Using machine-learning models for operational exception handling

A case study at IBM

Pim Schultz

October 2017

Supervised by:

IBM Amsterdam

J.P. Hazewinkel

University of Twente

Dr. M.C. Van der Heijden

Dr. A.B.J.M. Wijnhoven

Industrial Engineering and Management

Faculty of Behavioral Management and Social Sciences

Department of Industrial Engineering and Business Information Systems

"Essentially, all models are wrong,

but some are useful."

George Box, 1976

Executive summary

This master thesis project has been conducted at IBM Amsterdam. The goal of this thesis research was to analyze the possibilities of automating the exception handling process of operational service parts planning. In particular, we have analyzed the possibilities of using machine-learning and cognitive computing to predict a planner's actions in the exception handling process.

Motivation

Operational service parts planning at IBM is done via a mix of automated systems and human interventions. Currently a large number of orders are made automatically via Servigistics, a software solution used by the operational service parts departments. While Servigistics automates the standard service part orders, it does not automatically resolve exceptional service part situations. When confronted with an exceptional situation, the system alerts a human planner whom is expected to solve the issue. IBM argues that some of these exceptions could be automated by a cognitive computing system. If parts of the exception handling process can be automated, the planners will able to focus their efforts on the remaining cases. The goal of this research is thus defined as: *"Can the efficiency of the operational exception handling process be increased with cognitive computing systems?"*

Results and conclusions

Since the majority of the data used in the exception handling process is stored in structured databases and the desire of IBM for a system to automate parts of the exception handling process instead of advising employees, we have therefore created a traditional machine-learning model instead of a cognitive computing model.

A proof-of-concept model has been created to predict the planner's actions when presented with a specific exception type, concerned with projected inventory shortage, by the Servigistics system. Inputs given to the model consisted of basic service part information available to the operational planner. Outputs of the system were limited to the potential order types and collected via the Servigistics log files. Linking this input and output data over the months of May to September, 579 cases were used as input for the model. On average, the model correctly predicted the order type in 57% of the cases. Given the model's low accuracy, we would advise IBM to either conduct further research how to improve these models, or to use these models as a second opinion for the planner.

For the creation of a quantitative performance measuring tool for the impact of operational service part decisions, three different formulas have been proposed. When applying these measurement methods to 161 decisions made by the human exception planners, we found the inventory position returning to

normal in 46% of the cases during the evaluation period. If the inventory position did not return to its normal position during the evaluation, it does not need to be a direct result of a planner's decision as the dynamic nature of the service parts management environment makes it hard to predict the impact of decisions. Instead of dismissing these cases as incorrect, we advise the service parts operations department of IBM to inspect a selection of these cases for possible improvements of their exception handling process.

Recommendations for further research

While this research has shown that machine-learning models might be able to increase the efficiency of the exception handling process, further research will need to be undertaken to provide a conclusive answer. Based on the conclusions and limitation of this master research we therefore propose the following areas of interest for future research:

- Improvements to the machine-learning model. To improve the model from this research we would suggest adding more diverse cases to the model, increasing the number of inputs, or automating the input collection process. Increasing the diversity of cases and the number of inputs allows the model to easier distinguish between cases, thus improving its accuracy. Automating the input collection process of the model would enable it to store the contextual information at the moment of the planner's decision thus increasing the accuracy of the inputs.

- Extensions of the machine-learning model. The model proposed and tested in this research is specialized in the handling of a single exception type raised by Servigistics. Some of the other 39 exceptions raised by Servigsitics could be automated with a similar method. In selecting the exception for this research, we have mentioned that there are five other exceptions with similar input and output values. Future research could analyze the possibilities for machine-learning models on these exceptions.

- Cognitive systems in service parts management. While we created a traditional machine-learning model, we believe that there are options for cognitive computing systems in the service parts management environment at IBM. Cognitive systems could, for example, aid in the creation of new business cases by gathering information from unstructured sources such as contracts, emails or financial reports to indicate the profitability of the business case.

Table of Contents

List of figures

List of tables

Abbreviations list

3PL - Third Party Logistics

AI – Artificial Intelligence

CB – Central Buffer

CE – Customer Engineer

CPPS - Common Parts Processes and Systems

DSS – Decision Support System

EMEA – Europe, Middle-East and Africa

EXC – Excess stock level

GARS - Global Asset Recovery Services

IBM - International Business Machines Corporation

MOP – Must Order Point

MOR – Maximum Order up to point

PIMS - Parts Inventory Management System

PSS – Policy Safety Stock

SPO – Service Parts Operations

TSS - Technical Support Services

1 Introduction

This master thesis has been commissioned by IBM to research the possibilities of automating the exception handling processes in their operational service parts management. The research subject has been proposed as part of the research agreement within the ProSelo Next group and has been delegated to two students from different departments at the University of Twente, namely the department of Industrial Engineering & Business Information Systems and the department of Human Media Interaction. This chapter will be used to provide some background information on IBM, explain the ProSelo Next research agreement, give a small introduction to the field of service parts management and explain the collaboration between students in more detail.

1.1 IBM

Originally started as the Computing-Tabulating-Recording Company in 1911, the International Business Machines Corporation (IBM) has grown to become one of the biggest players in the tech industry as well as one of the most recognizable brands in the world (Brand Finance, 2017). Whereas IBM used to be big player in the consumer computer market in the 20th century, nowadays the focus is on the business to business market segment. Over the last decades IBM has shifted its focus from hardware sales to selling services, solutions and software products. IBM desires to be an organization that provides cutting edge, innovative, technology, and has therefore heavily invested in cloud and cognitive solutions. IBM expands its portfolio in these markets via inhouse developments such as the creation of their cognitive computing framework Watson, or by acquisition such as the takeover of the big data specialist The Weather Company in 2015 (IBM, 2016a).

1.2 ProSeLo Next

The research project Proactive Service Logistics for Advanced Capital Goods - the Next Steps (ProSeLo Next) is a follow up research initiative of the initial ProSelo project, which ran from 2010 to 2015. ProSeLo Next is a research agreement between three universities and nine companies: University of Twente, University of Tilburg, Erasmus University Rotterdam, ASML, Fokker Services, IBM, Marel Stork, Océ, Thales Nederland, Vanderlande, Gordian Logistic Experts, and the Stichting Service Logistics Forum (Basten, 2015). The research agreement will last until 2020 and consists of the following work packages:

1. Predictive maintenance and service logistics: By (remotely) monitoring assets and timely maintenance, efficient service logistics becomes possible. Interesting models have been developed for this, partly in the ProSeLo project; we aim to apply these in two pilot studies and

to use the results to improve the models.

2. Service business models: When an OEM or service provider assumes responsibility for maintenance and service logistics (i.e., the availability) of assets, both involved parties have their own interests that need to be aligned.

3. Service control towers: Coordinated management requires service control towers to be constructed. We focus on operational decision making in dynamically changing situations, as requested by companies in response to the need to quickly adapt to rapidly changing market requirements (NWO, 2016).

The research subject of this thesis is related to the third work package as we are researching the operational decision-making processes within the service parts control tower environment at IBM.

1.3 Cognitive computing

Cognitive computing is seen by many as the next step in the evolution of machine-learning, in particular as an extension of the artificial intelligence field (Chen, Argentinis, & Weber, 2016; Y. Wang, Zhang, & Kinsner, 2010). While there is no consensus on the definition of cognitive computing yet, in this thesis we will use the definition given by IBM: "*Cognitive computing refers to systems that learn at scale, reason with purpose and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment.*" (IBM, 2016a; Kelly, 2015). When these smart machines are introduced to a business their usage can be roughly classified in three categories: (i) they can be used to enhance existing human capabilities by, for example, assisting employees via historic decision feedback; (ii) they can take over tasks from the human workforce such as the usage of chatbots in the customer service industry; or (iii) they can be used in conjunction with the human work force to, for example, lift and relocate heavy objects (Dalton, Mallow, & Kruglewicz, 2015; Kelly, 2015). More information regarding machine-learning and cognitive systems can be found in 3.4.

1.4 Service parts management

Service parts management is defined as the activities required to ensure the right availability of service parts against minimal integral costs (Driessen, Arts, van Houtum, Rustenburg, & Huisman, 2015; Rustenburg, 2016; Slack, Brandon-Jones, & Johnston, 2013). Managing an organizations service parts portfolio requires operating in an environment with an ever-increasing complexity as:

- there is an increased focus on working capital,

- there are reductions in field knowledge and expertise,
- customers are demanding customized performance levels,
- products are becoming obsolete sooner due to fast product development cycles,
- reductions in the size of installed bases lead to more complicated demand patterns,
- organizational assets become increasingly complex due to technical advances limiting comprehensible failure recognition (Rustenburg, 2016).

When organizations are unsuccessful in their balancing act between costs and performance it may result in: lost revenues (e.g. manufacturing lines not being able to produce due to part failure), customer dissatisfaction (e.g. delays in transportation of goods such as parcels), or public safety hazards (e.g. failure of critical fail-safe parts) (Driessen et al., 2015).

Managing the service parts portfolio in this dynamic environment requires numerous planning and control options on all organizational levels. An indication of these options within the service parts management environment is given by Driessen et al. (2015). While the authors point out that this framework is not directly applicable to every organization, it can be used as a starting point for creating and analyzing maintenance service parts planning control systems. The first step of service parts management is the decision to include a service part in the existing portfolio, the so-called assortment management step. After the decision to include the part in the organization's portfolio, the expected demand and return distributions of the service part are calculated. When the demand for a part has been calculated, one or more suitable sources of new supply and/or repair shops must be located. These five steps together create the input for the inventory control systems and enable the organization to decide upon stocking and replenishment strategies. The order handling process uses all this data to decide upon acceptance, change, or rejection of an order, for releasing service parts mentioned in the order, and for handling the return orders of defective parts. Finally, the deployment process is responsible for the replenishment of the service parts inventories by either releasing procurement orders, repair orders, or lateral transshipment orders (inventory movements between locations on the same hierarchy level). The framework by Driessen et al. (2015) is shown in Figure 1.

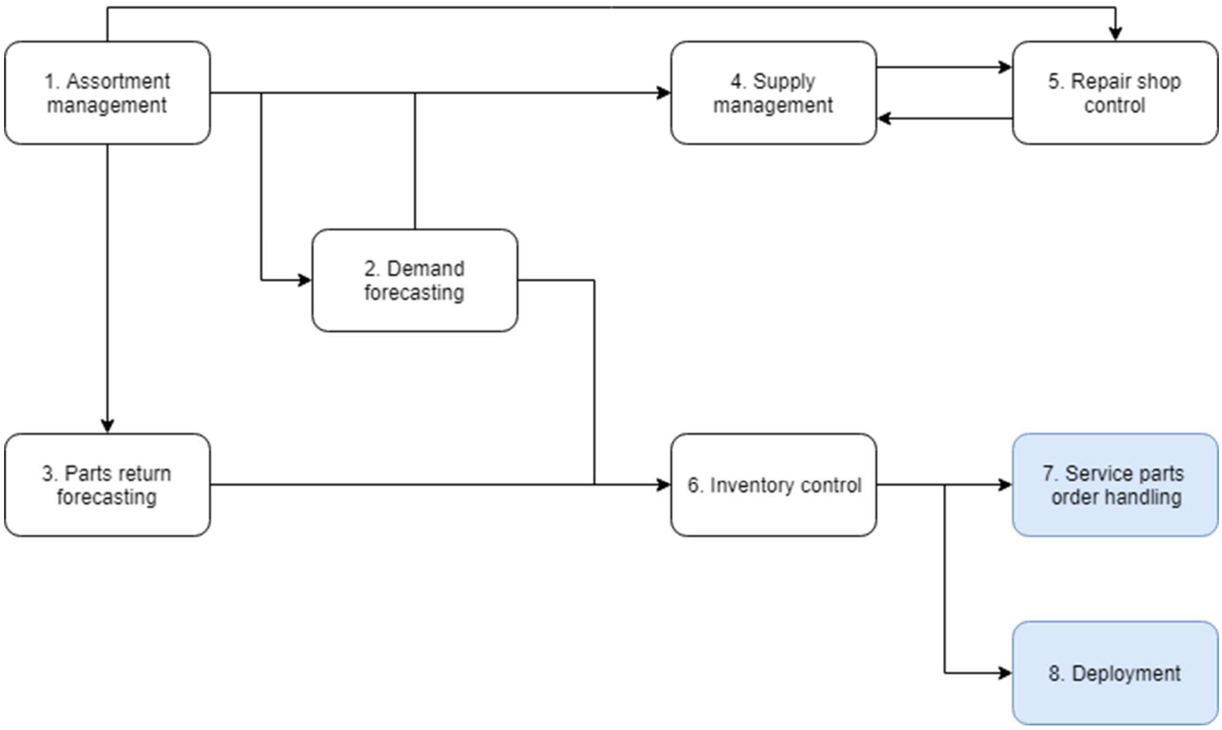

Figure 1: Decisions in maintenance logistics control

Given the focus of this thesis on the operational service parts management processes at IBM and the focus of the ProSelo Next project, we can state that this research will be concerned with last two parts of the framework: service parts order handling and deployment. A more detailed description of the operational service parts management environment is found in 3.1.

1.5 Collaboration

To gather insights in both the business and IT perspectives of the possibilities for a cognitive computing solution, IBM has proposed this thesis' research to two faculties at the University of Twente namely: Industrial Engineering & Management and Human Media Interaction. One student of each of these faculties will be granted the assignment and they will collaborate during their research period. For this thesis I, Pim Schultz, am the student from the faculty of Industrial Engineering & Management and will cooperate with Sofia Kyriazi from the faculty of Human Media Interaction. This collaboration will be done as follows: I will research the theoretical possibilities of implementing a cognitive solution within the operational service parts control tower environment and will be responsible for both the selection of data to include in the solution and the creation of a performance measurement tool. Sofia will delve into IBM's portfolio of existing tools and services to identify the most fitting ones for a possible solution as well as be responsible for the coding of the solution and the solution's user interface.

2 Research design

Within this part of the thesis we will describe IBM's motivation for the commissioning of this thesis, the research questions created to answer the overarching research subject, the scope for this thesis, the methodology used, the deliverables from this research, and provide an outline for the remainder of this thesis.

2.1 Motivation

For the past decade, IBM has been investing heavily in cognitive computing services such as natural language processing. IBM seeks to leverage this new technology to provide new solutions to existing clients, create solutions in new markets, and to improve its internal processes. An example of a solution for a new market is the implementation of cognitive systems in healthcare. Here, cognitive solutions can provide value to the medical staff by processing the large amounts of structured and unstructured data from medical records with a combination of natural language processing and machine-learning services. After the internalization of data, cognitive systems can assist healthcare professionals in their medical diagnoses. Initial test runs at several hospitals seem to be positive, and were seen as especially helpful in diagnosing patients that have been afflicted with a rare decease (Herper, 2017; IBM, 2016a, 2017; NG, 2016)

While these innovative technologies have helped IBM deliver more value to its customers, IBM would like to use these technologies to improve their own processes as well. One of these potential improvement areas identified by IBM is the exception handling process in the service parts operations (SPO) department. In recent years, most of the standard operational inventory management procedures have been automated but IBM still employs human planners to intervene when needed. These planners are prompted for action by the system when exceptional behavior of a service part is detected. Planners will then analyze the service part's situation and attempt to solve the issue.

However, there are three factors which may limit the efficiency of the operational exception handling process: (i) when operational planners perform well on their task, they are often promoted to higher positions, resulting in personnel turnover and brain drain at the operational level; (ii) a planner's service part portfolio changes every three months, resulting in planners having to relearn the intricacies of the service parts under their control; (iii) planners will generally not receive feedback on the impact of their decisions, resulting in planners not learning from their previous decisions.

To improve their operational exception handling process, IBM argues that if cognitive systems could be used to automate some exceptions, the operational planners would be able to focus their efforts on the truly exceptional cases. Additionally, cognitive systems could provide planner feedback on their earlier decisions. The goal of this research is thus defined as: *"Provide insights in the possibilities of applying cognitive computing to the operational exception handling process at IBM, in order to improve the efficiency of the exception handling process"*

2.2 Research questions

To provide an answer to this overall research goal and structure this research, we formulate five research questions. These research questions will answer the main research question of this thesis: "*How can cognitive systems improve the efficiency of the operational exception handling process at IBM?"* These research questions are:

1. What are the characteristics of the operational service parts environment?
    a. What are the characteristics of the decision maker in this environment?
    b. How can technologies be used to aid the decision maker in this environment?
2. What are the characteristics of machine-learning and cognitive computing?
    a. What are the differences between machine-learning and cognitive computing?
    b. What are the requirements for implementing these systems?
3. What are the characteristics of the operational exception handling processes at IBM?
    a. What information is available to the planners?
    b. What systems are used by the planners?
    c. What steps are taken in the current exceptions handling process?
4. How can a model support the exception handling process?
    a. What inputs should be given to the model?
    b. Which algorithm should be used by the model?
    c. How can the model's performance be measured?
5. How can we assess the quality of the exception handling process?
    a. What should be defined as a correct decision?

b. How can the impact of a decision be measured?

Following the answers to these research questions, a final chapter will be devoted to discussing and combining these answers into an overall conclusion.

2.3 Scope

This master research project will focus on the impact of the operational exception handling decisions related to clients serviced by IBM within the European, Middle-Eastern & African (EMEA) region. This region consists of 518 inventory warehouses which are generally replenished by parts from the central buffer location in Venlo, The Netherlands. Instead of tackling the exception handling problem for the entire EMEA region this thesis will focus on the exceptions raised at this central buffer location.

The argumentation for this scope definition is threefold: resources constraints, present knowledge base, and complexity reduction. First, given the time and resources available for this master thesis, handling the entirety of the worldwide service part data within this timeframe would be impossible due to its sheer volume, as the EMEA region alone has more than 1.4 million active service parts. Secondly, this research will be conducted from the IBM location at Amsterdam, the Netherlands, where most of the employees are responsible for the service parts management in the EMEA region and thus have extensive knowledge of this region's workings and intricacies.

Thirdly, since one of the goals of this research is to create a proof-of-concept model for the exception handling process within the limited timeframe available for this research, we want to reduce the complexity of the environment to be modelled. We accomplish this by choosing to focus this research on the central buffer location instead of the entire EMEA region, thus decreasing the complexity by eliminating the rules and regulations that exist in moving service parts between stock locations across the various countries. In addition to reducing the number of locations to be used in this research, the prototype model to be created during this research will be aimed at solving a single exception type. We believe that if the prototype model will be successful in automating a certain type of exception, similar design steps can be used to extend the model for other exceptions. More information regarding the selection of a suitable review reason for the initial model can be found in 5.2.1.

2.4 Methodology

Broadly speaking this thesis can be divided in three parts. First, we must understand the service parts management environment SPO operates in, the systems and methods they use, and the controls available to them. Second, after creating an overview of the current situation and environment, we can start our search for a fitting solution. Given the preference of IBM for a cognitive computing model, we will start our research in this direction. If, however, we discover a mismatch between the desired solution type and the environment it must be placed in, we will branch out to find other possible solutions. Finally, we will create a proof-of-concept model based on our research.

To systemically research the objectives of these three parts mentioned above, we have created five research questions, as mentioned in 2.2. To answer those five questions, we will use a combination of literature review, interviews, data review, expert opinions, and brainstorm sessions. These methods will use the following data sources: scientific literature, IBM employees, IBM databases & systems, IBM's internal research notes, and system documentation. The methodology and data sources to be used for each research question are shown in Table 1.

Table 1: Methodology broken down per research question

#	Subject	Methodology	Data sources
1	Characteristics of the service parts management environment.	Interviews Literature review	Scientific literature IBM (employees & databases)
2	Characteristics of machine-learning models and cognitive computing systems.	Interviews Literature review	Scientific literature IBM (employees)
3	Characteristics of the exception handling process at IBM.	Interviews Data review Expert opinions	IBM (employees & databases) Documentation
4	Model creation and testing.	Brainstorm Expert opinions Literature review	Scientific literature IBM (employees)
5	Review the actions made by operational planners.	Data review Expert opinions	IBM (employees)

2.5 Deliverables

The main deliverable of this research project will be a master thesis covering the following topics:

- An overview of the current operational service parts management process at SPO,
- A prototype model for the automation of the exception handling process,
- An indication of the possibilities and requirements of cognitive systems,
- Suggestions for the implementation of the model,
- Suggestions for further expansion of the model.

In addition to the thesis, I will hand over the source code of the model used for quality testing to IBM after the completion of this thesis.

2.6 Thesis outline

Within this thesis, we will use the five research questions stated earlier as a guideline for the chapter order. Chapter 3 will consist of the theoretical framework of this thesis and will contain the answer to the characteristics of the operational service parts environment, the characteristics of machine-learning, and the characteristics of cognitive systems. The following chapter will then provide an overview of the current situation at the SPO department of IBM and the information available to the operational planners. Then the characteristics of the prototype model and its performance, will be discussed in chapter 5. In chapter 6 we will propose a method of assessing the quality of the exception handling process. The last chapter is reserved for the overall conclusions, limitations and suggestions for further research. A breakdown of the research questions and their relevant chapter is given in Table 2.

Table 2: Thesis outline

	Research question	Chapter(s)
1	What are the characteristics of the operational service parts environment?	3
2	What are the characteristics of machine-learning and cognitive computing?	3
3	What are the characteristics of the operational exception handling processes at IBM?	4
4	How can a model support the exception handling process?	5
5	How can we assess the quality of the exception handling process?	6

3 Theoretical framework

Within this chapter, we will describe the results of the literature search. The goal of this literature research is threefold: (i) define the characteristics of the operational service parts management environment and the role of the operational decision maker; (ii) identify systems that can be used to automate the exception handling process; and (iii) describe the characteristics of a control tower environment. The insights gained from this literature research will be used to answer the first two research questions regarding the characteristics of the operational service parts environment and the characteristics of machine-learning models.

3.1 Operational service parts management

In any organization, we can distinguish between three levels of decision making: strategical, tactical, and operational. Strategic decisions relate to the long-term usage of the organizations assets to achieve its goals, e.g. deciding what markets to enter. Tactical decisions relate to the implementation of the strategic decisions such as deciding a part's inventory stock levels. Finally, operational decisions relate to the day-to-day decisions such as creating the order delivery schedule (Schmidt & Wilhelm, 2000; Slack et al., 2013). Applying these definitions to the service parts management environment, we find that operational decisions are mostly related to the daily management of the inventory positions of the various service parts. A task which, in a predictable environment, should be relatively easy to automate. The operational service parts management environment however, should not be classified as a predictable environment but as a dynamic decision-making environment according to Brehmer (1992). In this paper, he argues that a dynamic decision-making environment consists of four characteristics:

1. Reaching the goal requires a series of decisions. The system's controller will need to make numerous decisions to maintain in control of the system.
2. Decisions made within the system are not independent. A decision made by the system's controller will influence decisions made at a later point in time.
3. The decisions problem's state changes, both autonomously and because of the actions made by the system's controller.
4. Decisions need to be made in real time, as decision makers are forced to react to changes in the environment which are largely out of their control.

Based on these four characteristics we can conclude that the service parts management environment is a dynamic decision-making environment. The reasoning for this conclusion is as follows: (i) decisions made by controllers in the service parts operational environment have the goal of attaining a certain service level for the customers, requiring a series of decisions during the part's lifecycle; (ii) any decision made by a planner will also impact his future decisions, e.g. placing a large order today may lead to excess holding costs or scrapping parts later in the service part's lifecycle; (iii) the system's problem state will change both without input from the planner, e.g. a client order, and as a direct result of the planner's actions, e.g. a stock replenishment order; (iv) decisions are made in response to external events outside of a planner's control and failing to react in time could have major consequences on the system's health.

3.2 Human decision maker

Now that we have defined the service parts management environment as a dynamic one, we will delve deeper into the characteristics of the operational service parts planner's tasks. The tasks appointed to these planners can be distilled to six dimensions according to the general systems theory (Brehmer, 1992). These dimensions are: the complexity of the system, the feedback quality, the feedback delay, the rate of change, the relation between the process to be controlled and those used for control, and the extent to which the decision-making power can be delegated.

Complexity of a task differs per controller and should be defined as a relative concept that compares the capability of the controller to the complexity of the system. It is in this dimension that the differences between a human and a computerized controller are the most diverse. Human decision makers will have a hard time processing many items simultaneously and will thus define a system with many elements as a more complex system. Traditional computerized controllers will have difficulties controlling systems in which decisions must be made based upon interpretations of data instead of fixed rules, but cognitive computing systems will be able to circumvent this limitation (Dalton et al., 2015; Kelly, 2015; Lerch & Harter, 2001).

Quality and delay of feedback in the system are directly related to the variations in the observability of a system and thus to the consequences of the planner's actions. In addition to the quality and the delay of the feedback, the frequency of feedback also influences the performance of the user's decisions (Brehmer, 1992; Sterman, 1989). Receiving feedback too often, might lead to users becoming overwhelmed with feedback and becoming less selective in their processing of the feedback (Lurie & Swaminathan, 2009). Providing feedback to a decision maker in a dynamic environment in the form of information technology has been proven to help decision makers if the feedback is not solely focused on the decision outcomes. The most effective way to use information technology for the feedback process is by providing users with a feedforward based expert system detailing the steps a system expert would take in the situation

(Gibson, 2000; Gonzalez, 2005; Lerch & Harter, 2001).

The rate of change in a process refers to the possibility of the controller to delay decisions based on direct feedback from the system. In cases of a system with a high rate of change, decisions need to be made quickly to maintain in control, e.g. close-combat engagements resulting in the controller to make decision based on feedforward control. Systems with a slow rate of change are less dynamic as their decision series become less independent and more sequential in nature (Brehmer, 1992). Within the field of service parts management, the rate of change depends highly on the type of part. Some parts have a high turnover rate, so-called fast movers, creating a more dynamic planning problem whereas other parts, slow movers, may be sparingly demanded, resulting in a more deterministic planning environment.

As for the relation between the process to be controlled and the processes available for control, also referred to as the requisite variety within the system theory, decision makers within the service parts industry should be able to influence most factors related to inventory management (Godsiff, 2010). While some factors will remain out of their control, freak of nature events such as earthquakes, they could still hedge their bets to some degree against the impact of such events.

The delegation of decision making power within the system mostly refers to the ability to distribute the decision-making responsibilities to the lower system levels closest to the actual events. Brehmer (1992) states that, especially in system with delays, this may increase the system's efficiency. A good example of delegating the decision-making power downwards can be found in the military where the squad leader has full decision-making power during an operation.

3.3 Decision support systems

To aid the human decision maker in their complex task, organizations are increasingly often implementing technical solutions such as decision support systems (DSS). These systems aim to minimize the decision maker's cognitive errors and maximize their performance by aiding users in the different decision-making steps such as information collection, situation evaluation, alternatives generation, alternatives selection, and solution implementation. However, no matter the power of the support system, the final responsibility of the actions, remains with the human decision maker (Niu, Lu, & Zhang, 2009; Papathanasiou, Linden, & Ploskas, 2016). Even though decision support systems can provide planners with good suggestions, research has shown that planners will not always follow these directions. This can be explained by the belief that an increased mental effort in an employee's task should lead to an increase in performance compared to blindly following the system, especially in situations where employees are confident in their expertise. This issue can be remedied by increasing the person's trust in the system

(Fransoo & Wiers, 2008; Geertjes, 2014).

The decision support systems most commonly implemented by organizations can be classified as: model-driven, communication-driven, data-driven, document-driven and knowledge-driven. Model-driven DSS are based upon statistical optimization or simulation such as linear programming used for forecasting demands. Communication-driven DSS aid in the effective sharing of information between members of a group, an example of such a DSS would be an electronic meeting system. Data-driven DSS are often tools such as IBM's Cognos that aid in the collection and manipulation of large amount of internal and external company data. Document-driven DSS are specialized in processing and manipulating electronic documents to, for example, automate a financial approval process. Finally, knowledge-driven DSS apply information from human expertise in areas such as business procedures and rules, to generate decision suggestions for a specific domain (Niu et al., 2009; Power, 2016; Power, Sharda, & Kulkarni, 2007). A framework of the DSS types and their characteristics as proposed by Power (2016) is shown in Table 3. This framework will be used to classify the proposed model in 5.1.

Table 3: Decision support system framework

Decisions support system type	Dominant DSS component(s)	User(s)	Purpose(s)	Enabling technology
Communications-driven	Communications	Internal teams	Conduct a meeting	Bulletin boards
Data-driven	Database	Managers and staff, suppliers	Query a data warehouse	Relational databases
Document-driven	Document storage and management	Specialists and user groups	Search Web pages	Search engines
Knowledge-driven	Knowledge base, Artificial intelligence	Internal users, new customers	Management advice	Expert systems
Model-driven	Quantitative models	Managers and staff, new customers	Scheduling, forecasting	Linear programming

3.4 Machine-learning

Machine-learning systems are systems that improve automatically through experience. Improvements in this area have been rapid, as can be illustrated by the influx of machine-learning systems in health care, manufacturing, and education. The methods used for machine-learning can be roughly divided in three categories: supervised learning, unsupervised learning, and reinforcement learning (Jordan & Mitchell, 2015; Kubat, 2015).

Supervised learning is the most widely adopted machine-learning method and is able to link a system's outputs to inputs. Creating these links, also called mapping, requires a training set of labelled data. A supervised learning machine will then use algorithms such as decision trees and logistic regression to create the mapping. Links between the input and output pairs can be deterministic or probabilistic depending on the dataset and the requirements. This method can, for example, be used to predict the type of groceries bought depending on social economic factors.

Unsupervised learning enables machines to explore unlabeled data and find clusters within. Algorithms such as cluster analysis and market basket analysis are often used to identify segments within the given dataset. A usage example of this machine-learning method is the ability to identify a shop's consumer segments enabling it to more directly target these segments in its ads.

Reinforcement learning is a mix of both supervised and unsupervised learning as training samples given to the system will only provide the system with an indication of the correct input-output pair. Systems using reinforcement learning will therefore attempt to continuously search for the best possible outcome given a set of input values. This machine-learning method are often implemented in a general control-theoretic setting in which an agent is required to learn a control strategy for acting in an unknown dynamical environment. The aim of the agent is to choose a set of actions for any given state where he will maximize his expected returns. The agent can continuously update his best course of action, by incorporating the feedback given by the system on the agent's previous actions. Examples of reinforcement learning applications are self-driven cars and product delivery routing.

3.4.1 Cognitive computing

Cognitive computing is seen by many researchers the next step in the evolution of machine-learning and artificial intelligence (AI). The main difference between AI and cognitive computing can be found in the answering style of both systems. G. Rometty, IBM's CEO, describes this difference as follows: "*In an artificial intelligence system, the system would tell the user which course of action to take based on its analysis. In cognitive computing, the system provides information to help the user make that decision*" (Consortium, 2016; Dalton et al., 2015; Rometty, 2017). This statement embodies the vision of IBM to use cognitive computing as an intelligence amplification tool.

One thing most researchers agree on however, are the four characteristics of a cognitive system. To receive this classification a system must be adaptive, interactive, iterative, and contextual. *Adaptive* refers to the ability of the system to resolve ambiguity, tolerate unpredictability, and process dynamic data in, or near, real-time. This is sharp contrast with the plethora of programmable systems used today, which produce their output based on deterministic rules. *Interactive* cognitive systems allow its users to define their needs and wishes intuitively and comfortably. This dimension refers to both the interaction between user and machine, and the interaction of cognitive systems with other machines and services. The *iterative* dimension of a cognitive systems relates to the cyclical problem-solving methods incorporated by these systems. A cognitive system helps with the problem definition by prompting the user for more information or by collecting additional sources of input in case of statement ambiguity or incompleteness. Additionally, the system should remember earlier process interactions with the user to return the most appropriate information at the end of the interaction. Finally, the system should understand, identify and extract *contextual* information such as natural language syntax, regulations, and user profiles. It is this requirement that enables cognitive systems to ingest a larger diversity of data compared to traditionally programmed systems (Consortium, 2016; Dalton et al., 2015).

Cognitive systems generate their answers via a process known as hypotheses generation. Instead of using programmed rules to immediately a single, definitive, answer, cognitive systems generate multiple hypotheses. The cognitive system ranks the hypotheses based on their fit to the question context and its previous interactions. It then provides the user with the answer that has the highest probability of being the correct one, also known as the one with the highest confidence. Users are asked to provide feedback on the generated answer to allow the system to learn from this interaction. This feedback will be incorporated in the knowledge base for future reference (Ferrucci et al., 2010; Goksel Canbek & Mutlu, 2016; Kelly, 2015). An illustration of this hypotheses generation and answering is shown in Figure 2.

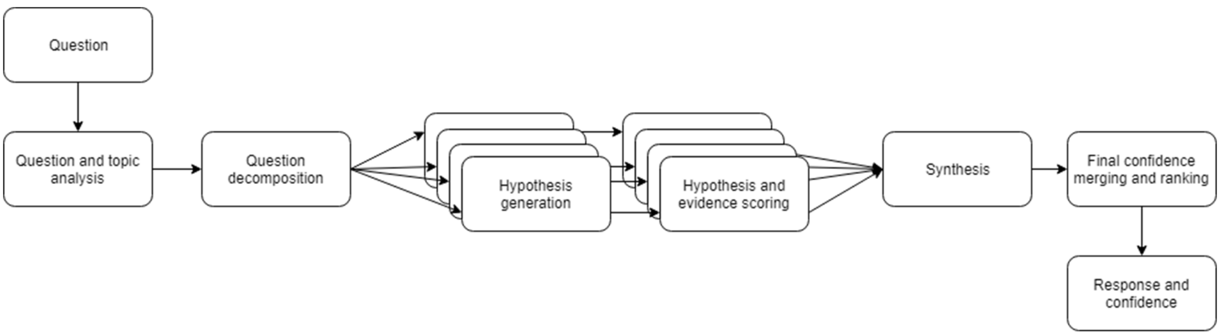

Figure 2: Cognitive system hypothesis generation

3.5 Control tower environment

Supply chains of organizations are the linkages of upstream and downstream flows of products services and information. Control towers are designed to visualize all these processes in a centralized environment. In addition to a simple visualization of the processes, the functions of a control tower are fivefold. It can be used by the supply chain's organizations for planning and routing, auditing and reporting, forecasting, event management, and decision making (Bhosle et al., 2011). To properly fulfill these five functions a supply chain control tower consists of five main layers. These are, from bottom to top: the supply chain business layer, the information perception layer, the information operation control layer, the information service platform layer, and the information manpower layer, as shown in Figure 3 (Shou-Wen, Ying, & Yang-Hua, 2013; Trzuskawska-Grzesińska, 2017).

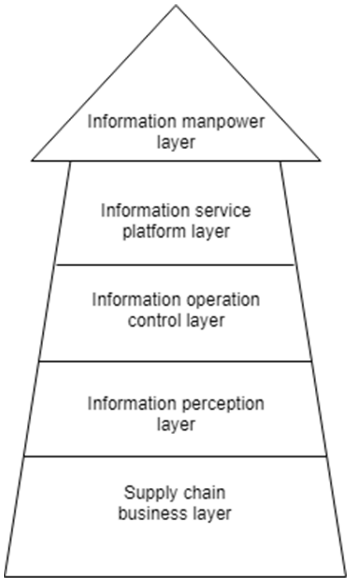

Figure 3: Control tower environment

Supply chain business layer

The supply chain business layer consists of all the processes required in a successful supply chain and thus is therefore modeled as the base of the supply chain control tower. Within the general supply chain control tower proposed by Shou-Wen et al. (2013) this layer contains the procurement, transportation, warehousing, loading and unloading, handling, distribution processing and packaging, distribution, and information service processes. Organizations that are typically present in this layer of the supply chain include: raw material suppliers, manufacturers, outsourcing logistics service providers, distributors, dealers and users.

Information perception layer

Within the information perception layer technologies such as the Internet of Things can be used for integrating information collection and transmission. Data gathered from the information perception layer is at the heart of the analytical processes and thought should be given to the collection methods used. This data can be collected and transmitted by the supply chain organizations themselves or they can hire an external organization to do so. Hiring an external organization instead of creating an in-house solution, could be used to alleviate asymmetrical trust levels between the organizations (Cetindamar, Çatay, & Basmaci, 2005; de Kok, van Dalen, & van Hillegersberg, 2015; Khurana, Mishra, & Singh, 2011).

Information operation control layer

In the information operation control layer, we find the supply chain information storage and the supply chain information control sections. The data originating from the supply chain business layer and gathered through the methods from the information perception layer are stored within this layer (Shou-Wen et al., 2013). Storing the data should be done in a repository that is accessible to all the members of the supply chain. Nowadays organizations are often using cloud-based data storage as cloud storage is easily scalable and can be acquired as a service, reducing the initial financial investment (Raj & Sharma, 2014). The data acquired from a product's journey through the supply chain, will be used to provide feedback on its quality after every step of the supply chain (Shou-Wen et al., 2013; G. Wang, Gunasekaran, Ngai, & Papadopoulos, 2016; L. Wang & Alexander, 2015).

Information service platform layer

Next, we have the information service platform layer and its purpose is threefold. Firstly, it is used to centrally store and dynamically update the information from the lower control tower layers to improve and maintain the required transparency. Secondly, it is used for the real-time monitoring of the information processes within the supply chain. Finally, it provides an insight in the whole supply chain quality issues and its feedback control by comparing the current processes to the supply chain metrics. The combination of real-time monitoring and transparency of the data, enables the supply chain partners to reduce their time to action in case of an unwanted event occurring at any point in the supply chain reducing the potential impact of these events (Heaney, 2014; Jüttner, Peck, & Christopher, 2003).

Information manpower layer

At the "top" of the supply chain control tower is the information manpower layer. Within this layer the human supply chain decision makers are located. They use the data gathered in the "lower" layers of the supply chain control tower to detect anomalies and intervene. This information can be used to control the design of the overall supply chain network on the strategical level, enable proactive planning of the supply chain operations and distributions on the tactical level, and enable real time functionality such as transportation insights and inventory tracking on the operational level (Shou-Wen et al., 2013).

3.6 Chapter summary

This chapter was written to create a better understanding of the theoretical side of the operational service parts management environment, to provide a basis the role of the operational decision maker, to research the characteristics of a machine-learning system, and to explore the control tower environment. The most important findings in this chapter are:

- The service parts management environment can be classified as a dynamic decision-making environment based on: reaching the goal requires a **series of decisions** to be made, these decisions are **dependent** on each other, the decision **problem state changing continuously**, and decisions needing to be made in **real time**.
- The characteristics of decision makers within this environment can be classified along six dimensions: the **complexity** of the system, the **feedback quality**, the **feedback delay**, the **rate of change**, the **relation between the process to be controlled and those used for control**, and the extent to which the **decision-making power can be delegated**.
- Decision support systems can be deployed to **aid** human decision makers in their task by **minimizing cognitive errors, collecting information, evaluating the situation, generating alternative actions**, and **implementing solutions**.
- Machine-learning can be used to **predict** outcomes, **uncover** segments within the data, or **improve** the system's decision-making processes over time. Such models can be divided in three categories: **supervised**, **unsupervised**, or **reinforcement** learning.
- Cognitive systems can be used to **enhance, replace**, or **cooperate** with the human workforce. The **strengths** of a cognitive system compared to a programmable system are its **adaptability**, their ability to process both **structured** and **unstructured** data, and the ability to **improve** over time by incorporating user **feedback**.
- The typical control tower environment consists of five layers: the **supply chain business layer**, the **information perception layer**, the **information operation control layer**, the **information service platform layer**, and the **information manpower layer**.

4 Status of IBM processes and systems

Within this part of the thesis, an overview will be given regarding the processes and systems at IBM relevant to this thesis. We will start with a general explanation of the service part operations (SPO) and its supply chain. Then, an overview of the systems used within this environment are explained. Finally, an overview of the planning and exception handling processes, and a more detailed explanation of the exceptions raised by the systems are given. By using this information, we aim to answer the third research question regarding the information currently available to the operational planning staff at IBM, and to illustrate the current status of SPO at IBM.

4.1 Service parts operations

The main responsibilities of IBM's SPO are the allocation, cost management, and inventory planning of the service parts for their customers. The service parts managed by SPO are only meant for the business-to-business market and the contractual agreements are highly customizable. For service parts regarding IBM machines, so called logo products, and consigned brands such as Lenovo and Cisco, SPO is responsible for all the previously mentioned processes. Other organizations however, only want SPO to be responsible for some of the processes such as the inventory planning. Whatever the responsibilities of SPO for the customer, SPO always has to carry out its operations in a highly variable and uncertain environment.

Within the SPO supply chain there are over 300 different suppliers and nearly 500 locations ranging from large central warehouses to small lockers. Last year, 2016, more than 420,000 orders were related to SPO activities across IBM's entire network, these orders ranged from simple electrical cables costing a couple cents to high tech servers costing thousands of dollars. IBM has been working to automate most of these orders via their Servigisitics Plan system for the past decade. Last year 77% of the all the SPO orders were handled by the system, so called hands-free orders, therefore freeing up large amounts of human capital. Human intervention, i.e. hands-on ordering, was only required in less than a quarter of the cases. A breakdown of the global hands-on orders is given in Table 4.

Table 4: Global hands-on orders breakdown

Order type	Orders	% of total orders
Warranty	44,181	45%
New Buy	32,507	33%
Repair	20,062	20%
Emergency	1,569	2%
New business opportunity	95	<1%
Excess, surplus and scrap	47	<1%
Total	**98,461**	**100%**

A bit more than a fourth (27%) of the global SPO orders where directly related to the EMEA region. The biggest difference between the global hands-on order breakdown and that of the EMEA region is the decrease in warranty and repair orders. This difference can be attributed to the other regions using PIMS (Parts Inventory Management System) instead of CPPS (Common Parts Processes and Systems) for their service parts management. Within PIMS, the parts available for repair are not linked to the Servigistics ordering system, thus requiring the operational planners to create a feed-ordering process via spreadsheet calculations.

A characteristic specific to the operational planning of the EMEA region, is that all the emergency orders are placed by local warehouses in the network on the region's central buffer, i.e. internal replenishment orders. Planners in this region cannot place emergency orders on suppliers as this option is not included in the contracts between the supplier and IBM. The order types are further broken down in Table 5.

Table 5: EMEA hands-on orders breakdown

Order type	Orders	% of total orders
New Buy	11,615	91%
Emergency	754	6%
Warranty	204	2%
Repair	181	1%
Excess, surplus and scrap	28	<1%
New business opportunity	8	<1%
Total	**12,790**	**100%**

4.1.1 Supply chain

The service parts department of IBM should be seen as the link between the service part suppliers and end users. Service parts can enter the supply chain via manufacturers, new buy vendors, reutilization or repair vendors, Global Asset Recovery Services (GARS) as remains from machine dismantling, or the open market after being tested by GARS. In almost all cases within the EMEA region, the supplier will send its products to the central buffer location in Venlo, the Netherlands. There are two exceptions: (i) a service part is bought via the open market, in this case parts are first send to GARS in Germany for quality checking; (ii) a service part is bought from a local supplier, in this case the service parts are directly send to one of the local stocking hubs.

Next to the large central buffer warehouses and the local stocking hubs, IBM also implemented small lockers filled with critical parts close to the customer. Deciding where and how much of the service parts to store depends on a combination of factors with the most important ones being: the expected failure rate of the part, the criticality of the part, the cost of the part and the demands of the customer. In case of expensive slow-moving parts, SPO may decide to stock the part only on the global buffer in Mechanicsburg, USA. When this part is needed, a transshipment from the global buffer to the region's central buffer is ordered. Lateral replenishments are almost non-existent in the EMEA region, due to legal difficulties in importing and exporting service parts across the countries of the region.

Finally, if parts are deemed obsolete, they will be marked as scrap and transported to a scrap vendor, thus leaving the network. An illustration of this supply chain environment is shown in Figure 4.

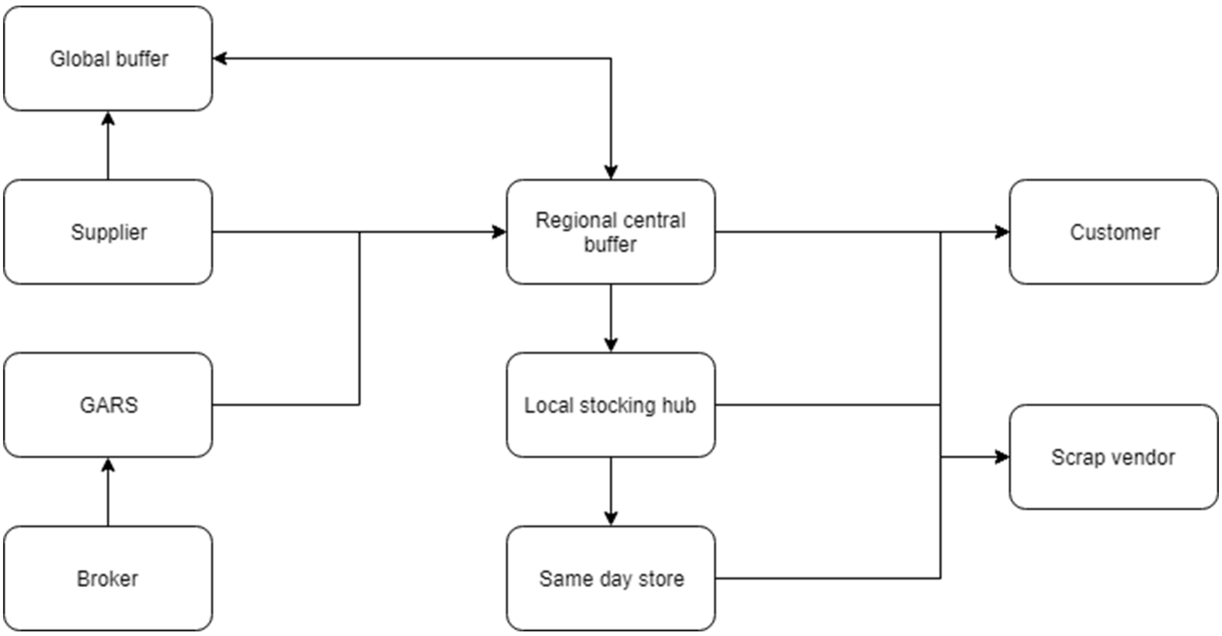

Figure 4: SPO EMEA Supply chain

4.1.2 Service parts management control tower

The current service parts management control tower environment at IBM has been dubbed control tower

2.0. At the base of its control tower are the operational data storage and datamarts within IBM. These processes store and organize all the information of the service parts management process, from the service part stock levels at location to the orders placed at the suppliers. The data from these datamarts is directly accessible and manipulatable by the systems used at SPO, CPPS and Servigistics, which are described on more detail in 4.1.3. An external organization, Entercoms, uses the information from IBM's datamarts to analyze the performances of the service parts management on a tactical level and provides insight into KPI's such as the order fulfillment rate per supplier. The information gained from the Entercoms analyses are then reported back to IBM and can be used by the managerial staff at SPO to improve IBM's. An illustration of the control tower 2.0 is shown in Figure 5.

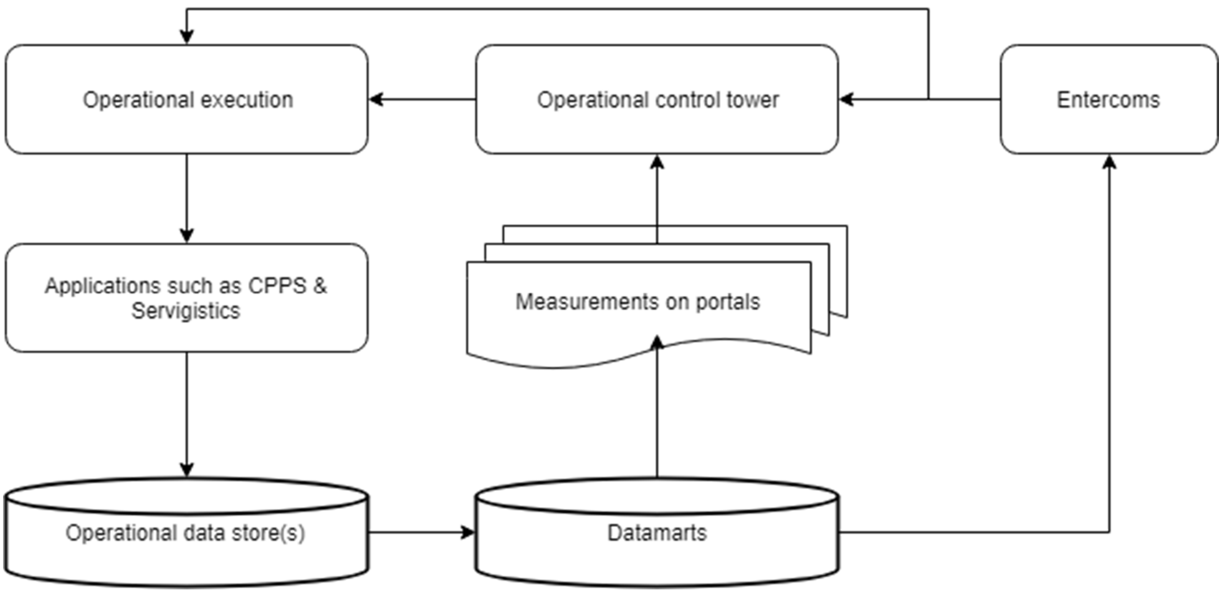

Figure 5: IBM's control tower 2.0

4.1.3 Customer machine failure situation

To provide a bit more background to the service parts process a short description of a failure of a machine at a customer location will be given. A customer machine failure situation usually starts with a call from the customer, the technician or the machine itself (if equipped with the required sensors) to the IBM call center in Sofia, Bulgaria. Upon receiving the call, a ticket will be created and a call center agent will start the diagnosis. After the diagnosis, a technician and a part (if needed) are scheduled to arrive at the customer site at the same time. Shipping of the part then done either directly to the customer site or to a location where the technician should pick it up such as a locker or warehouse. Upon arrival at the customer site, the technician will attempt to repair the machine. Parts that are defective will be returned to reutilization and parts that he did not use in the repair process are returned as good returns. The technician is incentivized to return unneeded parts quickly to the network with a financial compensation to prevent unnecessary ordering of parts by the planners or hoarding of parts by engineers. If the problem has been fixed, the call will be closed and the customer will be informed. If the engineer on site is unable to solve the problem, he will contact the call center again and the ticket will be transferred to a higher support level support, the engineering department or even procurement depending on the level of complexity and possible solutions. A flowchart of the customer failure situation is shown in Figure 6.

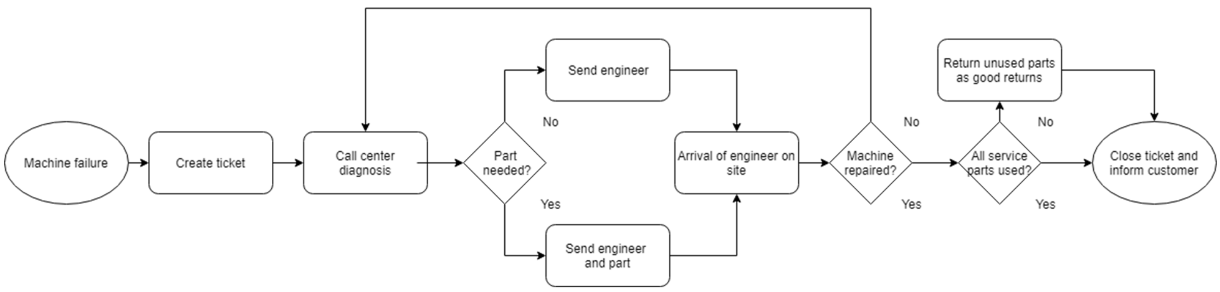

Figure 6: Customer machine failure flowchart

4.2 Systems

The planners responsible for the EMEA region at SPO mainly use CPPS and Servigistics Plan for operational planning and control. The largest differences between the two systems are seen in the interface and the relevance of the data. Servigistics has a more modern graphic interface, provides more information on the same screen and suggests planner actions but the service part information is limited to the inventory situation at the latest database update cycle. CPPS on the other hand, has a traditional terminal interface and requires queries to pull the information but has access to real-time stock levels at all locations. More details regarding these systems are given in the remainder of this chapter.

4.2.1 Servigistics Plan

Servigistics Plan, formerly known as XelusPlan, has been used by IBM since 2004 for operational service parts planning. Within this system planners have an overview of all the service parts under IBM's control. Servigistics has been designed to automatically handle most of the service parts management. It uses the input from IBM's databases to forecast customer demand, to compare demand to on hand stock balances and to recommend order quantities. When these orders, and their expected consequences, remain within the boundaries of the system's parameters, they will be automatically processed by Servigistics. If data from a service part falls outside of these boundaries however, the relevant part planner will be prompted for action via the appearance of a review reason in his work queue.

Operational planners tasked with exception handling within the Servigistics system will mostly be using the planner work queue, the planner worksheet and the order book modules for their activities. These three parts of the system will therefore be described in more detail below.

Planner work queue

Upon starting the Servigistics system a planner is shown his work queue. Within this queue, he is given a list of parts that need attention. He will be given information regarding the part name, the location, the review reason, the priority of the review reason and the time since the review reason has appeared for the part. This provides him with a quick indication of the part's status and helps him prioritize his daily work. Clicking on the service part within his queue redirects him to the planner work sheet.

Planner work sheet

The primary features of Servigistics used by the planners are located on the planner worksheet. The planner work sheet consists of five different windows, namely: the item data window, the schedules data workspace, the notepad, the review reason list, and the item family window. The worksheet in general tracks the results of the service part inventory plan over time and alerts the user when a planning rule has been crossed or is projected to be crossed. A more detailed description of the different windows can be found in Appendix A.

Order book

Users can open the order book in Servigistics to find a list of service part order information such as supplier, order quantity and order status. Next to listing this information the order book can be used to review, adjust or reconcile orders. Most of the orders required for the service parts management are placed by the system itself, however if human intervention is required in the ordering process, Servigistics will still attempt to aid the user by suggesting an order quantity and due date. This suggested quantity will be based upon the difference between the maximum inventory threshold and the expected stock level at the time of delivery.

4.2.2 CPPS

Before IBM implemented the Servigistics system, all service parts management for the EMEA region was done via CPPS. Other regions serviced by IBM use a similar system called PIMS but their capabilities are mostly the same. CPPS is a DB2 database system in which users can find real-time service parts information of the EMEA and China regions. Most of the information that can be found using CPSS can also be found through Servigistics as they interact with the same database, but the main advantage of using CPPS is access to the real-time information of the location inventory data. Users can only use CPPS for informational purposes however, since all the ordering processes must be done via the Servigistics interface. Accessing the information within CPPS is done by querying the database through a terminal interface or via programs such as QMF for Workstation.

4.3 Operational exception handling

After the description of the systems used by the operational planners we will describe the planning and exception handling processes in more detail. First, we will describe these two processes and thereafter we will describe the characteristics of the exceptions raised by Servigistics.

4.3.1 Planning

Most of the operational planning of the service parts has been outsourced to locations in Hungary and India. The planning of service parts starts in India, where the planners are responsible for the inventory levels at all the regional central buffer locations. Service parts planning for the other locations within the EMEA region, so called country planning, is done in Hungary. Finally, the employees at the Amsterdam location are responsible for supplier contact and approval of high impact decisions. The warehousing and transport for the service parts has been fully outsourced by IBM to third party logistic (3PL) providers.

The operational planning decisions made by the service parts planners at SPO are based upon the tactical levels set by the management. These tactical levels are based upon parameter such as: unit price, stocking cost, criticality, predicted usage and the service level required by the customer(s). Using these parameters, SPO has created four tactical stock levels, which are, from lowest to highest: the must order point (MOP), the policy safety stock (PSS), the maximum order up to point (MOR), and the excess stock level (EXC). While most organizations use safety stock levels as a buffer for unexpected increases in demand, SPO has two tactical levels situation below a service part's ideal level, the MOP and the PSS. The MOP should be seen as an indicator for immediate action, if a service parts inventory position drops below this line, operational planners should take immediate action to increase the inventory position. The PSS, usually higher than the MOP, is used as a guideline in the ordering process. Most orders within Servigistics are planned to arrive when the inventory position reached this level and, usually, the incoming order will increase the inventory position to the MOR level. The stock inventory position in a normal situation thus moves between the PSS and the MOR. Above the MOR, we have the EXC, above which service part stock is eligible for scrapping orders.

4.3.2 Exception handling

The exception handling process for a planner begins by opening his work queue in Servigistics. Here, he will be notified of the parts under his control that require attention. These parts will be linked to a review reason describing the exception. Clicking on the review reason within its work queue will take the planner to the item's work sheet. Now the planner's problem-solving process begins. Experienced planners will know what characteristics to check based upon the part's review reason, the basic training manual suggests planners to start their analysis by:

- Checking inventory
- Reviewing forecasts
- Reviewing open orders
- Checking actual demand
- Checking allocated quantities
- Checking for past-due orders
- Checking for large future orders

After their analysis, planners can either add, cancel or change orders, change forecasts, or contact other parties such as the procurement department. When the exception has been handled, planners are instructed to start working on the next item in their work queue thus starting the process anew.

4.3.3 Review reasons

Whenever the inventory position of a service part moves outside the predefined parameters or the characteristics of a service part change, Servigistics will alert the planner responsible for the service part by placing a review reason in his work queue. Planners will then attempt to resolve these review reasons as fast as possible, as ignoring review reasons might result in further escalations. Currently there are 40 triggerable review reasons within Servigistics and the complete list can be found in 0. Within this part of the chapter we will research what kind of actions can be taken by operational planners to handle these review reasons, how these review reasons are distributed among the service parts, and if there is a correlation between the set of review reasons linked to the same service part.

Looking at the actions required for resolving a review reason in the exception handling manual found in Appendix C, we state that these can be divided among five categories: planner review, informational, order, forecast, part settings and procurement. *Planner review* indicates a need for the planner to verify an action or change in the part characteristics. As an operational planner is not allowed to change all settings, e.g. part relationships, some of these cases will need to be referred to other users. *Informational* review reasons inform the planner of a part development and, often, do not require immediate action.

When the suggested action is *order* related, planners are advised to create, cancel, or adjust orders related to the service part. In case of a recommended *forecast* change, planners should critically analyze the current forecast method(s) used for the part and decide if a different forecasting method or adjusting the current forecast might be a better fit of the service part's movements. The *procurement* category indicates the need for the planner to contact the procurement department to solve the exception, e.g. there not being any supply sources for the service part while there is customer demand for the part. Most of the review reasons are concise enough to require a single type of action, but some of them have different recommended actions based on the situation, e.g. manually creating new orders or changing the forecast are both valid options to remedy a projected stock-out.

To give an overview of the spread of the review reasons over these five categories, we extracted a list of all the review reasons present at EMEA's central buffer location on 27-09-2017. This list contained 94,363 review reason entries spread over 72,146 service parts. From this analysis, we find that, while most recommended actions are in the order category, actions related to review and information appear the most in practice. For the review and information categories there exists a single review reason that is responsible for the lion's share of the appearances. Nearly 70% of the review actions are related to verification of machine configuration changes, whereas 77% of the informational review reasons are related to part with a weighted average cost (WAC) of 0. The other three categories have their appearances more evenly spread out between the different review reasons.

When we look at the suppliers related to the service parts up for review, we note that 47.958 of these review reasons belong to parts from a single supplier, Lenovo, a supplier with whom IBM will stop cooperating in the near future. Eliminating these parts from IBM's portfolio could therefore result in a potential reduction of roughly 51% in review reasons volume at this location. When removing the Lenovo parts from the analysis we find that, while there is a large reduction in volume, there appears to be no significant change in the distribution of the review reasons. A summary of this analysis can be found in Table 6. When looking at this table, one might note that the sum of the appearances is larger than the total number of review reasons, this is explained by the possibility of a single review reason having a set of actions spanning multiple categories.

Table 6: Review reason action categories

Action category	Review reasons in category	Appearances	Non-Lenovo appearances
Review	10	52,193	24,315
Information	5	37,374	19,183
Order	22	5,681	3,829
Forecast	8	3,091	2,377
Procurement	4	2,611	2,047

Next to analyzing the possible action categories of the review reasons, we are also interested in possible correlations between the appearance of review reasons. Since a single service part can have multiple review reasons attached to it, we analyze if there are certain sets of review reasons that appear often, as this could provide extra information for an operational planner. Using the same dataset of the action category analysis, we find that most of the parts, 72%, had a single review reason attached to them at the time of review. The highest number of review reasons attached to a single part in the dataset was eight with the average being 1.31. The number of review reasons attached to a service part is shown in Table 7.

Table 7: Frequency of review reasons per service part

Number of review reasons attached to part	Frequency
1	53,383
2	15,490
3	3,123
4	127
5	18
6	4
7	0
8	1

While most of the time a part will have a single review reason, in more than a quarter of the cases the operational planner is presented with multiple review reasons. While planners usually focus on solving the exception with the highest priority ranking, they might use the set of review reasons to support their decision. We have therefore analyzed the correlation between the appearance of these review reasons to see if some review reasons are likely to appear as a set. For this analysis, we used the Apriori algorithm within SPSS Modeler. The Apriori algorithm can discover association rules within the data in the form of: if (antecedent) then (consequent). An example of an Apriori associate rule could be: if a customer buys bread and ham in supermarket, there is a 70% chance he will also buy cheese. This probability that a consequent will appear if its antecedent(s) are present is called the rule confidence.

From our analysis, we find a relatively high rule confidence for some of the sets, but the confidence for the link between antecedent and consequent quickly diminishes. Looking at the ten sets with the highest confidence, we find that most of these sets consist of informational review reasons. The rule with the highest support level however, between R24 and R83, consists of two order related review reasons. A closer look at these two review reasons reveals that these are closely linked to each other as they are both related to a projected low inventory level with R24 indicating a projected stock out situation and R83 indicating a projected stock position below the must order point. Nevertheless, this analysis indicates that some sets of review reasons have a high probability of appearing together and that this information might be used in the operational exception handling process. A list of the ten rules with the highest confidence are shown in Table 8 and a complete list of review reason sets and their confidence values can be found in Appendix D.

Table 8: Review reason rules

Consequent	Antecedent	Confidence %
R24	R83	48.5
R67	R22 and R3	43.9
R22	R51 and R67	43.2
R67	R51 and R22	37.5
R22	R51	32.0
R67	R22	28.5
R3	R22 and R67	28.5
R67	R51	27.8
R3	R67	26.4
R9	R4	26.3

4.4 Chapter summary

Within this chapter, we have explained the workings of IBM's SPO department, we have outlined the current operational control tower environment, and we have described the systems and processes currently used in their operational planning activities. The goal of this chapter was twofold: (i) provide the reader with an overview of the current situation at SPO; (ii) answer the third research question regarding the characteristics of IBM's operational exception handling process. In summary, we can conclude the following:

- Operational planners at **SPO** use a combination of **Servigistics** and **CPPS** for their planning and exception handling, whereas the managerial staff leverages the **Entercoms** analyses to keep tabs on supplier's performance.
- By using Servigistics and CPPS, planners have access to **general part information** such as WAC and criticality, part **forecasts** such as demands and returns, part **substitutes**, **orders**, and a method of **communicating** between users via notes.
- Of all the orders made by SPO, **77%** made **automatically** by the Servigistics system. In case of a **manual order**, Servigistics will attempt to aid the planner by **recommending** an order quantity and due date.
- When a service part becomes problematic or is projected to become problematic, Servigistics will alert operational planners by creating **review reasons**. Most review reasons are simply **informational** or require the planner to **verify** a change in part characteristics. The remaining review reasons require either **order** management, changes in the **forecast** method or **contacting** the procurement department.
- Of the review reasons raised by Servigistics, we found that **72%** of the parts only have a **single review reason** linked to them. The other **28%** however, have between **two** and **eight** review reasons attached to them. We discovered that there are certain **sets of review reasons** that are likely to appear on the same service part, indicating that these sets could provide extra information to the operational planner.

5 Model

Within this part of the thesis we will explain the characteristics of our model. First, we will explain the type of decision support model we propose and our support for this model. Then, we will explain the model's scope and discuss the relevant data for its input and output. Based on the results, we will select the algorithm to be used. Within this chapter we therefore aim to answer the fourth research questions regarding the model characteristics and its performance.

5.1 Decision support classification

Given our findings in our literature research, our understanding of the operational service parts management environment at IBM, and IBM's motivation for the commissioning of this thesis, we believe that the proposed decision support system should be a mix between a knowledge-driven and a model-driven one. We come to this conclusion as follows: (i) the overall goal of the model is to use data from the present SPO knowledge base and combine this with the machine-learning and cognitive tools within IBM to automate the exception handling process, and these solution components are consistent with those of a knowledge-driven DSS; (ii) the decisions from the experts will be quantitatively scored and used to improve the model, which is consistent with the optimization of a model-driven DSS.

After defining the type of DSS, we must define what the most fitting model is for the exception handling at SPO. Given our theoretical research into the requirements of cognitive computing, we find that a cognitive computing solution is not a good fit for this task and our support for this statement is as follows: (i) there is little unstructured data present in the current operational service parts environment; (ii) the model does not require interaction with the end-user, and has therefore does not need to understand or respond in natural language; and (iii) given IBM's preference for quantitative feedback, there will be no room for active learning from users. Instead of using a cognitive computing model, we therefore propose to use a classic machine-learning model.

5.2 Model characteristics

Within this part of the chapter we will define the model's scope, data to be used as input and outputs of the model, and select a suitable machine-learning algorithm. The input and output data will be based upon the information available to the operational service parts planners. An initial dataset will be used to test the performance of the possible models.

5.2.1 Scope

As we have discovered from the interviews and documentation regarding exception handling at SPO, most review reasons have very specific triggers and decision steps. As this thesis is meant to provide an insight in the possibilities of automating the exception handling process, we will focus our model on a single exception type. Selecting this review reason was done with the following steps: (i) limit the possibilities to the review reasons tagged as highly important by the operational planning management, (ii) limit the possibilities by removing review reasons that cannot be linked to a planner's actions as they are merely informational or explicitly require human review, (iii) choose a set of review reasons from the remaining possibilities that have a similar set of actions, enabling easier generalization of the model, (iv) choose a review reason that can prevent an event from occurring as these may reduce the number of fire-fighting review reasons appearing as well, (v) finally from this set, choose a review reason that appears relatively often, thus insuring enough data, and has a relatively high priority level. Given these criteria, we have decided to use the review reason number 38, hereafter to be called R38, for the initial testing of the model. An added benefit of selecting the R38 for testing the model, is the similarity between R38 and three other review reasons: R5, R6, R7, R18 and R39. These six review reasons have large overlap between their set of actions and, we presume, also in their relevant inputs. We believe that, a working solution for handling a R38, will therefore be easily transferable to these other review reasons.

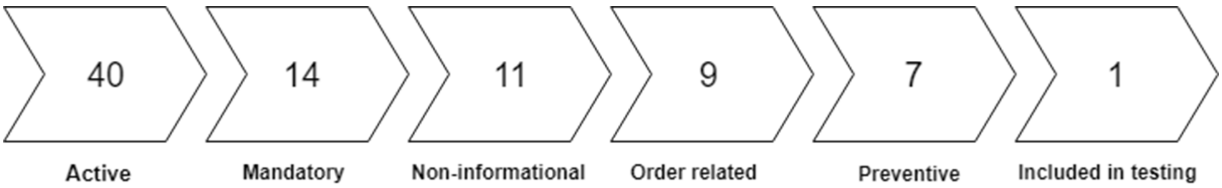

Figure 7: Review reason filtering process

Within Servigistics, an R38 is triggered when the existing orders would result in the service part's inventory level dropping below the PSS in a four-day period after the current date plus the lead-time, e.g. if we have a service part with a lead-time of five days, the analysis period for the R38 trigger will be between five and nine days into the future. This example is illustrated in Figure 8 with the period of analysis is visualized by a rectangle. Servigistics takes into account the different service parts lead-times for a normal, warranty or repair order during its analysis and shows the relevant lead-time to the planner in the description of the R38 in his planner worksheet. It will thus preventively signal the operational planner that, given the existing set of orders and the forecasted usage, the inventory levels will drop below the desired levels.

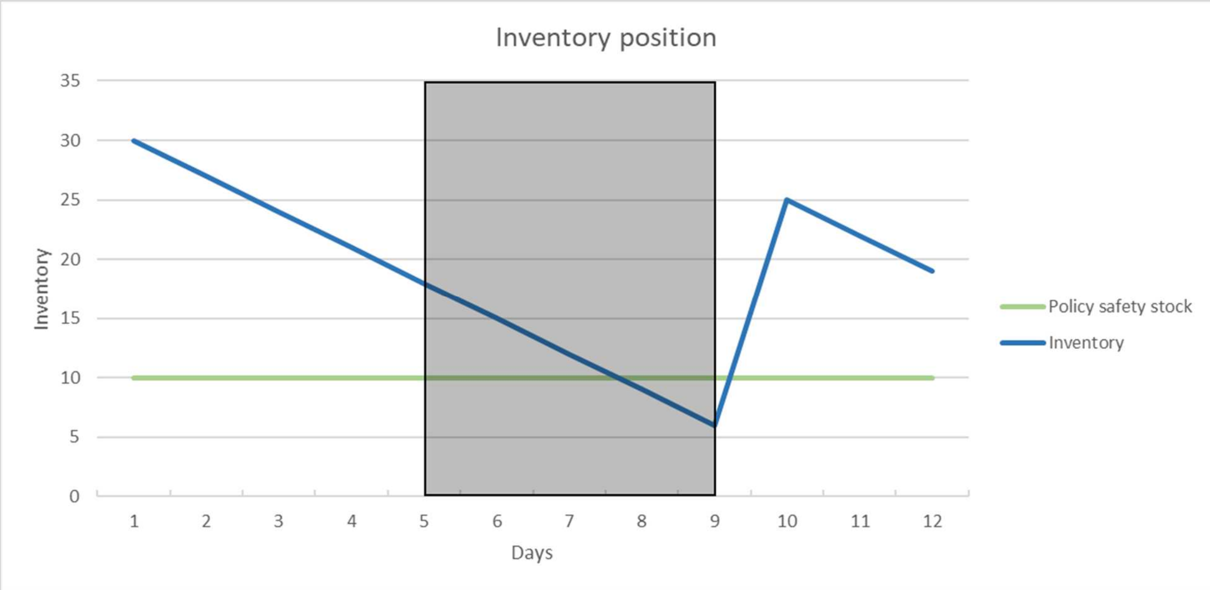

Figure 8: R38 evaluation period

When encountering a R38, the planner manual lists four steps to solve the problem: (i) review the ordering location, supplier and order type associated with the order increase message; (ii) resolve other review reasons blocking automatic order approval; (iii) if the order increase is concerned with repair or warranty, decide to wait for the returns or place an order on an alternate supply source; or (iv) change the ordering process flag. Since not all actions are logged in the database, e.g. planners are not reporting their findings after review, and are thus unavailable for our model, we interviewed the SPO employees to find the frequency of these actions. These interviews revealed that: (i) a planner's review rarely leads to solving the problem and should be seen as the start of the planner's analysis; (ii) changing the ordering process flag is rarely done and the option to do so is not available to every planner; and (iii) in the large majority of these cases, a R38 requires an order related action. We have therefore decided to include the possible ordering options for this location as feasible model outputs and these are described in 5.2.3. More detail on review reasons themselves can be found in 4.3.3, whereas the complete list of review reason triggers and advised solutions can be found in Appendix C.

5.2.2 Input

For the inputs, we include the basic service part information used by the operational planners at SPO. This information has been extracted from SPO's CPSS database. The service part information selected for the model's input included the following:

- Service part's ordering information regarding new part, warranty, and repair orders (service part's

WAC at the EMEA central buffer, unit price, contractual order lead-time, minimum order quantity, its location planning hierarchy, last time buy information, and division owner code)

- Service part's characteristics (vitality, category, substitute type, shelf life, end-of-service date, first stocking date, first appearance date in the SPO network, and the analyzer responsible for the part)
- Inventory information (inventory level at the EMEA central buffer, demand within the previous four-week period, and the number of parts available for repair and warranty at the repair and warranty locations)
- Tactical inventory planning information (excess level, maximum order up to point, policy safety stock, and the must order point)

The complete list of model inputs and their description can be found in Appendix E.

5.2.3 Output

The model should predict the action used by a planner to resolve the R38 given the situation. Since planners are only able to influence the service parts network through Servigistics, we use Servigistic's log files to extract the planner's actions. In this log file all changes made to the service parts network through Servigistics can be found as well as automated actions and responses. To link the actions made by the planner in response to the R38, we did the following: (i) scan the log file for instances in which the R38 disappeared from the planner's work queue; (ii) use the timestamp, planner code, and service part number to find the action taken by the planner; (iii) extract the timestamp, planner's code, service part number, and action from this log file. This information is then joined with the contextual information extracted from the CPPS database to create an input-action pair. As mentioned in 5.2.1, these actions are mostly order related and planners have the following options at their disposal:

- New buy order. These are purchase orders for service parts owned by IBM, placed on the contractual supplier.
- New business opportunity (NBO) order. These are new buy orders for service parts from external organizations for which IBM is responsible for the inventory management.
- Repair order. These are orders for service parts present at the repair location to be repaired and send to the regular stock locations.
- Warranty order. These are orders for service parts that are still in their contractual warranty time

to be send replaced and send to the regular stock locations.

- Internal demand order. These are hub-to-hub rebalancing orders where parts are ordered from another region's central buffer location
- Excess order. These are orders directly from the service part's manufacturer (excess) inventory.
- No order. This is done by the planner if he, for example, has knowledge of changes to future ordering and/or demand.

While this is the complete list of possible ordering types linked to this review reason, not all of these ordering types are available for all the service parts in the database. An operational planner can, for example, never create a repair order for a service part without a repair source.

5.2.4 Machine-learning algorithm

IBM has expressed preference for the usage of SPSS Modeler for the creation of the model as this software is owned by IBM. We will therefore only consider the machine-learning models available in version 18.1 of SPSS Modeler, the most up-to-date version at time of this research. Within SPSS modeler, there are three different categories of machine-learning models: classification, association, and segmentation. *Classification* relates to models that use one or more input values to predict the value of one or more output fields. *Association* models are used to find patterns in the data where one or more entities are associated with one or more other entities. *Segmentation* models are merely able to divide the data into segments, or records that have similar patterns of input fields, they are unable to process output fields. Since our model is intended to predict an action based on the problem context, it will use one of the classification models.

Within the classification category, we have identified seven candidates which fit our data. These seven models are: C5.0, Quick Unbiased Efficient Statistical Tree (QUEST), Bayesian networks, neural networks, Chi-squared Automatic Interaction Detection (CHAID), Classification and Regression (C&R), and random trees. A small summary of these models can be found in Appendix F.

We will analyze and compare the performance of these models on the data collected from May to September. Within this dataset we analyze the actions that lead to the removal of a R38 from the planner's work queue and eliminate those without stock inventory information, resulting in 597 records for the analysis. Creating machine-learning models from this dataset will require it to be partitioned in a training and a testing set. The training set is then used to generate the model's rules. These rules are then applied to the testing set, in which the model will provide output based upon these rules. The model's outputs are then compared to the real outputs present in the testing set. Performance of these models will be calculated based upon the ratio of correctly predicted outputs by the model compared to those present in the testing set, also called the accuracy of the model. If the model would, for example, predict a warranty order for two cases whereas in reality only one of these cases resulted in a warranty order, the accuracy of the model will be 50% (Dietterich, 1997; Hoffman & Bhattacharya, 2016; Jordan & Mitchell, 2015; Proper & Tadepalli, 2004).

To validate the accuracy of these models we will use k-fold cross-validation. When using this method, one partitions the entire dataset in *k* subsets which are approximately equal in size. Each of these subsets is used as the testing partition once, while the other subsets are used as the training partition. This method thus produces *k* models that are trained on a different configuration of training sets and tested on the remaining partition (Varian, 2014). We will then use the average accuracy of these *k* models for our model selection.

5.2.5 Results

Here we will discuss the results of the model type selection, the predictor importance of the model, the model's decision confidence values, and will analyze the model's performance in more detail. These results will be used as an indication of the usefulness of implementing a machine-learning model to the operational exception handling process. In the analyses of the results we define a prediction to be correct if the model generates the same output as the action taken by the planner for a specific case.

Model selection

Selecting the model to be used in the remainder of this research will be done via comparing the average accuracy of the seven candidate models with k-fold cross-validation. For this comparison we have decided to split the dataset into ten different subsets, as is common in the field of machine-learning (Varian, 2014). To divide the input-action pairs in these subsets, we have randomly assigned them a value between one and ten, resulting in ten different subsets with an average size of 60. The models in our analysis all used these same ten subsets, to provide a fair comparison, and every model used its default settings since we lack the time to do an exhaustive comparison.

The results from the analysis show that four models (C5.0, QUEST, CHAID, and C&R) had a similar average accuracy. While the performances of these models indicate that any of these four could be selected for further analysis, given the resources available to this research we will only use the C5.0 for the remainder of this research. This model has been selected because: (i) unlike the QUEST model, it is able to handle weighted input fields, an option that might be useful in future iterations of the model; (ii) the C5.0 models within SPSS provides us with the most decision tree and predictor importance values, enabling deeper model understanding. The performances of the models are summarized in Table 9, and a visualization of the SPSS model is shown in Appendix G.

Table 9: Model performances

Model	Min accuracy %	Max accuracy %	Average accuracy %
C5.0	49.3	68.0	56.9
QUEST	40.7	64.0	54.4
Bayesian networks	0.0	0.0	0.0
Neural networks	0.0	3.4	0.8
CHAID	45.8	72.0	55.6
C&R	42.4	59.7	55.4
Random trees	7.4	30.4	18.6

Model rule sets

The C5.0 models within SPSS produce both a rule set and a decision tree based upon the training data to classify the data. These rule sets and decision trees provide insight into the workings of the model and how the different input values are used to categorize the cases. An example of such a rule would be dividing the parts based upon their WAC value. In the first C5.0 model, this input value is used as a rule by dividing the cases into those related to service parts with a WAC value higher than €81.49 and those with a value equal to or lower than €81.49. For service parts in the second group, the first model predicts the planner to make a new buy order, a decision which is correct in 73% of the cases. The first group is then split multiple times into smaller groups by inputs such as analyzer code or the historic usage. The complete rule set of the first model is included in Appendix H.

Predictor importance

Within SPSS predictor importance values indicate the relative weight of each input of the cases in the training set when creating the model's ruleset. The sum of all predictors will be equal to one (1) and due to the differences in training and testing sets of the ten different C5.0 models, each model has the potential to have a different predictor importance distribution. SPSS modeler was able to generate these predictor importance values for nine of the ten C5.0 models. We will use the average values of these for this analysis. We find that, on average, the analyzer code linked to the service part plays a rather large role with an average predictor importance value of 0.47. Indicating that certain planners might make a lot of the same decisions, or that these planners are responsible for parts that often require a similar intervention. The other two predictors with an average importance of more than 0.05, are the actual on-hand balance of the service part at the central buffer location and the service part's WAC value, with average values of 0.13 and 0.07 respectively. Indicating these two inputs are relatively important for the creation of the model and should be included as inputs for future iterations of the model. A complete list of the predictor importance values per model can be found in Appendix I.

Decision confidence level

The next step in the analysis of the model's results is the analysis of the confidence level of its decisions made in the testing set. Within the C5.0 model the decision confidence is defined as the ratio between the number of correctly classified cases with similar input values in the training set and the total number of cases with the same input values in the training set. When the model has a set of different outcomes linked to this category of cases, it chooses the option with the highest decision confidence level as its prediction. If there are, for example, ten parts with similar input values in the training set and for seven of these cases a warranty order has been placed; the model will predict a warranty order with a decision confidence of 70% (7/10) when presented a similar case from the testing set. When implementing machine-learning models, practitioners usually limit the autonomous actions of such a model to actions above a certain confidence threshold. If the model predicts an outcome with a low confidence level, it may request user feedback before making a decision.

Within our testing set we find an average confidence level of 72% of the ten models. The highest average confidence level for one of the models was 74% and the lowest average confidence level for one of the models was 68%. While these are the model averages there are two aspects of the models we will analyze in more detail: the distribution of the confidence values and the difference in confidence values between correct and incorrect decisions. Analyzing the distribution of the average confidence level values we find that most of the decisions made in the testing set, have a confidence level between 70% and 80%. While these confidence levels are relatively high, this may be a result of there not being enough distinguishable inputs between cases or the small number of cases in the testing set. Looking at the differences in confidence values between correct and incorrect decisions we find an average mean confidence level of 74% for the correct decisions and an average mean confidence of 70% for the incorrect decisions, indicating that there are a significant number of incorrect decisions with a high confidence level. It might therefore be hard to properly choose a confidence level for automation for the exception handling process. A summary of the average confidence levels and their (cumulative) frequency are given in Figure 9, and a frequency table of all confidence values is given in Appendix J.

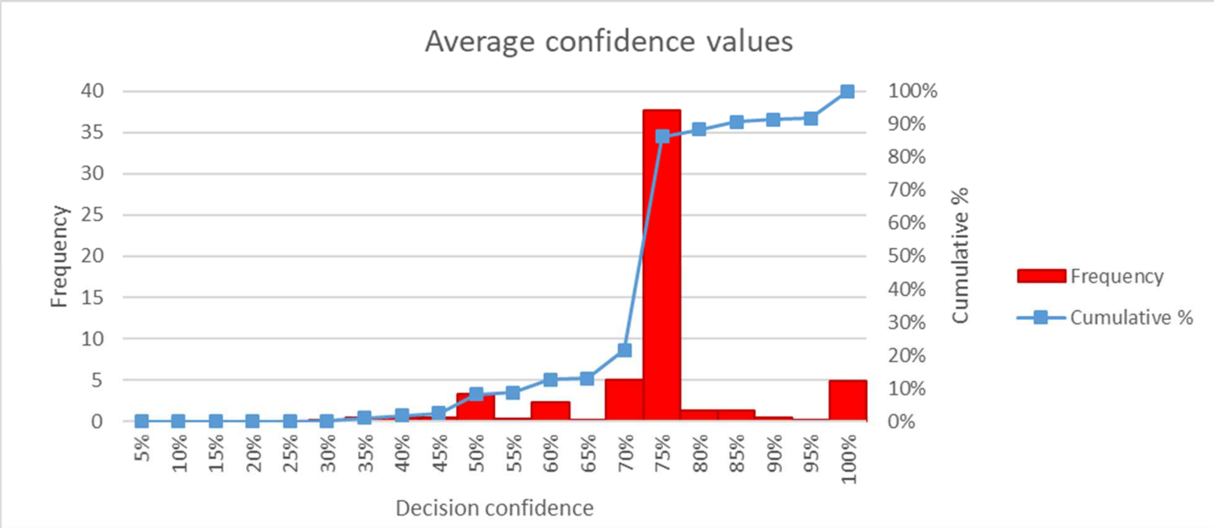

Figure 9: C5.0 average confidence values

Detailed analysis

While looking at the overall performance and confidence values of the decisions gives an indication of the overall usefulness of the model, these values might differ across the different service part characteristics. The model might, for example, be able to very accurately generate decisions for parts with a low criticality, but not for those with a high criticality. If this analysis reveals that there are certain categories of service parts for which the decision confidence values are relatively high, IBM might consider adopting this model for those categories and tweak the model for other categories.

For this analysis, we compare the accuracy and average decision level confidence for the parts over the following categories: the part division's owner, the order type lead-times, the part vitality, the part WAC, and the different outcomes (order types) suggested by the model. These categories have been chosen since these are the main characteristics of the parts defined by the employees at IBM. In addition to the part characteristics we have included the possible outcomes of the model in this analysis to see if the model is able to predict one of these outcomes with more certainty than the others. We will provide a summary of the results here and the complete results per category can be found in Appendix K. Some categories will have less cases as not all service parts can be included in every category, e.g. a service part without a contractual repair option will not appear in the repair order lead-times table.

From the results we find the model being relatively bad at predicting the correct order type for service parts with a contractual repair (29% accurate) or warranty (33% accurate) lead-time. Actions regarding service parts within these categories might therefore not have a very straight-forward solution or additional information is required by the model to handle these cases. The most accurately predicted cases by the model concerned services parts that are supplied by Lenovo, have a very short new buy order lead-time, have a low criticality, and have a WAC value between €1 and €100. Service parts with all off these characteristics accounted for 32 cases in the dataset and for 28 of these cases the model generated the same output as the planner's actual decision. For every case however, the model proposed a new buy order whereas the planners also proposed repair orders and internal demand orders. While not able to predict all the outcomes for these exceptions with 100% accuracy, service parts with these characteristics would be good candidates for automated exception handling.

5.3 Position in IBM's control tower

When IBM moves forward with a machine-learning approach in operational exception handling, it will need to be integrated in their control tower environment. We would argue that the current version of the model cannot be completely positioned in either the information service platform or the informational manpower layers of a control, but should be positioned as a connection between them. The current iteration of the model will require interaction between the operational execution layer as it is not yet able to autonomously handle parts of the exception handling process. When it is able to autonomously handle exceptions, it can be situation on the same layer as the decision makers.

When we translate this theoretical positioning to the control tower environment of IBM described in 4.1.2, we propose the model to be situation parallel to the CPPS and Servigistics applications. The biggest difference compared to the other two systems in this layer, would be the model having a two-way interaction with the operational execution, as it uses operational execution decisions as input and generates operational execution output. An illustration of the current model's placement within the IBM control tower environment is shown in Figure 10.

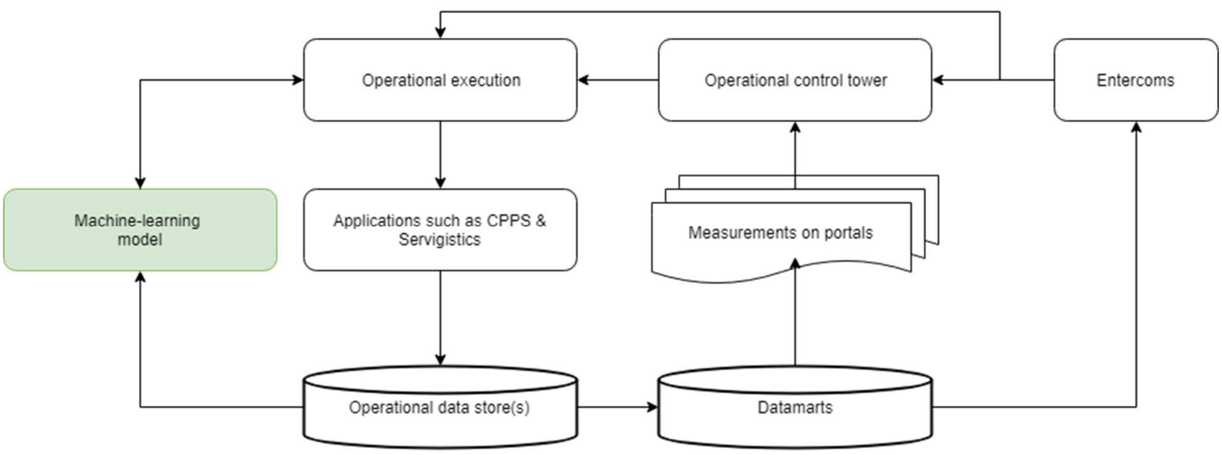

Figure 10: Model placement in IBM's control tower

5.4 Chapter summary

Within this chapter, the proposed model and its data requirements have been discussed. Furthermore, the prototype's scope, input, output and algorithm and their reasoning have been explained. Within this chapter we have thus attempted to answer the fourth research questions of this master thesis regarding the characteristics of the model and its performance in predicting the operational exception handling decisions. In summary the results from this chapter are:

- The model should be a mix between a **knowledge-driven** and a **model- driven** DSS. Instead of a cognitive computing model, a standard **machine-learning model** is more fitting to the SPO exception handling environment.
- A single review reason, **R38,** has been chosen for the initial test run of the prototype. This exception is related to predicted inventory shortages at least an order's lead-time into the future.
- Feasible outputs of the model are all **order related** with the different order types being: **new buy** (either for IBM owned stock or externally managed stock)**, repair, warranty, internal demand**, or **excess.**
- We have decided to use the **C5.0** algorithm for the model. This algorithm was able to, on average, **correctly classify 56.9**% of the testing set.
- Service supplied by **Lenovo** with a very **short new buy order lead-time**, a **low criticality**, and a **WAC** value between **€1** and **€100**, are best suited for automated exception handling with the model correctly predicting the outcome in **88%** of the cases.

6 Performance measurement

In this chapter we will propose a method for the quantitative performance measurement of a planner's action made in response to an R38 exception. Identifying actions which lead to the desired result will help IBM to provide feedback to the operational planners and to identify cases which can be valuable to include in future iterations of a machine-learning model. To create this quantitative performance measurement method, we will first define the ideal situation to be used as a reference, then define the timing of the evaluation, and finally propose a scoring formula.

6.1 Ideal situation

In an ideal world, the operational planners know when a part is needed, the quantity that is required, and all the lead times are deterministic. These characteristics would enable them to match the order arrivals with the exact moment of requirement, thus leading to minimal inventories. Sadly, this is not the case in the probabilistic world of service parts management. Organizations in this market will therefore need to have a certain number of parts on stock to be able to react to events. Setting this level too low might result in not being able to deliver parts when needed, whereas setting this level too high might result in the organization tying up capital in inventory that might not be used during the part's lifetime. Operational planners are therefore responsible for keeping the inventory levels between these levels.

Within IBM, service part inventory is measured against the following tactical levels, from lowest to highest: the must order point (MOP), the policy safety stock (PSS), the maximum order up to point (MOR), and the excess level (EXC). A more detailed explanation of these levels can be found in 4.3.1. An operational planner within IBM will be instructed to keep the stock between the PSS and the MOR. In general, the system will automatically place orders such that when the inventory drops to the PSS, a replenishment arrives which increases the stock level to the MOR. The ideal situation of such an ordering policy is stylized in Figure 11.

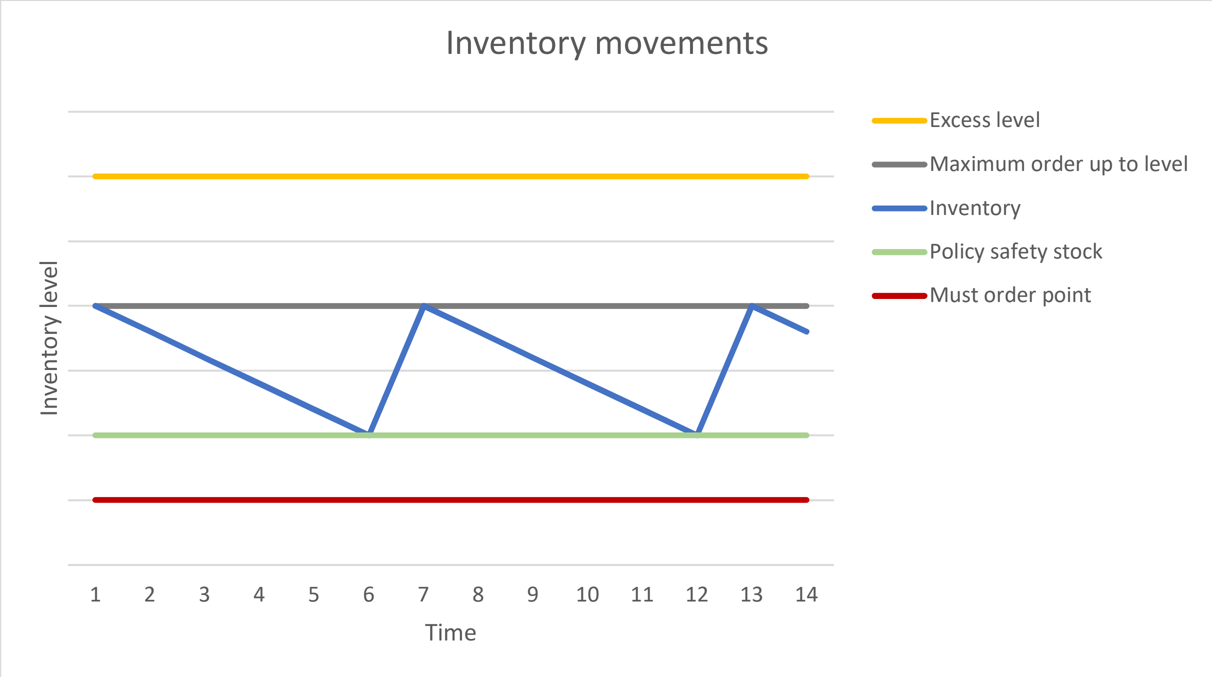

Figure 11: Inventory ideal situation

6.2 Evaluation timing

Since most of the decisions made by the planners will not have an immediate impact on the physical inventory levels, we must decide the period for review where the planner's impact can be measured. Since the planner is prompted for action via the appearance of a review reason and a review reason's trigger is linked to a specific analysis period, we propose that the impact evaluation should also be linked to the review reason's analysis period. To measure the impact of an action regarding the R38, we therefore choose an evaluation period starting a lead-time into the future from the appearance date of the review reason, and stop our evaluation period after four additional days have passed, thus looking at a five-day period. During this lead-time, other events may occur that influence the part's health but these will be linked to different review reasons and should thus not be taken into review.

6.3 Scoring

Now that we have defined the ideal reference situation and the evaluation period of the performance measurement, we will propose a quantitative performance scoring formula at the end of this chapter. Before we do so, we will define the characteristics of the scoring formula itself. From interaction with the employees and theory we conclude that, in most cases, it is not an immediate disaster if the stock position moves below the PSS, as this is a tactically chosen level to absorb risk. It becomes increasingly problematic however, if the stock position moves further away from these tactical levels. The scoring method should be able to capture this increase in risk and punish the stock position increasingly harsh the further it moves from these tactical levels. Additionally, we found that, SPO would like the outcomes of the quantitative scoring method to be comparable between the different service parts, and that there is a desire for this quantitative measurement to capture the differences in services parts, e.g. the service part's vitality.

If we only use an increasing scoring method based upon the distance from the ideal band, the scoring method would only consider the inner most tactical levels, the PSS and the MOR as there is no extra punishment for crossing the outer most tactical levels, the MOP and the EXC. Thus, we propose that the scoring method should use an extra variable to indicate the area of the inventory position at that point in time. These areas of a service parts inventory are: below the MOP (A), between the PSS and the MOP (B), between the MOR and the PSS (C), between the MOR and the EXC (D), or above the EXC (E), as shown in Figure 12.

Figure 12: Inventory stock positions

Furthermore, these area variables allow the quantitative scoring method to distinguish between moving towards a stock out scenario, i.e. below the ideal band, and moving towards excess stock, i.e. above the ideal band. These directions may be weighed differently for different service parts based upon characteristics such as vitality or unit cost. High vitality parts may for example be punished harder if they move towards a stock out situation as these are critical to the customer's operations while high unit cost parts may be punished harder when moving towards excess stock as this can lead to large scrapping costs. In general, interviews with the employees indicate that, while moving away from the ideal zone would always be undesirable, SPO would rather have excess stock than low stock inventory levels, as they prefer a potential reduction in revenue to a potential reduction in customer satisfaction levels.

Finally, the scoring method should be useable for any service part and the results should be comparable across service parts. This means that an expensive, slow moving part with a low PSS would preferably use the same scoring method as an inexpensive, faster moving part with a high PSS. The scoring method will thus need to be scaled for the performances of the interventions to be comparable. Stock moving below the PSS can easily be scaled as the ratio between the current stock level to the quantity between the PSS and a stock-out situation. Scaling the scoring above the MOR requires more thought, as an infinite amount of inventory can theoretically be bought. We have decided to set the maximum for upward deviation as the EXC plus the width of the ideal band. We reason that since having large amounts of excess stock is a rare occurrence, the probability of having excess stock plus a lead-time worth of stock would be extremely rare. Since the width of the ideal band can be defined as the difference between the MOR and the PSS, we can simplify this as follows: $EXC + (MOR - PSS) - MOR = EXC - PSS$

In summary, we find the following requirements: (i) scoring should take increase based on the deviation from the ideal zone; (ii) the tactical levels should be considered; (iii) the scoring method should consider which side of the ideal zone the inventory position is; and (iv) the scoring method should use a ratio to compare scores between different service parts. We propose three different scoring methods to be discussed in further detail: a linear formula, a quadratic formula, and an exponential formula. To make comparison between these methods easier, all three will start their penalty score at a value of one (1) if the inventory positions is outside the ideal zone. The scoring methods for these three are given in Equation 1, Equation 2, and Equation 3 respectively.

$$P_t = \begin{cases} W_a * (1 + \left(\frac{PSS - Q_t}{PSS}\right)), when\ 0 < Q_t < MOP \\ W_b * \ (1 + \left(\frac{PSS - Q_t}{PSS}\right))\ , when\ MOP \leq Q_t < PSS \\ 0\ , when\ PSS \leq\ Q_t \leq\ \ MOR \\ W_d * (1 + \left(\frac{Q_t - MOR}{EXC - PSS}\right)), when\ MOR < Q_t \leq EXC \\ W_e * (1 + \left(\frac{Q_t - MOR}{EXC - PSS}\right)), when\ EXC < Q_t \end{cases}$$

Equation 1: Linear penalty score

$$P_t = \begin{cases} W_a * (1 + \left(\frac{PSS - Q_t}{PSS}\right))^2, when\ 0 < Q_t < MOP \\ W_b *\ (1 + \left(\frac{PSS - Q_t}{PSS}\right))^2\ , when\ MOP \leq Q_t < PSS \\ 0\ , when\ PSS \leq\ Q_t \leq\ \ MOR \\ W_d * (1 + \left(\frac{Q_t - MOR}{EXC - PSS}\right))^2, when\ MOR < Q_t \leq EXC \\ W_e * (1 + \left(\frac{Q_t - MOR}{EXC - PSS}\right))^2, when\ EXC < Q_t \end{cases}$$

Equation 2:Quadratic penalty score

$$P_t = \begin{cases} W_a * e^{\left(\frac{PSS - Q_t}{PSS}\right)}, when\ 0 < Q_t < MOP \\ W_b *\ e^{\left(\frac{PSS - Q_t}{PSS}\right)}\ , when\ MOP \leq Q_t < PSS \\ 0\ , when\ PSS \leq\ Q_t \leq\ \ MOR \\ W_d * e^{\left(\frac{Q_t - MOR}{EXC - PSS}\right)}, when\ MOR < Q_t \leq EXC \\ W_e *\ e^{\left(\frac{Q_t - MOR}{EXC - PSS}\right)}, when\ EXC < Q_t \end{cases}$$

Equation 3: Exponential penalty score

The variables within these equations are defined as follows: *P* = penalty score at time t, W_i is the weight for area i, and Q_t is the inventory quantity at time t. As an illustration to the proposed penalty formula we have created the graph shown in Figure 13. In this graph, the penalty scores given to an inventory position ranging from 0 to 130 are given. For this graph, we have used the following tactical levels: must order point of 30, policy safety stock of 50, maximum order up to level of 80, and an excess stock level of 100. The weights were give as follows: W_a of 2, W_b of 1, W_c of 1, and a W_d of 1.5, to indicate the preference for excess stock over the possibility of a stock out situation. As one can see from this graph, values further from the center band are punished more harshly, a clear distinction is seen when moving from one tactical level to another, and inventory positions within the ideal zone are not punished at all.

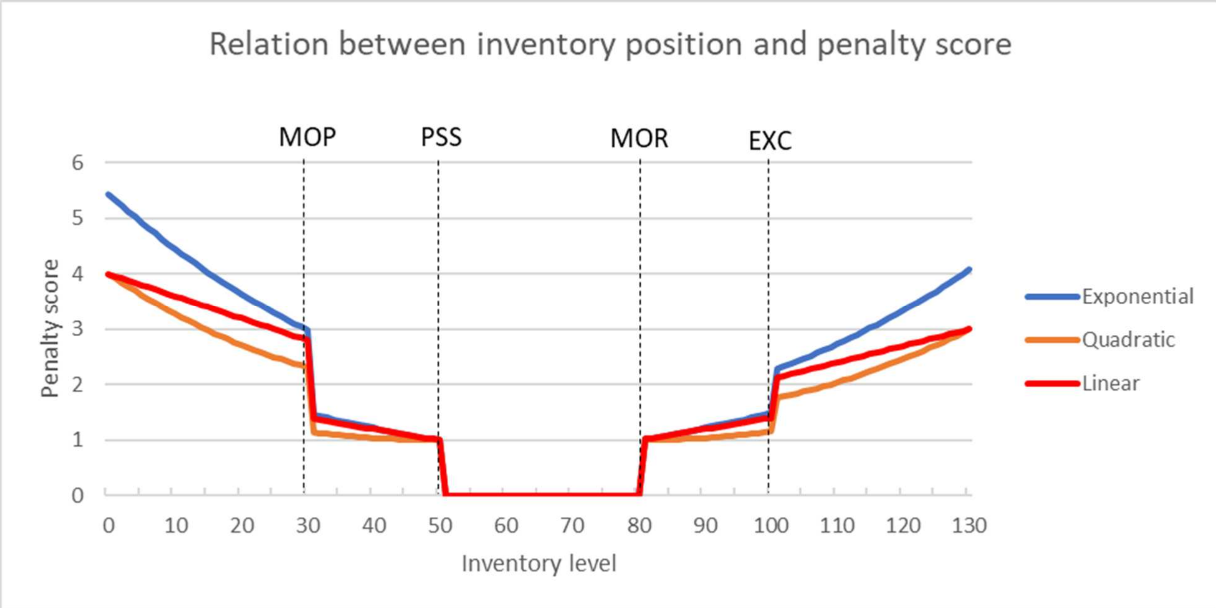

Figure 13: Penalty score example

Now let us give an example of this scoring method from the dataset we have. Let us assume an operational service parts planner opens his work queue on 27-07-2017 and sees that a R38 triggered at that same date. He looks at the planner worksheet related to the service part and finds the following values: there are currently 268 parts in stock, the lead-time for a new buy order is 14 days, and the service part inventory level is indeed expected to be below the PSS of 375 in 14 days. He looks at the existing orders and notices that there is an existing new buy order in the order book with a due date of 09-08-2017 and decides to send an expedite request for an additional 13 parts to the supplier.

To measure the quality of his decision, we apply the penalty score discussed earlier to the review reason's trigger period. In this example, given the lead-time of 14 days and trigger date of 27-07-2017, we record the penalty scores starting from 10-08-2017 to 14-08-2017. For this example, we have chosen the area weights as 2, 1, 1, and 1.5 for inventory levels in area A, B, D, and E respectively. When we track the inventory position, we can see that the part starts in an unhealthy situation as it is even below the MOP of 300. Over time we note that the inventory level of the service part increases to a healthier level and during the trigger period of the R38 is situation between the PSS and the MOR. Looking at the review reason trigger period, we see that the planner's decision resulted in a healthy part and the decision was thus a correct one. The inventory position of the part and the penalty score per day is given in Table 10. When looking at this table we stress that we have included the penalty score for every day up to the triggering period for illustrative purposes, but only the penalty score of the days in the triggering period themselves will be included in the performance evaluation of the model.

Table 10: Example scoring

Date	Stock quantity	Exponential penalty score	Quadratic penalty score	Linear penalty score
27-07-17	268	2.660	2.160	2.572
28-07-17	262	2.703	2.182	2.602
29-07-17	262	2.703	2.182	2.602
30-07-17	254	2.762	2.208	2.645
31-07-17	245	2.829	2.240	2.693
01-08-17	242	2.851	2.252	2.709
02-08-17	369	1.016	1.000	1.016
03-08-17	363	1.033	1.001	1.032
04-08-17	357	1.049	1.002	1.048
05-08-17	357	1.049	1.002	1.048
06-08-17	356	1.052	1.003	1.051
07-08-17	361	1.038	1.001	1.037
08-08-17	361	1.038	1.001	1.037
09-08-17	457	0.000	0.000	0.000
10-08-17	**452**	**0.000**	**0.000**	**0.000**
11-08-17	**445**	**0.000**	**0.000**	**0.000**
12-08-17	**445**	**0.000**	**0.000**	**0.000**
13-08-17	**453**	**0.000**	**0.000**	**0.000**
14-08-17	**453**	**0.000**	**0.000**	**0.000**

6.4 Performance measurement results

We ran the performance measurement algorithm mentioned in 6.3 on the set of 597 input-action pairs we have regarding handling the R38 exception. From this set, 161 actions were gradable, as the period of evaluation, the lead-time of the relevant order type, fell within the period of our database extracts (May to September 2017). For the evaluation of these input-action pairs, we have decided to use a weight of one (1) for all four areas as a starting point. When analyzing the data, we found that not all inventory information of every service part is available for the entire evaluation period. Some daily inventory extracts made for this research became corrupted, resulting in information losses. Due to the structure of IBM's database, these values are not recoverable as these databases do not store historic inventory positions. For 11 of the gradable pairs, we have access to the complete five-day period of evaluation, for 109 of the gradable pairs we have information for at least three of the five evaluation days. We will however, still use the entire data-set to provide an example of the distribution of the performance measurement results. Inclusion of pairs for the next iteration of the model can be limited to those pairs with at least a certain number of data points.

The performance measurement indicated that in 74 out of the 161 cases (46%), the inventory position of the part was within the ideal zone for at least one of the evaluation days whereas in 66 (41%) cases the inventory position was within the ideal zone for the entire period. These results show that despite the best efforts of the operational planners, the inventory position of the service part is within the ideal zone in less than half of the cases. Part of this result can be explained by the dynamic and unpredictable nature of the service parts management environment but further analysis of the incorrect decisions could reveal areas of improvement for IBM’s exception handling process. Additionally, a planner’s performance could be measured by these scores and they can be given feedback on these less than optimal cases to improve their future decisions. The descriptive statistics of the different scoring methods are shown in Table 11. The scores of all 161 cases can be found in Appendix L.1, box plots visualizing the distributions, without the inclusion of the extreme outliers for readability, are shown in Appendix L.2.

Table 11: Average score descriptive statistics

	Exponential	Quadratic	Linear
Mean score	33.21	4.48	1.26
Minimum	0.00	0.00	0.00
Q1	0.00	0.00	0.00
Q2	1.19	1.39	1.17
Q3	2.26	3.52	1.82
Maximum	4761.67	148.28	9.47

6.5 Case inclusion

By using this performance measurement method, we can select the cases to be included in the next iteration of the model, since including more cases with desired results will improve the value of the model’s predictions. However, when selecting which cases to include, we must make a trade-off between the quantity and the quality of the cases to include. Ideally, we would like to model to only make perfect decisions, and as such would only like to include actions that resulted in the best possible result. Due to the dynamic nature of service parts management however, some bad actions may produce good results and vice versa. This is especially a problem when then the training set is rather small as these cases will be averaged out in a larger dataset. Additionally, if only the best cases will be fed to the model it will take a long time to generate a large training set. Thus, we propose that some less than optimal cases be selected for inclusion in the earlier stages of the model to provide enough information for the model.

We can roughly divide the evaluation of the actions in three groups: (i) a group where the inventory position was within the ideal zone for the *entire* duration of the evaluation period; (ii) a group where the inventory position was within the ideal zone for *one to four* days of the evaluation period; and (iii) a group where the inventory position was *not* within the ideal zone at any point in the evaluation period. We will discuss the proposed results of the actions in more detail and have included a flowchart for the inclusion decisions based solely on a decisions quantitative score in Figure 14.

Actions that resulted in the inventory position of the part being the first group would always be included in the next iteration of the model as these have led to the desired situation. Inclusion of data in the second group will require a more detailed approach since if the inventory position of the part was within the ideal zone at a point during the evaluation period, the action might still have the desires result. If for example, the planner made an order to resolve the R38 based on the contractual lead-time and the order incurred delays, the inventory position may be outside the ideal zone for the first two days but inside the ideal zone for the last three. This action could then still be classified as a correct action as the operational planner is not in control of supplier delays and the desired result has eventually been reached. The reverse however, is also possible, an inventory position that starts within the ideal zone and moves outside during the evaluation period, or a combination of moving inside and outside of this zone. In these cases, we propose that cases where the inventory position is within the ideal zone at the end of the evaluation period always be included whereas in the other cases the average over the evaluation period is calculated. Based upon this average performance measurement, IBM will have to decide whether to include or exclude the case for the next iteration of the model.

Finally, we have the case in which the inventory position is outside the ideal zone for the entire duration of the evaluation period. Again, we propose to use the average score over the evaluation period for this decision to capture the impact of the decision over the entire evaluation period. Here, IBM will need to decide how big the deviation from the ideal zone can be for the decision to still be somewhat acceptable. It might still be worth it for IBM to manually review the extreme outliers as these can provide extra information regarding their inventory management process and the decision made by the operational planners. To illustrate this, we use the most extreme outlier found in 6.4. Here the service part inventory position during the evaluation period was around 4,000, far above the excess level of 300 for this service part. We found that this inventory position was the result of a last buy order meant to support the related machines for the coming three years. If this order would not have been made, planner's decision would have resulted in an inventory position within the ideal zone. Furthermore, analyzing this case revealed that while a last-buy order had been planned, not all of the relevant database indicators had been updated. This lack of contextual information could have resulted in the planner making this decision, instead of considering the large last-buy order coming in. Thus, analyzing extreme values could prove useful for IBM to uncover anomalies in their service parts operational management.

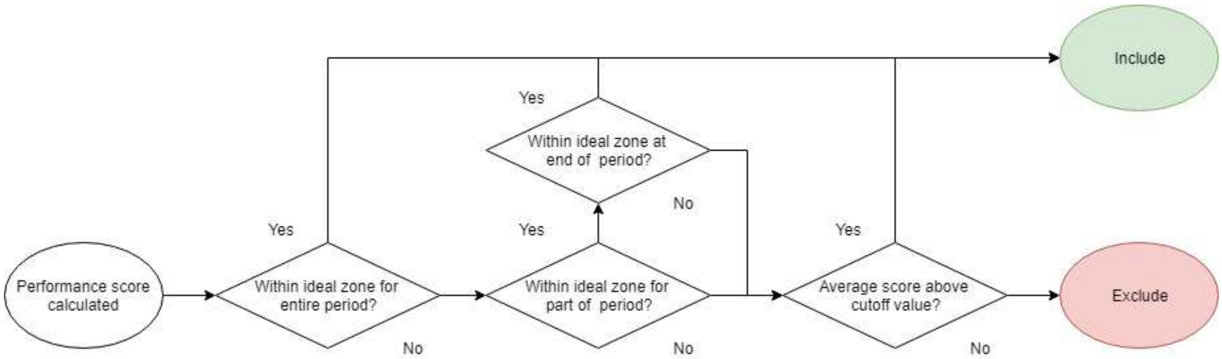

Figure 14: Case inclusion flowchart for quantitative measurement

After the cases for inclusion have been selected, they will need to be added to the model. Generally, the inclusion of new cases for a machine-learning model can be done via two methods: on-line or off-line learning. In on-line learning, whenever all the relevant information of a new case has been selected, it will be presented to the existing model and the model will use this information to adjust its rule sets. The other learning method, off-line learning, collects the new cases in a separate database, merges the old and the new cases at a later point in time, and creates a new model from scratch (Mohri, Rostamizadeh, & Talwalkar, 2012).

Selecting either of these two options requires IBM to think about their method of inclusion. If IBM wants to only use quantitative measurement for the selection of new cases, i.e. a hands-off approach, both methods would be viable. On the other hand, if IBM wants to use a combination of quantitative and qualitative measurement, the off-line learning method would be preferred as this would allow more efficient use of human capital. Presenting a batch of actions to be analyzed by an operational decision maker once per month for example, would allow him to focus on the task at hand instead of analyzing these cases immediately after their individual evaluation periods have ended.

6.6 Chapter summary

Within this part of the thesis we have proposed a method for accessing and measuring the exception handling performance. This method for performance measurement has been applied to 161 decisions made by the operational planners in response to an R38 exception. These results from this analysis can be used by IBM to generate feedback on the planner's decisions and to identify cases which can be used to improve a machine-learning model. In summary, the following were discussed in this chapter:

- The performance measurement should be evaluated during a **five-day period**, starting a lead-time after the R38 has entered the planner's work queue. We have proposed three different formulas for the performance measurement: **linear**, **exponential**, and **quadratic**. We would advise IBM to use the exponential formula as its penalty increase fits best with SPO's requirements.
- Scoring the impact of the actions made by the operational planners revealed that in **46%** of the cases the inventory position was **within the ideal zone** during the evaluation period. Cases with less than optimal results can be used as **learning** material for the planners and to provide managerial staff **insights** into the performance of the operational exception handling process.
- New cases can be added to the existing model either via **on-line** or **off-line** learning methods. If IBM wants to include cases based solely on their **quantitative** score, we would recommend the **on-line** method. Whether, case inclusion based on **both quantitative and qualitative** scoring would be better suited for the **off-line** learning method.

7 Conclusions and recommendations

Within this part of the thesis, we will discuss the answers to the research questions, and provide an overall conclusion to the main research question of: "*How can cognitive systems improve the efficiency of the operational exception handling process at IBM?"* Furthermore, we will provide IBM with recommendations for further research and a suggestion for implementation of a machine-learning model in IBM's control tower environment.

7.1 Conclusions

We will first provide answers to the five research questions central to this thesis before we provide an answer to the overall research questions.

What are the characteristics of the operational service parts environment?

This research revealed that the operational service parts management environment can be classified as a dynamic decision-making environment since reaching goals in this environment will require a series of interdependent decisions to be made in real-time during an everchanging problem state. The task of operational planners within this environment should be seen as complex due to the many interactions between items, without the ability to delay decisions, and with long feedback delays. To aid decision makers, organizations are use decision support systems to aid planners by minimize cognitive errors, providing situation specific information, generating possible actions, and implementing solutions.

What are the characteristics of machine-learning and cognitive computing?

Machine-learning, and by extension cognitive computing, enable organizations to create systems that improve through experience. These systems can be used to predict outcomes by linking historical inputs to historical actions via supervised learning. The biggest differences between traditional machine-learning models and cognitive computing are the capability of cognitive computing systems to use unstructured data, the ability to interacted with a user via natural language, and to provide the user with reasoning behind its answers.

What are the characteristics of the operational exception handling processes at IBM?

The operational exception employees responsible for the EMEA region use two systems, CPPS and Servigistics. Through these systems, they have access to general part information such as part WAC and criticality, different forecasts such as usage and returns, insights into the order book and the option to adjust or place orders, and a communication interface between other users. When prompted by an exception, planners are advised to start their problem analysis by checking the inventory, forecasts, orders, and demand of the service part. Based upon this information and their training they will attempt to solve the issue.

How can a model support the exception handling process?

For our proof-of-concept model we have used the basic service part information available to the operational planner as input. We have compared the performances of seven different machine-learning algorithms. The C5.0, QUEST, CHAID, and C&R algorithms performed similarly and we have chosen the C5.0 algorithm for our model. The model was used to predict the planner's actions for a single exception related to a projected inventory shortage. It was able to predict the planner's action in 59% of the cases, with an average confidence value of 72%.

How can we assess the quality of the exception handling process?

Our performance measurement focusses on the inventory position of a service part. The ideal situation is defined as a service part inventory between the MOR and the PSS. Since the planner makes a decision based upon the analysis period of the review reason, we propose to use the same analysis period to measure the impact of the decision. The impact of the operational exception handling decisions is then measured by penalizing deviations from this ideal situation during the evaluation period. For this purpose, we have proposed three different formulas as shown in 6.3.

By answering these five research questions we aimed to provide an answer to the overall research question regarding cognitive systems in the operational exception handling environment. Our research showed that, while technically possible, implementing a cognitive computing system in the current control tower environment of IBM would be ill-fitted. None of the strengths of a cognitive system compared to the traditional machine-learning models would be optimally used. We come to this conclusion due to the lack of unstructured data and there being no requirement for interaction between the user and the system in a natural language format.

7.2 Recommendations for further research

Given the performance of our prototype model we would advise IBM to further research the possibilities of a machine-learning models for the operational service parts exception handling. As this research provides an indication of the possibilities, we have identified areas for further research. Given the desire of IBM to use cognitive computing, we also propose an area within SPO where cognitive computing could be applied. These suggestions for research can be undertaken internally by IBM, be commissioned to external researchers, or be used as a basis for student research projects.

Improvements to the machine-learning model.

To improve the model from this research we would suggest adding more diverse cases to the model, increasing the number of inputs, or automating the input collection process. Increasing the diversity of cases and the number of inputs allows the model to easier distinguish between cases, thus improving its accuracy. Automating the input collection process of the model would enable it to store the contextual information at the moment of the planner's decision thus increasing the accuracy of the inputs.

Expansion of the model scope

The model proposed and tested in this research is specialized in the handling of a single exception type raised by Servigistics. Some of the other 39 exceptions raised by Servigsitics could be automated with a similar method. In selecting the exception for this research, we have mentioned that there are five other exceptions with similar input and output values (R5, R6, R7, R18, R39). Future research could analyze the possibilities for machine-learning models on these exceptions.

Cognitive computing systems in service parts management

If IBM wants to use cognitive systems to its internal processes, we would advise them to research the possibilities of using them in areas where much of the data is unstructured and the system has room to learn. One such an area of interest within SPO is the creation of business cases for new customers. Cognitive systems would be able to quickly gather information from unstructured sources such as contracts, emails or financial reports to indicate the profitability of the business case. Additionally, they can be used to aid the current employees by providing them with the information from earlier, similar business cases.

8 References

Basten, R. (2015). *Proactive Service Logistics Next has been awarded a one milion euro grant*. Retrieved 31-05-2017, from https://www.tue.nl/universiteit/faculteiten/industrial-engineering-innovation-sciences/01-01-1970-nieuws/17-11-2015-proactive-service-logistics-next-has-been-awarded-a-one-milion-euro-grant/

Bhosle, G., Kumar, P., Griffin-Cryan, B., van Doesburg, R., Sparks, M., & Paton, A. (2011). *Global Supply Chain Control Towers*. Retrieved 02-07-2017, from https://www.capgemini.com/resources/global-supply-chain-control-towers/

Brand Finance. (2017). *Global 500*. Retrieved 31-05-2017, from http://brandfinance.com/images/upload/global_500_2017_locked_website.pdf

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica, 81*(3), 211-241.

Cetindamar, D., Çatay, B., & Basmaci, O. S. (2005). Competition through collaboration: insights from an initiative in the Turkish textile supply chain. *Supply Chain Management: An International Journal, 10*(4), 238-240. doi:10.1108/13598540510612686]

Chen, Y., Argentinis, J. E., & Weber, G. (2016). IBM Watson: how cognitive computing can be applied to big data challenges in life sciences research. *Clinical therapeutics, 38*(4), 688-701.

Consortium, C. C. (2016). *Cognitive Computing Definition*. Retrieved 20-09-2017, from https://cognitivecomputingconsortium.com/definition-of-cognitive-computing/

Dalton, R., Mallow, C., & Kruglewicz, S. (2015). *Disruption ahead: Deloitte's point of view on IBM Watson*. Retrieved 20-07-2017, from

de Kok, T., van Dalen, J., & van Hillegersberg, J. (2015). *Cross-Chain Collaboration in the Fast Moving Consumer Goods Supply Chain*. Eindhoven: Eindhoven University of Technology.

Dietterich, T. G. (1997). Machine-learning research. *AI Magazine, 18*(4), 97.

Driessen, M., Arts, J., van Houtum, G.-J., Rustenburg, J. W., & Huisman, B. (2015). Maintenance spare parts planning and control: a framework for control and agenda for future research. *Production Planning & Control, 26*(5), 407-426.

Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A. A., . . . Prager, J. (2010). Building Watson: An overview of the DeepQA project. *AI Magazine, 31*(3), 59-79.

Fransoo, J. C., & Wiers, V. C. S. (2008). An empirical investigation of the neglect of MRP information by production planners. *Production Planning & Control, 19*(8), 781-787. doi:10.1080/09537280802568163

Geertjes, F. (2014). *A first step towards fully automated spare parts planning systems: An empirical study on ordering behaviour.* (Master), University of Twente, Enschede.

Gibson, F. P. (2000). Feedback Delays: How Can Decision Makers Learn Not to Buy a New Car Every Time the Garage Is Empty? *Organizational Behavior and Human Decision Processes, 83*(1), 141-166. doi:http://dx.doi.org/10.1006/obhd.2000.2906

Godsiff, P. (2010). Service systems and requisite variety. *Service Science, 2*(1-2), 92-101.

Goksel Canbek, N., & Mutlu, M. E. (2016). On the track of Artificial Intelligence: Learning with Intelligent Personal Assistants. *International Journal of Human Sciences, 13*(1), 592-601. doi:10.14687/ijhs.v13i1.3549

Gonzalez, C. (2005). Decision support for real-time, dynamic decision-making tasks. *Organizational Behavior and Human Decision Processes, 96*(2), 142-154.

Heaney, B. (2014). *Supply Chain Visibility and Segmentation: Control Tower Approach*. Retrieved 03-04-2017, from https://www.inttra.com/Assets/Documents/aberdeen-paper---supply-chain-visibility-and-segmentation.pdf?msc=anlypg

Herper, M. (2017). *MD Anderson Benches IBM Watson In Setback For Artificial Intelligence In Medicine*.

Retrieved 26-09-2017, from https://www.forbes.com/sites/matthewherper/2017/02/19/md-anderson-benches-ibm-watson-in-setback-for-artificial-intelligence-in-medicine/#414a56f13774
Hoffman, B., & Bhattacharya, R. (2016). *Machine Learning, Deep Learning 101*. Retrieved 19-09-2017, from https://www.ibm.com/developerworks/library/l-machine-learning-deep-learning-trs/
IBM. (2016a). *IBM Annual Report*. Retrieved 20-07-2017, from https://www.ibm.com/annualreport/2016/images/downloads/IBM-Annual-Report-2016.pdf
IBM. (2016b). *IBM SPSS Modeler 18.0 Modeling Nodes*. Retrieved 22-10-2017, from ftp://public.dhe.ibm.com/software/analytics/spss/documentation/modeler/18.0/en/ModelerModelingNodes.pdf
IBM. (2017). *IBM Watson Health*. Retrieved 11-09-2017, from https://www.ibm.com/blogs/watson-health/watson-and-cancer-get-the-facts/
Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science, 349*(6245), 255-260.
Jüttner, U., Peck, H., & Christopher, M. (2003). Supply Chain Risk Management: Outlining an Agenda for Future Research. *International Journal of Logistics : Research & Applications, 6*(4), 197-210.
Kelly, J. E. (2015). *Computing, cognition and the future of knowing: How humans and machines are forging a new age of understanding*. Retrieved 10-04-2017, from https://www.research.ibm.com/software/IBMResearch/multimedia/Computing_Cognition_WhitePaper.pdf
Khurana, M. K., Mishra, P. K., & Singh, A. R. (2011). Barriers to Information Sharing in Supply Chain of Manufacturing Industries. *International Journal of Manufacturing Systems, 1*(1), 9-29.
Kubat, M. (2015). *An introduction to Machine Learning*: Springer International Publishing.
Lerch, F. J., & Harter, D. E. (2001). Cognitive support for real-time dynamic decision making. *Information systems research, 12*(1), 63-82.
Lurie, N. H., & Swaminathan, J. M. (2009). Is timely information always better? The effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes, 108*(2), 315-329. doi:10.1016/j.obhdp.2008.05.005
Mohri, M., Rostamizadeh, A., & Talwalkar, A. (2012). *Foundations of machine learning*: MIT press.
NG, A. (2016). *IBM's Watson gives proper diagnosis after doctors were stumped*. Retrieved 17-07-2017, from New York: http://www.nydailynews.com/news/world/ibm-watson-proper-diagnosis-doctors-stumped-article-1.2741857
Niu, L., Lu, J., & Zhang, G. (2009). Decision Making and Decision Support Systems. In *Cognition-Driven Decision Support for Business Intelligence: Models, Techniques, Systems and Applications* (pp. 3-18). Berlin, Heidelberg: Springer Berlin Heidelberg.
NWO. (2016). *Pro-active service logistics for capital goods - the next steps (ProSeLoNext)* Retrieved 31-05-2017, from https://www.nwo.nl/onderzoek-en-resultaten/onderzoeksprojecten/i/92/14292.html
Papathanasiou, J., Linden, I., & Ploskas, N. (2016). *Real-World Decision Support Systems: Case Studies* (Vol. 37). Switzerland: Springer International Publishing
Power, D. J. (2016). Computerized Decision Support Case Study Research: Concepts and Suggestions. In J. Papathanasiou, N. Ploskas, & I. Linden (Eds.), *Real-World Decision Support Systems: Case Studies* (pp. 1-13). Cham: Springer International Publishing.
Power, D. J., Sharda, R., & Kulkarni, U. (2007). Understanding Decision Support Systems for Global Enterprises. In U. Kulkarni, D. J. Power, & R. Sharda (Eds.), *Decision Support for Global Enterprises* (pp. 3-14). Boston, MA: Springer US.
Proper, S., & Tadepalli, P. (2004). *Scaling Average-reward Reinforcement Learning for Product Delivery.* Paper presented at the American Association for Artificial Intelligence.
Raj, S., & Sharma, A. (2014). *Supply chain management in the cloud* from
Rometty, G. (2017) *Bloomberg Businessweek Debrief with Ginni Rometty/Interviewer: M. Murphy*. Sooner

Than You Think at Cornell Tech, Bloomberg.
Rustenburg, J. W. (2016). Planning Services: A Control Tower Solution for Managing Spare Parts. In H. Zijm, M. Klumpp, U. Clausen, & M. t. Hompel (Eds.), *Logistics and Supply Chain Innovation: Bridging the Gap between Theory and Practice* (pp. 239-259). Cham: Springer International Publishing.
Schmidt, G., & Wilhelm, W. E. (2000). Strategic, tactical and operational decisions in multi-national logistics networks: a review and discussion of modelling issues. *International Journal of Production Research, 38*(7), 1501-1523.
Shou-Wen, J., Ying, T., & Yang-Hua, G. (2013). Study on Supply Chain Information Control Tower System. *Information Technology Journal, 12*(24), 8488-8493.
Slack, N., Brandon-Jones, A., & Johnston, R. (2013). *Operations Management*: Pearson.
Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science, 35*(3), 321-339.
Trzuskawska-Grzesińska, A. (2017). Control towers in supply chain management – past and future. *Journal of Economics and Management, 27*, 114-133. doi:10.22367/jem.2017.27.07
Varian, H. R. (2014). Big Data: New Tricks for Econometrics. *The Journal of Economic Perspectives, 28*(2), 3-27.
Wang, G., Gunasekaran, A., Ngai, E. W. T., & Papadopoulos, T. (2016). Big data analytics in logistics and supply chain management: Certain investigations for research and applications. *International Journal of Production Economics, 176*, 98-110. doi:10.1016/j.ijpe.2016.03.014
Wang, L., & Alexander, C. A. (2015). Big Data Driven Supply Chain Management and Business Administration. *American Journal of Economics and Business Administration, 7*(2), 60-67. doi:10.3844/ajebasp.2015.60.67
Wang, Y., Zhang, D., & Kinsner, W. (2010). *Advances in cognitive informatics and cognitive computing*. Berlin: Springer-Verlag.

Appendix A Planner worksheet

The first part of the planner worksheet is the item data window, and provides planners with general information regarding the service part. Here the user can choose the aggregation level of their analysis, e.g. a single location or the consolidated stock for the entire region. The user is also able to change some of the part settings that influence the degree of automation such as: automatic allocation, eligibility for automatic forecast selection, releasing a forced worksheet, disabling automatic planning for the part, automating transshipments and automatic excess recall. Lastly, the item data window provides the user with basic part information such as the part forecast settings, the part criticality, and the part's WAC.

The next part of the planner worksheet is the schedules data workspace. In this part of the worksheet planners are shown periodic stock level data, either in a table or a graph. Users can switch between different tabs showing information regarding the part forecast, the levels and orders of parts in the field, a graphical representation of the part movements, the inventory levels of the part, stock analysis per period, the stock planning levels e.g. safety stock, and a summary of the part analysis. The operational service parts planners responsible for the central buffer locations will see the consolidated demand for the entire region at the central buffer locations.

The review reason list on the planner work sheet contains all the exceptions raised by the system for this specific service part. When an exception has been found, the review reason will appear both in the planner's work queue and in the review reason list for the relevant service part. Failing to respond to a single review reasons may lead to other review reasons being triggered in the future. An example being a backorder being created due to a planner ignoring the forecasted inventory shortage. The number of review reasons linked to a specific service part is not limited but planners are advised to keep this number as low as possible. In the case of a service part having multiple review reason, the planner will focus on the one with the highest priority first while considering the others during its decisions process. As understanding review reasons is vital to this research, they will be described in more detail in 4.3.1.

Next, we have the notepad area. This area is used to communicate between planners regarding their actions for the service part and to show historic actions taken. Planners write their notes in the specific region sheet and these notes are stored in the system linked to the service part. Entries in this area can provide extra information regarding the planner's thought process and could be helpful for new planners to learn what actions others have taken in the past.

The last module of the planner worksheet is the item family window. This window gives planner a quick overview of the on-hand balance of the service part and any related parts such as substitutes. Within Servigistics a service part is either a prime service part or an alternate part. A prime service part is the "active" part, i.e. the part that is used in the forecast and against which orders are placed. Alternate parts can be substitutable parts or earlier models. The demand history and stock balances of alternate parts are considered in the systems' calculations of the prime's forecast.

Within the item family window, a user can see the substitution relation between the service parts. Most of the substitutable service parts within IBM are linked in a predecessor successor relationship. Within this relationship, the predecessor is the "older" version of the part and the successor is the "new" replacement. During the service parts management processes, planners try to deplete the supply of older parts first before moving on to the new parts, if this is in accordance with the customer. When a service part is able to fulfill the requirements of one part but not the other way around, the first part is called a forward substitutable for the second part. Lastly, there are parts within the service parts portfolio of IBM that are matrix substitutable. These are parts that have equal form, fit and function but differ in price. Deciding which part to use in this case will depend on factors such as the part lead time, the availability of the stock and the client's preference. An example of a matrix substitution would be if a customer requires a 300 GB HDD but IBM has no such HDD on stock. In this case planners, could send a more expensive model, assuming it fits the machine, such as 600 GB HDD instead. While a planner can send the 600 GB HDD instead of the 300 GB HDD, he can't do it the other way around as a 300 GB HDD is not able to perform up to the level of a 600 GB one.

Appendix B List of review reasons

ID	Name	Priority
R1	Forced Review A-Item	43
R2	Forced Worksheet	1
R3	Zero Unit Cost	39
R4	Backorder Quantity	2
R5	Order Decrease Inside L/T	8
R6	Order Change Outside L/T	44
R7	Order Increase Inside L/T	7
R8	Recommends Zero On order	45
R9	Behind Sched. Over Policy	16
R10	Block Passing PSS Shortfall	79
R11	Predecessor Can Be Unregistered	80
R14	No Forecast, but Activity This Period	18
R15	Change Forecast Setting	3
R16	Forecast Tracking Error	4
R17	History Value Out of Range	5
R18	Order Change Inside L/T	46
R19	Order on Sumcode 2 Part	17
R20	Scheduled Demand Exceeds Forecast	47
R21	Inventory on Sumcode 2 Part	48
R22	Configuration Change	34
R23	Critical Review Item	49
R24	Projected Stock Out	22
R25	Stocked Out	19
R26	Below Must Order Point	20
R27	Stock Below/at Minimum	21
R28	BOM Configuration Changed	50
R29	IPLS Order Received	27
R31	Part Needs to Be Allocated	13
R32	Part Needs to Be Transshipped	32
R33	Incorrect Receipt Exists	37
R34	Failed to Auto Approve Orders	6
R35	Audit Transaction Logged	51
R36	Large Future Order	52
R37	Unsatisfied Allocation Requirement	14
R38	Order Increase Outside L/T	10
R39	Order Decrease Outside L/T	11
R40	Excess Inventory	53
R41	Quantity Outside Location Hierarchy	24
R42	Order Overdue	38

R43	Zero Option Quantity	54
R44	Opportunity to Maximize Repair	55
R45	Transship Instead of Buy	31
R46	Order Reschedule Inside L/T	9
R47	Reprioritize Orders ILT	41
R48	Projected Stock Below MOP	42
R49	Order Type Enabled Mismatch	56
R50	Underutilized Repair Stock	57
R51	Order Note Received	28
R53	Return Orders Exceed Capacity	58
R54	Newbuy Orders Exceed Capacity	59
R55	Stop Repair Line	60
R56	Repair Line Contract Review	61
R57	Repair Line Safety	62
R58	Above POS Limit	63
R59	Below POS Limit	64
R60	New Order	65
R61	Trapped Surplus Inventory	36
R62	Allocating into Trapped Surplus Location	66
R63	Opportunity to use sumcode 2 inventory	67
R64	Repair Stock May Be Moved	68
R65	Orders Exist Beyond EOL Due Date	69
R66	EOL Quantity Not Calculated	68
R67	Good As It Gets Date Changed	33
R68	Returned Stock Overused	25
R69	Opportunity to use Secondary Repair	71
R70	Returns Order with No Repair-From Part	40
R71	Theater Collaboration	35
R72	Part May Be Recalled	30
R73	Above POS Limit With OHB	76
R77	Consider Croston's Forecasting Method	70
R78	Order Requires Legal Entity Approval	29
R79	Order Reschedule Outside L/T	12
R80	Projected Inventory Above XS Point	15
R81	Supply Constraint Over Consumed	26
R82	Order Change within Approval Horizon	72
R83	Projected Inventory Below MOP	23
R84	Local Projected Stock Out	73
R85	Local Projected Below MOP	74
R86	Failed to Reach EOL Inventory Target	75
R100	Unsatisfied Internal Demand Recall Requirements	77

Appendix C Review reasons causes and actions

Description	Probable Cause	Suggested Actions
Forced Worksheet (R2)	Triggered when the 'F' – Force Worksheet part flag is ON indicating a desire to review the part on a regular basis regardless of other active review reasons. Force Worksheet part flags can be set by location in the item data section of the planner worksheet.	The analyzer sets the Force Worksheet flag when they want to periodically review specific parts. Analyzers may review and action whatever data they desire based on the specific reasons they forced the part into their work queue. Turn the Force Worksheet flag off when the part no longer needs periodic review.
Zero Unit Cost (R3)	Triggered when the WAC at a location is explicitly set to zero. This will not happen on new parts because the default WAC is set to $0.01 in XelusPlan.	This RR is information only.
Backorder Quantity (R4)	Triggered when there are open customer orders that are past their due date. Backorders that trigger this RR are: • Customer orders that are past due and not yet filled. Backorders that already have stock allocated to them will be excluded. • Low priority backorders filled after replenishment of internal locations (US Only). These low priority orders are identified via a lower priority emergency code (Code X and B) in PIMS. Reference customer backorders are captured but will not trigger this RR. These are backorders that have been referred to another geography or to a supplier. They will display against the location where the original order was registered.	Review your open order position and determine if there are enough parts on order and/or if expediting or rescheduling is necessary. Review other locations to determine if inventory from another location can be transshipped to satisfy the backordered demand.
Order Decrease ILT (R5)	Triggered when the recommended orders result in an order decrease inside lead-time (ILT). The lead-time for evaluation is specific to part, source, customer and order type. The RR trigger will only occur if decrease freeze days are less than lead-time. Since most parts are set to auto approve, this RR will	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order decrease message. Review the 'Net Change' schedule in the MAIN schedule book, Analysis page, to see the magnitude of the recommended changes ILT.

	only trigger if auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation.	Resolve RRs blocking automatic order approval and then approve orders manually. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Order Increase ILT (R7)	Triggered when the recommended orders result in an order increase inside lead-time (ILT). The lead-time for evaluation is specific to part, source, customer and order type. The RR trigger will only occur if increase freeze days are less than lead-time. Since most parts are set to auto approve, there are only two reasons for this RR to trigger: • Auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation. • The order is a repair or warranty order recommended against forecast.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order increase message. Review the 'Net Change' schedule in the MAIN schedule book, Analysis page, to see the magnitude of the recommended changes ILT. Resolve RRs blocking automatic order approval and then approve orders manually. If the order increase is against repair or warranty forecast, decide to wait for the returns or place an order on an alternate source of supply. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Behind Schedule/ Over Policy (R9)	Triggered when the Overdue order (backlog) quantity for the part exceeds the recommended order quantity in an ordering location for the current period, and the maximum inventory level reached during the current period is greater than the XS Point threshold. Backlog orders are never candidates for order change recommendations.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the backlog order. Determine if canceling or reducing the overdue (backlog) order(s) is an option. If a repair or warranty order is changed, update the order "Comm Flag" to 'D' – Direct Interface if it should not wait for the weekly interface to SAP.
No Forecast, but Activity This Period (R14)	Triggered when the executive forecast over the next two years is zero and there is recurring demand quantity in the current period.	Drill-Down on the RR notes to view the location with the activity data, the forecasted location and the type of forecast (demand, repair or

	When this is evaluated for repair or warranty return forecasts it is triggered when there is zero forecast over the next two years, but a repair or warranty return quantity exists in the current period.	warranty). For new parts consider using planner adjustments to estimate the requirements over the next few periods. For old parts, you may just approve this RR with no action if you are bleeding off inventory through end of service.
Change Forecast Setting (R15)	Triggered when an alternate forecast setting is performing more accurately (has a lower mean square error (MSE)) than the executive forecast format for 1 (P9) consecutive periods. The difference between the executive and alternate forecast must also be greater than 1 piece. This RR is evaluated for all forecast schedules: demand, repair and warranty. Since most parts are set to auto forecast, this RR should only be triggered on parts where auto forecasting is disabled (the 'E' – Eligible for Auto Forecast part flag is off).	Drill-Down on the RR notes to view the forecasted location and the type of forecast (demand, repair or warranty) requiring a change to forecast settings. Open the forecast settings window to select the alternate forecast setting that is performing better than the executive. The lower MSE ratio equates to a better performing forecast setting.
Forecast Tracking Error (R16)	Triggered when the back cast is biased high or low when compared to actual historical activity for 3 (P18) consecutive periods. The accuracy is measured using a smoothed error tracking signal (SETS) calculation. The number of standard deviations used to determine the boundaries for outliers is controlled by P19 – SETS Sigma. This RR is evaluated for all forecast schedules: demand, repair and warranty.	Drill-Down on the RR notes to view the forecasted location and the type of forecast (demand, repair or warranty) exhibiting the forecast tracking error. Open the forecast settings window to see if an alternate forecast is performing marginally better than the executive. Select a better performing forecast setting if one exists. If you are seeing this RR a lot and changing forecast settings does not typically clear the RR then notify the super user. Work may need to be done to tune the SETS forecast parameters.
History Value Out Of Range (R17)	Triggered when the recurring demand summed over primes and alternates falls outside the demand filter trip limits. These limits are defined as a 2.0 (P4)	Drill-Down on the RR notes to view the forecasted location and the type of forecast (demand only) experiencing an outlier.

	standard deviations above and below the back cast for the most current prior period. The square root of the forecast mean square error (MSE) is one standard deviation. This will only be triggered for demand outliers relative to the demand forecast. (not on warranty or repair forecasts)	Users associated with PIMS locations will manually enter demand corrections in the "Manual Demand Correction" schedule. Enter a negative quantity to reduce recurring demand and increase non-recurring demand, leaving net change to total demand zero. CPPS will feed demand adjustments daily A user may also enter a planner (forecast) adjustment if the outlier demand is known to occur over a number of future periods.
Order on Sum Code 2 Part (R19)	Triggered if an open order exists for a sum code 2 alternate in the prime/alternate chain. Sum code 2 alternates are down level parts in an EC substitution chain with a disposition of scrap, re-label or return to source.	Validate the prime/alternate relationship. Contact your WWDC representative if the substitution chain is incorrect. Typically sum code 2 inventory is obsolete so the action should be to cancel the order or change the ordered part number to the up-level/prime part.
Configuration Change (R22)	Triggered when a new part has been added to XelusPlan or the prime/alternate chaining relationship has been modified as follows: • Prime part reclassified as an alternate • Alternate part reclassified as a prime • Alternates for a giving prime part have changed If a condition changes for an alternate item that is a member of the prime-alternate chain, the system sets this review reason on the prime. Prime/alternate chains are used for the following substitutions in WWDC: • Fully compatible EC • Fully compatible Matrix • Equivalent	Validate the prime/alternate relationship. Contact your WWDC representative if the substitution chain is incorrect or contact the Reutilization PA is the SKU data is incorrect. If it's a new part, verify that the corresponding part data such as location hierarchy, supplier and auto order flags are correct.

	• SKU Inventory	
Projected Stock Out (R24)	Triggered when a stock out is projected in the next two years (i.e. projected inventory goes negative) in a network-ordering sub-network, and the order recommendations are constrained such that the stock out cannot be avoided. XelusPlan order recommendation constraints may be one or more of the following: • Increase freeze days • No valid part/sources for any order type that is enabled. • No order types enabled that have valid part/sources. • "Good as it Gets" effective dates • Capacity or Supply constraints • All three of the following conditions must be satisfied for this review reason to be triggered: • The planning inventory, calculated over the prime/alternate chain, must be greater than zero. • The sum of the Total Requirements (T1) two years out must be greater than zero. • The projected inventory calculated over the prime/alternate chain for the ordering sub-network must be negative on one or more days in the future.	Verify the executive forecast setting is correct. Determine what ordering constraint is prohibiting the system from placing orders sooner. Investigate whether an exception to this constraint can be negotiated with the appropriate supplier. (e.g. reschedule orders inside freeze or GAIG date) If there are no orders or order recommendations investigate new sources of supply.
Stocked Out (R25)	Triggered when the network planning inventory, calculated over the prime/alternate chain, is equal to or less than zero and the sum of Total Requirements in the network over the next two years is greater than zero. Planning inventory is the sum of the following inventories less receipt corrections: • New Stock • Stock Reserved/Allocated to a Customer	Verify the executive forecast setting is correct. Determine if there are existing orders or recommended orders outside freeze that can be expedited and received sooner. If there are no orders or order recommendations investigate new sources of supply. If this is a new part, review the order

	Order • Kit Stock (EMEA Only) • Dock to Stock (US Only) • Stock On Loan • Limited Use Stock: • Used Class Stock • Non-Certified Repair Stock • Stock Trapped due to the fact that is it below the Minimum Recall Value • New and Limited Use In transit Stock	recommendations and validate there is a sufficient quantity on order.
Below Must Order Point (R26)	Triggered when the network-planning inventory (T18) for a part in an ordering sub-network is below the computed T63 Must Order Point (MOP) in the current period. This review reason is triggered only if both the on-hand balance and MOP quantities are non-zero.	Verify the executive forecast setting is correct. Determine if there are existing orders or recommended orders outside freeze that can be expedited and received sooner. If there are no orders or order recommendations investigate new sources of supply. If this is a new part, review the order recommendations and validate there is a sufficient quantity on order.
Stock Below/at Minimum (R27)	Not Yet Active	
IPLS Order Received (R29)	Triggered when a new order through IPLS from a Plant is sent to the analyzer for review and acceptance. (TX400) While the order is in a placed/queued status, the RR will be active on the part and the RR disable days will control whether the user gets the part in their work queue for this RR again.	The analyzer can make changes to order quantities and dates. Based on due date logic they will have a minimum of 5 days to take action before the order is sent out for shipment. No order changes will be accepted from s Plant. There will be a business process that forces the manufacturing planner to negotiate order changes with the ITS analyzer. The analyzer will make appropriate changes to orders in XelusPlan.
Part Needs To Be Allocated (R31)	This review reason is triggered when there is a recommended replenishment order from parent (or	The Allocation List window will show the recommended allocation moves in the

	grandparent) to a child location. This move will not be recommended unless the child location's allocation on-hand balance is below its order point. This review reason will only be triggered if the parent location has sufficient stock to fill the recommended order. Most parts will be configured for automatic allocation and there are no review reasons that will block automatic allocation so this review reason should appear infrequently.	Recommended Due Out column for the shipping location and in the Recommended Due In column for the receiving location. Generating the recommended allocation orders clears this review reason. At the top of the Allocation List window there are three checkboxes for Allocation, Transship and Excess Recall. Make sure that the 'Release using Allocation' box is checked and hit the Generate button in the upper left corner of the Allocation List window. Once the allocation orders are generated the review reason will be cleared. The recommended order quantity may be changed prior to generation by highlighting the shipping location and clicking on the Source button in the upper left corner of the Allocation list window.
Part Needs To Be Transshipped (R32)	This review reason is similar to 'Part Needs to be Allocated' but the recommended move order is from a location other than the receiving location's primary replenishment location. The review reason will be triggered when there is a valid transship link between the shipping location and the receiving location. This review reason will only be triggered if the shipping location has sufficient stock to fill the recommended order. Most parts will be configured for automatic transshipment and there are no review reasons that will block automatic transshipment so this review reason should appear infrequently.	The Allocation List window will show the recommended transship moves in the Recommended Due Out column for the shipping location and in the Recommended Due In column for the receiving location. Generating the recommended transship orders clears this review reason. At the top of the Allocation List window there are three checkboxes for Allocation, Transship and Excess Recall. Make sure that the 'Release using Transship' box is checked and hit the Generate button in the upper left corner of the Allocation List window. Once the transship orders are generated the review reason will be cleared. The recommended order quantity may be changed prior to generation by highlighting the shipping location and clicking on the Source button in the upper left corner of the Allocation

		list window.
Forelog Exists (R33)	When more inventory has been received than was on the purchase order, the forelog exists RR will be triggered.	If inventory is received that cannot be matched to a purchase order the buyer and planner have to decide if they agree to receive it. If they agree the inventory will be received, the analyzer should create an order with the 'I' – Incorrect Receipt Comm Status that will be sent to SAP to receive against.
Failed to Auto Approve Orders (R34)	Triggered if another RR is active that is configured to block auto order OR if an order for the same orderbook is in a "change cycle". Orders in a change cycle, for example, are those in a placed/sent status waiting for confirmation from SAP.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the auto approval failure message. Resolve RRs blocking automatic order approval and then approve orders manually. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP. If the auto approval is blocked by orders in a change cycle, approve the part and wait for the order communication loop to complete.
Unsatisfied Allocation Requirement (R37)	This review reason is triggered when a child location cannot be replenished to its Order Point (allocation threshold 2). The location's parent (and its parent if it is a bypass location) does not have sufficient inventory to bring the child up to its order point. For this review reason to be triggered there would be insufficient local inflow in the look-ahead period to bring the location's inventory above the order point and there would be insufficient transship opportunities.	Check the review reason notes to find the locations that cannot be satisfied. The Allocation List window will provide information about the order points, the available inventory at the location and at its parent and any incoming supply orders in the look-ahead period. The analyzer may be able to manually transship stock from a location that is not configured to automatically share inventory with the location in need. Manual transship orders may be created in Add Orders window in the Order List window. The analyzer may be able to expedite re-supply orders due in to the location or due in at the parent location.

Order Increase OLT (R38)	Triggered when the recommended orders result in an order increase outside lead-time (OLT). Recommendations on days more than 4 days OLT are not considered. The lead-time for evaluation is specific to part, source, customer and order type. Since most parts are set to auto approve, there are only two reasons for this RR to trigger: • Auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation. • The order is a repair or warranty order recommended against forecast.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order increase message. Resolve RRs blocking automatic order approval and then approve orders manually. If the order increase is against repair or warranty forecast, decide to wait for the returns or place an order on an alternate source of supply. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Order Decrease OLT (R39)	Triggered when the recommended orders result in an order decrease outside lead-time (OLT). Recommendations on days more than 4 days OLT are not considered. The lead-time for evaluation is specific to part, source, customer and order type. Since most parts are set to auto approve, this RR will only trigger if auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order decrease message. Resolve RRs blocking automatic order approval and then approve orders manually. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Quantity Outside Loc/Hier (R41)	This review reason will be displayed when there is planning data in the current period or some future period for a location that is not attached to the part's location hierarchy. Data types that will trigger this review reason are: • T13 - Executive Forecast • T18 – On-hand Balance • T19 - Total Backorders • T48 - Proposed Orders • T25 - Returns Forecast • T28 - Repair On-hand Stock	Notify your super user to investigate this problem. The super user will most likely need to make a change to the appropriate location hierarchy to include the location.
Order Overdue	Triggered when the due date on a new buy or returns	This RR is information only.

(R42)	type order is more than 1 day past due. All of the following order criteria must be true to trigger this review reason: • The order due date + 1 day must be less than or equal to the current date. • An outstanding balance must exist on the order (i.e. the order is not fully received). • The order must not be closed	The buyer will manage overdue orders.
Tranship Instead of Buy (R45)	New Buy or Return type order recommendations are greater than zero within lead time for the primary source and tranship is also recommended with the ordering location as the customer. The lead-time for evaluation is specific to part, source, customer and order type. This review reason is only evaluated at ordering locations.	The review reason notes will specify which locations satisfy the conditions of this review reason. The specific order recommendations may be viewed in the Order List window. The tranship recommendations may be reviewed in the Allocation List window. The analyzer should decide which recommendations to accept and should generate orders accordingly.
Order Reschedule ILT (R46)	Triggered when an order recommendation exists to reschedule an order with an original due date OR rescheduled due date inside lead-time (ILT). The lead-time for evaluation is specific to part, source, customer and order type. The RR trigger will only occur if freeze days are less than lead-time. Since most parts are set to auto approve, this RR will only trigger if auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order reschedule message. Review the 'Net Change' schedule in the MAIN schedule book, Analysis page, to see the magnitude of the recommended changes ILT. Resolve RRs blocking automatic order approval and then approve orders manually. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Reprioritize Orders ILT (R47)	Triggered when at least one order recommendation exists to cancel or decrease an order against a lower	Drill-Down on the RR notes to see the ordering location where orders are being reprioritized.

	priority order type to place or increase an order on a higher priority order type inside lead-time. The orders may exist on either primes or sum code 1 alternates. The lead-time for evaluation is specific to part, source, customer and order type of the existing order associated with the lower priority order type. Since most parts are set to auto approve, there are only two reasons for this RR to trigger: • Auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation. • The order is a repair or warranty order recommended against forecast.	Resolve RRs blocking automatic order approval and then approve orders manually. If an order increase is against repair or warranty forecast, you may decide to wait for the returns or leave the existing lower priority order in place. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Order Note Received (R51)	This review reason is triggered when an order note is received on an order confirmation record from SAP.	A column in the Order Lists window indicates the existence of a note. Select the order by clicking on the order record and the Order Note In and Order Note Out buttons in the Order List will become active. Use the Order Note In button to view the note. This review reason is triggered by the feed of the order note. Once it has been approved the review reason will not appear again until another note is received.
Trapped Surplus Inventory (R61)	Triggered when planning inventory (T18) in a trapped location exceeds PSS (Policy Safety Stock) + 3 POS (periods of supply). For most locations that are not forecasted, this will simply be triggered if planning inventory is greater than PSS. This RR will also be triggered is limited use inventory in a limited use trapped location exceeds PSS + 3 POS. When you have nested trapping locations as you do in both LA and AP, note that trapped surplus is not summed, it is calculated at each trapping location.	This RR is information only. Drill down on the RR notes to view the location(s) with trapped surplus inventory. Manually initiate a transshipment of excess inventory out of the trapping location if that is possible. Do not initiate a transshipment of Limited Use Excess Inventory. This inventory cannot cross country borders.

Good As It Gets Date Changed (R67)	The Good As It Gets Date is used as a temporary freeze period when a buyer knows that the due date for an order cannot be improved or an analyzer knows there is a temporary problem with a supplier. All orders from this supplier due before the GAIG date are affected. The buyer may set the Good As It Gets date on individual orders in SAP. This will be communicated to Xelus as an order confirmation record. Analyzers may also set a Source Good As It Gets date on a part/source in the Sources window of the Planner Worksheet. XelusPlan will calculate the effective Good As It Gets date by taking the most future of the Source date set by an analyzer and the dates set on any orders fed from SAP. The effective Good As It Gets date is displayed in the Sources window. This review reason is triggered if the effective Good As It Gets date is changed by a feed from SAP.	The analyzer should review the supply plan for the location in question. If the due date on the order setting the new Good As It Gets date will cause supply shortages other review reasons should also be active. Projected Below Must Order Point or Projected Stock Out review reasons are likely to be triggered. The analyzer should take appropriate steps to resolve those review reasons. This review reason is triggered by the feed of the order confirmation. Once it has been approved the review reason will not appear again until the Good As It Gets date changes again.
Returned Stock Over Used (R68)	Triggered when the unyielded quantities on repair or warranty orders exceed the available AFR + WIP. This RR will be evaluated for each SKU/ordering location or FRU/ordering location in EMEA. The system will never recommend this condition. It may be caused by AFR stock updates or changes made upon order confirmation after communication with WMS.	Drill-Down on the RR notes to see the ordering location and order type where AFR stock is over used. Validate the AFR quantities are correct between Xelus and WMS. Manually decrease or cancel the repair or warranty orders. These orders are firm so the system will never recommend changing them. All repair and warranty orders are set to firm after interfacing with SAP since communication has been done to the warehouse to pick/pack/ship the AFR.
Returns Order with No Repair-From Part (R70)	Triggered when the repair-from part number on a repair or warranty order is blank.	Analyzers should notify the super users if this condition exists because repair-from part number is required on repair and warranty

		orders. There may be a feed or integration issue that needs to be resolved.
Theater Collaboration (R71)	This is triggered when an analyzer in another theater sends a collaboration message via the Notepads.	The review reason notes show the receiving location. The analyzer responsible for this location should open the appropriate Notepad to receive the intended message. The Notepads can be accessed through the View dropdown menu and the Item Data selection.
Part May be Recalled (R72)	This review reason is triggered when excess recall has been configured and there is inventory to be moved using the excess recall rules. Excess recall runs after all other allocation and transshipment moves have been recommended. If stock-on-hand at a location exceeds the allocation excess threshold and a recall location has been assigned to the inventory location this review reason will be triggered. The review reason notes will show the locations with excess stock to be recalled. Most parts will be configured for automatic excess recall and this review reason will not appear. There are no review reasons that will block automatic excess recall.	The Allocation List window will show the recommended recall moves in the Recommended Due Out column for the shipping location and in the Recommended Due In column for the recall location. This window also displays the assign excess recall location for each stocking location, if a recall location has been configured. In addition, the Allocation List window shows the value of the allocation excess threshold for each location. Generating the recommended recall move orders clears this review reason. At the top of the Allocation List window there are three checkboxes for Allocation, Transship and Excess Recall. Make sure that the ‘Release using Excess Recall’ box is checked and hit the Generate button in the upper left corner of the Allocation List window. Once the recall orders are generated the review reason will be cleared.
Order Requires LE Approval (R78)	This review reason is tripped when there is an allocation or re-supply order that meets the criteria for legal entity approval. Allocation orders require legal entity approval if the order’s source and customer locations have different legal entities. The value of the order must exceed the legal entity approval threshold for the receiving location. Re-supply orders require legal entity approval if the	Open the Order List window and select the Legal Entity Approval toolbar button. Review the orders that require approval and approve the orders or modify the proposed quantity and then approve the orders.

	legal entity of the user responsible for planning the part is different than the legal entity of the customer location. The value of the order must exceed the legal entity approval threshold for the receiving location.	
Order Reschedule OLT (R79)	Triggered when an order recommendation exists to reschedule an order whose original due date AND rescheduled due date are both outside lead-time. The lead-time for evaluation is specific to part, source, customer and order type. Recommendations for rescheduled orders whose original and rescheduled due date are more than 4 days OLT are not considered. Since most parts are set to auto approve, this RR will only trigger if auto order approval was blocked from automatically approving the order(s). The Failed to Auto Approve Orders RR will accompany this RR in this situation.	Drill-Down on the RR notes to see the ordering location, supplier, and order type associated with the order reschedule message. Resolve RRs blocking automatic order approval and then approve orders manually. Change the order "Comm Flag" to 'D' – Direct Interface on repair and warranty orders that should not wait for the weekly interface to SAP.
Projected Inventory Above XS Point (R80)	Triggered when the projected inventory in an ordering sub-network goes above XS Point in the next two years and ordering constraints are prohibiting decrease or cancel recommendations. This will not trigger if the XS Point is exceeded due to an order maximized to a threshold above XS Point or due to an EOL buy order. This also will not trigger if the OR inventory position is above XS point and there are no future orders causing projected inventory to go above XS point. XelusPlan order recommendation constraints may be one or more of the following: • Decrease freeze days • "Good as it Gets" dates • Minimum, module or EOQ order quantities This RR may also trigger if allocation orders are due into a location that will put that location into an excess position.	Drill-Down on the RR notes to see the location projected to exceed excess point. Validate that the XS Point quantities are a reasonable level above Maximum Inventory Level (MAX). Determine what ordering constraint is prohibiting the system from canceling or reducing orders. Investigate whether an exception to this constraint can be negotiated with the appropriate supplier. (e.g. change orders inside freeze, violate minimum order change parameters or violate MOQ, module or EOQ quantities). If the excess point is being exceeded by an allocation order due in, validate the stocking levels on the part are correct. If all data is correct and no changes can be made, approve the part and the remaining above XS

		point condition.
Newbuy Supply Constraint Over Consumed (R81)	Triggered a when the advertised supply constraint for ESZ, GARS or H2H (internal demand) is over consumed in a period. The proposed order quantity that is not yet confirmed will be compared to the proposed available supply to evaluate this RR. The system will never recommend this condition. It may be caused by updates to supply constraint quantities or manual changes made to orders during order approval or manual order entry.	Drill-Down on the RR notes to see the ordering location and order type where supply constraint is over consumed. Manually decrease or cancel the orders so as to not over consume the supply constraint.
Projected Inventory Below MOP (R83)	Triggered when the projected inventory in a network ordering sub-network goes below Must Order Point (MOP) in the next two years because one or more ordering constraints are preventing orders from being recommended. XelusPlan order recommendation constraints may be one or more of the following: • Increase freeze days • No valid part/sources for any order type that is enabled. • No order types enabled that have valid part/sources. • “Good as it Gets” dates • Capacity or Supply constraints All three of the following conditions must be satisfied for this review reason to be triggered: • The planning inventory, calculated over the prime/alternate chain, must be greater than zero. • The sum of the Total Requirements (T1) two years out must be greater than zero. • The projected inventory calculated over the prime/ alternate chain for the ordering sub-network must be below MOP on one or more days in the future.	Verify the executive forecast setting is correct. Determine what ordering constraint is prohibiting the system from placing orders sooner. Investigate whether an exception to this constraint can be negotiated with the appropriate supplier. (e.g. reschedule orders inside freeze or GAIG date) If there are no orders or order recommendations investigate new sources of supply.

Appendix D Review reason sets

Rule	Consequent	Antecedent	Support %	Confidence %	Rule Support %
1	R24	R83	1.62	48.46	0.78
2	R67	R22 and R3	6.08	43.86	2.67
3	R22	R51 and R67	2.19	43.17	0.95
4	R67	R51 and R22	2.53	37.47	0.95
5	R22	R51	7.90	32.01	2.53
6	R67	R22	32.76	28.54	9.35
7	R3	R22 and R67	9.35	28.53	2.67
8	R67	R51	7.90	27.77	2.19
9	R3	R67	38.83	26.38	10.25
10	R9	R4	1.08	26.29	0.28
11	R22	R3 and R67	10.25	26.03	2.67
12	R67	R3	41.35	24.77	10.25
13	R22	R67	38.83	24.08	9.35
14	R3	R22	32.76	18.56	6.08
15	R3	R4	1.08	16.62	0.18
16	R22	R3	41.35	14.71	6.08
17	R26	R83	1.62	10.38	0.17
18	R51	R22 and R67	9.35	10.13	0.95
19	R22	R61	1.89	9.45	0.18
20	R51	R83	1.62	8.06	0.13
21	R38	R4	1.08	7.86	0.09

22	R51	R22	32.76	7.71	2.53
23	R25	R4	1.08	7.47	0.08
24	R38	R83	1.62	7.29	0.12
25	R42	R4	1.08	6.96	0.08
26	R83	R4	1.08	6.44	0.07
27	R22	R83	1.62	6.09	0.10
28	R80	R4	1.08	5.93	0.06
29	R51	R67	38.83	5.65	2.19
30	R4	R83	1.62	4.29	0.07
31	R67	R61	1.89	4.25	0.08
32	R24	R4	1.08	4.00	0.04
33	R81	R83	1.62	3.86	0.06
34	R42	R83	1.62	3.60	0.06
35	R42	R51	7.90	2.55	0.20
36	R17	R83	1.62	2.49	0.04
37	R22	R4	1.08	2.45	0.03
38	R51	R61	1.89	2.20	0.04
39	R67	R83	1.62	1.80	0.03
40	R14	R4	1.08	1.68	0.02
41	R83	R51	7.90	1.65	0.13
42	R26	R4	1.08	1.55	0.02
43	R80	R83	1.62	1.54	0.03
44	R42	R51 and R22	2.53	1.54	0.04

45	R3	R51 and R67	2.19	1.39	0.03
46	R9	R51	7.90	1.39	0.11
47	R16	R83	1.62	1.37	0.02
48	R83	R51 and R22	2.53	1.37	0.04
49	R17	R4	1.08	1.16	0.01
50	R51	R4	1.08	1.16	0.01
51	R80	R51	7.90	1.11	0.09

Appendix E Model input

Input	Description
Actual on-hand balance	Quantity of parts actually in stock at this location.
Analyzer code	Identifier code of the analyzer linked to the action
Birth date	First appearance date of the part in IBM's database
Critical stock level	Quantity of a part which are reserved for a location's own demands.
Critical part indicator	Code which specifies whether a part is either systematically or manually specified to be in a critical situation regarding the need versus supply for a specific location
Division owner code	Code which represents the product platform associated with the part
End of service date	End of service date for the part
Excess stock level	Quantity of parts which are allowed to be held at this location. Stock above this level is considered excess.
First stock date	First stocking date of the part within one of IBM's locations
Last time buy indicator	Code which specifies whether or not a last time buy process was performed. A last time buy (order) is a formal commitment to purchase goods outside fo the (CI 144) transfer of sourcing responsibility activities.
Location hierarchy	Indicates the location planning hierarchy
Maximum stock level	Quantity of parts which are normally allowed to be in stock and above which cancellations of orders are allowed.
New buy order lead-time	New buy order lead-time
New buy order minimum order quantity	New buy order minimum quantity
New buy order part price	New buy order unit cost
On-hand repair balance	Number of service parts available at the repair location
On-hand warranty balance	Number of service parts available at the warranty location
Order review inventory position	Quantity of parts in stock used during the order review process. For service stock this attribute is on a aggregate level.
Part commodity name	The name of the part's class such as power unit
Part substitute type	Code which indicates what type of substitution is applicable for the part.
Part vitality	The vitality code of the service part. Ranges from 1 to 5, with 1 being vital to the machine and 5 optional accessories. Also known as criticality.

Part WAC	Part WAC value for the EMEA CB location, converted to euro if originally in other currency
Planning equivalent stock level	Quantity of a part which are supposed to be in stock in a hierarchy.
Policy safety stock	Policy safety stock value for the central buffer location
Repair order lead-time	Repair order lead-time
Repair order minimum order quantity	Repair order minimum quantity
Repair order part price	Repair order unit cost
Reorder point	Quantity of a part which, if stock becomes lower than this, triggers replenishment
Reorder point aggregated	Sum of the reorder points of all locations in a lower hierarchy level than the central buffer
Shelf life indicator	Indicates whether the part has a limited shelf life
Usage previous period	Sum of the service part quantity used in the previous four-week period
Warranty order lead-time	Warranty order lead-time
Warranty order minimum order quantity	Warranty order minimum quantity
Warranty order part price	Warranty order unit cost

Appendix F SPSS model types

All descriptions of the models are taken from the IBM SPSS Modeler 18.0 Modeling Nodes manual (IBM, 2016b).

Appendix F.1 C5.0

This node uses the C5.0 algorithm to build either a decision tree or a rule set. A C5.0 model works by splitting the sample based on the field that provides the maximum information gain. Each subsample defined by the first split is then split again, usually based on a different field, and the process repeats until the subsamples cannot be split any further. Finally, the lowest-level splits are reexamined, and those that do not contribute significantly to the value of the model are removed or pruned.

Note: The C5.0 node can predict only a categorical target. When analyzing data with categorical (nominal or ordinal) fields, the node is more likely to group categories together than versions of C5.0 prior to release 11.0.

C5.0 can produce two kinds of models. A decision tree is a straightforward description of the splits found by the algorithm. Each terminal (or "leaf") node describes a particular subset of the training data, and each case in the training data belongs to exactly one terminal node in the tree. In other words, exactly one prediction is possible for any particular data record presented to a decision tree.

In contrast, a rule set is a set of rules that tries to make predictions for individual records. Rule sets are derived from decision trees and, in a way, represent a simplified or distilled version of the information found in the decision tree. Rule sets can often retain most of the important information from a full decision tree but with a less complex model. Because of the way rule sets work, they do not have the same properties as decision trees. The most important difference is that with a rule set, more than one rule may apply for any particular record, or no rules at all may apply. If multiple rules apply, each rule gets a weighted "vote" based on the confidence associated with that rule, and the final prediction is decided by combining the weighted votes of all of the rules that apply to the record in question. If no rule applies, a default prediction is assigned to the record.

Example: A medical researcher has collected data about a set of patients, all of whom suffered from the same illness. During their course of treatment, each patient responded to one of five medications. You can use a C5.0 model, in conjunction with other nodes, to help find out which drug might be appropriate for a future patient with the same illness. Show me

Requirements: To train a C5.0 model, there must be one categorical (i.e., nominal or ordinal) Target field, and one or more Input fields of any type. Fields set to Both or None are ignored. Fields used in the model must have their types fully instantiated. A weight field can also be specified.

Strengths: C5.0 models are quite robust in the presence of problems such as missing data and large numbers of input fields. They usually do not require long training times to estimate. In addition, C5.0 models tend to be easier to understand than some other model types, since the rules derived from the model have a very straightforward interpretation. C5.0 also offers the powerful boosting method to increase accuracy of classification.

Appendix F.2 QUEST

QUEST—or Quick, Unbiased, Efficient Statistical Tree—is a binary classification method for building decision trees. A major motivation in its development was to reduce the processing time required for large C&R Tree analyses with either many variables or many cases. A second goal of QUEST was to reduce the tendency found in classification tree methods to favor inputs that allow more splits, that is, continuous (numeric range) input fields or those with many categories.

- QUEST uses a sequence of rules, based on significance tests, to evaluate the input fields at a node. For selection purposes, as little as a single test may need to be performed on each input at a node. Unlike C&R Tree, all splits are not examined, and unlike C&R Tree and CHAID, category combinations are not tested when evaluating an input field for selection. This speeds the analysis.
- Splits are determined by running quadratic discriminant analysis using the selected input on groups formed by the target categories. This method again results in a speed improvement over exhaustive search (C&R Tree) to determine the optimal split.

Requirements: Input fields can be continuous (numeric ranges), but the target field must be categorical. All splits are binary. Weight fields cannot be used. Any ordinal (ordered set) fields used in the model must have numeric storage (not string).

Strengths: Like CHAID, but unlike C&R Tree, QUEST uses statistical tests to decide whether or not an input field is used. It also separates the issues of input selection and splitting, applying different criteria to each. This contrasts with CHAID, in which the statistical test result that determines variable selection also produces the split. Similarly, C&R Tree employs the impurity-change measure to both select the input field and to determine the split.

Appendix F.3 Bayesian networks

The Bayesian Network node enables you to build a probability model by combining observed and recorded evidence with "common-sense" real-world knowledge to establish the likelihood of occurrences by using seemingly unlinked attributes. The node focuses on Tree Augmented Naïve Bayes (TAN) and Markov Blanket networks that are primarily used for classification.

Bayesian networks are used for making predictions in many varied situations; some examples are:

- Selecting loan opportunities with low default risk.
- Estimating when equipment will need service, parts, or replacement, based on sensor input and existing records.
- Resolving customer problems via online troubleshooting tools.
- Diagnosing and troubleshooting cellular telephone networks in real-time.
- Assessing the potential risks and rewards of research-and-development projects in order to focus resources on the best opportunities.

A Bayesian network is a graphical model that displays variables (often referred to as nodes) in a dataset and the probabilistic, or conditional, independencies between them. Causal relationships between nodes may be represented by a Bayesian network; however, the links in the network (also known as arcs) do not necessarily represent direct cause and effect. For example, a Bayesian network can be used to calculate the probability of a patient having a specific disease, given the presence or absence of certain symptoms and other relevant data, if the probabilistic independencies between symptoms and disease as displayed on the graph hold true. Networks are very robust where information is missing and make the best possible prediction using whatever information is present.

Requirements: Target fields must be categorical and can have a measurement level of Nominal, Ordinal, or Flag. Inputs can be fields of any type. Continuous (numeric range) input fields will be automatically binned; however, if the distribution is skewed, you may obtain better results by manually binning the fields using a Binning node before the Bayesian Network node. For example, use Optimal Binning where the Supervisor field is the same as the Bayesian Network node Target field.

Example: An analyst for a bank wants to be able to predict customers, or potential customers, who are likely to default on their loan repayments. You can use a Bayesian network model to identify the characteristics of customers most likely to default, and build several different types of model to establish which is the best at predicting potential defaulters.

Appendix F.4 Neural networks

A neural network is a simplified model of the way the human brain processes information. It works by simulating a large number of interconnected processing units that resemble abstract versions of neurons.

The processing units are arranged in layers. There are typically three parts in a neural network: an input layer, with units representing the input fields; one or more hidden layers; and an output layer, with a unit or units representing the target field(s). The units are connected with varying connection strengths (or weights). Input data are presented to the first layer, and values are propagated from each neuron to every neuron in the next layer. Eventually, a result is delivered from the output layer.

The network learns by examining individual records, generating a prediction for each record, and making adjustments to the weights whenever it makes an incorrect prediction. This process is repeated many times, and the network continues to improve its predictions until one or more of the stopping criteria have been met.

Initially, all weights are random, and the answers that come out of the net are probably nonsensical. The network learns through training. Examples for which the output is known are repeatedly presented to the network, and the answers it gives are compared to the known outcomes. Information from this comparison is passed back through the network, gradually changing the weights. As training progresses, the network becomes increasingly accurate in replicating the known outcomes. Once trained, the network can be applied to future cases where the outcome is unknown

Appendix F.5 CHAID

CHAID, or Chi-squared Automatic Interaction Detection, is a classification method for building decision trees by using chi-square statistics to identify optimal splits.

CHAID first examines the crosstabulations between each of the input fields and the outcome, and tests for significance using a chi-square independence test. If more than one of these relations is statistically significant, CHAID will select the input field that is the most significant (smallest p value). If an input has more than two categories, these are compared, and categories that show no differences in the outcome are collapsed together. This is done by successively joining the pair of categories showing the least significant difference. This category-merging process stops when all remaining categories differ at the specified testing level. For nominal input fields, any categories can be merged; for an ordinal set, only contiguous categories can be merged.

Exhaustive CHAID is a modification of CHAID that does a more thorough job of examining all possible splits for each predictor but takes longer to compute.

Requirements: Target and input fields can be continuous or categorical; nodes can be split into two or more subgroups at each level. Any ordinal fields used in the model must have numeric storage (not string).

Strengths: Unlike the C&R Tree and QUEST nodes, CHAID can generate nonbinary trees, meaning that some splits have more than two branches. It therefore tends to create a wider tree than the binary growing methods. CHAID works for all types of inputs, and it accepts both case weights and frequency variables.

Appendix F.6 C&R

The Classification and Regression (C&R) Tree node is a tree-based classification and prediction method. Similar to C5.0, this method uses recursive partitioning to split the training records into segments with similar output field values. The C&R Tree node starts by examining the input fields to find the best split, measured by the reduction in an impurity index that results from the split. The split defines two subgroups, each of which is subsequently split into two more subgroups, and so on, until one of the stopping criteria is triggered. All splits are binary (only two subgroups).

Pruning. C&R Trees give you the option to first grow the tree and then prune based on a cost-complexity algorithm that adjusts the risk estimate based on the number of terminal nodes. This method, which allows the tree to grow large before pruning based on more complex criteria, may result in smaller trees with better cross-validation properties. Increasing the number of terminal nodes generally reduces the risk for the current (training) data, but the actual risk may be higher when the model is generalized to unseen data. In an extreme case, suppose you have a separate terminal node for each record in the training set. The risk estimate would be 0%, since every record falls into its own node, but the risk of misclassification for unseen (testing) data would almost certainly be greater than 0. The cost-complexity measure attempts to compensate for this.

Example: A cable TV company has commissioned a marketing study to determine which customers would buy a subscription to an interactive news service via cable. Using the data from the study, you can create a stream in which the target field is the intent to buy the subscription and the predictor fields include age, sex, education, income category, hours spent watching television each day, and number of children. By applying a C&R Tree node to the stream, you will be able to predict and classify the responses to get the highest response rate for your campaign.

Requirements: To train a C&R Tree model, you need one or more Input fields and exactly one Target field. Target and input fields can be continuous (numeric range) or categorical. Fields set to Both or None are ignored. Fields used in the model must have their types fully instantiated, and any ordinal (ordered set) fields used in the model must have numeric storage (not string).

Strengths: C&R Tree models are quite robust in the presence of problems such as missing data and large numbers of fields. They usually do not require long training times to estimate. In addition, C&R Tree models tend to be easier to understand than some other model types--the rules derived from the model have a very straightforward interpretation. Unlike C5.0, C&R Tree can accommodate continuous as well as categorical output fields.

Appendix F.7 Random trees

The Random Trees node is a tree-based classification and prediction method that is built on Classification and Regression Tree methodology. As with C&R Tree, this prediction method uses recursive partitioning to split the training records into segments with similar output field values. The node starts by examining the input fields available to it to find the best split, which is measured by the reduction in an impurity index that results from the split. The split defines two subgroups, each of which is then split into two more subgroups, and so on, until one of the stopping criteria is triggered. All splits are binary (only two subgroups).

Random Trees adds two features compared to C&R Tree:

- The first feature is bagging where replicas of the training dataset are created by sampling with replacement from the original dataset. This action creates bootstrap samples that are of equal size to the original dataset, after which a component model is built on each replica. Together these component models form an ensemble model.
- The second feature is that, at each split of the tree, only a sampling of the input fields is considered for the impurity measure.

Requirements: To train a Random Trees model, you need one or more Input fields and one Target field. Target and input fields can be continuous (numeric range) or categorical. Fields that are set to either Both or None are ignored. Fields that are used in the model must have their types fully instantiated, and any ordinal (ordered set) fields that are used in the model must have numeric storage (not string). If necessary, the Reclassify node can be used to convert them.

Strengths: Random Trees models are robust when you are dealing with large data sets and numbers of fields. Due to the use of bagging and field sampling, they are much less prone to overfitting and thus the results that are seen in testing are more likely to be repeated when you use new data.

Appendix G SPSS Model

Here we provide an overview of the nodes used in SPSS Modeler 18 to generate the different machine-learning models. We have included the k-folding validation model used for the C5.0 model, the validation for the other six models was done using a similar setup.

Appendix G.1 Model overview

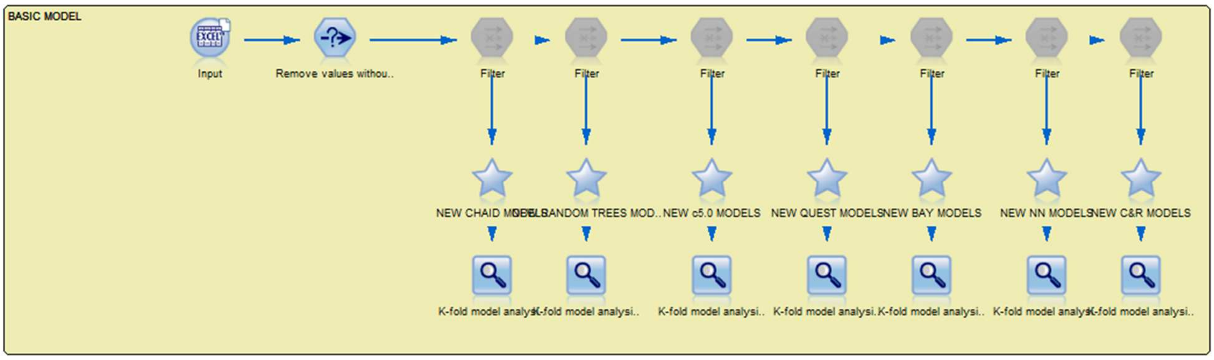

Appendix G.2 C5.0 K-folding validation model

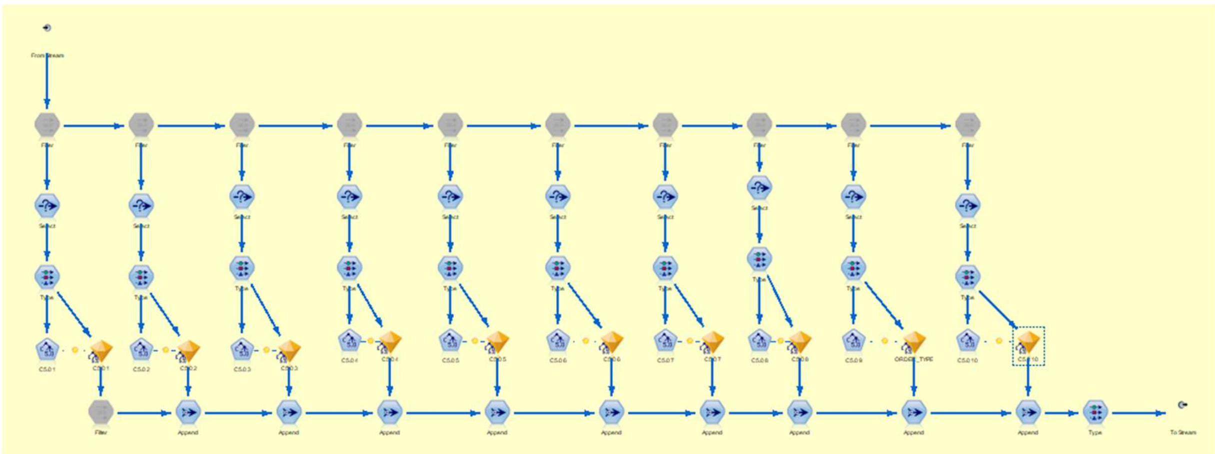

Appendix H C5.0 first model rule set

PART_WAC_CONV <= 81.490 [Mode: New Buy] => New Buy

PART_WAC_CONV > 81.490 [Mode: New Buy]

ANALYZER_COD in ["02M" "07Q" "08F" "08Z" "09E" "09J" "09Y" "0G7" "0W9" "14" "23" "4" "99"] [Mode: New Buy] => New Buy

ANALYZER_COD in ["03D" "07A" "07F" "07K"] [Mode: -] => -

ANALYZER_COD in ["04A"] [Mode: New Buy]

N_PART_PRICE <= 120 [Mode: -] => -

N_PART_PRICE > 120 [Mode: New Buy] => New Buy

ANALYZER_COD in ["05C" "0C9" "0I3" "0J7" "0K7" "0U7" "10"] [Mode: Internal Demand] => Internal Demand

ANALYZER_COD in ["07B" "07C" "07D" "07E" "07H" "07J" "07N" "07V" "07Y" "09A" "09D" "09G" "09K" "0D7" "12" "8"] [Mode: New Buy] => New Buy

ANALYZER_COD in ["07I"] [Mode: Repair] => Repair

ANALYZER_COD in ["08L"] [Mode: New Buy]

USAGE_PERPERIOD_Sum <= -109 [Mode: -] => -

USAGE_PERPERIOD_Sum > -109 [Mode: New Buy] => New Buy

ANALYZER_COD in ["09M"] [Mode: ESZ]

PART_SUB_TYPE in [""] [Mode: ESZ]

PART_WAC_CONV <= 126.960 [Mode: ESZ] => ESZ

PART_WAC_CONV > 126.960 [Mode: Allocation] => Allocation

PART_SUB_TYPE in ["A" "AE" "C" "CE" "E" "EE" "G1B" "GE"] [Mode: ESZ] => ESZ

PART_SUB_TYPE in ["G1D"] [Mode: -] => -

ANALYZER_COD in ["0C6"] [Mode: Warranty]

PART_WAC_CONV <= 155.520 [Mode:] =>

PART_WAC_CONV > 155.520 [Mode: Warranty] => Warranty

ANALYZER_COD in ["0G3"] [Mode: Repair]

PART_SUB_TYPE in [""] [Mode: Repair]

LOC_HIER in [""] [Mode: -]

ORDREV_INVENT_POS <= 22 [Mode: Internal Demand] => Internal Demand

ORDREV_INVENT_POS > 22 [Mode: -] => -

LOC_HIER in ["A1"] [Mode: Allocation]

PART_WAC_CONV <= 207.780 [Mode:] =>

PART_WAC_CONV > 207.780 [Mode: Allocation] => Allocation

LOC_HIER in ["A2"] [Mode: Repair]

EXC_STOCK_LVL <= 52 [Mode: Repair] => Repair

EXC_STOCK_LVL > 52 [Mode: Internal Demand] => Internal Demand

LOC_HIER in ["A4" "B1" "L1"] [Mode: Repair] => Repair

PART_SUB_TYPE in ["A" "CE" "E" "EE" "G1B" "G1D"] [Mode: Repair] => Repair

PART_SUB_TYPE in ["AE"] [Mode: -] => -

PART_SUB_TYPE in ["C" "GE"] [Mode: New Buy] => New Buy

ANALYZER_COD in ["0M7"] [Mode: Warranty] => Warranty

ANALYZER_COD in ["0V7"] [Mode: New Buy]

PART_SUB_TYPE in [""] [Mode: New Buy]

N_PART_PRICE <= 191 [Mode: New Buy] => New Buy

N_PART_PRICE > 191 [Mode: Internal Demand] => Internal Demand

PART_SUB_TYPE in ["A" "C" "CE" "E" "EE" "G1D"] [Mode: New Buy] => New Buy

PART_SUB_TYPE in ["AE"] [Mode: New Buy] => New Buy

PART_SUB_TYPE in ["G1B"] [Mode: Allocation] => Allocation

PART_SUB_TYPE in ["GE"] [Mode: Repair] => Repair

ANALYZER_COD in ["0W7"] [Mode: Internal Demand]

PART_COMMOD_NAME in [""] [Mode: Internal Demand] => Internal Demand

PART_COMMOD_NAME in ["A COVER" "ADAPTER" "AMD" "AOP" "BASE COVER" "BATTERY" "BEZEL" "CABLE" "CABLES" "DONGLE CBL" "ECAT" "EMECHS" "FOOT" "HDDFAN BKT" "I/O-CARDS" "I/O-DEVICES" "KBD US BL" "KEYBRDS" "KYBD FR" "LCD BEZEL" "MECH ASM" "MECHANICAL" "MECHS" "MONITOR" "NB KYB" "NB_KYB" "ODD BEZEL" "OEM-BOX" "OTH-EQUIP" "PCI CARD" "PLNRWNTPM" "PLNTPMNOK" "PLRNOKNTPM" "POWER" "PWR_SUPPLY" "S" "SSD TRAY" "STAND" "STORAGE" "STRIP COVR" "THERMAL"] [Mode: Internal Demand] => Internal Demand

PART_COMMOD_NAME in ["OTH-BOMS"] [Mode: Repair] => Repair

ANALYZER_COD in ["21"] [Mode: -]

PART_SUB_TYPE in [""] [Mode: -] => -

PART_SUB_TYPE in ["A" "AE" "E" "EE" "G1B" "G1D"] [Mode: -] => -

PART_SUB_TYPE in ["C"] [Mode: Warranty] => Warranty

PART_SUB_TYPE in ["CE"] [Mode: Internal Demand] => Internal Demand

PART_SUB_TYPE in ["GE"] [Mode: Internal Demand]

USAGE_PERPERIOD_Sum <= 7 [Mode: Internal Demand]

CRITICAL_STOCK_LVL <= 1 [Mode: New Buy] => New Buy

CRITICAL_STOCK_LVL > 1 [Mode: Internal Demand] => Internal Demand

USAGE_PERPERIOD_Sum > 7 [Mode: New Buy] => New Buy

ANALYZER_COD in ["24"] [Mode: New Buy]

PART_SUB_TYPE in [""] [Mode: New Buy]

PART_COMMOD_NAME in [""] [Mode: New Buy]

DIV_OWNER_COD in ["13" "26" "44" "48" "4S" "71" "75" "9R" "LN" "LQ" "LT" "LU" "LX" "MN" "MP" "MT" "Z1"] [Mode: New Buy] => New Buy

DIV_OWNER_COD in ["2D"] [Mode: New Buy]

ORDREV_INVENT_POS <= 11 [Mode: New Buy]

ACTUAL_ONHAND_BAL <= -1 [Mode: New Buy] => New Buy

ACTUAL_ONHAND_BAL > -1 [Mode: Internal Demand] => Internal Demand

ORDREV_INVENT_POS > 11 [Mode: -]

W_LEAD_TIME <= 58 [Mode: Allocation] => Allocation

W_LEAD_TIME > 58 [Mode: -]

PART_WAC_CONV <= 267.040 [Mode: ESZ] => ESZ

PART_WAC_CONV > 267.040 [Mode: -] => -

DIV_OWNER_COD in ["2K"] [Mode: New Buy]

ONHAND_REPAIR <= 213 [Mode: Repair] => Repair

ONHAND_REPAIR > 213 [Mode: New Buy] => New Buy

PART_COMMOD_NAME in ["A COVER" "ADAPTER" "AMD" "AOP" "BASE COVER" "BATTERY" "BEZEL" "CABLE" "CABLES" "DONGLE CBL" "ECAT" "EMECHS" "FOOT" "HDDFAN BKT" "I/O-CARDS" "I/O-DEVICES" "KBD US BL" "KEYBRDS" "KYBD FR" "LCD BEZEL" "MECH ASM" "MECHANICAL" "MONITOR" "NB KYB" "NB_KYB" "ODD BEZEL" "OEM-BOX" "OTH-EQUIP" "PCI CARD" "PLNRWNTPM" "PLNTPMNOK" "PLRNOKNTPM" "PWR_SUPPLY" "S" "SSD TRAY" "STAND" "STRIP COVR" "THERMAL"] [Mode: New Buy] => New Buy

PART_COMMOD_NAME in ["MECHS"] [Mode:] =>

PART_COMMOD_NAME in ["OTH-BOMS"] [Mode: New Buy]

ACTUAL_ONHAND_BAL <= 5 [Mode: New Buy] => New Buy

ACTUAL_ONHAND_BAL > 5 [Mode: Internal Demand] => Internal Demand

PART_COMMOD_NAME in ["POWER"] [Mode: New Buy] => New Buy

PART_COMMOD_NAME in ["STORAGE"] [Mode: Warranty]

EXC_STOCK_LVL <= 476 [Mode: Warranty] => Warranty

EXC_STOCK_LVL > 476 [Mode:] =>

PART_SUB_TYPE in ["A" "C" "E" "G1B" "G1D"] [Mode: New Buy] => New Buy

PART_SUB_TYPE in ["AE"] [Mode: Warranty] => Warranty

PART_SUB_TYPE in ["CE" "EE"] [Mode:] =>

PART_SUB_TYPE in ["GE"] [Mode: Allocation]

PART_VIT in [1.000] [Mode: Allocation] => Allocation

PART_VIT in [2.000] [Mode: -] => -

PART_VIT in [3.000 4.000 5.000] [Mode: Allocation] => Allocation

ANALYZER_COD in ["26"] [Mode: New Buy]

SHELF_LIFE_IND = Y [Mode: Repair] => Repair

SHELF_LIFE_IND = [Mode: New Buy]

R_PART_PRICE <= 178 [Mode: New Buy] => New Buy

R_PART_PRICE > 178 [Mode: -] => -

ANALYZER_COD in ["5"] [Mode: New Buy]

ORDREV_INVENT_POS <= 6 [Mode: New Buy]

PART_COMMOD_NAME in [""] [Mode: New Buy] => New Buy

PART_COMMOD_NAME in ["A COVER" "ADAPTER" "AMD" "AOP" "BASE COVER" "BATTERY" "BEZEL" "CABLE" "CABLES" "DONGLE CBL" "ECAT" "EMECHS" "FOOT" "HDDFAN BKT" "I/O-CARDS" "I/O-DEVICES" "KBD US BL" "KEYBRDS" "KYBD FR" "LCD BEZEL" "MECH ASM" "MECHANICAL" "MECHS" "MONITOR" "NB KYB" "NB_KYB" "ODD BEZEL" "OEM-BOX" "OTH-EQUIP" "PCI CARD" "PLNRWNTPM" "PLNTPMNOK" "PLRNOKNTPM" "POWER" "PWR_SUPPLY" "S" "SSD TRAY" "STAND" "STORAGE" "STRIP COVR" "THERMAL"] [Mode: New Buy] => New Buy

PART_COMMOD_NAME in ["OTH-BOMS"] [Mode: -] => -

ORDREV_INVENT_POS > 6 [Mode: Repair]

R_PART_PRICE <= 69 [Mode: Repair]

R_PART_PRICE <= 54 [Mode: Repair] => Repair

R_PART_PRICE > 54 [Mode: Internal Demand] => Internal Demand

R_PART_PRICE > 69 [Mode: Repair]

PLNEQ_SAFSTOCK_LVL <= 2 [Mode: New Buy] => New Buy

PLNEQ_SAFSTOCK_LVL > 2 [Mode: Repair] => Repair

Appendix I Predictor importance

In these tables, the sum of the indicator performances per model will not always equal one as SPSS modeler shows a maximum of ten predictor importance values per model. SPSS Modeler was unable to produce the predictor importance values of model 1.

Predictor	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Actual on-hand balance	N/A	0.07	0	0.02	0	1
Analyzer code	N/A	0	0.59	0.62	1	0
Birth date	N/A	0	0	0	0	0
Critical part indicator	N/A	0	0	0	0	0
Critical stock level	N/A	0.07	0	0.02	0	0
Division owner code	N/A	0	0	0	0	0
End of service date	N/A	0	0	0	0	0
Excess stock level	N/A	0	0	0	0	0
First stock date	N/A	0	0	0	0	0
Last time buy indicator	N/A	0.07	0	0.02	0	0
Location hierarchy	N/A	0.07	0	0	0	0
Maximum stock level	N/A	0.07	0	0	0	0
New buy order lead-time	N/A	0	0	0	0	0
New buy order minimum order quantity	N/A	0	0	0	0	0
New buy order part price	N/A	0	0	0	0	0
On-hand repair balance	N/A	0	0	0	0	0
On-hand warranty balance	N/A	0	0	0	0	0
Order review inventory position	N/A	0.07	0	0	0	0
Part commodity name	N/A	0.07	0	0	0	0
Part substitute type	N/A	0	0.41	0	0	0
Part vitality	N/A	0.07	0	0	0	0
Part WAC	N/A	0	0	0	0	0
Planning equivalent stock level	N/A	0.07	0	0.02	0	0
Policy safety stock	N/A	0	0	0	0	0
Reorder point	N/A	0	0	0.02	0	0
Reorder point aggregated	N/A	0	0	0	0	0
Repair order lead-time	N/A	0	0	0	0	0
Repair order minimum order quantity	N/A	0	0	0	0	0
Repair order part price	N/A	0	0	0.02	0	0
Shelf life indicator	N/A	0	0	0.02	0	0
Usage previous period	N/A	0	0	0	0	0
Warranty order lead-time	N/A	0.07	0	0.02	0	0

Warranty order minimum order quantity	N/A	0	0	0	0	0
Warranty order part price	N/A	0	0	0.02	0	0

Predictor	Model 7	Model 8	Model 9	Model 10	Average
Actual on-hand balance	0	0.07	0.03	0	0.13
Analyzer code	0.28	0.3	0.52	0.92	0.47
Birth date	0	0	0	0	0.00
Critical part indicator	0	0	0	0	0.00
Critical stock level	0	0.07	0	0	0.02
Division owner code	0.01	0	0	0	0.00
End of service date	0	0	0	0	0.00
Excess stock level	0.01	0	0.03	0	0.00
First stock date	0	0	0	0	0.00
Last time buy indicator	0	0.07	0.03	0	0.02
Location hierarchy	0	0	0	0	0.01
Maximum stock level	0	0	0	0	0.01
New buy order lead-time	0.01	0	0	0	0.00
New buy order minimum order quantity	0	0.07	0	0	0.01
New buy order part price	0	0	0	0	0.00
On-hand repair balance	0	0	0	0	0.00
On-hand warranty balance	0	0	0	0	0.00
Order review inventory position	0.08	0.07	0.03	0	0.03
Part commodity name	0.15	0	0	0	0.02
Part substitute type	0	0	0	0	0.05
Part vitality	0	0	0	0	0.01
Part WAC	0.35	0.07	0.12	0.08	0.07
Planning equivalent stock level	0	0	0	0	0.01
Policy safety stock	0	0	0	0	0.00
Reorder point	0	0	0	0	0.00
Reorder point aggregated	0	0	0	0	0.00
Repair order lead-time	0	0	0.03	0	0.00
Repair order minimum order quantity	0	0	0	0	0.00
Repair order part price	0.01	0	0.03	0	0.01
Shelf life indicator	0	0.07	0	0	0.01
Usage previous period	0.01	0.07	0	0	0.01
Warranty order lead-time	0	0.07	0.03	0	0.02
Warranty order minimum order quantity	0	0	0	0	0.00
Warranty order part price	0.01	0	0.03	0	0.01

Appendix J Frequency table of model confidence values

Confidence level	1	2	3	4	5	6	7	8	9	10
0.05	0	0	0	0	0	0	0	0	0	0
0.1	0	0	0	0	0	0	0	0	0	0
0.15	0	0	0	0	0	0	0	0	0	0
0.2	0	0	0	0	0	0	0	0	0	0
0.25	0	0	0	0	0	0	0	0	0	0
0.3	0	0	0	0	1	0	0	0	0	0
0.35	1	2	0	0	0	0	0	1	1	0
0.4	0	0	0	0	0	0	0	0	4	0
0.45	0	1	2	0	0	0	0	1	0	0
0.5	2	4	5	2	1	5	6	1	2	5
0.55	0	3	0	0	0	0	0	0	0	0
0.6	0	2	0	3	2	5	3	3	3	2
0.65	0	0	1	0	0	0	0	1	0	0
0.7	0	2	2	4	9	1	9	10	5	8
0.75	36	37	37	45	39	33	30	37	33	50
0.8	3	1	0	0	0	6	3	0	0	0
0.85	2	0	0	3	2	2	0	2	0	2
0.9	0	0	3	0	0	0	0	0	1	0
0.95	0	1	0	0	0	0	1	0	0	0
1	5	3	3	6	5	5	5	4	6	6

Appendix K Model performance per category

Here the results of the model are shown broken down per category. Note, that the sum of number of cases per category might be less than the total number of cases, 579. This difference occurs because not every category encompasses all cases, e.g. a service part might have no repair lead-time.

Appendix K.1 Division owner

Division owner	Average confidence level	Number of cases	Correct cases
PSG-Netfinity	0.774	3	2
Costco Corp.	0.725	2	1
PSG	0.725	1	0
Telepix	0.688	1	1
RSS	0.795	44	29
RISC	0.708	32	6
Anixter	0.696	43	12
Tape systems	0.708	13	2
xSeries	0.738	2	1
HVP	0.850	17	8
Lenovo New	0.727	275	210
Lenovo EPG Server	0.638	13	6
Lenovo current	0.718	3	3
Lenovo New	0.723	14	13
Lenovo x86 products	0.680	111	33
MVS Networking	0.717	20	12
Netezza	1.000	1	0

Appendix K.2 New buy order lead-time

New buy lead-time duration (work days)	Average confidence level	Number of cases	Correct cases
<5	0.723	252	196
5-10	0.831	7	4
10-20	0.757	24	16
20-90	0.709	132	62
>90	0.755	71	41

Appendix K.3 Warranty order lead-time

Warranty order lead-time (work days)	Average confidence level	Number of cases	Correct cases
<5	N/A	0	0
5-10	0.667	9	1

10-20	0.711	73	30
20-90	0.677	94	27
>90	0.774	46	16

Appendix K.4 Repair order lead-time

Repair order lead-time (work days)	Average confidence level	Number of cases	Correct cases
<5	N/A	1	1
5-10	0.753	9	3
10-20	0.727	80	27
20-90	0.662	81	21
>90	0.734	35	7

Appendix K.5 Part vitality

Part vitality	Average confidence level	Number of cases	Correct cases
1	0.724	370	224
2	0.712	114	64
3	1.000	6	4
4	0.710	5	2
5	0.741	34	25

Appendix K.6 Order type

Order type	Average confidence level	Number of cases	Correct cases
Warranty	0.747	29	14
Repair	0.697	34	11
New buy	0.723	440	292
ESZ	0.922	9	4
NBO	N/A	10	0
Internal demand	0.743	37	10

Appendix K.7 WAC

WAC (euro)	Average confidence level	Number of cases	Correct cases
<=1	0.728	20	10
1-10	0.727	99	74
10-100	0.724	241	174
100-1000	0.720	198	72
>1000	0.693	39	10

Appendix L Performance measurements

Appendix L.1 Scoring of all 161 cases

Action number	Exponential scoring average	Exponential scoring minimum	Quadratic scoring average	Quadratic scoring minimum	Linear scoring average	Linear scoring minimum	Record count
1	3.104	2.876	4.541	4.229	2.130	2.056	5.000
2	2.969	2.917	4.360	4.287	2.088	2.070	4.000
3	4.207	4.148	5.937	5.869	2.437	2.423	3.000
4	4.246	4.206	5.983	5.937	2.446	2.437	3.000
5	0.000	0.000	0.000	0.000	0.000	0.000	1.000
6	2.114	1.509	3.027	1.993	1.725	1.412	3.000
7	1.875	1.875	2.652	2.652	1.629	1.629	2.000
8	0.000	0.000	0.000	0.000	0.000	0.000	2.000
9	0.000	0.000	0.000	0.000	0.000	0.000	3.000
10	1.702	1.696	2.347	2.335	1.532	1.528	2.000
11	0.305	0.000	0.360	0.000	0.300	0.000	4.000
12	1.822	1.822	2.560	2.560	1.600	1.600	3.000
13	3.388	1.822	4.660	2.560	2.100	1.600	2.000
14	0.000	0.000	0.000	0.000	0.000	0.000	4.000
15	0.000	0.000	0.000	0.000	0.000	0.000	3.000
16	2.197	0.000	3.228	0.000	1.556	0.000	4.000
17	2.197	0.000	3.228	0.000	1.556	0.000	4.000
18	0.000	0.000	0.000	0.000	0.000	0.000	3.000
19	1.948	1.948	2.778	2.778	1.667	1.667	5.000
20	1.948	1.948	2.778	2.778	1.667	1.667	4.000
21	0.553	0.000	0.605	0.000	0.550	0.000	2.000
22	7.389	7.389	9.000	9.000	3.000	3.000	2.000
23	0.769	0.000	0.871	0.000	0.762	0.000	3.000
24	1.649	1.649	2.250	2.250	1.500	1.500	4.000
25	1.512	1.091	1.982	1.181	1.391	1.087	3.000
26	1.154	1.154	1.306	1.306	1.143	1.143	3.000
27	40.793	20.086	21.400	16.000	4.600	4.000	5.000
28	1.355	1.314	1.700	1.620	1.303	1.273	3.000
29	1.303	1.234	1.598	1.465	1.263	1.211	2.000
30	1.199	1.199	1.397	1.397	1.182	1.182	1.000
31	0.000	0.000	0.000	0.000	0.000	0.000	2.000
32	1.025	1.025	1.049	1.049	1.024	1.024	3.000
33	1.199	1.199	1.397	1.397	1.182	1.182	4.000
34	3.903	3.903	5.578	5.578	2.362	2.362	1.000
35	0.000	0.000	0.000	0.000	0.000	0.000	4.000
36	0.000	0.000	0.000	0.000	0.000	0.000	1.000

37	0.000	0.000	0.000	0.000	0.000	0.000	2.000
38	0.000	0.000	0.000	0.000	0.000	0.000	4.000
39	0.376	0.000	0.419	0.000	0.374	0.000	3.000
40	0.000	0.000	0.000	0.000	0.000	0.000	5.000
41	4.318	4.318	6.065	6.065	2.463	2.463	5.000
42	3.136	3.136	4.592	4.592	2.143	2.143	3.000
43	0.000	0.000	148.278	136.111	-12.167	-12.667	4.000
44	2.718	2.718	4.000	4.000	2.000	2.000	3.000
45	4.860	4.698	6.659	6.488	2.580	2.547	4.000
46	4.860	4.698	6.659	6.488	2.580	2.547	4.000
47	0.000	0.000	0.000	0.000	0.000	0.000	3.000
48	5.294	5.294	7.111	7.111	2.667	2.667	5.000
49	1.105	1.105	1.210	1.210	1.100	1.100	4.000
50	0.000	0.000	0.000	0.000	0.000	0.000	2.000
51	0.000	0.000	0.000	0.000	0.000	0.000	3.000
52	0.000	0.000	0.000	0.000	0.000	0.000	4.000
53	0.000	0.000	0.000	0.000	0.000	0.000	3.000
54	0.000	0.000	0.000	0.000	0.000	0.000	4.000
55	0.000	0.000	0.000	0.000	0.000	0.000	3.000
56	0.000	0.000	0.000	0.000	0.000	0.000	2.000
57	0.000	0.000	0.000	0.000	0.000	0.000	3.000
58	0.000	0.000	0.000	0.000	0.000	0.000	2.000
59	0.000	0.000	0.000	0.000	0.000	0.000	2.000
60	0.000	0.000	0.000	0.000	0.000	0.000	4.000
61	2.264	2.171	3.301	3.151	1.817	1.775	3.000
62	2.226	2.171	3.241	3.151	1.800	1.775	4.000
63	0.000	0.000	0.000	0.000	0.000	0.000	2.000
64	0.532	0.000	0.563	0.000	0.531	0.000	2.000
65	0.787	0.000	0.823	0.000	0.786	0.000	4.000
66	1.284	1.284	1.563	1.563	1.250	1.250	3.000
67	8.080	7.898	9.542	9.404	3.089	3.067	3.000
68	1.246	1.194	1.488	1.386	1.219	1.177	3.000
69	171.204	171.204	37.735	37.735	6.143	6.143	2.000
70	1.181	1.181	1.361	1.361	1.167	1.167	4.000
71	20.086	20.086	16.000	16.000	4.000	4.000	3.000
72	20.086	20.086	16.000	16.000	4.000	4.000	3.000
73	2.663	2.607	3.918	3.835	1.979	1.958	2.000
74	0.517	0.000	0.534	0.000	0.517	0.000	2.000
75	1.122	1.108	1.244	1.216	1.115	1.103	2.000
76	2.718	2.718	4.000	4.000	2.000	2.000	4.000
77	1.101	1.069	1.201	1.138	1.096	1.067	4.000

78	2.718	2.718	4.000	4.000	2.000	2.000	4.000
79	0.000	0.000	0.000	0.000	0.000	0.000	3.000
80	1.061	1.061	1.121	1.121	1.059	1.059	3.000
81	0.000	0.000	0.000	0.000	0.000	0.000	3.000
82	1.091	1.091	1.181	1.181	1.087	1.087	4.000
83	1.091	1.091	1.181	1.181	1.087	1.087	4.000
84	1.221	1.221	1.440	1.440	1.200	1.200	2.000
85	2.399	2.399	3.516	3.516	1.875	1.875	3.000
86	12.182	12.182	12.250	12.250	3.500	3.500	2.000
87	0.000	0.000	0.000	0.000	0.000	0.000	4.000
88	1.105	1.105	1.210	1.210	1.100	1.100	3.000
89	0.000	0.000	0.000	0.000	0.000	0.000	4.000
90	0.000	0.000	0.000	0.000	0.000	0.000	3.000
91	1.186	1.186	1.370	1.370	1.171	1.171	2.000
92	1.186	1.186	1.370	1.370	1.171	1.171	2.000
93	0.000	0.000	0.000	0.000	0.000	0.000	2.000
94	16.253	15.094	14.337	13.796	3.786	3.714	2.000
95	14.053	14.053	13.270	13.270	3.643	3.643	1.000
96	18.701	18.701	15.434	15.434	3.929	3.929	3.000
97	1.396	1.396	1.778	1.778	1.333	1.333	1.000
98	0.000	0.000	0.000	0.000	0.000	0.000	3.000
99	1.023	0.000	1.361	0.000	0.952	0.000	3.000
100	8.166	8.166	9.610	9.610	3.100	3.100	4.000
101	8.166	8.166	9.610	9.610	3.100	3.100	3.000
102	2.226	2.226	3.240	3.240	1.800	1.800	2.000
103	1.137	1.137	1.273	1.273	1.128	1.128	2.000
104	2.849	2.663	4.188	3.919	2.046	1.980	4.000
105	2.849	2.663	4.188	3.919	2.046	1.980	4.000
106	1.185	1.087	1.368	1.174	1.167	1.083	4.000
107	1.133	1.133	1.266	1.266	1.125	1.125	3.000
108	1.110	1.087	1.220	1.174	1.104	1.083	4.000
109	1.173	1.137	1.344	1.273	1.159	1.128	5.000
110	1.161	1.143	1.321	1.286	1.149	1.134	4.000
111	1.158	1.158	1.316	1.316	1.147	1.147	3.000
112	1.314	1.314	1.620	1.620	1.273	1.273	5.000
113	1.354	1.354	1.698	1.698	1.303	1.303	3.000
114	1.199	1.199	1.397	1.397	1.182	1.182	2.000
115	1.340	1.314	1.672	1.620	1.293	1.273	3.000
116	0.000	0.000	0.000	0.000	0.000	0.000	3.000
117	11.124	11.124	11.622	11.622	3.409	3.409	2.000
118	1.778	1.733	2.481	2.403	1.575	1.550	2.000

119	12.507	12.507	12.435	12.435	3.526	3.526	4.000
120	2.050	2.050	2.951	2.951	1.718	1.718	3.000
121	2.050	2.050	2.951	2.951	1.718	1.718	3.000
122	1.025	1.025	1.050	1.050	1.025	1.025	2.000
123	0.000	0.000	0.000	0.000	0.000	0.000	2.000
124	1.479	1.472	1.936	1.922	1.391	1.386	3.000
125	1.066	1.066	1.132	1.132	1.064	1.064	1.000
126	1.022	1.022	1.043	1.043	1.021	1.021	4.000
127	1.066	1.066	1.132	1.132	1.064	1.064	5.000
128	6.638	5.017	8.160	6.827	2.839	2.613	3.000
129	10.882	10.882	11.472	11.472	3.387	3.387	4.000
130	10.882	10.882	11.472	11.472	3.387	3.387	3.000
131	10.316	6.927	10.657	8.617	3.237	2.935	3.000
132	1.331	1.331	1.653	1.653	1.286	1.286	4.000
133	1.037	1.037	1.073	1.073	1.036	1.036	2.000
134	1.331	1.331	1.653	1.653	1.286	1.286	4.000
135	1.035	1.014	1.070	1.028	1.034	1.014	4.000
136	1.249	1.249	1.494	1.494	1.222	1.222	5.000
137	1.118	1.118	1.235	1.235	1.111	1.111	3.000
138	1.249	1.249	1.494	1.494	1.222	1.222	2.000
139	2.919	2.718	4.280	4.000	2.067	2.000	3.000
140	2.226	2.226	3.240	3.240	1.800	1.800	2.000
141	2.886	2.886	4.242	4.242	2.060	2.060	2.000
142	1.916	1.916	2.723	2.723	1.650	1.650	5.000
143	0.000	0.000	0.000	0.000	0.000	0.000	3.000
144	0.000	0.000	0.000	0.000	0.000	0.000	4.000
145	0.000	0.000	0.000	0.000	0.000	0.000	2.000
146	0.000	0.000	0.000	0.000	0.000	0.000	2.000
147	0.000	0.000	0.000	0.000	0.000	0.000	4.000
148	0.000	0.000	0.000	0.000	0.000	0.000	4.000
149	0.000	0.000	0.000	0.000	0.000	0.000	1.000
150	1.064	1.064	1.129	1.129	1.063	1.063	3.000
151	0.545	0.000	0.591	0.000	0.543	0.000	2.000
152	1.284	1.284	1.563	1.563	1.250	1.250	1.000
153	1.284	1.284	1.563	1.563	1.250	1.250	1.000
154	1.396	1.396	1.778	1.778	1.333	1.333	3.000
155	1.396	1.396	1.778	1.778	1.333	1.333	4.000
156	8.839	8.524	10.105	9.878	3.179	3.143	4.000
157	0.000	0.000	0.000	0.000	0.000	0.000	2.000
158	0.000	0.000	0.000	0.000	0.000	0.000	2.000
159	1.306	1.306	1.604	1.604	1.267	1.267	3.000

160	4761.673	4761.673	89.650	89.650	9.468	9.468	2.000
161	0.000	0.000	0.000	0.000	0.000	0.000	3.000

Appendix L.2 Boxplot of performance measurements

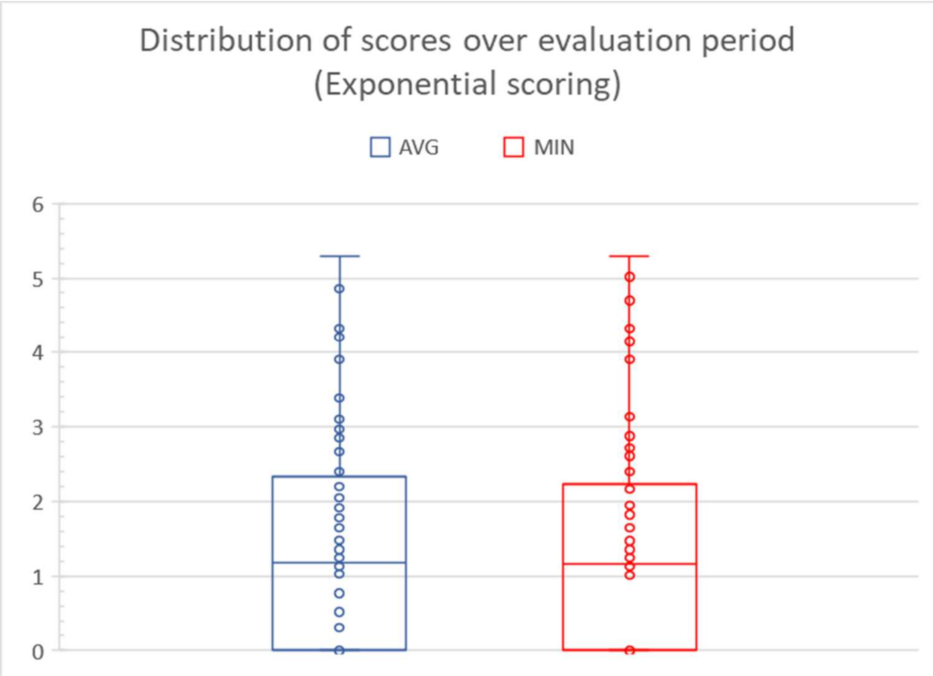

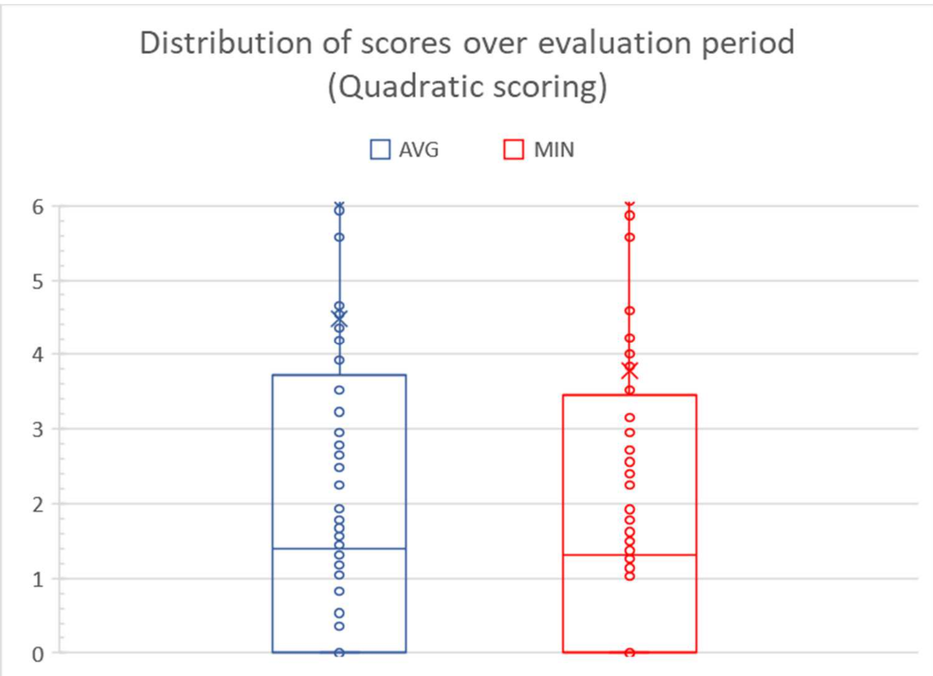

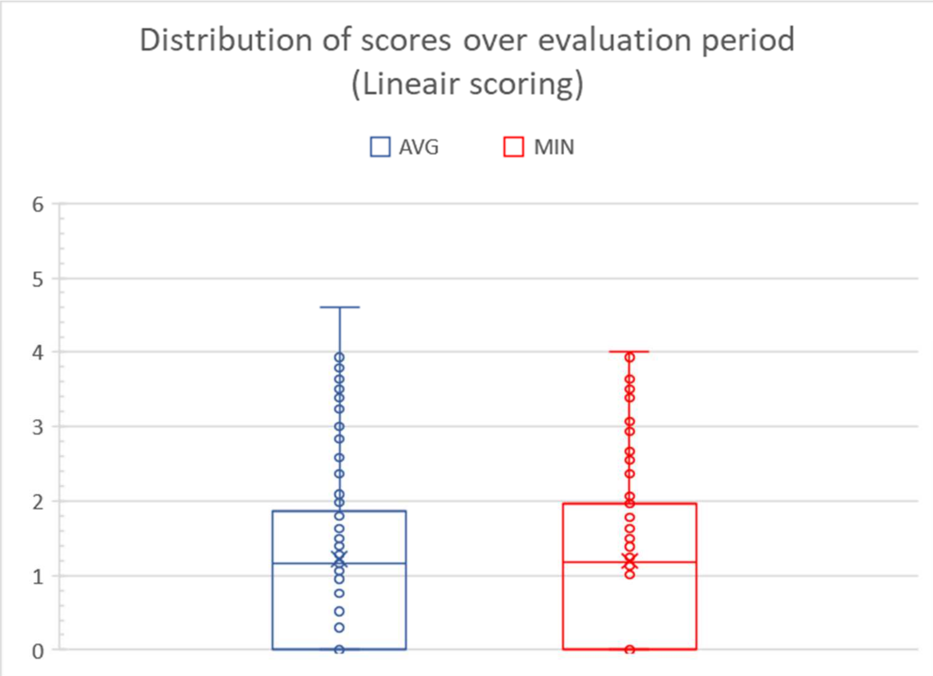
Distribution of scores over evaluation period
(Lineair scoring)
AVG
MIN
6
5
4
3
2
1
0