners but could also be used by a maintenance service provider to pre-purchase necessary spare parts and to increase the chance of getting the order when the demand arises.
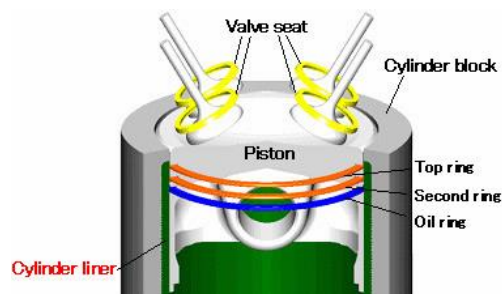
Figure 26 – Cylinder liner on a cross section of a cylinder.

In the previous instantiation, I discussed remote monitoring of the overall engine activity, but a forecasting model based on such a monitoring system could ideally rely on measures observing the wear of spare parts. Subsequently, by regression, the historical values could be extended into the future, estimating when a certain threshold value of wear is reached. As an example, a cylinder liner is a part of an engine installed as an inner wall of a cylinder, holding a piston in one

plane and providing minimum friction when the piston is in motion (see Figure 26).

Currently, a decision to replace a cylinder liner is taken based on manual inspection and measurement of its thickness. An engineer needs to dismantle the engine and measure, usually with a calliper, the exact thickness of the liner. The spare-part manufacturer provides a safety threshold value below which the part must be replaced. As an example, if average wear exceeds 18 mm, the wear starts to increase exponentially and, likely, the warranty for the cylinder will void. Monitoring the current cylinder liner wear and using the measurement to predict when it will exceed the safety threshold will not only enhance the quality of the service provider's prediction but would also allow the customer to avoid the costs associated with frequent engine inspections.

Data preparation (Step 2b)

To execute the evaluation function and calculate the limits on the cost of forecast error after implementing the system, the cost of potential investment and the baseline cost of forecast error associated with the state-of-the-art guide need to be established.

In the previous instantiation, I established the cost of the forecast error to be the sum of the cost of under-forecasting and the cost of over-forecasting (see the 'goal definition' section in the previous instantiation). The first measure was lost sales potential and was set to be a multiplication of the quotation volume of under-forecasted goods and the factor of quotations rejected due to a lack of availability (see goal definition in Instantiation 1 for details), and the average CM. Based on historical data, I can obtain both the factor estimate

(the average difference in quotation convergence between parts in stock and out of stock for materials with forecasts) as well as the volume of quotations for items that were out of stock and the CM for SKUs with forecasts. For over-forecasting costs, the over-forecasted volume was multiplied by factors related to lost opportunity costs and depreciation factors. The accounting factors provided previously are still valid, and the over-forecasted volume can be measured precisely as the stock value associated with the forecasted parts is still available for sales at the end of the month, corrected by the safety stock volume. The advantage of using this measure over a simple measure of the value of over-forecasted quantities is that the chosen notion better reflects the 'unnecessary' stock value, which is the intention of the formula.

Defining data sources for investment cost

When discussing investment costs, one piece of information needs to be shared; MAN Diesel & Turbo has included a sensor-based monitoring system project in its long-term company vision, and many steps towards its implementation have been taken. Based on this implementation project (see details in the section below), the implementation cost of the project, including the cost of establishing satellite communication between the central infrastructure and the monitored nodes, and the cost of the central infrastructure can be estimated to be 2.5 M€. The cost of data acquisition modules is not included in the calculation because, if sold with the complete new installation, the module cost is negligible compared to the complete engine cost, while the module also provides additional local monitoring functionality, justifying its cost.
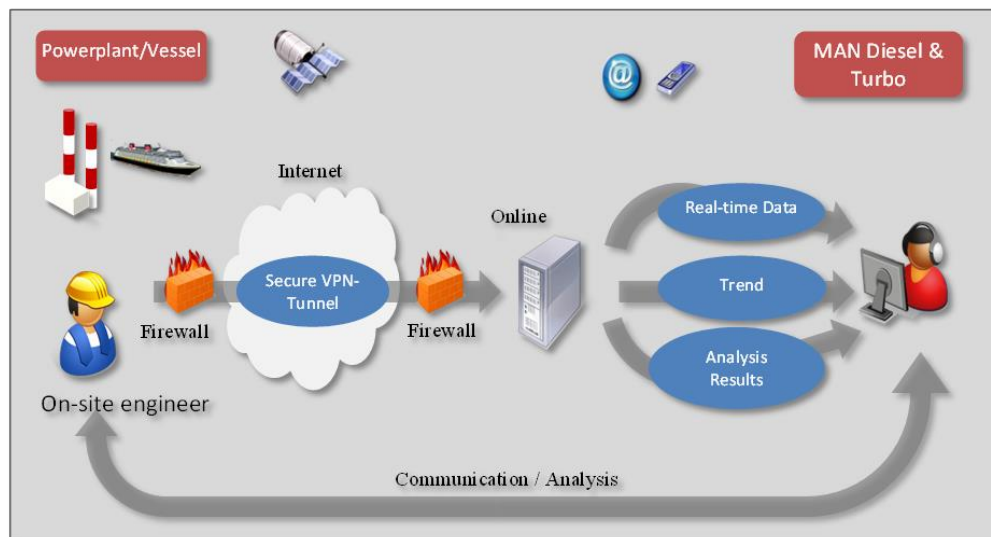
Figure 27 – MAN Diesel & Turbo Online Services project outline.

Model implementation (Step 2c)

A sensor-based data acquisition model has been included in all engines built since 2010, and in most of those built in the twenty-first century, enabling local sensor data inspection. In addition, MAN Diesel & Turbo employs a satellite communication system to systematically pool data from connected vessels sailing all over the world and fixed power stations. An important comment at this point is that, due to privacy concerns, customers need to agree to activate the transmission module for the data to be accessed by MAN Diesel & Turbo. A large central infrastructure for storing collected data and an elaborate graphical interface to facilitate data inspection and analytics was built. The system is currently running for a sample of 4-stroke MAN Diesel & Turbo engines, designed, and supported by the MAN Diesel & Turbo Augsburg headquarters, while all the predictions and operational data in the first instantiation were extracted based on 2-stroke engines, supported by the MAN Diesel & Turbo Copenhagen headquarters. As a result, due to technical and market differences between 2- and 4-stroke engines,

Model details

123

there is no possibility of using data collected from the sensors to produce predictions comparable with previous results.

Justify/evaluate: Cost of prediction error evaluation (Step 3a)

$$COST_{FE_{BEFORE}} - 6.1 \text{ M€}, COST_{INV} \sim 2.5 \text{ M€}$$

$$COST_{FE_{AFTER}} < 6.1 - \frac{2.5}{5}$$

$$COST_{FE_{AFTER}} < 5.6 \text{ M€} \tag{3}$$

Quant. results evaluations Based on the current operational values, including the stock-out percentage and surplus inventory value, the overall estimated cost of over-forecasting is 1.8 M€, and the cost of under-forecasting is 4.3 M€, adding up to 6.1 M€. With an investment cost estimate of 2.5 M€ (see previous section), the requested payback period of five years can be achieved with a forecast cost improvement of 8% (from 6.1 to 5.6 M€), assuming no additional maintenance cost for the new solution. This assumption is quite realistic, especially when only considering costs in the order of magnitude of other values included in this calculation.

Contextual systematic bias identification (Step 3b)

The focus of the current system development has been on establishing communication and data storage functionality with relatively less attention paid to the actual sensor technology responsible for the monitoring. The initial business case was to be able to compose monthly condition reports summarising overall engine performance in the past month. This application requires somewhat simple operational measures, such as the number of hours of

operation or the temperature or oil pressure at key engine points. Although these measurements can be proxies of spare-part wear, they are not straightforward to use for forecasting the wear of spare parts or for preventing breakdowns, especially if the sensor measures are not complemented by the detailed context of the engine activity.

As a result, spare-part prediction data usability, understood as 'convenient and practicable for use' (Merriam-Webster, 2015b), is considered low. As an example, excess wear of piston rings results in high friction between the rings and the cylinder liner, increasing the temperature inside the cylinder. Nevertheless, using the cylinder temperature as a predictor of ring wear is not feasible, as the temperature can vary during normal operation, depending on operational conditions (such as the time of operation or properties of fuel and lubrication). The dramatic increase of friction and temperature is a sign of excess wear that has already been experienced, at which point it is too late to use the information proactively.

Another systematic issue identified when evaluating the usage of sensor monitoring systems for spare-part demand predictions is limited context awareness. Without knowing what exactly happened to the engine, it might be difficult to interpret observed measurements, especially when the quantities measured are linked indirectly to the phenomenon of interest or are not unequivocal. General operational measures, such as temperature or oil pressure, might be significantly influenced by maintenance activities, such as switching to a different fuel or changing oil, or by replacing some spare parts or dismantling the engine. It is analogous to drawing conclusions about issues with heating based on a decrease in the room temperature when we are not aware someone opened a window.

Evaluation of the bias in the knowledge base (Step 3c)

Context awareness in the literature

The issue of context and context awareness for technological applications has been the focus of interest within the research community for many years. Historically, the main research stream investigated contexts for context-aware computing (Chen & Kotz, 2000) and intelligent environments (Coen, 1998). Although the importance of context awareness, using context to provide relevant information (Abowd et al., 1999), has been recognised in those areas, in the case of predicting quantities based on proxies, context awareness is even more important. It is a required condition, not a 'nice-to-have' feature. With the multiplicity of scenarios for issues and potential debugging routines, a flexible remote monitoring interface needs to be in place to collect relevant information on demand.

Iteration 2.2 – Flexible Remote Monitoring Interface: Red Helmet

Develop/build: Model definition (Step 2a)

Identifying the need for a flexible remote monitoring interface was the foundation for creating the idea of the Red Helmet project. The vision of the project was that any local engineer in a remote location in need of diagnosing a MAN Diesel & Turbo engine could use a monitoring device built as a helmet equipped with a camera, a microphone, and other necessary sensors connected to MAN Diesel & Turbo headquarters, allowing their experienced engineers to see and hear the engine as if they were at the location. In this way, they would be able to guide the on-site engineer to diagnose and solve any complex problem with the installation almost immediately, reducing extremely expensive potential downtime costs of an engine on a running vessel.

Remote monitoring in MAN context

This interface could also prove to be useful as an extension of the continuous multi-sensor monitoring system, as it is able to provide collected data in a context that was previously missing. If some continuously monitored quantities reported values outside of a normal operational range, this kind of interface could allow almost immediate and inexpensive intervention. In addition, for spare-part forecasting, this kind of interface could allow following up on installations with key sales forecasted, enabling the enhancement of predictive analytics via the expert knowledge of on-site engineers.

Objective for remote monitoring

Data preparation (Step 2b)

Data extraction logic for estimating the current forecasting cost is the same as for the continuous multi-sensor monitoring system (Step 2b). To estimate the cost of

127

investment, I will use the cost of development of the Red Helmet project acquired during the prototype project and multiply it by the number of engine installations currently under the responsibility of the MAN Diesel & Turbo headquarters in Copenhagen, which is obtained from the business DWH and is set to 19,462 for 2015. An additional 100 K€ will be reserved for creating a central infrastructure, similar to the one for the multi-sensor monitoring project. An estimate of the prototype cost was approximately 400€ for producing about 10,000 units. Using these assumptions, the cost of project implementation is calculated as 7,884,800€, which I round up to 8 M€.

Model implementation (Step 2c)

In terms of forecasting spare parts, Red Helmet infrastructure could be used as an on-demand remote monitoring interface that could remotely inspect key installations, returning precise estimates if a given spare part were to be replaced in the subsequent couple of months. This estimate would be informed by the knowledge of wear patterns and part replacement of the engineers from the headquarters, the data collected during the inspection, and the contextual information of recent activities concerning the installation provided by the on-board engineer.

Justify/evaluate: Cost of prediction error evaluation (Step 3a)

$$COST_{FE_{BEFORE}} - 6.1 \text{ M€}, COST_{INV} \sim 8 \text{ M€}$$

$$COST_{FE_{AFTER}} < 6.1 - \frac{8}{5}$$

$$COST_{FE_{AFTER}} < 4.5 \text{ M€} \qquad (4)$$

Based on the current operational values, including the stock-out percentage and surplus inventory value, the overall estimated cost of over-forecasting is 1.8 M€ and the cost of under-forecasting is 4.3 M€, adding up to 6.1 M€. With an investment cost estimate of 8 M€, the requested payback period of five years can be achieved with a forecast cost improvement of 25% (from 6.1 to 4.5 M€). That significant reduction of forecast error cost might be especially difficult to achieve when including the additional maintenance cost of the new solution, assuming that 10% of IBs will be monitored through the interface (about 2000 engines) and that one remote diagnostic session costs 100€, including human costs and satellite communication costs. Thus, the actual forecast error cost after implementing the system needs to be reduced by a further 200 k€, translating to a 30% improvement (6.1 to 4.3 M€).

Quant. results evaluation

Contextual systematic bias identification (Step 3b)

The strength of the presented solution is that the interplay of objective (audio and video stream and additional sensor readings) and subjective (interaction with on-site engineer) inputs could compensate for any systematic bias present in any of the individual inputs. Nevertheless, the clear downside of this approach is the very high cost, both for the implementation of the system and for maintenance and data acquisition after the system is implemented.

Understanding weaknesses of the approach

4.1.4 Process evaluation and conclusion (Step 4)

Examining the multiple iterations and instantiations, I presented four predictive IS using the following sensor technologies: a phase-out sensor, an activity sensor, a continuous multiple-sensor monitoring system, and the Red Helmet. The process of designing the four artefacts, despite being specific to a single case and a specific well-defined problem, uncovered some aspects of designing IS predictive artefacts that can be generalised into guidelines, often referred to as design principles in the literature. In the following paragraphs, I will discuss those aspects and the foundations for that generalisation.

The concept of design principles is not new in IS. The exact definition of the design principles that we followed is the one provided by Chandra et al. (2015). According to their work, a design principle must be defined with the materiality (the meaning of the exact system design aspect), action to be recommended, and the boundary condition defining the scope of the guideline. They recommended using a template phrased as follows: 'Provide (…) in order to (…) given that (…)'. In the section below, we followed all those recommendations.

The first aspect we generalise from our observations is related to the development of the Croston method with a phase-out component IS model artefact, and it is the design decision to which we greatly attribute its success, using an already established sensor/digital trace-based system and data already collected for a different purpose.

This approach not only greatly lowers the cost and simplifies the implementation process but also ensures a certain level of

DQ. If the data are already in use, they are being reviewed, and some potential quality issues might be spotted and corrected. The risk factor for the investment is also lowered significantly. At the time of design, the exact format and structure of the data are available, helping to avoid any surprises with the practical limitations of sensors. The generalisability of these observations beyond the discussed case led us to formulate the first design principle for designing predictive IS: *provide systems with access to an existing infrastructure to lower infrastructure-related implementation cost and complexity given that such a feasible infrastructure exists.*

While developing an activity-sensor-based predictive system, our focus was on DQ and data estimation methods. Our *Developing* experience shows that, when sample size and DQ problems *guideline 2* are predicted upfront, they can be mitigated, for example, by redesigning sensor infrastructure or data projection, capturing essential relationships from a small data sample, and encapsulating them into an estimation function (as in our implementation of the activity-sensor design). The generalisability of these observations beyond the discussed case led us to formulate the second design principle for designing predictive IS: *provide systems with a DQ measurement mechanism to mitigate issues as soon as possible given that predictive algorithms based on the data in question are in use.*

The conclusion from investigating a multi-sensor continuous monitoring system is that it is crucial when using sensor technologies to be clear about which is the most interesting variable data to sense and how this sensing activity should be implemented. Our example has shown that, when a system measures a phenomenon of interest through a proxy

measure, such as measuring the wear of piston rings through the temperature around the rings, the process of analysing data collected might be both significantly and unnecessarily hindered. The generalisability of these observations beyond the discussed case led us to formulate the third design principle for designing predictive IS: *provide systems with carefully selected sensors to optimally choose between artefact predictive ability and complexity.*

The insight related to implementing a remote monitoring interface revolves around the cost of the forecast error. The design proved to be quite expensive, both in terms of initial implementation and the running cost of the system. The benefits from having the system in place must outweigh the cost, but this can only happen in an environment where the cost of prediction is very high. In general, generic solutions, such as our baseline approach or those presented in the literature review section, will frequently be cheaper to implement, as they do not require investment in new sensor infrastructure.

This leads us to suggest that sensor-based predictive IS will usually be feasible for environments with high prediction error cost, typically characterised by high uncertainty. We have formulated that observation as the fourth design principle for designing predictive IS: *provide systems with context-aware predictive features to improve the prediction quality only given that the cost of prediction error is high enough to outweigh the cost of implementation.*

All design principles are presented in Figure 28 – Design principles for designing predictive information systems*.*

> **Design Principle 1: Access to Existing Infrastructure**
> Provide systems with access to an existing infrastructure to lower infrastructure-related implementation costs and complexity given that such a feasible infrastructure exists.
>
> **Design Principle 2: Data Quality Measurement**
> Provide systems with a DQ measurement mechanism to mitigate issues as soon as possible given that predictive algorithms based on the data in question are already in use. Be mindful when choosing what to sense and how to sense it.
>
> **Design Principle 3: Sensor-Problem fit**
> Provide systems with carefully selected sensors to optimally choose between the predictive ability and complexity of the artefact.
>
> **Design Principle 4: Context-aware Predictive Systems Feasibility**
> Provide systems with context-aware predictive features to improve the prediction quality given only that the cost of prediction error is high enough to outweigh the cost of implementation.

Figure 28 – Design principles for designing predictive information systems.

Looking more closely into the financial feasibility of sensor-enabled predictive IS, some further generalisations can be made based on the evaluated cases. Summarising what was mentioned in the previous paragraph and in the fourth guideline for designing predictive IS in the context of financial feasibility, this can be considered the difference between the cost of investment and the benefit provided by the investment. Clearly, the latter is difficult to estimate precisely before the system is in place. Thus, I chose the cost of *Introducing the feasibility matrix* investment to be the primary dimension in my feasibility analysis. Moreover, revisiting the third guideline, the system can provide a financial benefit when the sensors provide added value to the prediction, and this can be best achieved when the phenomenon of interest is measured directly. The ability to observe meaningful and unambiguous measures directly can be an asset during the design phase, as advocated in Guideline 2, so this could be used as a meaningful dimension to investigate the system being designed. The concluding 2x2 matrix for pre-assessing the financial feasibility of predictive IS with sensor technologies is presented in Figure 29.

Figure 29 – Matrix for pre-assessing financial feasibility of predictive information systems with sensor technologies.

Evaluating the previously presented sensor-based solutions based on the matrix designs from Instantiation 1, the phase-out sensor and activity sensor will fall into the low investment cost section of the matrix (the left-hand side). The activity sensor, due to the need for estimation and the indirectness of the data values, will fall into the medium feasibility quadrant, while the phase-out sensor design will be classified as belonging in the high feasibility quadrant.

Evaluating developed artefacts with the feasibility matrix

The designs presented in the second instantiation will, on the other hand, fall into the right-hand side of the matrix. The multi-sensor continuous monitoring system, as described in the previous section, will fall into the low feasibility quadrant due to using wear proxy measures, while the flexible monitoring system falls into the medium feasibility quadrant. This lens can also be useful to visualise the possibilities of increasing the financial feasibility of a system. The aim of the predictive IS designer is to move the system up and towards the left.

4.2 Chapter Summary

In this chapter, I introduced the empirical work conducted in the project's industrial partner environment. The work was conducted in two phases. In the first phase, which was also a first instantiation of the DISPA framework, I implemented three IT artefacts suited to the requirements of the case. The initial implementation was based on the solution found in the literature and was used as a baseline. The following two IT artefacts, the phase-out sensor forecasting system and the activity-sensor forecasting system, were iterative improvements of the baseline, developed in this research project. Considering their superiority to the baseline in the case context, they can be considered this project's contribution to spare-part forecasting literature.

The downside of the improvement in prediction quality when using custom sensor-based solutions is shown as a tight coupling of the solution and the problem, leading to high development costs, a lack of standardisation, and increased solution complexity. In addition, throughout the second phase and the second instantiation of DISPA, I presented the evaluation mechanism that DISPA offers prior to artefact development. As a result of the two iterations, I instituted design principles for designing predictive IS and a matrix for pre-assessing the financial feasibility of the predictive IS with sensor technologies to assist designers in the design process and to help designers evaluate the design environment prior to beginning design activities.

Chapter summary

5 Discussion and Conclusion

This chapter discusses the research findings in the context of the state-of-the-art research and evaluates its contributions through the lens of relevance and rigour. Throughout the chapter, I primarily discuss and evaluate the DISPA framework, as it is considered the main contribution of this work. I argue for the practical utility and relevance of DISPA, which I consider its greatest strength, and evaluate the process of designing it against the design science guidelines of a rigorous design process.

5.1 Project Contribution Summary

Structuring the contributions by the level to which they can be generalised, there are three levels defining the scope of usability of the contributions. The outputs of my work that can be applied to any predictive analysis with sensors were classified as generic. Contributions that are specific to an after-sales business were classified as belonging to the after-sales class, and those specific to the MAN Diesel & Turbo company context were categorised as case specific. As a result, the main generic contributions arise in the form of the DISPA framework and components facilitating its instantiation (the guidelines and flow chart). The framework is also extended by the matrix for pre-assessing the financial feasibility, classified as generic contributions.

The spare-part forecast method review and the original methods designed and developed in the scope of the project (activity sensor and phase-out sensor methods) are valid, reproducible contributions. As they concentrate around the business-specific applications, however, they are optimised for the after-sales predictive analytics business domain.

Lastly, the exemplifications of the DISPA instantiation, including exemplifications of case-specific parameters, such as the evaluation function, can be considered contributions specific to the given case environment, even though there is some generic value in exemplifying the model instantiation as a tool to facilitate the future diffusion of the model. The contributions and their classifications are presented in Table 16.

Table 16 – Contributions structured by the level of generalisability.

Generic Contributions	After-sales-specific Contributions	Case-specific Contribution
DISPA framework	Forecast method evaluation	Case-specific forecast method evaluation (DISPA instantiation)
Design principles for designing predictive information systems	Activity-sensor based forecast method	Evaluation function specific to the case function
Matrix for pre-assessing financial feasibility	Phase-out sensor forecast method	
Flow chart presenting a process of instantiation of the DISPA framework		

5.2 Designing Predictive Analytics in Practice: Relevance of DISPA Framework

After investigating the case of the MAN Diesel & Turbo predictive sensor base system and a few sensor-based predictive systems described in the introduction, it becomes apparent that effectively embedding sensor-mediated context into an IS artefact is not a trivial task. The difficulty arises from the fact that, due to heavy context dependency, the actual actions that are performed in the process are difficult *DISPA as a* to generalise. Integrating sensor input into a system for *facilitating* predicting spare-part use patterns in a diesel engine will *framework* entail completely different analysis and actions than using customer GPS tracking information to enhance the efficiency of ads displayed on a mobile phone, although both problems belong to the class of predictive analytics with sensors. The claim of the DISPA model relevance relies on the assumption that it makes the challenging task of designing and developing such a system easier. We consider how this kind of claim can be constructed and advocated.

The DISPA framework is not the only lens that can be applied to the class of problems enhancing predictive analytics with sensor data; DISPA itself is based on two such models. The value of the framework and the value of the research that produced it must be determined based on how, at least in *Generalisability* some context, DISPA can prove to be (more) useful, providing *and utility as a* additional guidelines or conceptual perspectives facilitating *trade-off* the design process beyond what alternative frameworks could deliver.

Due to the complexity of the task in question and its specificity to a context, all usable models are process oriented, as activity-based abstractions would not be

generalisable across cases. The level of abstraction of the model is a key design decision to consider. More general frameworks are applicable in many more contexts but provide a rather loose guideline, while specific models have more limited scope but can be more specific in the way they inform their users.

On that scale, by limiting the scope to designing predictive analytics, the DISPA framework can give its users the benefit of more specific guidelines in comparison to generic DSR. Furthermore, by building on Shmueli and Koppius's (2011) steps for predictive empirical research, it can encapsulate the knowledge of predictive models into the multi-iteration IS design paradigm, providing guidelines on the highest level within the scope of designing predictive IS. Summarising, the claim of the DISPA framework utility is built on its balance between generality and specificity, allowing DISPA to facilitate a design theory of predictive systems.

5.3 DISPA Design as an Exemplary Design Science Research Project

In this dissertation, I introduced a framework facilitating the process of rigorously designing predictive IS. The model was evaluated based on a case study that showed that the framework can provide useful guidelines to develop environment-specific sensor-based predictive models that can outperform state-of-the-art predictive methods in a given environment. Generalising this observation, in the absence of a one-size-fits-all solution, customised, context-specific ways of creating predictive designs will be gaining popularity, especially when considering the inevitable growth of the IoT and sensor technologies. For these approaches, this model can provide both a structure and a rigorous guideline, as it has proven in the example case.

DISPA as a facilitating framework for problem specific predictive analytics

To evaluate whether the research activities followed the requirements for effective DSR, we revisit the design principles for design science in information system research introduced by Hevner et al. (2004; see Figure 30).

Project execution of guideline 1

Looking at the overall project activities in relation to the fulfilment of Guideline 1, multiple artefacts have been developed. In the most literal DSR approach, during Instantiation 1, three IT artefacts for spare-part forecasting have been developed, namely, the Croston forecasting system (Iteration 1), the phase-out sensor forecasting system (Iteration 2) and the activity-sensor forecasting system (Iteration 3). The prototypes presented in Instantiation 2 also comply with the IT artefact definition. However, as they were developed in conjunction with the project activities, I am not presenting them as the outcome of this research process. I merely evaluate them using my theoretical lens of cost and

profitability. Nevertheless, the main artefact of the project is the DISPA framework.

Guideline	Description
Guideline 1: Design as an Artifact	Design-science research must produce a viable artifact in the form of a construct, a model, a method, or an instantiation.
Guideline 2: Problem Relevance	The objective of design-science research is to develop technology-based solutions to important and relevant business problems.
Guideline 3: Design Evaluation	The utility, quality, and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods.
Guideline 4: Research Contributions	Effective design-science research must provide clear and verifiable contributions in the areas of the design artifact, design foundations, and/or design methodologies.
Guideline 5: Research Rigor	Design-science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact.
Guideline 6: Design as a Search Process	The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment.
Guideline 7: Communication of Research	Design-science research must be presented effectively both to technology-oriented as well as management-oriented audiences.

Figure 30 – Design science research guidelines (Hevner et al., 2004).

The criterion of relevance, made explicit in the second guideline, has also been established. The project was initiated by a company, the industrial partner, to address a real-life problem with a potential for strong bottom line *Project* improvement. Throughout Chapter 2, I introduced relevant *execution of* recent activities in academia and industry, showing that the *Guideline 2* class of problems I researched can be found in current scholarly works. Furthermore, I found reliable academic sources within IS that view sensor-based content (Chen et al., 2012) and predictive analytics (Shmueli & Koppius, 2011) as emerging research opportunities, suggesting that the class of predictive analytical applications using sensor technologies will only increase in popularity, providing opportunities for utilising the DISPA framework and proving that the problem faced by MAN Diesel & Turbo was not company specific.

Regarding design evaluation, which is the topic of Guideline 3, an explicit evaluation method was established at every step of the design research activities. As the DISPA framework enforces objective evaluation by introducing the evaluation function as part of the evaluation step, all artefacts developed using DISPA correctly will comply with Guideline 3. In the research activities undertaken, the evaluation function was introduced either as a cost of forecast error (in the first instantiation) or the baseline implementation cost requirement (in the second instantiation). In addition, the quantitative evaluation performed using a specific evaluation function (DISPA) introduces a qualitative evaluation aimed at identifying any systematic prediction biases that can be later addressed by upcoming design improvement activities. As a result, the implemented design evaluation approach is a combination of observational and analytical approaches, as classified by Hevner et al. (2004). For the DISPA artefact, the section titled 'Framework validation: DISPA elements in the existing literature' investigates how elements and functions of the framework are already present in the current literature, establishing its validity.

Project execution of guideline 3

The project provides clearly articulated research contributions, as requested by Guideline 4. I consider the DISPA framework to be the main contribution, and it can be considered a contribution to the methodological aspect for a class of problems dealing with predictive analytics and sensors. In addition, design principles for designing predictive IS and a matrix for pre-assessing the financial feasibility of predictive IS with sensor technologies complements the toolkit for the designer of predictive IS.

For researchers dealing with forecasting methods, several forecasting methods using sensor data, especially the

approach using activity sensors, can be considered a promising contribution to the spare-part forecasting literature. Finally, the project provides contributions to practice, as it introduces a practical framework to conduct predictive analytics with sensor data, which is a real-life gap introduced in the motivations section (Section 1.2) that pushed MAN Diesel & Turbo to become involved in this research project.

Research rigour has been applied throughout the project, as required by the fifth guideline. All the developed artefacts are based upon recent, relevant literature. The DISPA methodology, despite being only introduced in this work, is consistent with well-accepted methodological approaches developed by IS field authorities, such as Peffers et al. (2007), Hevner et al. (2004), and Shmueli and Koppius (2011). As a result, I can argue that the artefacts developed use both the knowledge base and sound research methodology. Furthermore, the predictive methods introduced comply with the requirements for predictive analytics introduced by Shmueli and Koppius (2011), including sample size, separation of training and test sets and data preparation requirements. Artefact performance matrices, referred to in this work as the evaluation function, were developed to be consistent with the case environment via close cooperation with multiple stakeholders in the company to ensure they encapsulated meaningful evaluation criteria.

Regarding presenting design as a search process, as suggested in Guideline 6, the DISPA model directly encapsulates Herbert Simon's (1996) generate/test cycle. Through specific evaluation, including quantitative analysis aiming to exhibit systematic biases that can directly lead to artefact improvement, the gap between the generate and test

steps is shortened and more explicit, facilitating the *satisficing* process (Simon, 1996). Anchoring the heuristic development in the case context, as explained in the previous paragraph, also makes assessing the measurement of the 'goodness' of the solution easier.

Finally, to ensure clear communication of the research, according to the seventh guideline, this thesis was structured according to the publication schema for DSR study, co-developed by the main author of the guidelines developed to facilitate 'communicating the new ideas to the stakeholder community' (Gregor & Hevner, 2013 p. 349).

Project execution of guideline 7 In addition to this thesis, the main pillars of this work have been presented in different academic outlets (Furtak, Avital, & Ulslev Pedersen, 2015; Pedersen et al., 2013; Ragnarsdóttir et al., 2012). The general foundations of this project, with the DISPA framework, have been presented at the ECIS conference (Furtak et al., 2015), and efforts towards publishing the research findings from this project in a yet more prestigious outlet have been undertaken. In addition to academic knowledge dissemination, communication efforts towards industry have also been undertaken. MAN Diesel & Turbo, the industrial partner in the project, is regularly updated with the research findings.

5.4 Validation and Evaluation in a Design Process Combining Physical and Abstract System Design

As argued in the previous section, although this design project can be considered an exemplar of the DSR, there is one element that distinguishes it from most DSR projects. The DSR can be used for the design of all IS artefacts, including

both physical systems and abstract models. In fact, in this project both types of IS artefacts have been designed.

Initially, the DISPA framework was designed by combining relevant contributions from the existing knowledge base (Hevner et al., 2004; Shmueli & Koppius, 2011) and the business needs from the environment (the need to design predictive systems using IoT/sensors) for a well-defined class of problems. This design activity constitutes the first iteration of the design project (one that is abstract); the designed artefact is a physical system framework development. This part follows all the specifics of the design science project, but it could also be conducted as a non-design paradigm project, with the literature review pointing towards a research opportunity addressed by the contribution.

This research project as a combination abstract and physical design project

The result of the abstract part, namely, the DISPA framework, is used as a methodology for the following physical system design to address the instance of the problem in the case context. The DISPA instantiation was carried out by defining the problem-specific evaluation function and contextualising the feedback mechanism for evaluation in the context environment. After the first instantiation, aiming at designing a better forecasting solution for the case company, another instantiation was conducted, this time with the objective to assess other initiatives to enable designing an IT system that could outperform the previous design solution in the case context. The complete process is presented in Figure 31.
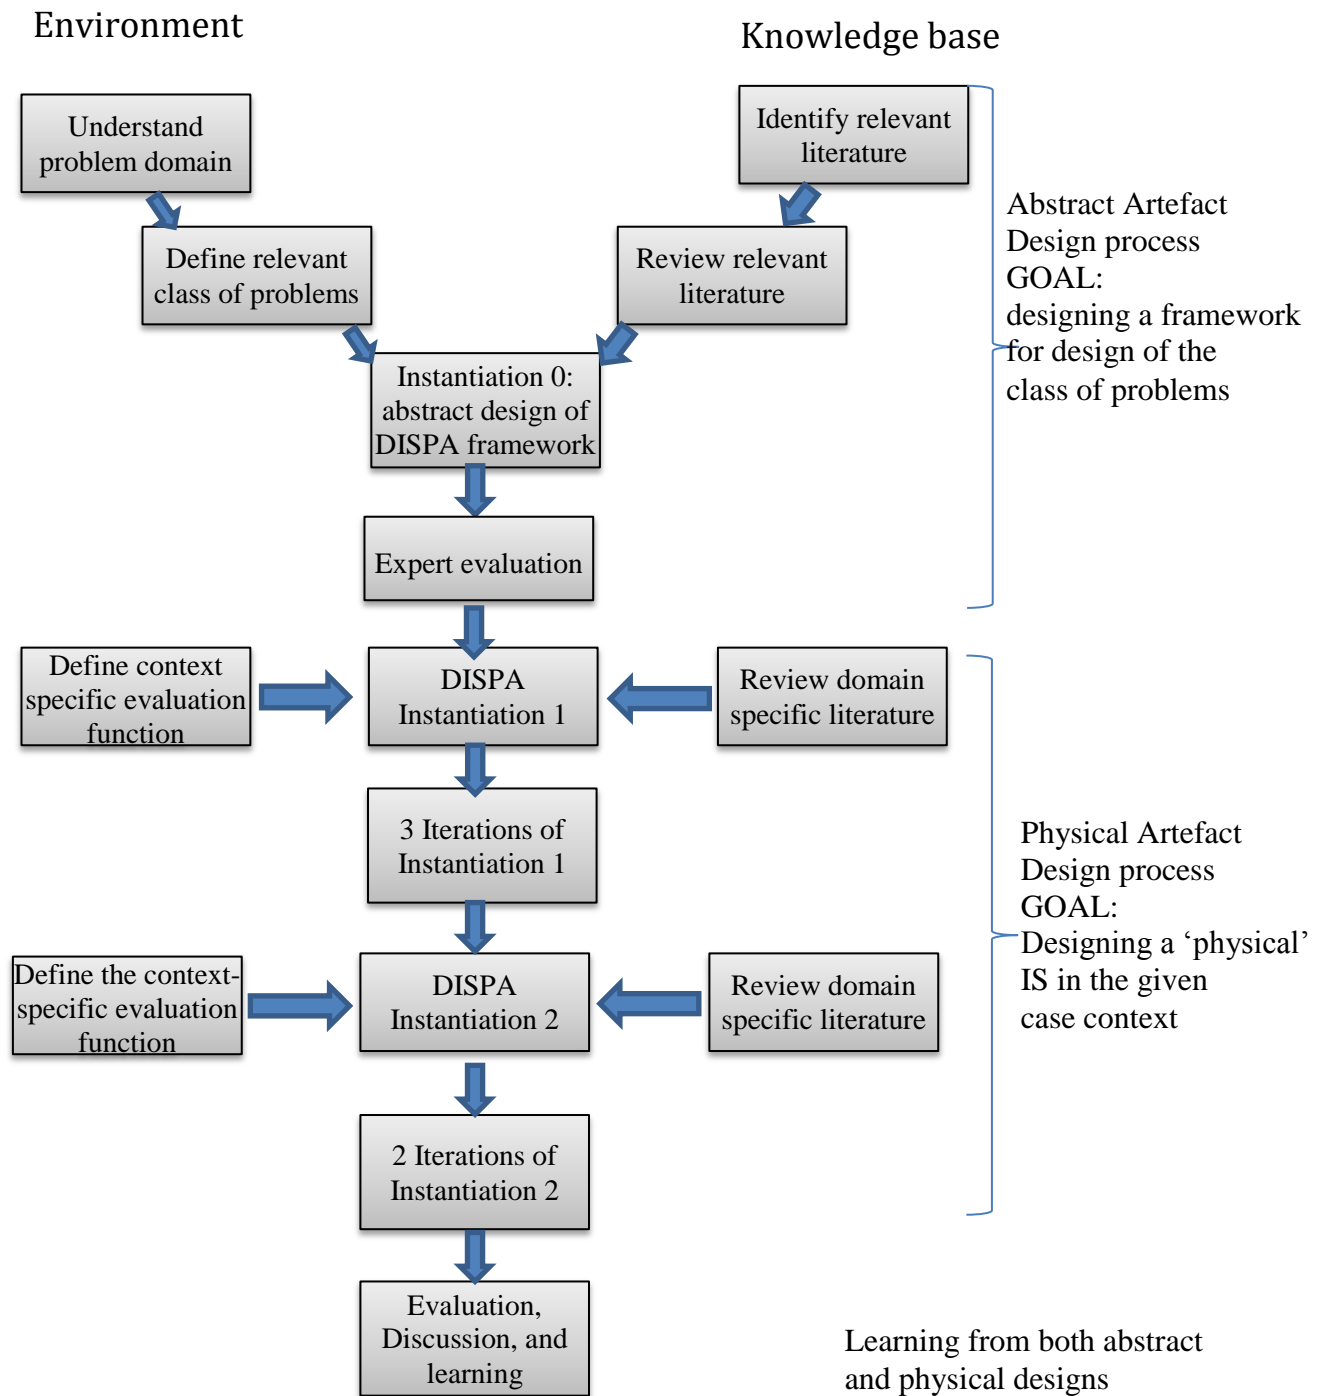
Figure 31 – Overall research process.

The combination of the abstract and physical artefact design introduces additional challenges in design evaluation and validation. According to March and Smith, the evaluation is

'the process of determining how well artefacts performs' (1995, p. 254). For physical designs, this performance is usually well established. In fact, any rigorous DSR project must discuss the quality and properties for evaluation criteria of the designed system. In this research project, the evaluation function for all instantiations has been defined, so that for evaluation of every physical system artefact, there was an easy-to-use quantitative evaluation tool.

For abstract designs aiming at designing a framework, the 'performance' of the artefacts is more difficult to assess. The challenge with the evaluation of an abstract process framework is that the subject of the evaluation will be the instantiation of the framework, not the framework itself. In this research project, I argue that DISPA is a framework facilitating the design of predictive systems with sensors. If I aimed to use scientific method to prove it, I would need to compare a 'performance' of two identical researchers working in two identical organisations on designing predictive IS, which is a situation that is neither realistic nor cost-effective. Even then, the evaluation would be an evaluation of the instantiation of the artefacts, but assuming conditions are the same, the differences in the instantiations would need to be linked to using or not using the framework.

Evaluation in abstract system design projects

To determine the best method to evaluate an abstract design process leading to DISPA development, I examined exemplary DSR designing abstract frameworks and analysed their validation strategies. As a result, I identified the two most common validation patterns. The first strategy is to illustrate the framework utility with field studies, exemplified by Pries-Heje and Baskerville (2008). The advantage of this strategy is that it not only provides insight into the model validity and utility but also provides guidelines and examples for the

Strategies in abstract system design projects

researcher who would like to use the framework in the future. The downside of this approach is that it only exhibits the utility of the framework when executed by the framework's author.

The alternative strategy is to illustrate and evaluate the framework by using it to analyse pre-existing studies. An exemplar of this approach is the work by Venable et al. (2016) on evaluation and the renowned DSR paper by Hevner et al. (2004) when validating the DSR guidelines. This approach allows us to pinpoint elements of the framework in existing literature, helping to some extent to address the issue of the framework utility that is limited to its author. On the other hand, the ex-post nature of this approach might be considered artificial and less informative for researchers planning to use the framework.

Validation in this research project

In this research project, I chose to use both validation approaches for the abstract design process. The process outcome, the DISPA framework, has been used to utilise the framework utility with two field studies, namely, Instantiations 1 and 2. The ex-post analysis of existing literature has also been performed, and its results are presented in the following section. The evaluation of the physical system design activities has been done using evaluation criteria embedded in the DISPA instantiation steps.

5.4.1 Generalising DISPA utility from the single case study: DISPA elements in the existing literature

The presented research project suffers from a major weakness; I build a case of framework utility based on its fit to a single problem domain, addressed by the author of the framework. This section aims to determine whether the framework can be used by researchers other than the author through a review of the existing literature, investigating the extent to which the existing contributions in forecasting for spare parts are consistent with the framework structure.

To conduct this analysis, I examined several spare-part forecasting contributions, marking sections that code the steps and sub-steps from the model (such as goal definition, model definitions, and quantitative method evaluation) and created a paper structure based on the presence and order of the steps discovered. Furthermore, the structure of the papers was compared and generalised into schemas describing predictive analytics, grouping publications that were similar regarding implementing the DISPA steps in the way in which the paper presented the predictive models.

Assessing utility of DISPA by other than the author

One important observation is that the DISPA framework describes a process of designing a predictive IS artefact and that this process does not necessarily need to be directly reflected in the paper structure. Nevertheless, finding similarities to the DISPA-structured process in the structure of the article would allow me to claim that the framework is usable and, to some extent, that the concept behind it has, in fact, already been used for some time.

The DISPA framework process schema

After analysing a sample of papers, I identified three schemas for describing predictive analytics: nominal, single-design,

and abstract. Nominal schema follows all the steps specified in the DISPA framework while introducing multiple design iterations with empirical evaluations. An example of exactly this kind of reporting style in the predicting spare-part demand domain is that of Teunter and Duncan (2009), but numerous other works follow an extremely similar approach (Hong, Koo, Lee, & Ahn, 2008; Leitner & Leopold-Wildburger, 2011). The schema is presented in Figure 32 without any modifications.

Another identified approach is the abstract reporting schema. In this approach, the deviation from the DISPA is that, in Steps 2b and 2c, data preparation and model implementation are implicit. This schema is exemplified in papers such as those by Dolgui and Pashkevich (2008) and Ghodrati et al.

Abstract reporting schema (2012). As the papers present the results of the model evaluation, the model must have been implemented, but there is no explicit description of how the abstract model was instantiated, nor are there any details concerning how the data were pre-processed before being fed into the model. The cause of the lack of this information might be an assumption that there is no difference between the performance characteristics of the abstract model and the implemented method. The schema is presented in Figure 32, and the deviation from DISPA is marked with a blue rectangle.

Single iteration schema The third identified reporting schema is reporting a single iteration. In the single iteration, after the design is presented, it is evaluated with the process of designing the artefact. Only one design is proposed, although a state-of-the-art baseline is typically presented, and the new design is evaluated against it. In the evaluated sample of work presenting predictive analytics, this schema was by far the most popular, with such examples as those by Hua, Zhang, Yang, and Tan, (2007),

Louit et al. (2011), and Tibben-Lembke and Amato (2001). Viewing this reporting process through the DISPA lens, this approach results in merging Steps 3 and 4 (see Figure 32 and the deviations marked in green).
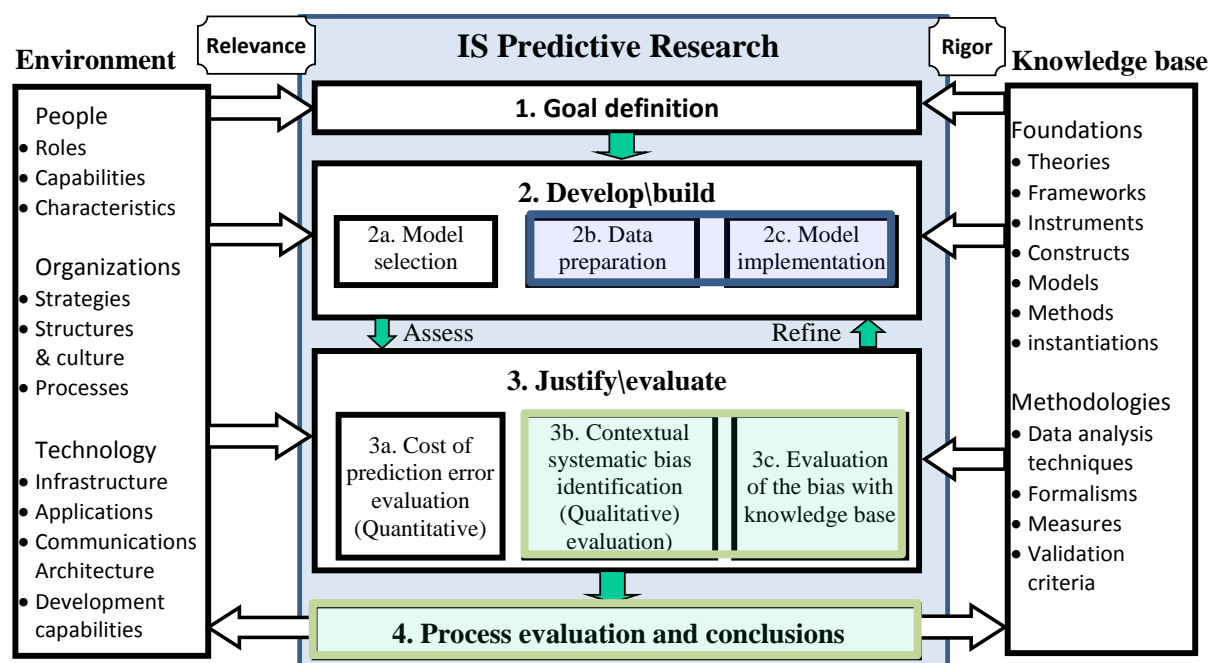
Figure 32 – DISPA model with deviation introduced by discovered predictive analytics reporting schemas (blue rectangles mark skipped steps in an abstract schema, while green marks steps unified in the single iteration approach).

One more common pattern observed was inverting Steps 1 and 2a, thus presenting the new model first and then the quantitative evaluation measure to assess it. This ordering seems to be introduced in the reporting to stress the main contribution of the paper, as it would be uncommon to generate a design outperforming a state-of-the-art solution without knowing the performance measure. Overall, presenting only one iteration does not necessarily mean that only one design was developed, as the authors could have simply decided to skip the process that led them to their final model and only present the output.

Concluding the section regarding the existing predictive analysis literature and their relation to the DISPA framework, the elements of the framework are well represented in the literature. In a sense, this is not surprising, as the framework itself is based on two popular models and scholarly contributions. Therefore, using them will use elements of the DISPA. My argument is that, although the DISPA has only been instantiated once in one context and only by the author of the framework, it can still be used by other scholars in other contexts. In addition to this line of argumentation, I provide a guideline regarding how to instantiate DISPA in any context in Appendix B.

5.5 Theory in Information Systems and the Design Research Community

5.5.1 Theory in information systems

The notion of theory is a crucial one to any academic discipline, as it defines the epistemological foundations of a discipline, stating how knowledge is acquired in a community. From this perspective, the IS research stream is no different. There has been a lot of interest in the ontology (what?), epistemology (how to develop it?), and use of theory in IS literature. I present some of the views on what theory is in Table 17.

Ontology of theory in IS as a debate

Table 17 – Different views on theory in the literature.
'A formal statement of the rules on which a subject of study is based or of ideas that are suggested to explain a fact or event or, more generally, an opinion or explanation'. (Cambridge Dictionary, 2016)
'A system of assumptions, principles, and rules of procedure devised to analyze, predict, or otherwise explain the nature or behavior'. (Chatterjee, 2015; American Heritage Dictionary, 2016)
'Statements that say how something should be done in practice'. (Cushing, 1990; Davis & Olson, 1984; Gregor, 2006)
'Statements providing a lens for viewing or explaining the world'. (Gregor, 2006; Orlikowski & Robey, 1991; Walsham, 1995)
'Statements of relationships among constructs that can be tested'. (Davis, 1989)
'Nets cast to catch what we call 'the world' to rationalize, to explain and to master it. We endeavor to make the mesh even finer and finer'. (Popper, 2005)

The main rationale behind presenting these multiple perspectives on theory is to show that the term is used in IS with different meanings, without a single common agreement. One elementary reason for this confusion is the partially empirical nature of IS research. According to Lewin, nothing is so practical as a good theory, but the presence of empirical research amplifies diversity in the epistemological attitude to research and hence theory in general. In his seminal work, Simon had a very eloquent way of phrasing the challenge of 'empirical theory':

The quest of practical theory

> How could one construct an empirical theory?
>
> I thought I began to see in the problem of artificiality an explanation of the difficulty that has been experienced in filling engineering and other professions with empirical and theoretical substance distinct from the substance of their supporting sciences. Engineering, medicine, business, architecture and painting are concerned not with the necessary but with the contingent—not with how things are but with how they might be—in short, with design. (1996, p. 12)

To harmonise and embrace this diversity of 'theories', Gregor (2006) introduced a taxonomy of theory types in IS research (Figure 33). The taxonomy defines five distinct types of theories in IS, presented as Type I to V. Type I is purely descriptive, while Type II is an explanatory theory, extending the descriptive element by providing a clarification concerning how or why the theory holds. Type III is a predictive theory, while Type IV combines the explanatory and predictive components. The mentioned four theory types are present it most of other academic disciplines and are based on the domain of natural science. The same cannot be said about the Type V theory, which is a theory of design and

Gregor's taxonomy of IS theory (2006)

154

action. This theory type is reserved for the science of the artificial, as Simon (1996) called it, so the domain of manufactured Type V theories defines how to do things, defining 'principles of form and function' (Walls et al., 1992).

With such an open definition, distinct from the other four types of definitions, a valid question to ask is what makes such a contribution to knowledge 'good'. In the literature, such criteria include the utility, novelty, and persuasiveness of claims that it is efficient (March & Smith, 1995) as well as, especially for models, the criteria of completeness, simplicity, consistency, ease of use, and quality of results (Hevner et al., 2004). One more interesting dimension pointed out by Simon (1996) is 'interestingness'.

What constitutes a good design theory?

Theory Type	Distinguishing Attributes
I. Analysis	Says what is. The theory does not extend beyond analysis and description. No causal relationships among phenomena are specified and no predictions are made.
II. Explanation	Says what is, how, why, when, and where. The theory provides explanations but does not aim to predict with any precision. There are no testable propositions.
III. Prediction	Says what is and what will be. The theory provides predictions and has testable propositions but does not have well-developed justificatory causal explanations.
IV. Explanation and prediction (EP)	Says what is, how, why, when, where, and what will be. Provides predictions and has both testable propositions and causal explanations.
V. Design and action	Says how to do something. The theory gives explicit prescriptions (e.g., methods, techniques, principles of form and function) for constructing an artifact.

Figure 33 – Taxonomy of theory types in information systems research (Gregor, 2006).

Gregor (2006) supplemented her typology by specifying seven distinct components of theory: four mandatory components and three optional components, dependent on the type. The mandatory components are representation, constructs, statement of relationships, and scope. Additionally, a theory could have casual explanations, a testable hypothesis, or prescriptive statements. The theory components by Gregor are presented in Figure 34.

Theory components by Gregor (2006)

As I am most interested in theory in the design paradigm. I would like to single out Type V from the four others. Using the lens of components, the design and action theory type has a prescriptive element, defining how something should be done and the scope (in what kind of situations it should be done). An interesting debate starts when we examine in detail the means of representation, constructs, and statements of relation in the design research.

Theory Component (Components Common to All Theory)	Definition
Means of representation	The theory must be represented physically in some way: in words, mathematical terms, symbolic logic, diagrams, tables or graphically. Additional aids for representation could include pictures, models, or prototype systems.
Constructs	These refer to the phenomena of interest in the theory (Dubin's "units"). All of the primary constructs in the theory should be well defined. Many different types of constructs are possible: for example, observational (real) terms, theoretical (nominal) terms and collective terms.*
Statements of relationship	These show relationships among the constructs. Again, these may be of many types: associative, compositional, unidirectional, bidirectional, conditional, or causal. The nature of the relationship specified depends on the purpose of the theory. Very simple relationships can be specified: for example, "x is a member of class A."
Scope	The scope is specified by the degree of generality of the statements of relationships (signified by modal qualifiers such as "some," "many," "all," and "never") and statements of boundaries showing the limits of generalizations.
Theory Component (Components Contingent on Theory Purpose)	Definition
Causal explanations	The theory gives statements of relationships among phenomena that show causal reasoning (not covering law or probabilistic reasoning alone).
Testable propositions (hypotheses)	Statements of relationships between constructs are stated in such a form that they can be tested empirically.
Prescriptive statements	Statements in the theory specify how people can accomplish something in practice (e.g., construct an artifact or develop a strategy).

Figure 34 – Theory components by Gregor (2006).

5.5.2 Design theory and capturing design knowledge

Summarising the previous section, traditionally the notion of theory in IS is conceptually anchored in natural science and the natural world domain. There seems to be a conceptually different domain dealing with manufactured things, which are not natural but artificial, hence the phrase 'science of the artificial'. Research shaping an environment is conceptually different than describing, explaining, or predicting an environment, which arguably requires different requirements *IS theory vs.* for the theory construction process (Chatterjee, 2015; Cross, *Design* 2001; Gregor, 2006; Simon, 1996). The difference is *Theory* pronounced in the Type V theory in the typology by Gregor (2006), shown in Figure 33. This systematic difference triggers a discussion regarding what components (like those in Figure 34) and goals should be used in design science theory.

In recent literature, we witnessed an interesting discussion regarding how design knowledge can be expressed as theory. Some researchers argue that the benefits of design research can be enjoyed without explicit theory formulation (Gregor & Jones, 2007), but the majority agree that there is a need for *Design* 'more formal knowledge of shape and configuration' (Cross, *knowledge* 2001; Hevner et al., 2004; Gregor & Johnes, 2007; Kuechler & *and theory* Vaishnavi, 2005). A challenge in formalising design knowledge and a distinction between craft and research is eloquently phrased by Cross:

> We must not forget that design knowledge resides in products themselves: in the forms and materials and finishes which embody design attributes. Much everyday design work entails the use of precedents or previous exemplars – not because of laziness by the

designer but because the exemplars actually contain knowledge of what the product should be. This is certainly true in craft-base design: traditional crafts are based on the knowledge implicit within the object itself of how best to shape, make and use it. This is why craft-made products are usually copied very literally from one example to the next, from one generation to the next. (2001, p. 4).

Craft vs research perspective on design The point by Cross (2001) is that there are different requirements for design as a craft level and as a research discipline, a point nicely summarised by Gregor and Jones: 'A craft can proceed with copying of a design artefact by one artesian after another. A discipline cannot' (2007, p. 312). The debate on what craft and what research are follows the same lines as what is and is not theory. A very strong voice in that discussion is a classic paper by Sutton and Staw from 1995, defining what theory is not. The point that they make is that:

- References are not theory,
- Data are not theory,
- List of variables or constructs are not theory,
- Diagrams are not theory,
- Hypothesis (or predictions) are not theory, and
- A statistical model is not theory (Sutton & Staw, 1995).

Anatomy of a design theory (2007) To facilitate specifying design theory so that it can be communicated, justified, and developed cumulatively as a discipline, Gregor and Jones (2007) created the anatomy of a design theory. One of the main conclusions from this seminal work is the definition of the eight components of an IS design theory, presented in Figure 35. Six components are defined as a minimum requirement while two are advised to increase the credibility of the theory.

Component		Description
Core components		
1)	Purpose and scope (the *causa finalis*)	"What the system is for," the set of meta-requirements or goals that specifies the type of artifact to which the theory applies and in conjunction also defines the scope, or boundaries, of the theory.
2)	Constructs (the *causa materialis*)	Representations of the entities of interest in the theory.
3)	Principle of form and function (the *causa formalis*)	The abstract "blueprint" or architecture that describes an IS artifact, either product or method/intervention.
4)	Artifact mutability	The changes in state of the artifact anticipated in the theory, that is, what degree of artifact change is encompassed by the theory.
5)	Testable propositions	Truth statements about the design theory.
6)	Justificatory knowledge	The underlying knowledge or theory from the natural or social or design sciences that gives a basis and explanation for the design (kernel theories).
Additional components		
7)	Principles of implementation (the *causa efficiens*)	A description of processes for implementing the theory (either product or method) in specific contexts.
8)	Expository instantiation	A physical implementation of the artifact that can assist in representing the theory both as an expository device and for purposes of testing.

Figure 35 – Eight components of an information system design theory.

I believe that it is very interesting to compare the seven theory components in IS by Gregor (2006) in Figure 34 with eight components of an IS design theory co-authored by Gregor (Gregor & Jones, 2007) in Figure 35. My general assumption is that IS design theory is a specific case of an IS theory, which means that the later classification must comply with the initial one, possibly being more specific on some aspects. The comparison is presented in Table 18.

Table 18 – Comparison of theory components in all IS theory and IS design theory; descriptions as provided by the authors; comparison added.

Structural Components Common to All Theory in IS (optional components in italic)		Core Components of an Information System Design Theory	
Component	Definition (Gregor, 2006)	Component	Description (Gregor & Jones, 2007)
Scope	The scope is specified by the degree of generality of the statements of relationships (signified by modal qualifiers such as 'some', 'many', 'all', and 'never') and statements of boundaries showing the limits of generalizations.	Purpose and scope	'What the system is for', the set of meta-requirements or goals that specifies the type of artefact to which the theory applies and in conjunction also defines the scope or boundaries of the theory.
Constructs	These refer to the phenomena of interest in the theory (Dubin's 'units'). All of the primary constructs in the theory should be well defined. Many different types of constructs are possible: for example, observational (real) terms, theoretical (nominal) terms and collective terms.	Construct	Representations of the entities of interest in the theory.
Means of representation	The theory must be represented physically in some way: in words, mathematical terms, symbolic logic, diagrams, tables or graphically. Additional aids for representation could include pictures, models, or prototype systems.	Principle of form and function	The abstract 'blueprint' or architecture that describes an IS artefact, either product or method/intervention.
Testable propositions (hypotheses)	*Statements of relationships between constructs are stated in such a form that they can be tested empirically.*	Testable propositions	Truth statements about the design theory.
Statements of relationship	These show relationships among the constructs. Again, these may be of many types: associative, compositional, unidirectional, bidirectional, conditional, or causal. The nature of the relationship specified depends on the purpose of the theory. Very simple relationships can be specified: for example, 'x is a member of class A'.	Justificatory knowledge	The underlying knowledge or theory from the natural or social or design sciences that gives a basis and explanation for the design (kernel theories).
		Artefact mutability	The changes in state of the artefact anticipated in the theory, that is, what degree of artefact change is encompassed by the theory.

Comparing IS theory and Design Theory components

The top three generic IS theory components, scope, construct, and means of representation, all find a close match in the specific design theory components, phrased in a more functional way ('requirements' instead of 'relationships' or engineering-rooted 'blueprint' as a means of representation). The testable proposition component is set to be mandatory in design theory, while it is only optional in generic IS theory, contingent on theory purpose and linked with theory types dealing with prediction, such as Types III and IV in Gregor's

typology in Figure 33. An interesting observation is that setting this component as mandatory in design theory, on top of the principle of the form and function component, makes design theory both design and action (Type V) and predictive (Type III or IV) according to the typology. In practice, that creates a requirement for design theory both to explicate how to do something and to predict something.

Another interesting observation can be made when we analyse mandatory components that are not yet discussed. There seems to be no explicit match between statements of relationships on generic IS theory and justificatory knowledge and artefact mutability in design theory. We reflect again that the design theory is an IS theory, meaning that the equivalent of the statement of relationships component must be present in design theory.

Implicit requirement for design theory to be predictive

Starting with this assumption and supposing that it is not included in the previously discussed four design theory components (purpose and scope, constructs, principles of form and function, and a testable hypothesis), we must conclude that a statement of relationships is included in justificatory knowledge and artefact mutability. This could easily be the case if we assume that an artefact is the unit of analysis and that we describe relationships between different forms and versions of a mutable artefact that is explained and informed by some justificatory knowledge (similarly to how figure 36 looks like). In that scenario, the testable propositions could also be about the artefact, with its functional performance being better or worse if certain design steps are taken.

Artefact as a unit of analysis

In the context of this research project, this is how I define the components of the implemented design theory, the framework for DISPA. Its main artefact is a process model that defines how to define case-specific artefacts. This method of using an artefact as a unit of analysis presents relationships between multiple instances of artefacts, showing artefact mutability, which is informed by justificatory knowledge used to cross-evaluate constructed artefact instances. A more detailed analysis of contributions of this project considering the notion of theory in IS and in design research is presented in the following chapter.

DISPA as an exemplar of design research with the artefact as a unit of analysis

5.6 DISPA as an Information Systems Design Theory

In this work, we introduced a toolkit framework facilitating the process of rigorously designing information systems with predictive analytics (DISPA). The framework consists of a design process model and a set of guidelines and assessment tools. It was evaluated based on a case study that showed that it can provide useful guidelines to develop environment-specific sensor-based predictive models that can outperform state-of-the-art generic predictive methods in an environment. Generalising this observation, we predict that, in the absence of a one-size-fits-all solution, custom context-specific predictive designs will be gaining popularity, especially when considering the almost inevitable growth of the IoT and the popularity of sensor technologies and digital traces. For these approaches, this framework can provide both a structure and a rigorous guideline, as it has proven in the example case.

DISPA as a facilitating framework for problem specific predictive analytics

Our framework constitutes an IS design theory according to the definitions of Walls et al. (1992), Gregor and Jones (2007), and Kuechler and Vaishnavi (2012). According to Walls et al. (1992), an IS design theory is a set of primary prescriptive statements describing how a class of artefacts should behave and how they can be constructed. Additionally, Gregor and Jones (2007) put a requirement for IS design theory, later confirmed by Kuechler and Vaishnavi (2012) to include 'justificatory knowledge', knowledge or theory from natural, social, or design science, explaining the design principles (also referred to as kernel theories).

DISPA as an IS design theory

Overall, Gregor and Jones (2007) proposed eight components that constitute IS design theory, out of which six are

mandatory. We claim that the presented framework includes all eight components, fulfilling Gregor and Jones's (2007) requirements for an IS design theory (see Table 19).

Table 19 – DISPA as an information system design theory.

Component	Description provided by Gregor and Jones (2007)	Reflection in the Framework
Purpose and scope	*'What the system is for', the set of meta-requirements or goals that specifies the type of artifact to which the theory applies and in conjunction also defines the scope, or boundaries, of the theory.*	The framework is intended to facilitate the design process for IS artefacts using context-aware predictive analysis.
Constructs	*Representations of the entities of interest in the theory.*	Examples are steps of developing and evaluating IS artefacts, the environment, and the knowledge base.
Principle of form and function	*The abstract 'blueprint' or architecture that describes an IS artifact, either product or method/intervention.*	The process description for DISPA describes the method/intervention.
Artefact mutability	*The changes in state of the artifact anticipated in the theory, that is, what degree of artifact change is encompassed by the theory.*	The process model encapsulates the iterative character of the model by Hevner et al. (2004), encouraging artefact mutability based on the evaluation.
Testable proposition	*Truth statements about the design theory.*	The process leads to improving artefact performance by exploring and evaluating other promising solutions and choosing the best one.
Justificatory knowledge	*The underlying knowledge or theory from the natural or social or design sciences that gives a basis and explanation for the design (kernel theories).*	The framework relies on related design theories (Hevner et al., 2004; Shmueli & Koppius, 2011) and the domain predictive analytics knowledge (Bacchetti & Saccani, 2012).
Principles of implementation	*A description of processes for implementing the theory (either product or method) in specific contexts.*	The process model explicitly defines steps for implementations of the theory.
Expository instantiation	*A physical implementation of the artifact that can assist in representing the theory both as an expository device and for purposes of testing.*	The instantiation based on the case of MAN Diesel & Turbo is provided.

The DISPA framework has a well-defined purpose and scope, which is facilitating the development of context-specific artefacts using predictive analytics. The process model of DISPA has well-defined constructs, understood as steps in the process of designing the artefacts, such as the model definition or quantitative error evaluation. The diagram of the process model is part of the blueprint of the theory, as is the feasibility matrix and the table with guidelines. The process of developing and evaluating artefacts, which the process model defines, includes artefact mutability as an essential part. *DISPA's* There is a prediction component, operating under the *fulfilment of* assumption that using evaluation to discover pitfalls of one *Gregor and* version of the artefact will lead to development of new *Jones's* artefacts that will perform better in the evaluation step, *criterion* effectively improving the performance of the designed solution. The framework is informed by justificatory knowledge and enables inserting such knowledge into the design process it facilitates. As it is a process model, it explicitly defines how to instantiate it in the specific context, providing a framework tutorial to make it easier for future users as well as an example that is an expository instantiation of the theory.

Consistent with Gregor and Jones (2007), the paper is also in agreement with the recent work by Mandviwalla (2015) on *DISPA's* design theory. Our project falls into the Type VIII DSR project *consistency* type, where a new artefact is generated, and a kernel theory *with other* is appropriated. It is also consistent with an 'ultimate' goal *definitions of* that Rossi, Henfridsson, Lyytinen, and Siau (2013) put *design* forward for theorising in DSR; we deliver 'an organized set of *theory* principles' for a certain well-defined domain.

5.7 Post-hoc Description of Interactions in the Case Company and the Knowledge Base

Introducing map of critical interactions in the environment

Any research project conducted in a practical setting is influenced by the case environment. To explicate the effect of the holistic case environment and individual actors in the case environment on the project results, I created a post-hoc map of the critical interactions with the relevant actors in the case company. The map of the critical interactions is presented in Figure 36. The actors relevant for the interactions in the case company are described in the introduction, in the section titled 'Engaged scholarship in practice: Practical project setup' (Section 1.2.3).

Understanding problem domain

The research project started with a set of activities aimed to understand the problem domain of predictive analytics in sensors in the context of the case company. Already, management, mostly from spare-part sales and the supply chain areas, were involved to deliver the initial set of requirements and expectations for the project. The initial step of the project was to understand the context of these requirements and the overall problem in the company. In parallel to the investigations in the environment, similar problems outside of the case context were investigated to identify the relevant literature. This investigation led to the definition of the relevant class of problems and enabled the presentation of the problem at hand as an example of that class, with particular characteristics explicitly discussed. The literature previously used for establishing a rough overview was then thoroughly reviewed.
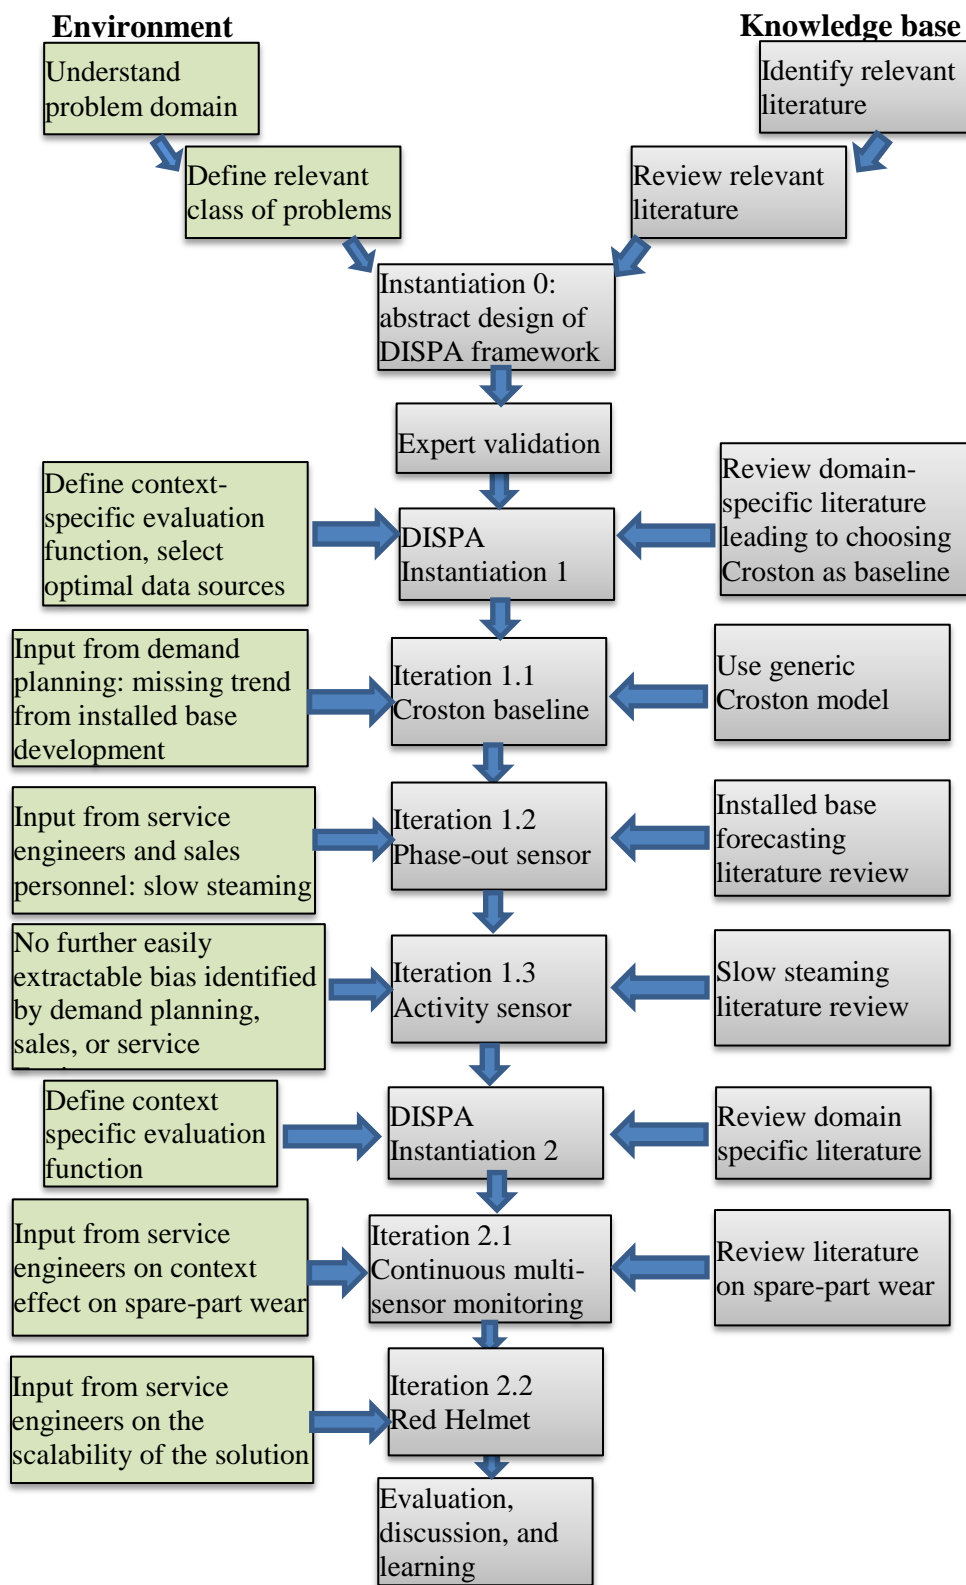
Figure 36 – Overview of critical interactions with the case environment and the knowledge base in the project period.

After understanding the problem context, both in the case company and in the academic literature, the next step was to establish a methodology that could facilitate designing an optimal artefact to address the challenge at hand. This step, the abstract design of the DISPA framework, was probably *Establishing* the most time-consuming and the one I backtracked to most. *methodology* After the framework was concluded, its instantiations followed, with a definition of the instantiation-specific evaluation function, identification of relevant data sources, and establishing meaningful business logic that can be used in the artefact evaluation process.

For Instantiation 1, the evaluation function was set to counterweight the cost of over-forecasting and under-forecasting, and the business logic and relevant data were aimed at investigating the financial effect of stock-outs and the cost of excess stock. A starting point was also established *Instantiation* using the literature to establish a state-of-the-art starting *1, baseline* point for the investigations. For the first iteration, the *and phase-* selected baseline was the Croston method, as a method *out sensor* common to academic research and a de-facto practitioner standard. The evaluation of the baseline was carried out using the evaluation function and insight from demand planning, suggesting the introduction of a link to IB development, leading to development of a phase-out sensor model.

The quantitative evaluation of this model was again performed using the previously developed evaluation function, and its results were further informed by sales force *Instantiation* and service engineers, who pointed out the changes on the *1, activity* market introduced by the 'slow steaming' behaviour. This *sensor* insight led to developing a model estimating the intervals between replacement in actual engine activity, not in time, called the 'activity sensor' model. The artefact was further

evaluated quantitatively, but the qualitative attempts to identify further bias that could be removed did not provide further results.

The activities were repeated for the second instantiation. A relevant evaluation function was designed, and it was used to sequentially evaluate innovative designs that were further *Instantiation* improved by the quantitative insight from service engineers. *2 and* The project was concluded be the overall discussion of the *discussion* process and results, and its findings were communicated both to the case company and academic community.

5.8 Implications

5.8.1 Internal and external validity of results

Design science research and a lot of prescriptive research in general does not deal with variable cause and effect, making it less feasible to statistically assess the validity of the results. Nevertheless, regarding the internal validity of the DSR activities relies on the assessment of the research output, the artefact, and the research process and its rigour. The assessment of the DISPA framework is done through its validation (section DISPA framework validation) and by ensuring its consistency with state-of-the-art literature (DSR and Shmueli and Koppius's models).

Rigor in DISPA's development process

The rigour of the process is built upon following the DSR methodology and ensuring research consistency with DSR guidelines (see previous section). The general purpose of internal validity is trying to avoid systematic bias. As a vehicle of conducting research, DISPA takes the same objective through directly incorporating activities to identify and correct for any systematic bias as a part of the validation of a model.
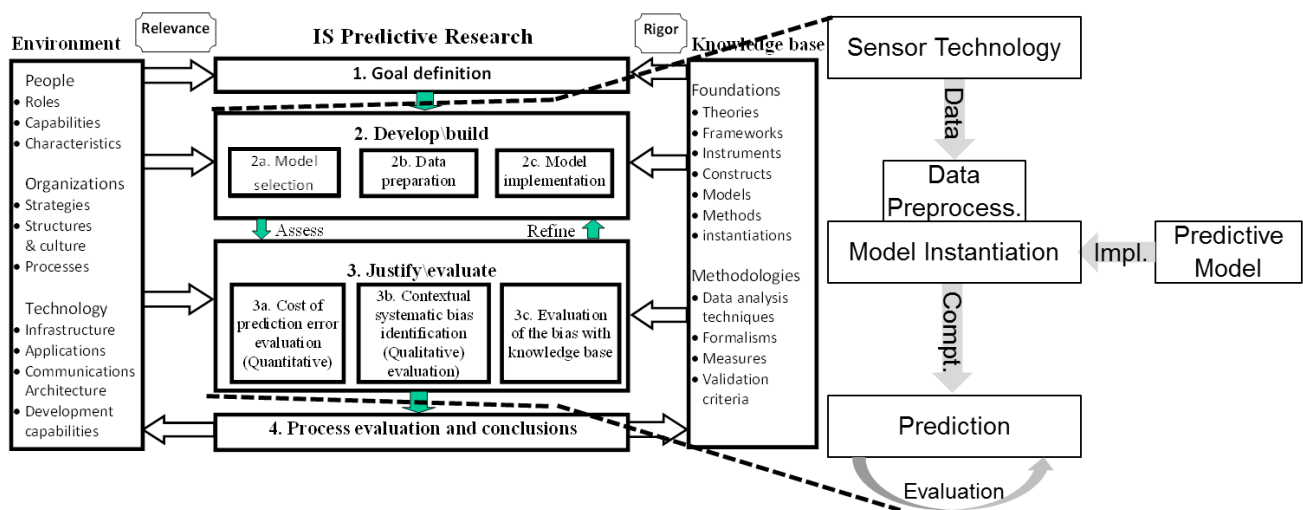
Figure 37 – DISPA steps as a reflection of activates present in every project dealing with predictive analytics with sensors.

Regarding external validity, I argue that, by basing the DISPA steps on the properties of including sensors in predictive analytics (including data preparation steps and considerations of model and implementation distinctions). More specifically, as presented in Figure 37, the actual sub-steps in develop/build are the natural consequence of implementing an abstract model using data sourced from sensor technology. Using a predictive model requires selecting a predictive model. Implementing it requires preparing (pre-processing) data to feed into the model. After the model is executed to make the exercise meaningful, there always must be an evaluation, specified here in a way that can facilitate finding potential improvements. In the DISPA framework, this cyclic process is wrapped around the goal definition (proceeding step) and the overall process evaluation (succeeding step).

DISPA steps as activities present in any predictive project with sensors

5.8.2 Overview of implications for theory and practice

We live in the age of data, when almost every step we take generates digital traces through sensors, social media, or

transactional systems. The ability to use this data, not only to observe current snapshots but to discover patterns that can be used for prediction of upcoming activities becomes a necessity for more organisations, and the actual applications of predictive analytics using it become increasingly complex. This project fits into this landscape because it reviews the activities, trends, and requirements around predictive analytics with sensor data from both academic and practical perspectives.

The world as a network of digital traces

The value of the project lies in its practical implications. I believe that the DISPA framework with the design principles for designing predictive information systems and the matrix for pre-assessing financial feasibility can be a vehicle facilitating the process of designing predictive IS using sensors and big data in the case-specific organisational context. I discuss how a framework dedicated specifically for developing predictive IS can be more informative and helpful for designers than generic design science, suggesting the initial process and starting steps for implementing a needed artefact. Embedding context through sensor data into organisational IS is complex and always very specific to a particular case. Providing a process description for those activities not only facilitates the activities but also enhances the diffusion of predictive IS in the industry.

DISPA as a practical framework

On the theoretical end, the DISPA framework can be considered a stub of a design theory for designing predictive IS. Although it is too early to discuss design theory as such, the only path to it is through the synthesis of practice and the knowledge base leading to the development of design prescriptions, a process referred to as the development of mid-range design theories (Chatterjee, 2015). I recognise that my efforts are a modest approach in the spectrum of theory

DISPA as a step towards design theory of predictive systems

development, but I see it as an iterative process and a milestone towards a more complete theory (see Figure 38).
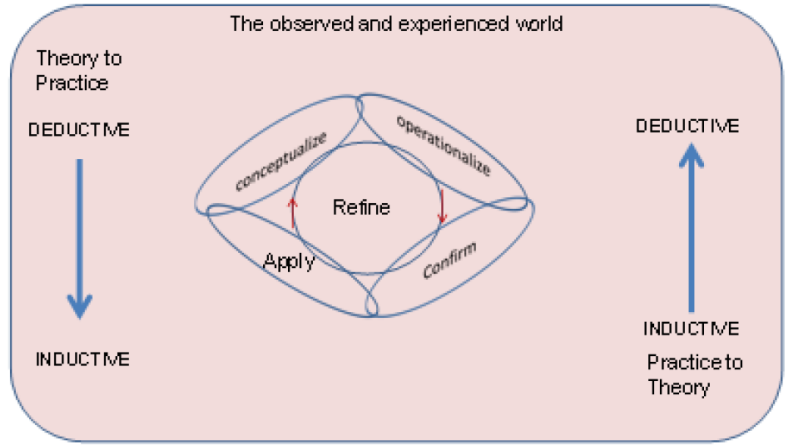
Figure 38 – Iterative theory building process (Chatterjee, 2015 p. 4).

One aspect of IS theory building, especially valid for mid-range theory development (Chatterjee, 2015), is that it is cyclic and includes (often multiple) transition between deductive and inductive reasoning. The process of theory development in this project has not been different. It started with combining the established conceptualisations of knowledge (DSR framework and predictive analytics model) into the DISPA framework design to facilitate its operationalisation. Then, the framework was operationalised in the case environment and multiple iterations of confirming and applying the inductive hypothesis have been conducted, leading to abstracting some of the case findings into conceptual domain-specific findings. As a result, knowledge that is valid in a scope of certain class of problems (predictive analytics with sensors) has been developed and tested, and a process for its development has been established. This is the theoretical contribution of this thesis.

DISPA as an output of an iterative theory building process

5.9 Concluding Remarks

Despite the increasing relevance of predictive analytics in IS research, the community has not devoted a great deal of attention to the issue until recently. Little attention has been directed towards issues related to forecasting. My study

DISPA as a step towards popularising predictive analytics in IS shows that the paradigms used in IS research, particularly in design science, can provide a useful lens for the analysis of environments characterised by a high degree of uncertainty and can provide promising solutions for challenges embedded in such environments. I intend my work to be a step towards addressing this shortcoming and hope that it initiates further efforts, both in predictive analytics in general and demand forecasting in volatile and uncertain environments.

DISPA as a contribution to academia and practice This work provides a contribution to IS research, and in particular to DSR, by introducing a model for DISPA that can serve as a method for developing IS artefacts. The framework constitutes an IS design theory consistent with the established definitions from the literature (Gregor & Jones, 2007; Kuechler & Vaishnavi, 2012; Walls et al., 1992). In addition, the paper introduces and systematically evaluates several spare-part forecasting methods, which can be considered a contribution to OR literature. Finally, as the model is sufficiently detailed to be instantiated in a real-life setting in the same way in which it was used in the case setup, the paper provides a contribution to industry and practice.

Future work should focus on the validation of the model in new environments by collecting data from further case studies. MAN Diesel & Turbo is a large company and is a market leader in its target market; thus, the utility of the DISPA model as experienced in the MAN Diesel & Turbo case

might not be as easily demonstrated in a smaller firm. More specifically, it would be interesting to see whether the contextual evaluation of systematic biases in other settings (Step 3b) could have significant depth to provide insight that is as useful as in MAN Diesel & Turbo's case. A flow chart presenting an instantiation process of DISPA has been presented to facilitate replication in another environment.

Testing new sensor technologies can also provide significant improvement in predicting sensor outputs. As discussed in the section about the continuous multi-sensor monitoring, the presence of noise signals unrelated to the variables intended to be monitored leads to difficulty in prediction. The development of a new generation of sensors, such as engine-monitoring ultra-sonic or non-invasive pressure sensors, can yield different insight when instantiating the DISPA framework in such environments.

6 References

Abowd, G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., & Steggles, P. (1999). Towards a better understanding of context and context-awareness. *Handheld and Ubiquitous Computing,* 304-307.

Ågerfalk, P.,J. (2013). Embracing diversity through mixed methods research. *European Journal of Information Systems, 22*(3), 251-256. doi:10.1057/ejis.2013.6

Agarwal, R., & Lucas Jr, H. C. (2005). The information systems identity crisis: Focusing on high-visibility and high-impact research. *MIS quarterly*, 381-398.

Aggarwal, R., & Singh, H. (2013). Differential influence of blogs across different stages of decision making: The case of venture capitalists. *MIS Quarterly, 37*(4), 1033-A3.

Ahamed, Z., Inohara, T., & Kamoshida, A. (2013). The servitization of manufacturing: An empirical case study of IBM corporation. *International Journal of Business Administration, 4*(2)

Aken, J. E. v. (2004). Management research based on the paradigm of the design sciences: The quest for field-tested and grounded technological rules. *Journal of Management Studies, 41*(2), 219-246.

Almeder, R. (2007). Pragmatism and philosophy of science: A critical survey. *International Studies in the Philosophy of Science, 21*(2), 171-195.

Ang, S., & Straub, D. W. (1998). Production and transaction economies and IS outsourcing: A study of the U.S. banking industry. *MIS Quarterly, 22*(4), 535-552.

Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. *Computer Networks, 54*(15), 2787-2805.

Baars, H., Gille, D., & Strüker, J. (2009). Evaluation of RFID applications for logistics: A framework for identifying, forecasting and assessing benefits. *European Journal of Information Systems, 18*(6), 578-591. doi:10.1057/ejis.2009.32

Babüroglu, O. N., & Ravn, I. (1992). Normative action research. *Organization Studies, 13*(1), 019-34.

Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega, 40*(6), 722-737. doi:http://dx.doi.org/10.1016/j.omega.2011.06.008

Baines, T., Lightfoot, H., Benedettini, O., & Kay, J. M. (2009). The servitization of manufacturing; A review of literature and reflection on future challenges. *Journal of Manufacturing Technology Management, 20*(5), 547-567. doi:10.1108/17410380910960984

Bansal, G., Sinha, A. P., & Zhao, H. (2008). Tuning data mining methods for cost-sensitive regression: A study in loan charge-off forecasting. *Journal of Management Information Systems, 25*(3), 315-336.

Barnett, N. J., Parry, G., Saad, M., Newnes, L. B., & Goh, Y. M. (2013). Servitization: Is a paradigm shift in the business model and service enterprise required? *Strategic Change, 22*(3), 145-156. doi:10.1002/jsc.1929

Bartezzaghi, E., Verganti, R., & Zotteri, G. (1999). Simulation framework for forecasting uncertain lumpy demand. *International Journal of Production Economics, 59*(1), 499-510.

Baskerville, R., & Wood-Harper, A. T. (1998). Diversity in information systems action research methods. European Journal of information systems, 7(2), 90-107.

Baskerville, R. (2008). What design science is not. *European Journal of Information Systems, 17*(5), 441-443.

Belcher, L. W., & Watson, H. J. (1993). Assessing the value of conoco's EIS. *MIS Quarterly, 17*(3), 239-253.

Bingi, P., Sharma, M. K., & Godla, J. K. (1999). Critical issues affecting an ERP implementation. *IS Management, 16*(3), 7-14.

Brookfield, S. (Ed.). (1987). *Developing critical thinkers,* Open University Press Milton Keynes.

Brookfield, S. D. (1995). Becoming a critically reflective teacher.

Brown, B., Chui, M., & Manyika, J. (2011). Are you ready for the era of 'big data'. *McKinsey Quarterly, 4*, 24-35.

Brown, R. M., Gatian, A. W., & Hicks Jr., J. O. (1995). Strategic information systems and financial performance. *Journal of Management Information Systems, 11*(4), 215-248.

Callegaro, A. (2010). *Forecasting methods for spare parts demand* (PhD thesis, Università degli Studi di Padova).

CFA Institute. (2015). Making sense of big data: Nate silver on the signal and the noise. *Enterprising Investor,*

Chatterjee, S. (2015). Writing my next design science research master-piece: But how do I make a theoretical contribution to DSR? *23rd European Conference on Information Systems 2015, Munster, Germany, may 26-29, 2015,*

Chen, G., & Kotz, D. (2000). A survey of context-aware mobile computing research. *Dartmouth Computer Science Technical Report, TR2000*(381)

Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly, 36*(4), 1165-1188.

Coen, M. H. (1998). Design principles for intelligent environments. *AAAI/IAAI,* 547-554.

Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the aftermarket. *Harvard Business Review, 84*(5), 129.

Cole, R., Purao, S., Rossi, M., & Sein, M. (2005). Being proactive: Where action research meets design research. *ICIS 2005 Proceedings, Paper 27*

Colen, P., & Lambrecht, M. (2013). Product service systems: Exploring operational practices. *The Service Industries Journal, 33*(5), 501-515.

Cook, D. J., & Das, S. K. (2004). Smart environments: Technology, protocols and applications. () John Wiley & Sons, Inc. doi:10.1002/047168659X.ch1

Corrêa, H. L., Ellram, L. M., Scavarda, A. J., & Cooper, M. C. (2007). An operations management view of the services and goods offering mix. *International Journal of Operations & Production Management, 27*(5), 444-463. doi:10.1108/01443570710742357

Cross, N. (2001). Design/science/research: Developing a discipline. *Fifth Asian Design Conference: International Symposium on Design Science, Su Jeong Dang Printing Company, Seoul, Korea,*

Clemen, R. T., & Winkler, R. L. (1993). Aggregating point estimates: A flexible modeling approach. *Management Science, 39*(4), 501-515.

Cushing, B. E. (1990). Frameworks, paradigms, and scientific research in management information systems. *Journal of Information Systems, 4*(2), 38-59.

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly, ,* 319-340.

Davis, G. B., & Olson, M. H. (1984). *Management information systems: Conceptual foundations, structure, and development* McGraw-Hill, Inc.

Dekker, R., Pinçe, C., Zuidwijk, R. A., & Jalil, M. N. (2010). On the use of installed base information for spare parts logistics: A review of ideas and industry practice. *Erasmus University Rotterdam, Econometric Institute,* (EI 2010-73)

Deshpandé, R., Farley, J. U., & Webster Jr., F. E. (1993). Corporate culture customer orientation, and innovativeness in japanese firms: A quadrad analysis. *Journal of Marketing, 57*(1), 23-37.

Dey, A. K. (2001). Understanding and using context. *Personal and Ubiquitous Computing, 5*(1), 4-7.

Dickson, G. W., Patridge, J. L., & Robinson, L. H. (1993). Exploring modes of facilitative support for GDSS technology. *MIS Quarterly, 17*(2), 173-194.

Dolgui, A., & Pashkevich, M. (2008). Demand forecasting for multiple slow-moving items with short requests history and unequal demand variance. *International Journal of Production Economics, 112*(2), 885-894.

Duan, W., Gu, B., & Whinston, A. B. (2009). Informational cascades and software adoption on the internet: An empirical investigation. *MIS Quarterly, 33*(1), 23-48.

Duhigg, C., & Bradsher, K. (2012, January 21, 2012). How U.S. lost out on iPhone work. *New York Times*

Fang, E., Palmatier, R. W., & Steenkamp, J. E. M. (2008). Effect of service transition strategies on firm value. *Journal of Marketing, 72*(5), 1-14. doi:10.1509/jmkg.72.5.1

Furtak, S., Avital, M., & Ulslev Pedersen, R. (2015). Sensing the future: Designing predictive analytics with sensor technologies. *23rd European Conference on Information Systems (ECIS) 2015,*

Galliers, R. D., & Land, F. F. (1987). Viewpoint: Choosing appropriate information systems research methodologies. *Communications of the ACM, 30*(11), 901-902.

Gebauer, H., & Fleisch, E. (2007). An investigation of the relationship between behavioral processes, motivation, investments in the service business and service revenue. *Industrial Marketing Management, 36*(3), 337-348. doi:10.1016/j.indmarman.2005.09.005

Gebauer, H., & Friedli, T. (2005). Behavioral implications of the transition process from products to services. *Journal of Business & Industrial Marketing, 20*(2), 70-78. doi:10.1108/08858620510583669

George, J. F., Carlson, J. R., & Valacich, J. S. (2013). Media selection as a strategic component of communication. *MIS Quarterly, 37*(4), 1233-A4.

Ghodrati, B., Benjevic, D., & Jardine, A. (2012). Product support improvement by considering system operating environment; A case study on spare parts procurement. *International Journal of Quality & Reliability Management, 29*(4), 436-450. doi:10.1108/02656711211224875

Ghodrati, B., & Kumar, U. (2005). Operating environment-based spare parts forecasting and logistics: A case study. *International Journal of Logistics Research and Applications, 8*(2), 95-105. doi:10.1080/13675560512331338189

Gill, T. G. (1996). Expert systems usage: Task change and intrinsic motivation. *MIS Quarterly, 20*(3), 301-329.

Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly,* , 611-642.

Gregor, S., & Jones, D. (2007). The anatomy of a design theory. *Journal of the Association for Information Systems, 8*(5), 312-335.

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly, 37*(2), 337.

Gummesson, E. (1999). *Qualitative methods in management research,* SAGE Publications.

Hellingrath, B., & Cordes, A. (2014). Conceptual approach for integrating condition monitoring information and spare parts forecasting methods. *Production & Manufacturing Research, 2*(1), 725-737.

Heskett, J. L. (1997). In Sasser W. Earl, Schlesinger Leonard A (Eds.), *The service profit chain, how leading companies link profit and growth to loyalty, satisfaction and value*. New York: Free Press.

Hevner, A., March, S., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75-105.

Hong, J. S., Koo, H., Lee, C., & Ahn, J. (2008). Forecasting service parts demand for a discontinued product. *IIE Transactions, 40*(7), 640-649.

Hong, W., Chan, F. K. Y., Thong, J. Y. L., Chasalow, L. C., & Dhillon, G. (2014). A framework and guidelines for context-specific theorizing in information systems research. *Information Systems Research, 25*(1), 111-136. doi:10.1287/isre.2013.0496

Hua, Z., Zhang, B., Yang, J., & Tan, D. (2007). A new approach of forecasting intermittent demand for spare parts inventories in the process industries. *Journal of the Operational Research Society, 58*(1), 52-61.

Hua, Z. S., Zhang, B., Yang, J., & Tan, D. S. (2007). A new approach of forecasting intermittent demand for spare parts inventories in the process industries. *Journal of the Operational Research Society, 58*(1), 52-61.

Ihde, G., Merkel, H., & Henning, R. (1999). In Baumgarten H., Ihde G. (Eds.), *Ersatzteillogistik. theoretische grundlagen und praktische handhabung* (3. Auflage ed.). München: Hussverlag.

Iivari, J. (2007). A paradigmatic analysis of information systems as a design science. *Scandinavian Journal of Information Systems, 19*(2), 5.

Jalil, M. N. (2011). *Customer information driven after sales service management: Lessons from spare parts logistics* (PhD).

Johns, G. (2006). The essential impact of context on organizational behavior. *Academy of Management Review, 31*(2), 386-408.

Johnson, P., & Duberley, J. (2000). *Understanding management research: An introduction to epistemology,* Sage.

Johnson, P., & Duberley, J. (2003). Reflexivity in management research. *Journal of Management Studies, 40*(5), 1279-1303.

Jonker, J., & Pennink, B. (2009). In SpringerLink (Online service) (Ed.), *The essence of research methodology, A concise guide for master and PhD students in management science*. Berlin, Heidelberg: SpringerLink Online service.

Kamis, A., Koufaris, M., & Stern, T. (2008). Using an attribute-based decision support system for user-customized products online: An experimental investigation. *MIS Quarterly, 32*(1), 159-177.

Karanasios, S., Thakker, D., Lau, L., Allen, D., Dimitrova, V., & Norman, A. (2013). Making sense of digital traces: An activity theory driven ontological approach. *Journal of the American Society for Information Science & Technology, 64*(12), 2452-2467. doi:10.1002/asi.22935

Kasper, G. M. (1996). A theory of decision support system design for user calibration. *Information Systems Research, 7*(2), 215-232.

Keil, M., Tan, B. C. Y., Wei, K., Saarinen, T., Tuunainen, V., & Wassenaar, A. (2000). A cross-cultural study on escalation of commitment behavior in software projects. *MIS Quarterly, 24*(2), 299-325.

Kuechler, W., & Vaishnavi, V. (2012). A framework for theory development in design science research: Multiple perspectives. *Journal of the Association for Information Systems, 13*(6), 395-423.

Lapide, L. (2012). Installed-base supply planning. *Supply Chain Management Review, 16*(3), 4.

Lee, H. L., Padmanabhan, V., & Whang, S. (1997). The bullwhip effect in supply chains1. *Sloan Management Review, 38*(3), 93-102.

Leitner, J., & Leopold-Wildburger, U. (2011). Experiments on forecasting behavior with several sources of information–A review of the literature. *European Journal of Operational Research, 213*(3), 459-469.

Lerch, F. J., & Harter, D. E. (2000). Cognitive support for real-time dynamic decision making. *Information Systems Research, 12*(1), 63-82. doi:10.1287/isre.12.1.63.9717

Louit, D., Pascual, R., Banjevic, D., & Jardine, A. K. S. (2011). Condition-based spares ordering for critical components. *Mechanical Systems and Signal Processing, 25*(5), 1837-1848. doi:http://dx.doi.org/10.1016/j.ymssp.2011.01.004

March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems, 15*(4), 251-266.

Maritime-Insight. (2013). *Shipbuilding outlook report.* ().

Martens, D., & Provost, F. (2014). Explaining data-driven document classifications. *MIS Quarterly, 38*(1), 73-A6.

Mathiassen, L. (2002). Collaborative practice research. *Information Technology & People, 15*(4), 321-345. doi:10.1108/09593840210453115

Mathieu, V. (2001a). Product services: From a service supporting the product to a service supporting the client. *Journal of Business & Industrial Marketing, 16*(1), 39-61. doi:10.1108/08858620110364873

Mathieu, V. (2001b). Service strategies within the manufacturing sector: Benefits, costs and partnership. *International Journal of Service Industry Management, 12*(5), 451-475. doi:10.1108/EUM0000000006093

Merriam-Webster. (2015a). Sensor in merriam-webster. Retrieved from http://www.merriam-webster.com/dictionary/sensor

Merriam-Webster. (2015b). Usability in Merriam-webster. Retrieved from http://www.merriam-webster.com/dictionary/usability

Minner, S. (2011). Forecasting and inventory management for spare parts: An installed base approach. In N. Altay, & L. A. Litteral (Eds.), *Service parts management* (pp. 157-169) Springer London. doi:10.1007/978-0-85729-039-7_8

Miranda, S. M., & Yong-Mi Kim. (2006). Professional versus political contexts: Institutional mitigation and the transaction cost heuristic in information systems outsourcing. *MIS Quarterly, 30*(3), 725-753.

Nahim, H., Younes, R., Nohra, C., & Ouladsine, M. (2015). Complete modeling for systems of a marine diesel engine. *Journal of Marine Science and Application, 14*(1), 93-104. doi:10.1007/s11804-015-1285-y

Nardi, B. A. (1996). *Context and consciousness: Activity theory and human-computer interaction* MIT Press.

Narock, T., Yoon, V., & March, S. (2012). On the role of context and subjectivity on scientific information systems. *Communications of the Association for Information Systems, 30*(1), 12.

NASDAQ. (2016). Crude oil price. Retrieved from http://www.nasdaq.com/markets/crude-oil.aspx?timeframe=10y

Neely, A. (2008). Exploring the financial consequences of the servitization of manufacturing. *Operations Management Research, 1*(2), 103-118. doi:10.1007/s12063-009-0015-5

Nissen, M. E., & Sengupta, K. (2006). Incorporating software agents into supply chains: Experimental investigation with a procurement task. *MIS Quarterly, 30*(1), 145-166.

Notteboom, T., & Cariou, P. (2013). Slow steaming in container liner shipping: Is there any impact on fuel surcharge practices? *International Journal of Logistics Management, 24*(1), 73-86. doi:10.1108/IJLM-05-2013-0055

Oliva, R., & Kallenberg, R. (2003). Managing the transition from products to services. *International Journal of Service Industry Management, 14*(2), 160-172. doi:10.1108/09564230310474138

Orlikowski, W. J., & Iacono, C. S. (2001). Research commentary: Desperately seeking the "IT" in IT research—A call to theorizing the IT artifact. *Information Systems Research, 12*(2), 121-134.

Orlikowski, W. J., & Robey, D. (1991). Information technology and the structuring of organizations. *Information Systems Research, 2*(2), 143-169.

Pedersen, R. U., Furtak, S. J., Häuser, I., Lauth, C., & Van Kranenburg, R. (2013). Mini smart grid@ copenhagen business school: Prototype demonstration. *Design science at the intersection of physical and virtual design* (pp. 446-447) Springer.

Pedersen, R. U., & Pedersen, M. K. (2014). Micro information systems: New fractals in an evolving IS landscape. In Claus-Peter Rückemann (Ed.), *Integrated information and computing systems for natural, spatial, and social sciences* (Information Resources Management Association (IRMA) ed., pp. 533-555) IGI global. doi:10.4018/978-1-4666-5125-8.ch022"

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45-77.

Ploesser, K., Recker, J., & Rosemann, M. (2011). Challenges in the context-aware management of business processes: A multiple case study. *ECIS,*

Ploesser, K. (2013). *A design theory for context-aware information systems*

Popper, K. (2005). *The logic of scientific discovery* Routledge.

Pries-Heje, J., & Baskerville, R. (2008). The design theory nexus. *MIS quarterly*, 731-755.

Punch, K. F., & Oancea, A. (2009). *Introduction to research methods in education* Sage.

Ragnarsdóttir, B., Ingadóttir, H., Akkerman, R., Hansen, K. R. N., Furtak, S., & Schwarck, G. (2012). Combined evaluation of forecasting methods and inventory models for intermittent demand: A case study. In R. W. Grubbström, & H. H. Hinterhuber (Eds.), *Pre-prints of the seventeenth international working seminar on production economics* (pp. 287-298) MEDIFAS: The Mediterranean Institute for Advanced Studies.

Rossi, M., Henfridsson, O., Lyytinen, K., & Siau, K. (2013). Design science research: The road traveled and the road that lies ahead. *Journal of Database Management (JDM), 24*(3), 1-8.

Rowe, G., & Wright, G. (1999). The Delphi technique as a forecasting tool: issues and analysis. *International journal of forecasting*, *15*(4), 353-375.

Saar-Tsechansky, M., & Provost, F. (2007). Decision-centric active learning of binary-outcome models. *Information Systems Research, 18*(1), 4-22. doi:10.1287/isre.1070.0111

Saunders, M., Lewis, P. and Thronhill, A., 2007. Research methods for business students. 3rd ed. London: Prentice Hall

Sein, M. K., Henfridsson, O., Purao, S., Rossi, M., & Lindgren, R. (2011). Action design research. *MIS quarterly,* 37-56.

Sengupta, K., & Te'eni, D. (1993). Cognitive feedback in GDSS: Improving control and convergence. *MIS Quarterly, 17*(1), 87-113.

Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. *MIS Quarterly, 35*(3), 553-572.

Shollo, A. (2013). *The role of business intelligence in organizational decision-making* (PhD).

Silver, E., Pyke, D. F., & Peterson, R. (1998). Inventory management and production planning and scheduling.

Silver, M. S. (1991). Decisional guidance for computer-based decision support. *MIS Quarterly, 15*(1), 105-122.

Silver, N. (2012). *The signal and the noise, the art and science of prediction*. London: Penguin Group.

Simon, H. A. (1996). *The sciences of the artificial*. Cambridge, Mass:

Sinha, A. P., & May, J. H. (1996). Providing design assistance: A case-based approach. *Information Systems Research, 7*(3), 363-387.

Song, J., & Zipkin, P. H. (1996). Evaluation of base-stock policies in multiechelon inventory systems with state-dependent demands. part II: State-dependent depot policies. *Naval Research Logistics (NRL), 43*(3), 381-396.

Speier, C., & Morris, M. G. (2003). The influence of query interface design on decision-making performance. *MIS Quarterly, 27*(3), 397-423.

Sun, H. (2013). A longitudinal study of herd behavior in the adoption and continued use of technology. *MIS Quarterly, 37*(4), 1013-A13.

Sutton, R. I., & Staw, B. M. (1995). What theory is not. *Administrative Science Quarterly, ,* 371-384.

Tashakkori, A., & Teddlie, C. (1998). *Mixed methodology: Combining qualitative and quantitative approaches,* Sage.

Teunter, R. H., & Duncan, L. (2009). Forecasting intermittent demand: A comparative study. *Journal of the Operational Research Society, 60*(3), 321-329. doi:10.1057/palgrave.jors.2602569

Teunter, R. H., Syntetos, A. A., & Zied Babai, M. (2011). Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research, 214*(3), 606-615. doi:http://dx.doi.org/10.1016/j.ejor.2011.05.018

The American Heritage Dictionary. (2016). "Theory" in The american heritage dictionary. Retrieved from https://www.ahdictionary.com/word/search.html?q=theory&submit.x=44&submit.y=13

The Cambridge Dictionary. (2016). "Theory" in The cambridge dictionary. Retrieved from http://dictionary.cambridge.org/dictionary/english/theory

Tibben-Lembke, R. S., & Amato, H. N. (2001). REPLACEMENT PARTS MANAGEMENT: THE VALUE OF INFORMATION. *Journal of Business Logistics, 22*(2), 149-164. doi:10.1002/j.2158-1592.2001.tb00008.x

Todd, P., & Benbasat, I. (1992). The use of information in decision making: An experimental investigation of the impact of computer-based decision aids. *MIS Quarterly, 16*(3), 373-393.

Turrof, M., Hiltz, S. R., Bahgat, A. N. F., & Rana, A. R. (1993). Distributed group support systems. *MIS Quarterly, 17*(4), 399-417.

Vaishnavi, V., & Kuechler, W. (2004). Design research in information systems.

Vaishnavi, V., and Kuechler, B. 2007. "Design Research in Information Systems," ISWorld (available at http://home.aisnet.org/displaycommon.cfm?an=1&subarticlenbr=279;accessed February 17, 2017).

Van de Ven, A. H. (2007). *Engaged scholarship : A guide for organizational and social research*. Cary, NC, USA: Oxford University Press.

Vandermerwe, S., & Rada, J. (1988). Servitization of business: Adding value by adding services. *European Management Journal, 6*(4), 314. doi:10.1016/0263-2373(88)90033-3

Venable, J. (2006). The role of theory and theorising in design science research. *Proceedings of the 1st International Conference on Design Science in Information Systems and Technology (DESRIST 2006),* 1-18.

Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: a framework for evaluation in design science research. *European Journal of Information Systems*, *25*(1), 77-89.

Verganti, R. (1997). Order overplanning with uncertain lumpy demand: A simplified theory. *International Journal of Production Research, 35*(12), 3229-3248.

Visnjic Kastalli, I., & Van Looy, B. (2013). Servitization: Disentangling the impact of service business model innovation on manufacturing firm performance. *Journal of Operations Management, 31*(4), 169-180. doi:10.1016/j.jom.2013.02.001

Walls, J. G., Widmeyer, G. R., & El Sawy, O. A. (1992). Building an information system design theory for vigilant EIS. *Information Systems Research, 3*(1), 36.

Walsham, G. (1995). Interpretive case studies in IS research: Nature and method. *European Journal of Information Systems, 4*(2), 74-81.

Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *Management Information Systems Quarterly, 26*(2), 3.

Welbourne, E., Battle, L., Cole, G., Gould, K., Rector, K., Raymer, S., . . . Borriello, G. (2009). Building the internet of things using RFID: The RFID ecosystem experience. *Internet Computing, IEEE, 13*(3), 48-55.

Wikipedia. (2016a). Bathtub curve. Retrieved from https://en.wikipedia.org/wiki/Bathtub_curve

Wikipedia. (2016b). Bullwhip effect. Retrieved from https://en.wikipedia.org/wiki/Bullwhip_effect

Winter, R. (2008). Design science research in europe. *European Journal of Information Systems, 17*(5), 470-475. doi:10.1057/ejis.2008.44

Wise, R., & Baumgartner, P. (1999). Go downstream: The new profit imperative in manufacturing. *Harvard Business Review, 77*(5), 133.

Woo, J., & Moon, D. S. (2014). The effects of slow steaming on the environmental performance in liner shipping. *Maritime Policy & Management, 41*(2), 176-191. doi:10.1080/03088839.2013.819131

Woollaston, V. (2013, 8 August 2013). First google, now apple's at it: New feature follows mobile users wherever they go and produces a location history using GPS. *Dailymail*

Xiao, B., & Benbasat, I. (2007). E-commerce product recommendation agents: Use, characteristics, and impact. *MIS Quarterly, 31*(1), 137-209.

Ye, L. R., & Johnson, P. E. (1995). The impact of explanation facilities on user acceptance of expert systems advice. *MIS Quarterly, 19*(2), 157-172.

Yin, J., Fan, L., Yang, Z., & Li, K. X. (2014). Slow steaming of liner trade: Its economic and environmental impacts. *Maritime Policy & Management, 41*(2), 149-158. doi:10.1080/03088839.2013.821210

Zhang, J., & Seidmann, A. (2010). Perpetual versus subscription licensing under quality uncertainty and network externality effects. *Journal of Management Information Systems, 27*(1), 39-38. doi:10.2753/MIS0742-1222270103

Zhang, P. (2013). The affective response model: A theoretical framework of affective concepts and their relationships in the ict context. *MIS Quarterly, 37*(1), 247-274.

7 Appendix A – Extended literature review on Decision Support Systems

7.1 Rationale and Background

Recently the capabilities of organizations to collect data being an input to decision making have increased rapidly. Access to new data could enable organizations to develop a new generation of Decision Support Systems (DSS), real-time, ubiquitous systems monitoring the subject of decision constantly through data collecting sensor-probes. Introducing these capabilities to DSS could change completely the cognitive attitude of decision makers towards the systems. Thorough analysis of the history and state-of the art of the IS literature is necessary in order to determine how currently accepted theories can cope with this emerging phenomenon.

7.2 Underlying Objective and Research Question

The objective of the study is to evaluate **how IS literature has perceived decision making (DM) process and DSS impact on that process**. In particular, the work will **determine how IS literature theorized about decision maker attitude towards DSS and the support for the decision it offers**. The answer to those questions can pave a path towards understanding how introducing new data sources to DSS can change the attitude of decision maker, including changing his cognitive bias and self-evaluative and self-affirmation concerns . In order to limit the scope of review to reasonable breadth IS literature will be understood as the AIS basket of 8.

7.3 Theoretical Underpinnings

The paper extends the frame of Webster and Watson (2002) in order to determine and evaluate state of the art literature on Decision making in IS body of knowledge.

7.4 Data collection

The object of the study, being all the articles published in the Basket of 8 IS journals having among their keywords "decision making" or "Decision support systems" , will be collected using EBSCO Business Source Complete database.

7.5 Data analysis

In order to answer the research questions all the extracted articles will be clustered based on their relation towards decision making and decision support systems. The overview of clusters will serve as overview of historical efforts of IS community towards researching decision making and DSS. Moreover, articles contributing to better understanding of decision making process and decision support systems will be isolated and thoroughly reviewed.

7.6 Research design

Clustering literature:

1. Extract articles from the basket of eight with matching subject terms (Decision making or Decision support systems) -> output 511 papers
2. Screen all documents and classify every one according to the clustering flowchart

For documents assign to most interesting classes (theorizing about Decision making and Theorizing about DSS) prepare a comparative overview (in progress).

7.7 Validity

(Ontological) assumptions:

- Only contributions in leading journals are significant enough to consider (limiting scope of review to basket of 8)
- Papers theorizing about DSS or DM will have among the subject terms "Decision support systems" or "Decision Making"

7.8 Result overview
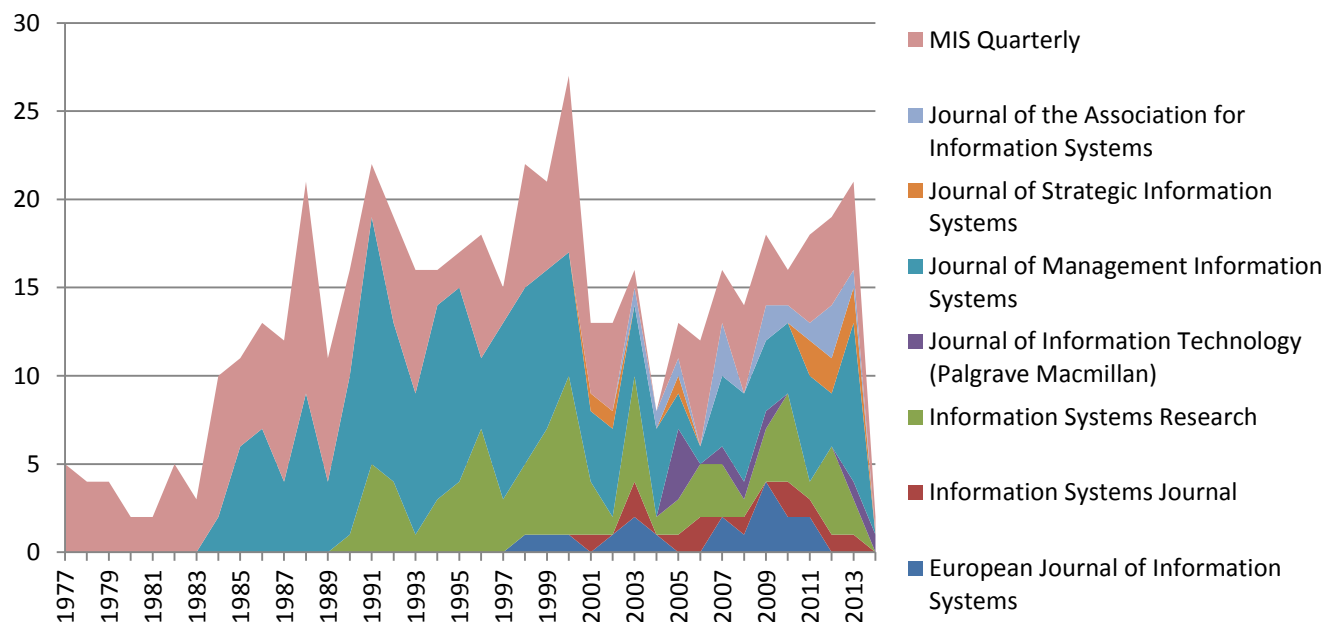
Figure 39 - The history of IS contributions about DM and DSS by journal

The initial search produced 511 papers. The initial picture, the history of IS contributions to DSS by journal, is presented in Figure 39. Overall, it can be seen that DSS has been very popular in the basket of eight, and the trend seems to have declined in the early 2000s, but picked up again in 2005.

I decided not to look into work published before 1990, on one hand to limit the number of articles and, on the other hand, not to explore trends that were older than 25 years. The articles were grouped into five segments, based on

the main object of theorizing: theories about decision making, about decision support systems, about information management or about a particular decision. Those not explicitly theorizing about any of the four were grouped as 'not related', and will not be presented in analysis below. After excluding 103 papers for being too old, 211 for not being related and nine because they were duplicate entries, I received 188 papers that were clustered as presented in Figure 40.
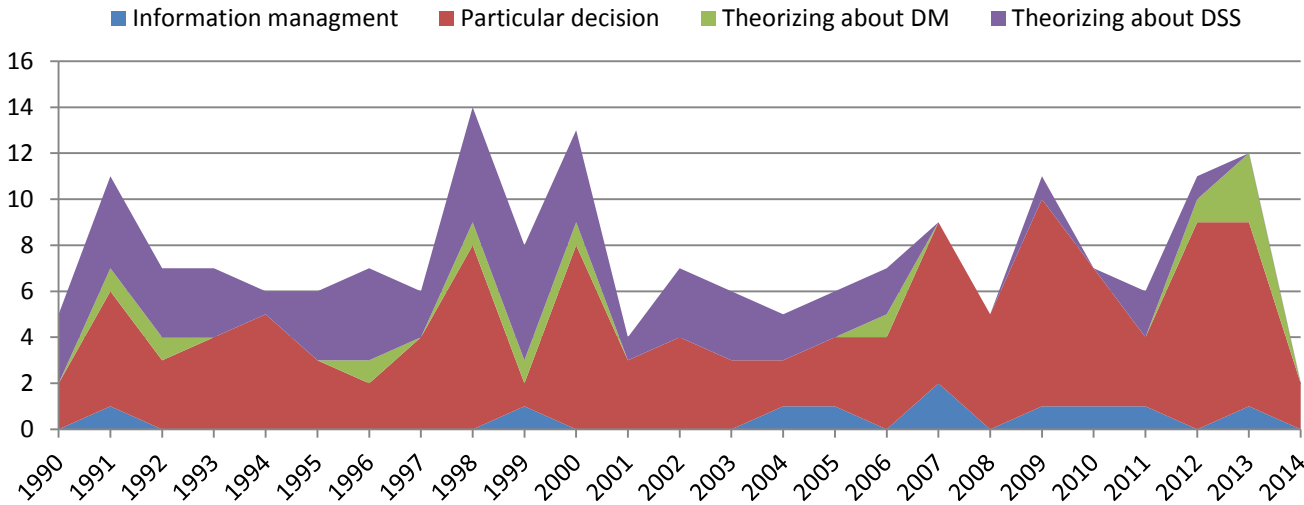
Figure 40 - DSS literature divided by the object of theorizing

Overall, investigating particular problems is by far the dominant. Theorizing about DSS seems to have a declining trend and theorizing about DM an inclining one, although the figures are so low that the generalization might be misleading. In spite of the fact that Decision Support Systems as publication keywords seem to be increasingly published under other topics, among most common of which are Business Intelligence, Business Analytics or Big Data, which should yield a less frequent occurrence of DSS keywords in published works, the sum of annual publications (especially when averaged over a few years) in the last 25 years does not seem to have changed dramatically. This observation might actually suggest that global popularity of decision support systems, regardless of the nomenclature, might be increasing.
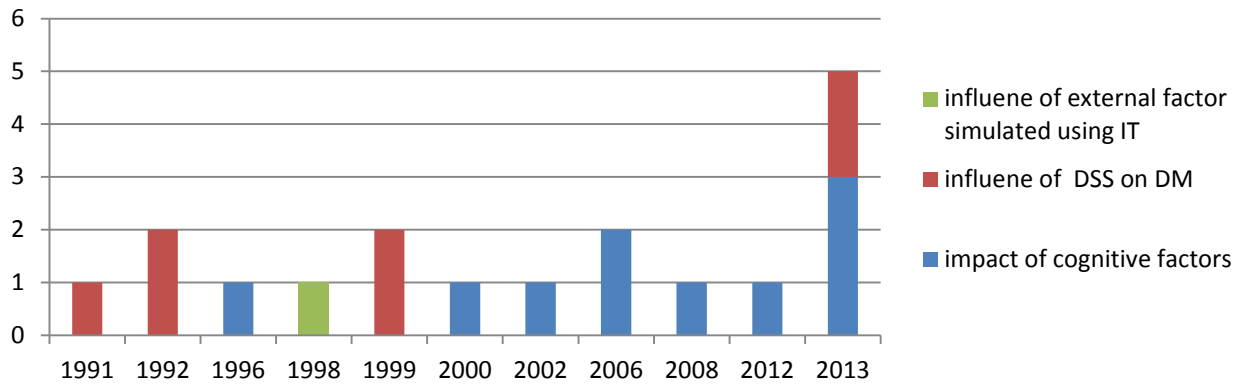
Figure 41 - Theorizing about Decision Making - cluster details

When we look closer into the clusters, more trends become apparent. When looking into the details of cluster grouping work theorizing about decision making (Figure 41 above), it seems that the recent IS efforts revolve around researching the impact of cognitive factors on decision making. With regard to work theorizing about Decision Support Systems (Figure 42, below), there seems to be a declining trend for research on the influence of particular factors in DSS performance. Overall, the research on DSS also seems to gradually attract less interest.
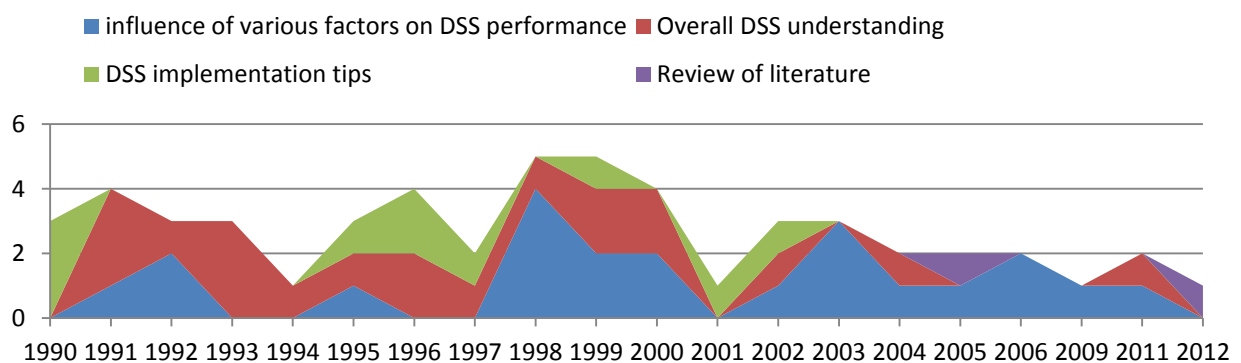
Figure 42 - Theorizing about Decision-support Systems - cluster details

In the context of this work, the cluster-grouping work theorizing about a particular decision seems to be the most interesting (see Figure 43 below). When we look into a detailed subject analysis, it becomes apparent that the publications on making particular decisions resonates with different terms

gaining traction in the field (the Group Support System in the mid-nineties, online purchase in the early 2000s, and outsourcing in 2005-2010), with some other topics (technology adoption, software piracy) running throughout the entire period.
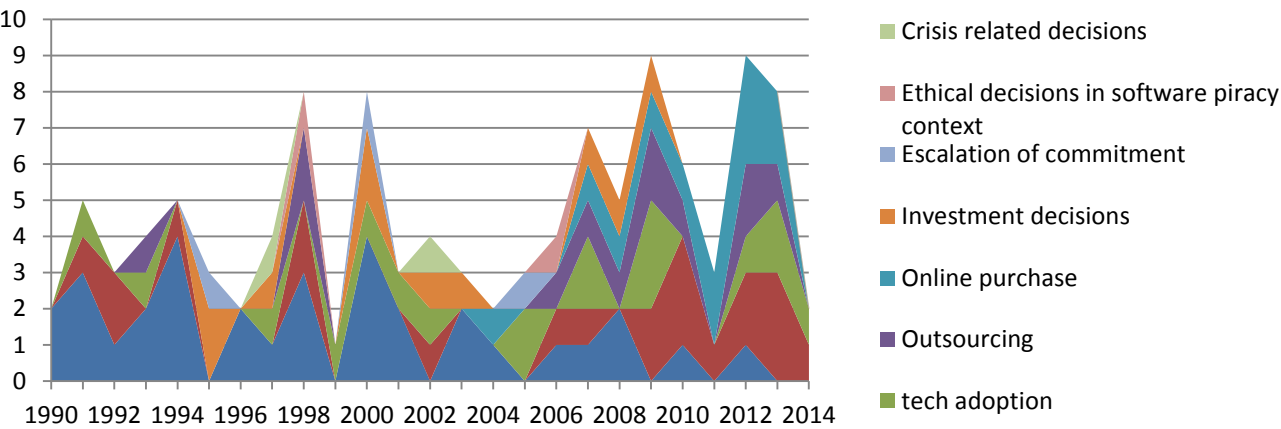
Figure 43 - Theorizing about a particular decision - cluster details

Overall, the number of publications concerned with particular decision support systems seems to have increased. This trend becomes even more apparent when we look into the percentage share of particular decisions concerning DSS in all DSS-related articles (see Figure 44, below). In the context of the previous section, this resonates well with solution-specific exploratory forecasting methods. Some of the evaluated DSSs had predictive functionality, such as systems dealing with investment decisions relying on predictions of external conditions and investment outcomes, although this was not presented as a main contribution for these papers. All these facts suggest that, although demand forecasting and, more broadly, predictive analytics are not very common in the DSS literature, the idea of context-specific solutions, such as that presented in the exploratory forecasting approach, seem very common and are still gaining in popularity.
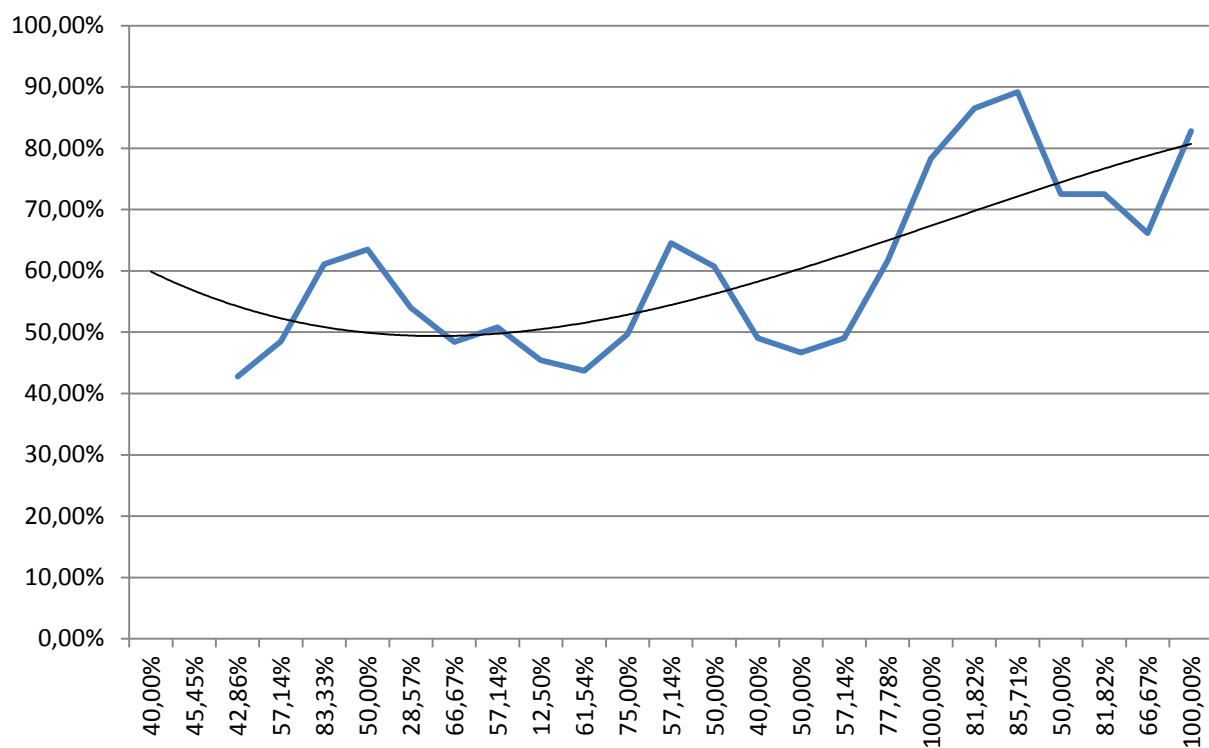
Figure 44 - Portion of the DSS literature articles theorizing about a particular decision

8 Appendix B- Framework tutorial – how to apply DISPA to a particular problem?

The foundation of the claim that DISPA is a useful contribution lies in the assumption that it can be used **by anyone trying to design a predictive Information System based on sensor data**. In order to facilitate the diffusion of the framework among the intended audience, I present a section describing exactly how to go about instantiating the framework in any sensor-based predictive context. The instructions are presented in the form of a step-by-step sequential tutorial that complements the previously presented graphical image of DISPA.

1. Outline in a quantitative and meaningful manner what defines the good IS system that you are trying to design (develop the system's purpose and a evaluation function to evaluate its fulfilment).
   A suggestion is to define it in terms of a monetary cost, but if other quantitative means are more meaningful, they can be applied.
2. Establish a qualitative evaluation process.
   Define a group of people that will use the IS artifact or who have expert knowledge about the environment in which the system operates. In addition, try to establish a questionnaire that could lead the group to determine exactly what happened in the system environment at the time and what kind of impact that could have on the system
3. Establish the benchmark.
   If an alternative state-of-the-art solution exists, either in the literature or in practice, evaluate it using the previously established evaluation function. If no known solution exists, use the first design developed as a benchmark.
4. Evaluate the benchmark, both quantitatively and qualitatively.
   Execute the evaluation function and qualitative evaluation process on the benchmark. Try to establish any systematic bias in the data used for quantitative evaluation (any reason that could decrease the performance of the evaluation function based on the selected dataset
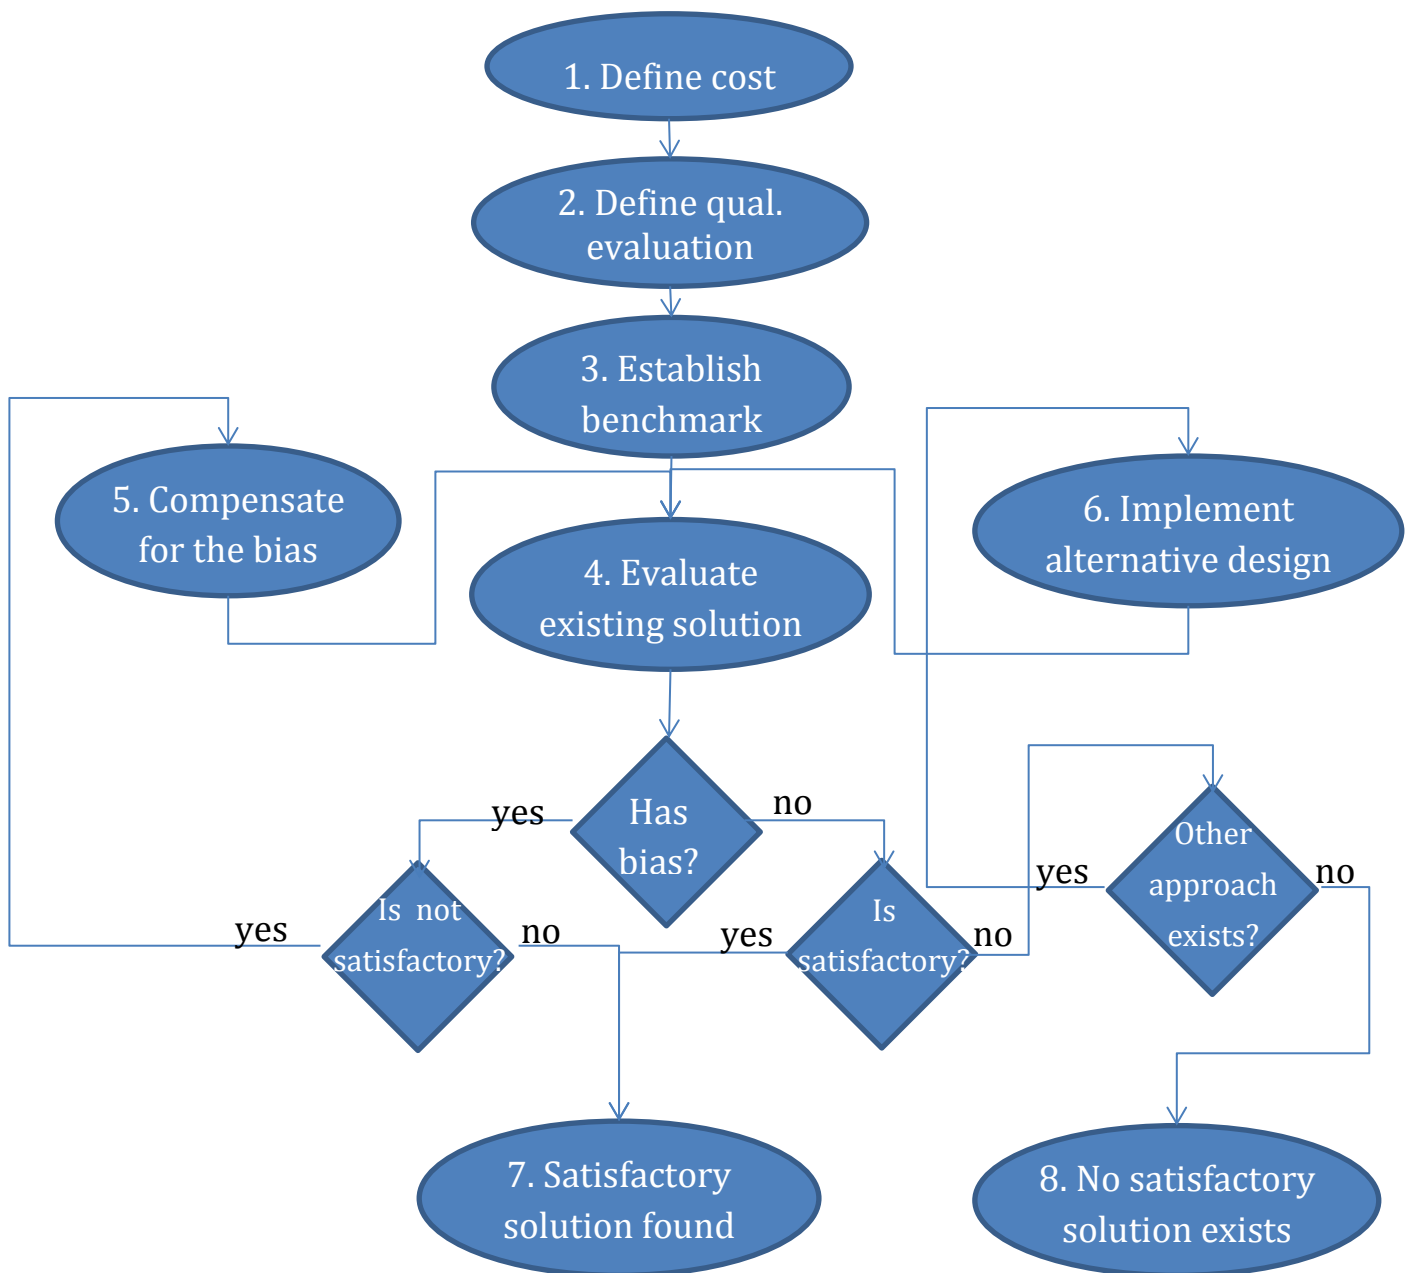
Figure 45 - Flow chart presenting a process of instantiation of the DISPA framework

5. If a potential systematic bias was discovered and the current solution is not satisfactory, re-design the system in order to control for the bias in the data and evaluate the new system (step 4).

Try to take countermeasures to limit the impact of the bias in the dataset under investigation. An example from the case previously presented can be extracting the impact of different engine usage patterns by measuring the interval between spare part replacement in engine usage (in running hours) rather than in time between replacements.

6. If no potential systematic bias was discovered and the current solution is not satisfactory, try to establish an alternative design that is not based on a previous solution and evaluate the new system (step 4).
Try to come up with a new design for the problem. This can be based on existing literature or, if no feasible solution exists in the literature, it could be a new, independent approach to solving the problem.

7. If after the evaluation step (4) the system performance is deemed to be satisfactory, the process terminates.

8. If establishing an alternative design (step 6) is not possible and no satisfactory system can be designed, evaluate whether the best performing system can be used or if no solution can be offered.

19. Thomas Lyse Hansen
Six Essays on Pricing and Weather risk in Energy Markets

20. Sabine Madsen
Emerging Methods – An Interpretive Study of ISD Methods in Practice

21. Evis Sinani
The Impact of Foreign Direct Investment on Efficiency, Productivity Growth and Trade: An Empirical Investigation

22. Bent Meier Sørensen
Making Events Work Or, How to Multiply Your Crisis

23. Pernille Schnoor
*Brand Ethos
Om troværdige brand- og virksomhedsidentiteter i et retorisk og diskursteoretisk perspektiv*

24. Sidsel Fabech
*Von welchem Österreich ist hier die Rede?
Diskursive forhandlinger og magtkampe mellem rivaliserende nationale identitetskonstruktioner i østrigske pressediskurser*

25. Klavs Odgaard Christensen
*Sprogpolitik og identitetsdannelse i flersprogede forbundsstater
Et komparativt studie af Schweiz og Canada*

26. Dana B. Minbaeva
Human Resource Practices and Knowledge Transfer in Multinational Corporations

27. Holger Højlund
*Markedets politiske fornuft
Et studie af velfærdens organisering i perioden 1990-2003*

28. Christine Mølgaard Frandsen
*A.s erfaring
Om mellemværendets praktik i en transformation af mennesket og subjektiviteten*

29. Sine Nørholm Just
*The Constitution of Meaning
– A Meaningful Constitution?
Legitimacy, identity, and public opinion in the debate on the future of Europe*

2005

1. Claus J. Varnes
Managing product innovation through rules – The role of formal and structured methods in product development

2. Helle Hedegaard Hein
*Mellem konflikt og konsensus
– Dialogudvikling på hospitalsklinikker*

3. Axel Rosenø
Customer Value Driven Product Innovation – A Study of Market Learning in New Product Development

4. Søren Buhl Pedersen
*Making space
An outline of place branding*

5. Camilla Funck Ellehave
*Differences that Matter
An analysis of practices of gender and organizing in contemporary workplaces*

6. Rigmor Madeleine Lond
Styring af kommunale forvaltninger

7. Mette Aagaard Andreassen
*Supply Chain versus Supply Chain
Benchmarking as a Means to Managing Supply Chains*

8. Caroline Aggestam-Pontoppidan
*From an idea to a standard
The UN and the global governance of accountants' competence*

9. Norsk ph.d.

10. Vivienne Heng Ker-ni
An Experimental Field Study on the

A case study of the Fashion and Design Branch of the Industrial District of Montebelluna, NE Italy

12. Mikkel Flyverbom
Making the Global Information Society Governable
On the Governmentality of Multi-Stakeholder Networks

13. Anette Grønning
Personen bag
Tilstedevær i e-mail som inter-aktionsform mellem kunde og med-arbejder i dansk forsikringskontekst

14. Jørn Helder
One Company – One Language?
The NN-case

15. Lars Bjerregaard Mikkelsen
Differing perceptions of customer value
Development and application of a tool for mapping perceptions of customer value at both ends of customer-suppli-er dyads in industrial markets

16. Lise Granerud
Exploring Learning
Technological learning within small manufacturers in South Africa

17. Esben Rahbek Pedersen
Between Hopes and Realities: Reflections on the Promises and Practices of Corporate Social Responsibility (CSR)

18. Ramona Samson
The Cultural Integration Model and European Transformation.
The Case of Romania

2007
1. Jakob Vestergaard
Discipline in The Global Economy Panopticism and the Post-Washington Consensus

2. Heidi Lund Hansen
Spaces for learning and working
A qualitative study of change of work, management, vehicles of power and social practices in open offices

3. Sudhanshu Rai
Exploring the internal dynamics of software development teams during user analysis
A tension enabled Institutionalization Model; "Where process becomes the objective"

4. Norsk ph.d.
Ej til salg gennem Samfundslitteratur

5. Serden Ozcan
EXPLORING HETEROGENEITY IN ORGANIZATIONAL ACTIONS AND OUTCOMES
A Behavioural Perspective

6. Kim Sundtoft Hald
Inter-organizational Performance Measurement and Management in Action
– An Ethnography on the Construction of Management, Identity and Relationships

7. Tobias Lindeberg
Evaluative Technologies
Quality and the Multiplicity of Performance

8. Merete Wedell-Wedellsborg
Den globale soldat
Identitetsdannelse og identitetsledelse i multinationale militære organisatio-ner

9. Lars Frederiksen
Open Innovation Business Models
Innovation in firm-hosted online user communities and inter-firm project ventures in the music industry
– A collection of essays

10. Jonas Gabrielsen
Retorisk toposlære – fra statisk 'sted' til persuasiv aktivitet

31. Fumiko Kano Glückstad
Bridging Remote Cultures: Cross-lingual concept mapping based on the information receiver's prior-knowledge

32. Henrik Barslund Fosse
Empirical Essays in International Trade

33. Peter Alexander Albrecht
Foundational hybridity and its reproduction
Security sector reform in Sierra Leone

34. Maja Rosenstock
CSR - hvor svært kan det være?
Kulturanalytisk casestudie om udfordringer og dilemmaer med at forankre Coops CSR-strategi

35. Jeanette Rasmussen
Tweens, medier og forbrug
Et studie af 10-12 årige danske børns brug af internettet, opfattelse og for-ståelse af markedsføring og forbrug

36. Ib Tunby Gulbrandsen
'This page is not intended for a US Audience'
A five-act spectacle on online communication, collaboration & organization.

37. Kasper Aalling Teilmann
Interactive Approaches to Rural Development

38. Mette Mogensen
The Organization(s) of Well-being and Productivity
(Re)assembling work in the Danish Post

39. Søren Friis Møller
From Disinterestedness to Engagement Towards Relational Leadership In the Cultural Sector

40. Nico Peter Berhausen
Management Control, Innovation and Strategic Objectives – Interactions and Convergence in Product Development Networks

41. Balder Onarheim
Creativity under Constraints
Creativity as Balancing 'Constrainedness'

42. Haoyong Zhou
Essays on Family Firms

43. Elisabeth Naima Mikkelsen
Making sense of organisational conflict
An empirical study of enacted sense-making in everyday conflict at work

2013
1. Jacob Lyngsie
Entrepreneurship in an Organizational Context

2. Signe Groth-Brodersen
Fra ledelse til selvet
En socialpsykologisk analyse af forholdet imellem selvledelse, ledelse og stress i det moderne arbejdsliv

3. Nis Høyrup Christensen
Shaping Markets: A Neoinstitutional Analysis of the Emerging Organizational Field of Renewable Energy in China

4. Christian Edelvold Berg
As a matter of size
THE IMPORTANCE OF CRITICAL MASS AND THE CONSEQUENCES OF SCARCITY FOR TELEVISION MARKETS

5. Christine D. Isakson
Coworker Influence and Labor Mobility Essays on Turnover, Entrepreneurship and Location Choice in the Danish Maritime Industry

6. Niels Joseph Jerne Lennon
Accounting Qualities in Practice Rhizomatic stories of representational faithfulness, decision making and control

7. Shannon O'Donnell
Making Ensemble Possible
How special groups organize for collaborative creativity in conditions of spatial variability and distance

26. Kristian Gylling Olesen
*Flertydig og emergerende ledelse i
folkeskolen
Et aktør-netværksteoretisk ledelses-
studie af politiske evalueringsreformers
betydning for ledelse i den danske
folkeskole*

27. Troels Riis Larsen
*Kampen om Danmarks omdømme
1945-2010
Omdømmearbejde og omdømmepolitik*

28. Klaus Majgaard
Jagten på autenticitet i offentlig styring

29. Ming Hua Li
*Institutional Transition and
Organizational Diversity:
Differentiated internationalization
strategies of emerging market
state-owned enterprises*

30. Sofie Blinkenberg Federspiel
*IT, organisation og digitalisering:
Institutionelt arbejde i den kommunale
digitaliseringsproces*

31. Elvi Weinreich
*Hvilke offentlige ledere er der brug for
når velfærdstænkningen flytter sig
– er Diplomuddannelsens lederprofil
svaret?*

32. Ellen Mølgaard Korsager
*Self-conception and image of context
in the growth of the firm
– A Penrosian History of Fiberline
Composites*

33. Else Skjold
The Daily Selection

34. Marie Louise Conradsen
*The Cancer Centre That Never Was
The Organisation of Danish Cancer
Research 1949-1992*

35. Virgilio Failla
*Three Essays on the Dynamics of
Entrepreneurs in the Labor Market*

36. Nicky Nedergaard
*Brand-Based Innovation
Relational Perspectives on Brand Logics
and Design Innovation Strategies and
Implementation*

37. Mads Gjedsted Nielsen
Essays in Real Estate Finance

38. Kristin Martina Brandl
*Process Perspectives on
Service Offshoring*

39. Mia Rosa Koss Hartmann
*In the gray zone
With police in making space
for creativity*

40. Karen Ingerslev
*Healthcare Innovation under
The Microscope
Framing Boundaries of Wicked
Problems*

41. Tim Neerup Themsen
*Risk Management in large Danish
public capital investment programmes*

2015

1. Jakob Ion Wille
*Film som design
Design af levende billeder i
film og tv-serier*

2. Christiane Mossin
*Interzones of Law and Metaphysics
Hierarchies, Logics and Foundations
of Social Order seen through the Prism
of EU Social Rights*

3. Thomas Tøth
*TRUSTWORTHINESS: ENABLING
GLOBAL COLLABORATION
An Ethnographic Study of Trust,
Distance, Control, Culture and
Boundary Spanning within Offshore
Outsourcing of IT Services*

4. Steven Højlund
*Evaluation Use in Evaluation Systems –
The Case of the European Commission*

TITLER I ATV PH.D.-SERIEN

1992
1. Niels Kornum
    Servicesamkørsel – organisation, øko-nomi og planlægningsmetode

1995
2. Verner Worm
    Nordiske virksomheder i Kina Kulturspecifikke interaktionsrelationer ved nordiske virksomhedsetableringer i Kina

1999
3. Mogens Bjerre
    Key Account Management of Complex Strategic Relationships An Empirical Study of the Fast Moving Consumer Goods Industry

2000
4. Lotte Darsø
    Innovation in the Making Interaction Research with heteroge-neous Groups of Knowledge Workers creating new Knowledge and new Leads

2001
5. Peter Hobolt Jensen
    Managing Strategic Design Identities The case of the Lego Developer Net-work

2002
6. Peter Lohmann
    The Deleuzian Other of Organizational Change – Moving Perspectives of the Human

7. Anne Marie Jess Hansen
    To lead from a distance: The dynamic interplay between strategy and strate-gizing – A case study of the strategic management process

2003
8. Lotte Henriksen
    Videndeling – om organisatoriske og ledelsesmæs-sige udfordringer ved videndeling i praksis

9. Niels Christian Nickelsen
    Arrangements of Knowing: Coordi-nating Procedures Tools and Bodies in Industrial Production – a case study of the collective making of new products

2005
10. Carsten Ørts Hansen
    Konstruktion af ledelsesteknologier og effektivitet

TITLER I DBA PH.D.-SERIEN

2007
1. Peter Kastrup-Misir
    Endeavoring to Understand Market Orientation – and the concomitant co-mutation of the researched, the re searcher, the research itself and the truth

2009
1. Torkild Leo Thellefsen
    Fundamental Signs and Significance effects A Semeiotic outline of Fundamental Signs, Significance-effects, Knowledge Profiling and their use in Knowledge Organization and Branding

2. Daniel Ronzani
    When Bits Learn to Walk Don't Make Them Trip. Technological Innovation and the Role of Regulation by Law in Information Systems Research: the Case of Radio Frequency Identification (RFID)

2010
1. Alexander Carnera
    Magten over livet og livet som magt Studier i den biopolitiske ambivalens