Sensing the Future

Designing Predictive Analytics with Sensor Technologies

Furtak, Simon; Avital, Michel; Ulslev Pedersen, Rasmus

Document Version
Final published version

Published in:
ECIS 2015 Proceedings

DOI:
10.18151/7217323

Publication date:
2015

License
Unspecified

Citation for published version (APA):
Furtak, S., Avital, M., & Ulslev Pedersen, R. (2015). Sensing the Future: Designing Predictive Analytics with Sensor Technologies. In J. Becker, J. vom Brocke, & M. De Marco (Eds.), *ECIS 2015 Proceedings* Article 54 Association for Information Systems. AIS Electronic Library (AISeL). https://doi.org/10.18151/7217323

Spring 5-29-2015

Sensing the Future: Designing Predictive Analytics with Sensor Technologies

Szymon Furtak
Copenhagen Business School, sf.itm@cbs.dk

Michel Avital
Copenhagen Business School, michel@avital.net

Rasmus U. Pedersen
Copenhagen Business School, rup.itm@cbs.dk

SENSING THE FUTURE: DESIGNING PREDICTIVE ANALYTICS WITH SENSOR TECHNOLOGIES

Completed Research

Szymon Furtak, Copenhagen Business School, Denmark, sf.itm@cbs.dk

Michel Avital, Copenhagen Business School, Denmark, michel@avital.net

Rasmus Ulslev Pedersen, Copenhagen Business School, Denmark, rup.itm@cbs.dk

Abstract

As digital technologies become prevalent and embedded in the environment, "smart" everyday objects like smart phone and smart homes have become part and parcel of the human enterprise. The ubiquity of smart objects that produce ever-growing streams of data presents both challenges and opportunities. In this paper, we argue that extending these data streams, referred to as "predictive analytics", provides a solid basis for the design and development of IS artefacts that can generate additional value. Subsequently, we introduce a model for Designing Information Systems with Predictive Analytics (DISPA), extending Design Science Research specifically towards predictive analytics. The model is evaluated based on a case study of MAN Diesel and Turbo, a leading designer of marine diesel engines. The case illustrates that the framework provides useful guidelines for developing environment-specific sensor based predictive models that can out-perform the traditional state of the art predictive methods especially in volatile and uncertain environments.

Keywords: Predictive Analytics, Design Science Research, Forecasting, Sensors.

1 Introduction

As digital technologies become prevalent and embedded in the environment, it makes more and more everyday objects smart – smart phones, smart cars, smart homes, and even smart clothes have become part and parcel of the human enterprise. In this context, the adjective "smart" denotes that an object is able to collect, process, and often communicate data with regard to its functionality and operating environment (Cook, Das 2004). Subsequently, all smart objects must be equipped with sensors that can collect various kinds of data. Although there are many examples of the successful utilization of current snapshots of such data, the identification of patterns from historical sensor data in order to make predictions is only now entering everyday applications. For example, GPS data on phones can provide a current location, but it is not currently able to guess where one is going. The main players on the mobile market, Google and Apple, are trying to close this gap, introducing services like Google Now and Apple frequent locations, collecting data with a similar functionality in mind, but usable applications based on predictions seem significantly more difficult to implement then those using a snapshot picture (Woollaston 2013). Forecasting an event upfront is especially important if there is a substantial cost associated with that event. With vast amounts of data collected for snapshot analysis and the main players clearly looking to extend it towards the future, it is apparent that more and more applications designed to benefit from predictive analysis of historical sensor data will be entering the market.

In this context, predictive analytics refers to empirical methods that are aimed at creating empirical predictions and assessing their quality (Shmueli, Koppius 2011 p. 554). Although the application of predictive analytics toolkits is not common in IS research (Shmueli, Koppius 2011), it is quite prevalent within the body of management science and in particular in Operations Research (OR), both as a method and as a topic, namely forecasting. Demand forecasting, one of the key processes in Supply Chain Management, offers generic forecasting algorithms that became a part of standard ERP systems implementations, making them easy to use in virtually any business. The growing interest in sensor technologies and their ramifications, like data volume and velocity or information processing capabilities, moves the process of demand forecasting to the front burner of current IS research.

In this paper, we introduce a framework that can be used to facilitate the process of rigorously designing predictive information systems. The framework draws on Design Science (Hevner, March et al. 2004) and the steps for building a predictive, empirical model (Shmueli, Koppius 2011), using a mixed method approach for design validation (Tashakkori, Teddlie 1998, Ågerfalk 2013). The framework is evaluated based on a case study of MAN Diesel and Turbo, a leading designer of marine diesel engines. The case illustrates that the framework provides useful guidelines to developing environment-specific sensor based predictive models that can out-perform the traditional state of the art predictive methods. The improvement of prediction quality will, however, come at the expense of a high level of coupling between the problem and the solution, making it financially feasible only in environments where the gain of forecast improvement overweighs the cost of solution implementation. This kind of environment would be characterized by, on one hand, a high level of volatility and uncertainty, and on the other, a high cost associated with forecasting error.

2 Theoretical Foundations

In this chapter we will review the state of the art literature in order to determine what kind of frameworks have already been developed that could be used for designing predictive information systems, including aspects of design specific to sensor technologies. Our requirements towards the framework are that it should be generic enough to allow the design of predictive information systems in various contexts but, as well, specific enough to provide meaningful guidelines to designers in the given context. The remaining part of this chapter will be two-fold: analysis of solutions proposed for predictive information systems in the case context, which is spare-part forecast, as well as a general investigation of the designing of predictive systems in IS.

2.1 Traditional approach to forecasting spare part demand

Given that predictive analytics had rarely been discussed in the IS discourse, either as a method or a subject (Shmueli, Koppius 2011), we looked in other management disciplines for similar problems. Our intention here is not to perform a systematic review of forecasting techniques in the management literature, but rather to provide a flavour of what constitutes a good forecasting method and how to determine which technique to use as a state of the art baseline. As within Operations Research forecasting has been studied thoroughly, multiple literature reviews are available, also for a specific application like spare part demand. Selected methods from Callegaro (2010) and Bacchetti and Saccani (2012) are presented in Table 1 below. The purpose of presenting this list is to illustrate the main difference between our intended approach and current OR research: scholars in Operations develop complex algorithms to predict next items in a (time) series based on the previous values. The main objective seems to be the transformation of historical data. Conceptually, it means predicting an output of the black-box only by analysing its previous outputs. Our general idea is that the prediction can be more informed, i.e. smarter, if it is based on an understanding of the activity within the black-box. An extended search of the literature shows that a limited number of contributions attempt to conceptually unwrap this black box (Ghodrati, Kumar 2005, Dolgui, Pashkevich 2008, Hellingrath, Cordes 2014), but their approaches are so problem-specific that we were not able to replicate them in our context as baseline.

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

Table 1 – Selected spare-part forecasting methods, model classification added by authors

To generalize, the selected models can be broadly sorted into three classes: (1) models that are based on a computing forecast as a single-dimensional aggregation of previous observations were classified into a "time series" cluster; (2) models that are based on computing separately demand magnitude and interval demand points and later combining them into the prediction, were clustered as "Croston-based"; and (3) models that are based on calculating a forecast value based on other properties of the previous value set, rather than the raw values, were grouped into a "stochastic" class. Our review shows that benchmarks of intermittent demand forecasting are inconclusive with regard to the relative performance of any of these models. Petropoulos et al. (2013) benchmark time series and Croston-based methods and conclude that their relative performance depends heavily on parameters used in the implementation. On the other hand, Kourentzes (2013) presents a study where a stochastic solution, namely Neural Networks, outperforms both time-series and Croston based algorithms. Finally, in the study of Teunter and Duncan (2009), time-series methods perform significantly worse than the two other classes, while there is no significant difference between two Croston-based methods and bootstrapping. In the absence of clear conclusions from scientific research, we looked at industry and discovered that a de facto standard is the Croston method: it is the only method specific to intermittent demand in standard SAP R/3 (SAP most current version) and it is explicitly recommended by SAP for products with intermittent demand (SAP 2013 p. 12). Consequently, we selected the Croston method as a starting point and the state of the art benchmark for all our future predictive methods.

2.2 Designing predictive analytics in IS

Shmueli and Koppius (2011) provide a model for the process of building a rigorous Predictive or Explanatory Empirical Model in IS (PoEEIS). Figure 1 represents the 8 steps suggested for building an empirical model. Comparing this framework to Hevner's Design Science Research (DSR) model reveals many similarities: the Evaluation, Validation step in PoEEIS matches almost exactly the Justify/Evaluate from the DSR, the 5 preceding steps of the PoEEIS model can be seen as a more detailed, application specific version of the Develop/Build of DSR. The original PoEEIS is not cyclic, as it considers one iteration of the predictive design, however when we envision multiple iterations, we can imagine cyclic arrows pointing from the Data Collection box to Evaluation and back, which is similar to the logic of DSR. In summary, we notice that the PoEEIS model seems to be a more specific guideline for a single Develop/Build-Justify/Evaluate cycle of Hevner's DSR model. This observation underlines the foundation of a merging of the two models to structure a multi-iterator design of predictive IS artefacts. For a more comprehensive review of the usage of predictive techniques in IS we recommend Shmueli and Koppius's review.

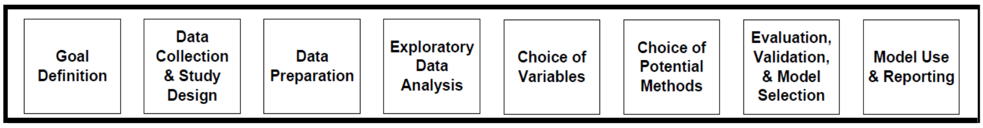

Figure 1 – Steps for building predictive empirical models (Shmueli, Koppius 2011)

2.3 Combining DSR with PoEEIS: Designing Information Systems with Predictive Analytics (DISPA)

In this section we will analyse how DSR (Hevner, March et al. 2004) and PoEEIS models can be utilised to devise a framework structuring a rigorous design process of IS artefacts using predictive models using sensor data. Hevner's DSR framework is intended to help design any IS artefacts while Shmueli's framework was developed having in mind building predictive empirical models. As our overall goal, designing IS artefacts using predictive models includes the goals of both of the frameworks, combining the two should fulfil our intentions, resulting in a more specific version of DSR towards predictive analytics.

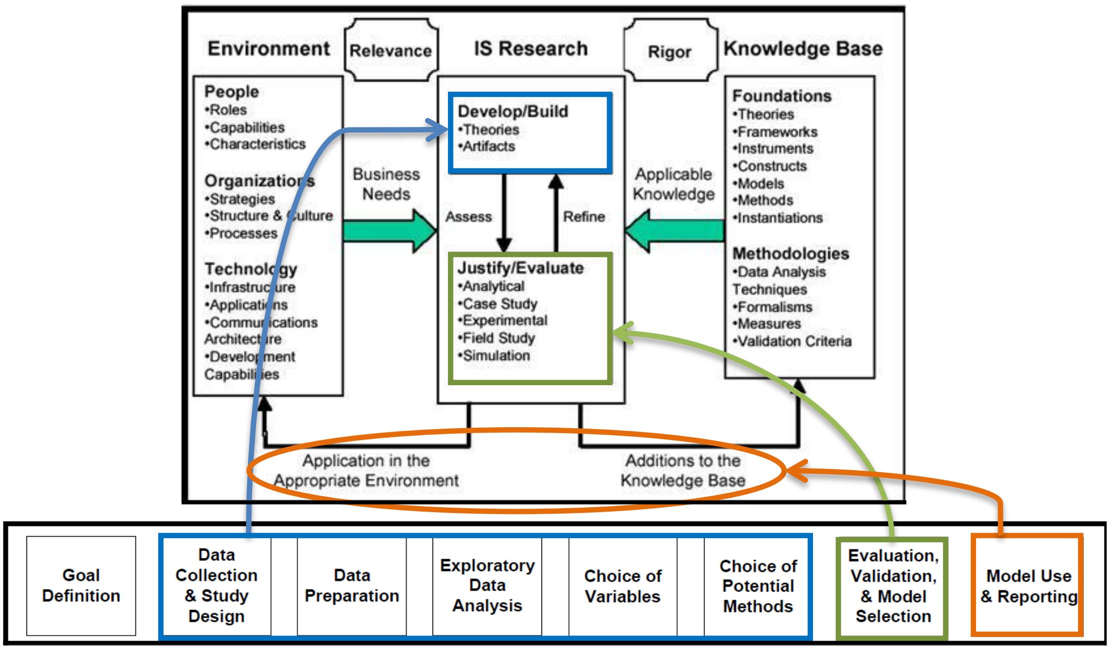

Figure 2 - Combining Hevner's and Shmueli's frameworks

As a first step to combining Hevner's and Shmueli's frameworks let us analyse how exactly the steps suggested by Shmueli and Koppius match those suggested by Hevner. The initial step, goal definition, understood as defining a purpose of the design process and properties constituting a good design for that purpose, does not have an explicit match in Hevner's framework. The following five steps, namely Data Collection & Study Design (initial design choices for the model such as simulation vs experiment study, data collection strategy, sample size), data preparation, exploratory data analysis and variable selection, and choice of a predictive method, seem to be conceptually included in the develop/build step. Nevertheless, in the particular context of IS artefacts using predictive models the output framework could potentially benefit from defining at least some of them more specifically as sub-steps. Evaluation, validation and model selection from Shmueli's model seem to correspond to Hevner's Justify/Evaluate. Finally, Model use matches Hevner's Application in Appropriate Environment and Reporting corresponds to Addition to the Knowledge Base. The graphical matching of the two frameworks can be seen in Figure 2 above.

Based on those observations we started to construct the model for designing IS artefacts using predictive analytics. As a starting point we decided to explicitly include the previously missing goal definition step. In this context the designer needs to answer questions at this step such as what exactly needs to be designed (including what needs to be predicted) and what makes a design for that purpose good. The next step to follow is Hevner's Develop/Build, but with sub-steps inspired by Shmueli's model. We observed that four steps from Shmueli's model (Data collection and study design, exploratory data analysis, choice of variables, choice of potential methods) are very tightly coupled, lacking the required flexibility in step ordering. To name some scenarios, the nature of available variables and methods impacts heavily on study design; or the choice of a method might change the choice of variable. To avoid this ordering struggle we suggest structuring the Develop\Build step in three sub-steps: Model definition, Data preparation and Model implementation. Specifically to sensor data, in the second sub-step (Data preparation) an investigation of match between sensor-measured quantities and predicted values should be specifically discussed.

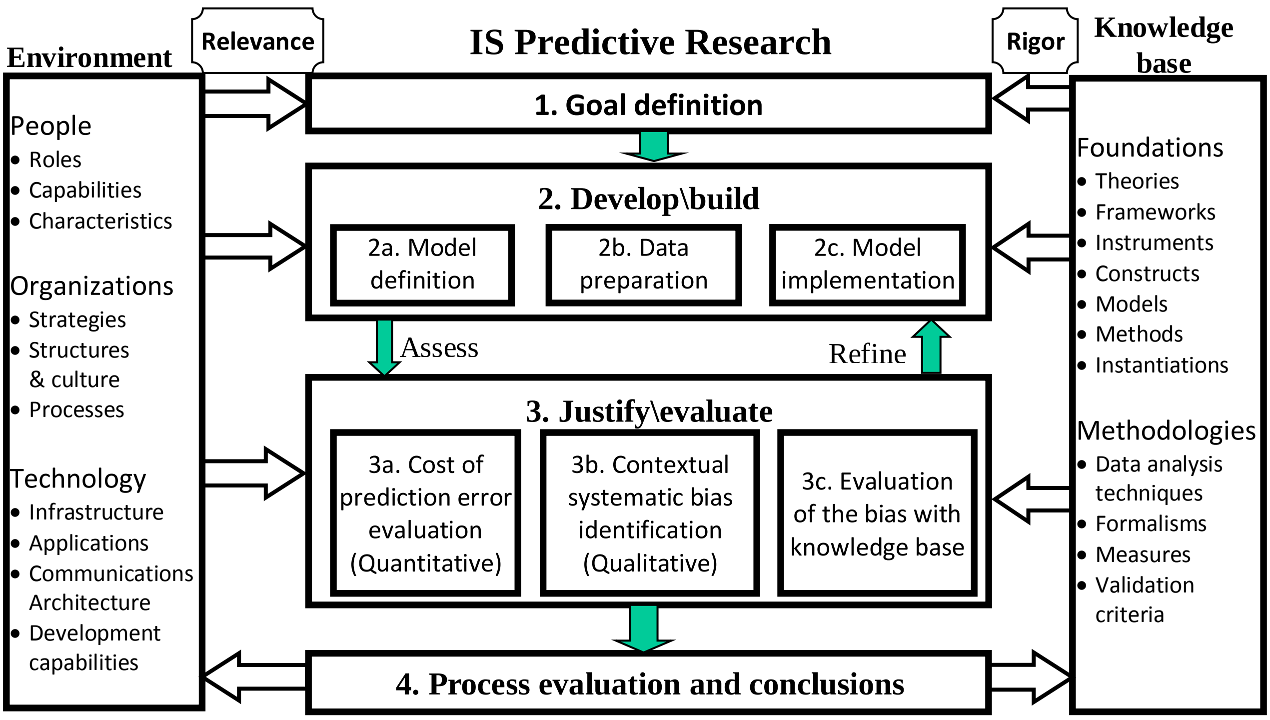

Figure 3- The model for Designing Information Systems with Predictive Analytics (DISPA)

Although both of the models specify the validation step as one of the keys to conducting a rigorous study we felt the need to further structure the validation process. We defined our intended validation process as an objective (and quantified) comparison of various models, but we also like it to extract insights on why different methods produce better or worse quantitative results, hoping to identify systematic biases that can be corrected for later. The combination of quantitative and qualitative elements pointed us towards the mixed method approach (Tashakkori, Teddlie 1998, Ågerfalk 2013). According to Ågerfalk repeating after Creswell, four central parts of the design of a mixed method study are (1) the sequence, (2) the relative priority and the stage of the project (3) stage when qualitative and quantitative components will be integrated, as well as the extent to which the components will be embedded in an overarching framework (4) (Creswell 2013, Ågerfalk 2013). With this guideline in mind we initially structured the validation as a quantitative and quantitative evaluation. Afterwards we decided to further specify the process: defining quantitative evaluation in terms of a single meaningful dimension, to facilitate comparisons and general understanding, and our choice was cost of prediction error. The purpose of the qualitative step is the analysis of the context of the study in order to identify and extract any systematic bias and to underline it the step was renamed accordingly. Additionally we discovered that in order to generalize the contextual findings of the qualitative step some validation with a general knowledge base might be necessary, which led to introducing the third sub-step. The framework is concluded with a Process evaluation and conclusion step. Specifically to sensor data, in this step an investigation clarifies if there is a more direct way to monitor predicted value. The final version of the framework is presented in Figure 3 above.

3 Illustration: Designing predictive analytics with sensors

3.1 Case overview

MAN Diesel and Turbo is the world market leader for large diesel engines for use in ships and power stations, and is one of the three leading suppliers of turbo machines. The roots of the company go back to 1758. In the years 1893-1897 Rudolf Diesel and MAN engineers developed the first diesel engine and in 1904 the company constructed its first steam turbine. According to latest Shipbuilding outlook report (Maritime-Insight 2013) MAN has designed about 70% of engines for active goods-carrying vessels which together with almost 90% of seaborne trade share in world trade (IHS Global Services 2009), which means that MAN engines propel more than half of world trade! Nowadays, MAN Diesel and Turbo does not build engines. The company's strategy concentrates on engineering-intensive engine design process and creating revenues from selling manufacturing licenses to third parties, as well as from the aftersales part of the engine business, namely offering spare parts and services.

The focus on aftersales introduces challenges to MAN's supply chain especially in the area of forecasting spare part and service demand. Aftersales-based business models usually involve a higher level of heterogeneity and product variation than initial sales environments, leading to higher levels of demand uncertainty and making demand predictions relatively more difficult (Teunter, Syntetos et al. 2011). Moreover in the marine business design changes introduced in manufacturing process are very common, typically due to local material availability or shipyards' manufacturing limitations, causing alterations in instantiations of the same design and additional variation of the installed base. Finally, the license-based business model implemented by MAN creates additional obstructions to aftersales activities, as it introduces the engine builder (MAN licensee) as an intermediary between MAN and the end customer (ship owner), that on the aftersales market is a MAN competitor. This setup limits information flow between customer and MAN, additionally hindering the forecasting.

The emphasis on demand forecasting in the case context is also introduced by the aftersales-oriented business model. In the aftersales environment the customer purchases spare parts and services based on two main criteria: availability and price. Availability is merely dependent on accurate demand prediction: if the demand is expected in advance, items or services can be ready at the time customer requests them, increasing sales probability without the cost of excess inventory. Moreover, procuring parts in advance (engine spare parts or elements necessary for performing additional engine services, like retrofit installations) enables stable production pipelines that lower overall procurement costs by avoiding rush orders and expensive rush transportation, helping to keep the price on levels acceptable for customers. In an environment characterized by a high installed base heterogeneity and high number of offered products and services, additionally underpinned by potentially incomplete information on product build and use, an effective forecasting process can be considered as difficult as it is important.

3.2 Introduction to framework instantiation

After having understood the case environment and the relevant knowledge base we are ready to instantiate the model in this context. In the initial step, goal definition, a thorough evaluation of the environment needs to be performed in order to determine what makes the predictive design suitable for the given context, and how to quantitatively measure the cost associated with prediction error. The following two steps, design\build and justify\evaluate, are executed iteratively for multiple designs under evaluation. We suggest starting with a state-of-the-art solution from the knowledge base in order to provide a baseline and to ensure the necessary grounding in previous academic work. The evaluation step should then evaluate the previously developed cost function as well as identify variables that are not monitored, introducing systematic bias that can be removed in following design iteration. When iteration cycles provide satisfactory output the environment and the installed base can be fed back with the newly designed predictive model and the insides acquired during the design process.

3.3 Designing spare part forecasting for marine heavy machinery industry

3.3.1 Instantiation 1 – currently implementable solution

Goal definition (including cost function) (1)

The goal of the empirical part of this paper is to design, develop, evaluate and continuously improve a system to predict the frequency of sales of a selected product in the given case context. Initially a state-of-the-art solution will be selected from the literature as a benchmark. The process of evaluation requires additional explicit structure: quantitative analysis will be performed in an experimental setting: data will be partitioned into learning and test-periods, predictions will be made for test periods based on parameters extracted from the learning sample and the prediction will be evaluated by an objective cost function. Two reliability tests will be repeated 3 times for 3 learning/test samples. Qualitative evaluation will follow, by collecting insights concerning systematic pros and cons of the chosen approach, as well as possibilities to improve it. Based on those suggestions, verified in existing literature, refinements leading to new designs will be made, that will then be finalized and implemented and will undergo the same systematic evaluation process. Linking qualitative feedback to quantitative results should enable evaluation not only of holistic solutions, but also their systematic properties.

The key to a legitimate quantitative evaluation of a design is a meaningful cost function. There are many standard measures for a prediction error, but because of their generic properties they are not able to capture very context-specific factors, such as asymmetric error cost. In order to cater for this diversity two separate cost functions, for over- and under-forecasting scenarios, are necessary. The actual cost associated with an under-forecasting situation occurs due to missed sales potential: as goods are not available when demand occurs some customers will decide to drop the order rather than wait for the items. The percentage of those customers can be determined by the difference in conversion ratio of quotes to orders (also referred to as hit rate) for in-stock quotes versus the stock-outs. In the case of an example component group, piston ring sales at MAN, in a scenario where goods are in-stock average hit rate will oscillate around 39 % but in the case of a stock-out only 30% of quotes would convert to orders, ceteris paribus, we would assume 9% of customers gave the purchase up due to lack of availability. In order to calculate lost profit, the average hit rate difference between in-stock quotes ($\mathbf{HR_{is}}$) and the stock-out hit rate factor ($\mathbf{HR_{so}}$) needs to be multiplied by under-forecasted volume $\mathbf{U_{vol}}$ (to compute sales volume missed due to stock-out) and, to convert sales turnover to EBIT profit, multiplied by average contribution margin **CM**.

The cost associated with over-forecasting can be divided into two categories: opportunity cost, also known as the cost of frozen capital, as well as cost of potential depreciation and scrap, both proportional to over-forecast volume ($\mathbf{OF_{VOL}}$). The opportunity cost is experienced since, for over-forecast, the investment in inventory was unnecessary and the money could be invested differently, bringing certain profit to the company. Most of the firms have some baseline working capital ratio to be used for such calculations. For MAN, in 2014 this working capital ratio (**OC%**) was set to 10% (internal) per year. As in our context sales are expected every month, over-forecasting in one month will lead to lower replenishment cost in the following month, so that the frozen capital cost will always be calculated for a single month. In case of cost of deprecation and scrap factor (**DF%**), this reflects a possibility that unsold inventory will not move for a period of time, leading the inventory to be written-down by a certain depreciation factor, or even, if parts are no longer sellable, to be written down completely and scrapped. For MAN, depreciation and scrap factor for 2014 is set to 5%. Putting all the parameters together, the cost of forecast error, $\mathbf{COST_{FE}}$, can be described as in formula 1 below:

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

Iteration 1 – state of the art solution (Croston)

Develop\build: Method selection (2a)

The initial method is selected based on the literature review (see section 2.1), where the Croston method was selected as the state of the art. It is a two-step approach, calculating a separately expected interval between demand points and the quantity of demand for every piece of equipment that MAN has ever sold spare parts to. Additionally, if there has not been any demand for more than 3 times the expected interval between sales, we expected no more sales from that installation. According to the prescription from the original solution (Croston 1972), the in-between sales interval and expected magnitude of the demand are set by calculating a simple exponential smoothing of the values registered in the past. More details about forecasting methods recommended for the case setup are described in the literature review section, as well as in the original paper (Croston 1972).

Data preparation (2b)

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

Method implementation (2c)

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

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

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

Table 2 - Cost of prediction error for baseline Croston method

Contextual systematic bias identification (3b)

The Croston method seems to provide a fairly good forecast, but it is very generic, so from the beginning we started looking for case specific information that could inform the prediction. The first promising idea coming from Demand Planners included the lifecycle information, by monitoring the installed base (equipment in use) with phase out sensors. Moreover, in the case of absence of equipment phase-out sensors, as marine diesel engines need to be legally supervised by a third party (a classification society), vessels and engines that are scrapped are regularly reported. This way all data regarding dead installations can be completely removed from both training and test sets.

Evaluation of the bias with knowledge base (3c)

Forecasting spare part demand using "installed base" (IB) information (knowledge on age and status of products and systems in use, and customer maintenance and replacement policies (Minner 2011)) has been a subject of recent academic research. The main stream of this research is concerned with optimizing inventory policies using detailed geographical information about customers and equipment (Jalil 2011, Ihde, Merkel et al. 1999, Song, Zipkin 1996). The academic efforts to develop a demand forecasting method exploring IB are limited: as Dekker et al point out: "scientific research on Installed base forecasting is limited and the term is pretty scarce in the operations literature" (Dekker, Pinçe et al. 2010 p. 2). The outline of the idea that demand forecast can be based on the installed base was drawn by Lapide (Lapide 2012), but the approach is so simplistic that it cannot be considered an applicable method. An extremely interesting theoretical forecasting framework is presented by (Minner 2011): the framework estimates the probability of spare-part sales for equipment of a certain age and based on age of equipment in the field, estimates total spare part demand. All those occurrences of IB use in the context of forecasting make it a promising candidate to include in the forecasting method.

Iteration 2 – Croston with phase-out sensor

Develop\build (2): Method selection (2a)

Based on the insights regarding an installed base from the case and the knowledge base, an enhancement for the Croston method will be developed, including a phase-out sensor output. For all pieces of equipment that are already not in use no forecast will be calculated. For all others, the same algorithm of Croston will be used. Notice that only data related to the phase-out of engines is used: although phase-in data is available, it is unusable for the Croston algorithm, as no historical sales are available.

Data preparation (2b)

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

Method implementation (2c)

The implementation is very similar to a previous Croston method, the only difference being that for engines not in use the forecast will always be set to 0. In order to ensure consistency, the data related to currently dead installations are removed from the learning data sets as well.

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

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

Table 3 – Cost of prediction error for the Croston method with IB phase-out component

Table 3 above shows the cost evaluation of the Croston method with IB phase-out component. It is important to note that compared with the baseline method the cost went down by more than 4%. The improvement is visible for 2 data partitioning scenarios (2012 and 2013) while for the last partition (2014) it remained practically constant.

Contextual systematic bias identification (3b)

Demand planners, service engineers and sales persons drew our attention towards a market diverse in the period of the study. As an effect of the lower demand on transportation services, being the result of the global financial crises, sailing patterns of most of the customers were said to be changed. Rather than travel with maximum frequency and speed the vessel management companies were said to concentrate on cost reductions, maximizing load per vessel and scarifying the time of a transport. Moreover, to optimize fuel consumption and a vessel's wear, ships would travel with the most efficient, rather than maximum, speed. This phenomenon in shipping industry is often referred to as slow steaming. As the global economy started to recover from the crises the situation started to go back towards the previous status quo. All those changes could potentially lead to a very significant change in demand for spare parts. This means that for example, if ordinarily a ship owner would replace a given spare part every 5 months, under the slow steaming scenario, assuming a vessel activity reduction by 20%, the matching period would be 6 months. In the context of the Croston method this suggests that data collected in the slow-steaming period has to be somehow "normalized" to be comparable to the previous observations.

Evaluation of the bias with knowledge base (3c)

Slow steaming has been widely recognized in recent shipping literature (Notteboom, Cariou 2013, Woo, Moon 2014, Yin, Fan et al. 2014). Three reasons for the popularity of slow steaming among ship managers are oversupply of shipping capacity, increase of fuel price and environmental pressure (Yin, Fan et al. 2014). According to Notteboom and Cariou, the strategy has been gradually implemented by the main liner shipping companies since 2008, very seriously affecting the dataset we are analysing. Those observations make apparent the potential of including activity sensor information, directly observing ship engine utilization patterns rather than inferring them from time intervals between replacements, into the predictive model.

Iteration 3 - Activity sensor

Develop\build: Method selection (2a)

In order to compensate for the bias in data caused by a changing market pattern in the period of study (application of slow steaming) the way of normalizing data periods based on actual engine activity needs to be introduced. The simplistic idea behind this approach is that if in a given period an engine was used 20% less than in an ordinary period, the calculation of intervals between spare part replacement would extend the expected lifetime of spare parts in that engine by 20%. In order to achieve this goal the unit of interval between replacements will be changed, from time (in months) to engine running hours with equivalent maximum revolutions. Intuitively, an engine can accomplish 1 running hour with maximum equivalent revolutions by either running 1 hour at full speed or 2 hours at half of maximum speed, and so on. In order to measure running hours with maximum revolutions, engine activity sensors were introduced. The prediction model will predict magnitude of sales the same way as the traditional Croston approach, but the interval between sales will now be predicted based on the engine's activity, rather than time. Conceptually, you can see this modification as a next step, after implementing a phase-out sensor – phase out monitors engine activity in a binary fashion ("in use" vs. "not in use"), while the activity sensor creates a more continuous scale of engine activity.

Data preparation (2b)

Running hours full-speed resolution equivalent was estimated based on engine application. Due to customer privacy protection policy MAN does not have access to all activity sensors actually mounted on the equipment (although they are installed on practically all engines). The values are estimated using the average value expected for an engine application: typically, according to MAN service engineers, an engine in a stationary plant will run close to all the time at full speed (about 8600 hours of full-speed resolution per year), the main engine on a ship around 6000 hours, while an auxiliary one will only run about 3000 hours. Furthermore, we introduced scaling factors, based on market behaviour according to the literature and MAN experts for a given year (slow steaming scenario popularity grew from 2008 and peaked in 2011 and from then engine average activity stabilized at levels slightly below norms from before 2008). Running hour values are also extrapolated to the future, using the same estimation logic. The values are extracted for the same periods as historic sales data, reaching back to 2008. . Because of the estimation factor the match between sensor data usage and measured variable is not perfect– clearly, a more optimal way would be to use “real” running hours measured on every engine.

Method implementation (2c)

The implementation is identical to the baseline Croston method for calculation of expected demand magnitude, but the expected interval before the next replacement is now calculated using the running hours readying: for every month with a replacement running hours output is looked up and this way the average running hours elapsing between replacements is obtained. The forecast is the extrapolation of that average, with the magnitude set using exponential smoothing, as in the original Croston.

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

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

Table 4 –Cost of prediction error for the Croston method with activity sensor output (using running hours)

Introducing activity sensor output has improved the baseline Croston cost by as much as 20% and the previously proposed Phase-out sensor implementation by 17%. The improvement is visible for 2 data partitioning scenarios (2012 and 2013) while for the last partition (2014) it remained practically constant.

Contextual systematic bias identification (3b)

The overall impressive improvement is achieved even in spite of the assumption of a rather rough estimation leading to a potential data quality issue: although engine application seems to be a good estimate of engine utilization, clearly there must be a variation within segments of the same engine application, so that the full potential of this solution could be achieved if estimated values were replaced by the real observations from in-situ installations. Nevertheless, as on one hand the realized forecast quality improvement is significant, and on the other hand further model development will require significant investment in infrastructure, at this point it is thus not feasible to run another design iteration.

Evaluation of the bias with knowledge base (3c)

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

Process evaluation and conclusion (4)

In summary, all the implemented sensor-based designs show prediction quality improvements when compared to the baseline Croston solution. Unfortunately, the quality comes at the price of complexity and specificity to a given environment. An initial Croston solution could be easily implemented for any data series. Sensor-based solutions require very specific additional information and the quality improvement they provide is gradually coupled more tightly with the application, and this tight coupling and specificity increase together with the increased prediction quality. Furthermore, the additional information comes from sensor installation (or its simulation), that needs to be pre-installed and this complexity introduces cost not present in the Croston scenario. Moreover, those observations suggest that sensor-enabled forecasting solutions would be financially feasible in environments where gain of forecast improvement overweighs the cost of solution implementation: this kind of environment would be characterized by, on one hand, high level uncertainty, and on the other, have a high cost associated with forecasting error.

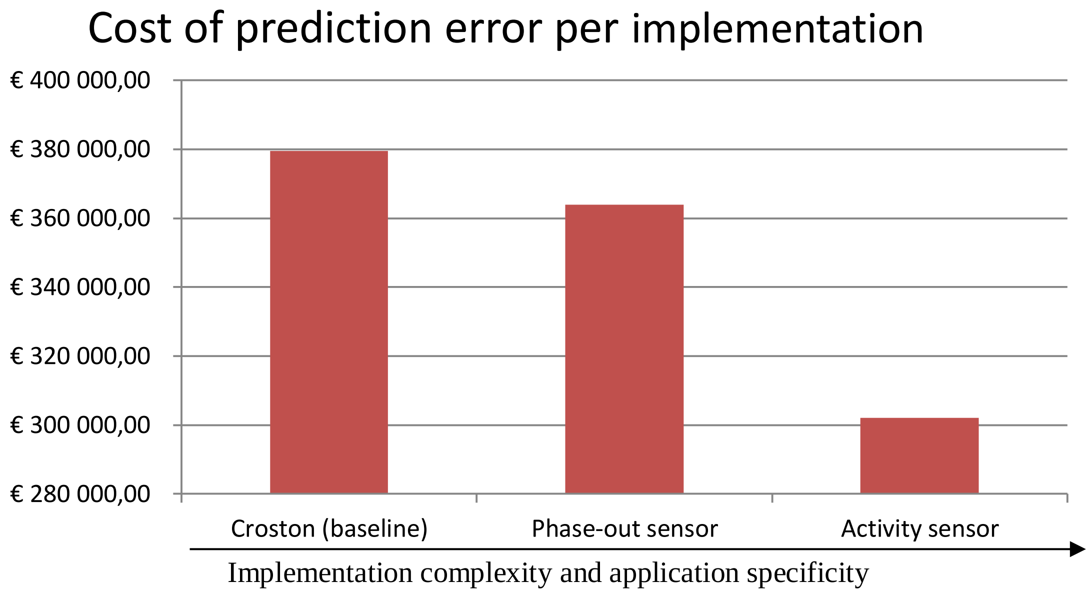

Figure 4 - Qualitative result summary

The qualitative output of four implemented designs is presented in Figure 4. Phase out sensor design improves the baseline prediction quality by 4%, while an activity sensor, on the other hand, beats the baseline by 20% and the Phase-out solution by 17%. This activity sensor example shows that together with including the additional dataset's data quality (DQ) problems can be faced, as the data might have been estimated or generated from a source that did not have that specific data usage in mind. In that case new data management routines should be implemented, leading to gradual DQ improvement.

4 Implications and Contributions

In this paper we introduced a framework facilitating the process of rigorously designing predictive information systems. The model was evaluated based on a case study that showed that the framework can provide useful guidelines to develop environment-specific sensor based predictive models that can out-perform, in a given environment, state of the art predictive methods,. Generalizing this observation, in the absence of a one-size-fits-all solution custom, context specific ways of predictive designs will be gaining popularity, especially when considering an inevitable growth of IOT and sensor technologies. For those approaches our model can provide both a structure and a rigorous guideline, as it has proved to do in the example case.

This paper provides a contribution to Information Systems Research and in particular Design Science Research by introducing a model for Designing Information Systems with Predictive Analytics (DIS-PA) that can serve as a method for developing IS artefacts. Additionally, the paper introduces and systematically evaluates a number of spare-part forecasting methods, which can be considered a contribution to Operations Research literature. Finally, as the model is detailed enough to be instantiated in a real-life setting in the same way that it was used in the case setup, the paper provides a contribution to industry and practice.

5 Conclusions

Despite the increasing relevance of predictive analytics in IS research, the community has not devoted great attention to the issue until recently. In particular, little attention has been directed at issues related to forecasting. Our study shows that paradigms used in IS research, in particular Design Science, can provide a useful lens for the analysis of environments characterized by a high degree of uncertainty, and provide promising solutions for challenges embedded in such environments. We intend our work to be a step towards addressing this shortcoming and hope that it initiates more efforts both in the area of predictive analytics in general and demand forecasting in volatile and uncertain environments.

Future work has to focus on validation of the model in new environments, by collecting data from further case studies. More specifically, it would be interesting to see if in other settings contextual evaluation of systematic biases (step 3b) can have the significant depth to provide insights as useful as in the case of MAN. Furthermore, to additionally structure the design process, written guidelines could be useful for the designer. Finally, as the designing process might turn out to be more expensive than the gain of forecast improvement, a framework to pre-asses an environment's suitability for a sensor-based predictive solution could be helpful for managers when deciding if similar projects should be undertaken.

References

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

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

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly (1970-1977), 23*(3), 289-303.

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

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

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

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

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

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

IHS Global Services. (2009). *IHS global insides.* ().

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

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

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

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

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

SAP. (2013). Product documentation - Supply chain planning and control. Retrieved from https://help.sap.com/saphelp_byd1302/en/KTP/Software-Components/01200615320100003379/ SAP_BBD/Print_Files/PDF_FILES/SupplyChainPlanningAndControl_BA_en.pdf

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

Song, J., & Zipkin, P. H. (1996). Evaluation of base-stock policies in multiechelon inventory systems with state-dependent demands. part II: State-dependent depot policies. *Naval Research Logistics (NRL), 43*(3), 381-396. doi:10.1002/(SICI)1520-6750(199604)43:3<381::AID-NAV5>3.0.CO;2-5

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

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

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

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

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