International Connectedness and Local Disconnectedness

MNE Strategy, City-regions and Disruption

Lorenzen, Mark; Mudambi, Ram; Schotter, Andreas

Document Version
Final published version

Published in:
Journal of International Business Studies

DOI:
10.1057/s41267-020-00339-5

Publication date:
2020

Sensing the Future: A Design Framework for Context-Aware Predictive Systems

Michel Avital
Copenhagen Business School, michel@avital.net

Samir Chatterjee
Claremont Graduate University, profsamir1@gmail.com

Szymon Furtak
Copenhagen Business School, szymonfurtak@hotmail.com

Journal of the Association for Information Systems (2023) **24**(x), xxx-xxx
doi: 10.17705/1jais.xxx**xx**

RESEARCH ARTICLE

ISSN 1536-9323

Sensing the Future: A Design Framework for Context-Aware Predictive Systems

Michel Avital[1], Samir Chatterjee[2], Szymon Furtak[3]

[1]Copenhagen Business School, Denmark, michel@avital.net
[2]Claremont Graduate University, USA, profsamir1@gmail.com
[3]Copenhagen Business School, Denmark, szymonfurtak@hotmail.com

Abstract

Sensors embedded in smart objects, smart machines, and smart buildings produce ever-growing streams of contextual data that convey information of interest about their operating environment. Although an increasing number of industries embrace the utilization of sensors in routine operations, no clear framework is available to guide designers who aim to leverage contextual data collected from these sensors to develop predictive systems. In this paper, we applied the Design Science Research methodology to develop and evaluate a general framework that helps designers to build predictive systems utilizing sensor data. Specifically, we developed a framework for designing context-aware predictive systems (CAPS). We then evaluated the framework through its application in MAN Diesel & Turbo, which served as a case company. The framework can be generalized into a class of demand-forecasting problems that rely on sensor-generated contextual data. The CAPS framework is unique and can help practitioners make better-informed decisions when designing context-aware predictive systems.

Keywords: Design framework, Systems design, Sensor data, IoT data, Predictive analytics, Forecasting, Design Science Research.

[Senior editor name] was the accepting senior editor. This research article was submitted on [manuscript submission date] and went through [number of revisions] revisions.

1 Introduction

Context-aware computing emerged as a term with the introduction of ubiquitous computing (Weiser, 1991). *Context* here refers to any information that describes the situation or environment of an entity, such as a person, place, or object (Abowd et al., 1999). Sensors are best equipped to capture context by leveraging situational and environmental information to offer timely, situated, and usable content, functions, and experiences (Perera et al., 2013).

As digital technologies continue to evolve, we see everyday objects – smartphones, cars, homes, and even clothes – getting embedded with sensor technologies that respond to physical stimuli and generate contextual data (Aarts & Marzano, 2003; Abowd et al., 1999; Javaid et al, 2021; Uckelmann, Harrison, & Michahelles, 2011). Such objects can collect, process, and communicate data about their situation, functionality, and operating environment that can then be used for context-aware computing (Cook & Das, 2004; Shim et al., 2019). Although there are ample examples of successful utilization of sensor-generated contextual data, the identification of patterns in such data for the purpose of harnessing them to make predictions is still not commonplace in everyday applications (Sampath et al., 2019).

The data generated by sensors can be used to make sense of past events and also to predict future events. Although making predictions is an established practice in organizational decision making (Aldahiri, et. al., 2021; Deming, 2000; Kim et al., 2018), predictive analytics applications that utilize sensor data are a relatively recent phenomenon. **Predictive analytics** refers to the use of data to predict future trends and events. It uses historical data to forecast potential scenarios that can help inform strategic decisions. Predictive analytics using contextual data differ from traditional analytics methods in the way in which data are collected and used (Nardi, 1996). Predictive analytics using contextual data extract data from sensors, which are often custom-designed to obtain specific data types. These sensors can generate and collect a wide range of data, including readings on temperature, water levels, air moisture, fuel levels, electrical impulses, and a host of other metrics depending on how and where they are deployed and the needs at hand. The growing demand for additional trusted data sources to be utilized by predictive analytics systems has led to the development of new sensors and smart devices that can generate contextual data streams.

Applications of predictive analytics include demand forecasting, which is one of the key processes in Supply Chain Management (Lapide, 2012; Srogninis, 2021; Suma, 2021). An example is predicting the behavior of machines or devices used in real-life business scenarios, such as when manufacturers track machines or raw materials used during production (Chen, 2001). As more sensor-enabled smart devices enter the market, embedded software solutions and massively improved device connectivity continue to generate streams of sensor data. However, working with sensor data in industrial settings introduces challenges for system designers (Gungor et al., 2009; Marabelli et al., 2017; Pech et. al., 2021), and there is no framework or guideline to help design these systems. This paper offers a design framework and propositions that can help develop these applications and advance the science behind them. Specifically, we pursue the following research question:

RQ: How can we design sensor-based context-aware predictive systems in industrial-scale settings?

In this paper, we developed a framework that can facilitate the process of designing Context-Aware Predictive Systems (CAPS). We defined CAPS as information systems that can make predictions (e.g., device lifespan, energy used, current temperature, and longevity) based on contextual data. We developed the CAPS framework building on Design Science Research (DSR) (Hevner, March, Park, & Ram, 2004) and data-driven predictive modeling (Shmueli & Koppius, 2011). We then evaluated the CAPS framework through its application in MAN Diesel & Turbo, which served as a case company. Finally, we also derived propositions for future designers through reflection on feedback from the case company.

The CAPS framework contributes to the growing stream of domain-specific information systems in Design Science Research (e.g., Meth, Mueller, & Maedche, 2015; Müller-Wienbergen, Müller, Seidel, & Becker, 2011; Pries-Heje & Baskerville, 2008). The framework aims not only to shorten the design time and reduce the cost of context-aware predictive systems development projects, but also to improve the quality of such systems. The case of MAN Diesel & Turbo illustrates the usefulness of the CAPS framework for developing sensor-based predictive models.

2 Background and Prior Work

IS-related research has widely referred to the notion of "context." For example, context plays a vital role in application development (Kumar & Sharma, 2020), search engine design (Storey, Burton-Jones, Sugumaran, & Purao, 2008), and human-computer interaction (Nardi, 1996), as well as in sense-making (Narock, Yoon, & March, 2012) and general management (Johns, 2006). Although the exact characterizations of context seem to vary, the definition used in the introduction (Cook & Das, 2004) is rooted in the framing of Dey (2001), who defines context as "any information that can be used to characterize the situation of an entity."

As the fourth industrial revolution sets in, more and more objects are being embedded with sensors that generate contextual data, which is used for various applications (Okano, 2017). For example, factories and businesses operating in the manufacturing sector are taking advantage of IoT sensors and data collection (Suma, 2021). With the use of IoT sensors attached to factory machines or robots, measures of usage and lifespan such as temperature, vibration, and wear and tear can be tracked (Kim, Lee & Shin, 2018). These data, when fed into an analytics model or algorithm, can help predict when a machine will likely require maintenance or replacement. In this way, the manager can order parts or supplies ahead of time to avoid costly downtime or expensive machine repairs (Hellingrath & Cordes, 2014).

Another industrial example is monitoring the physical and chemical characteristics of water locally to provide a fine-grained map of water condition. New water distribution channels are equipped with IoT sensors that monitor water quality and potential contamination. Such advanced water monitoring systems help to control risks related to the spread of polluted water and diseases (Nikam & Pawar, 2016).

While forecasting has received limited attention in the IS discipline, it has been studied thoroughly in other management disciplines. Naturally, we do not aim to review all the *forecasting techniques* in the management literature here; rather, our goal is specifically to reveal what constitutes *a suitable forecasting method* and how to determine which technique to use as a state-of-the-art baseline. Notably, multiple literature reviews in the context of Operations Research (OR) cover forecasting thoroughly, including dedicated studies of spare-part demand. Table 1 presents a list of selected methods from Callegaro (2010) and Bacchetti & Saccani (2012).

The list highlights that the common approaches in OR research are derived from the development of complex algorithms to predict the next items in a (time) series based on the previous values. The primary objective centers on the transformation of historical data. Conceptually, this refers to predicting an output of the black box by analyzing only its prior outputs. In contrast, the literature on a more *informed* prediction that is based on an understanding of the activity within the so-called black box is limited and problem-specific (e.g., Hellingrath & Cordes, 2014).

Table 1. Selected Spare-Part Forecasting Methods (Based on Bacchetti & Saccani, 2012; Callegaro, 2010)

Model Classification	Method Name	Inputs	Description	Important Features
Time series Arithmetic average with optional additional 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)	- Stresses recent trends - Easy to compute
	Single exponential smoothing (SES)	- Historical sales data - Smoothing constant	Moving average of demand with smoothing constant	- Works with a few samples - Easy to compute
	Box-Jenkins method	- Historical data - Multiple constants	Moving average and auto-regression, selected alternatively based on historical error	- Captures complex trends and seasonality - Requires much historical data to perform well
	Grey prediction model	- Historical data	Adaptive time series approach using least-squares estimate as feedback to correct for the error	- Works under massive uncertainty to predict events like hurricanes
Croston-based[1] Two average values with exponential smoothing	Croston's method	- Historical sales data - Smoothing constants	Single exponential smoothing for both typical demand magnitude and typical periods between demand points	- Performs well with materials that have intermittent demand (many periods without demand)
	Syntetos-Boylan approx.	- Historical sales data - Smoothing constants	Extension of Croston method removing the positive bias	- Provides a statistically proven bias reduction resulting in lower forecast error
Stochastic Probabilistic	Bootstrap method	- Historical sales data - Limited number for resampling	A randomly chosen subset of historical samples (e.g., forecast for the next 3 periods is 3 randomly chosen periods from the past)	- Offers a probabilistic approach
	Neural networks	- Historical sales data - Neural network layout	Inference from the connection between input and output of the training set to estimate future values	- Offers a method tested in various areas as a predictor
	Order Over-planning (early sales)	- Historical sales data	Extrapolation of sales orders placed by each single customer instead of the overall demand to estimate future sales	- Caters for business where some customers use to purchase well in advance
	Failure rate analysis	- Equipment failure rates - Installed base data	Extracting expected lifetime and failure rates of components based on historical data and extrapolating the forecasted values based on installed base	- Caters for spare part and heavy machinery business - Requires good data about historical incidents and replacement

[1] Croston (1972).

The selected models can be broadly sorted into three classes: (1) models that are based on a computed forecast as a unidimensional aggregation of previous observations that are classified into a time series cluster (Callegaro, 2010); (2) models that are based on separately computing demand magnitude and interval demand points and later combining them into predictions that are clustered as Croston-based (Croston, 1972); and (3) models that are based on calculating a forecast value that is based on other properties of the previous value set, rather than the raw values, and that are grouped into a Stochastic cluster (Bartezzaghi, Verganti & Zotteri, 1999).

Our review shows that benchmarks of intermittent demand forecasting are inconclusive about 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. However, Kourentzes (2013) presents a study where a stochastic solution – namely, Neural Networks – outperforms both time series and Croston-based algorithms. Finally, Teunter & Duncan (2009) find that time-series methods perform significantly worse than the other two classes, while there is no significant difference between the two Croston-based methods and stochastic bootstrapping.

In the absence of clear conclusions from the related scientific research, we acknowledge that the Croston-based methods are the *de facto* traditional standard. It is the only method class specifically aimed at intermittent demand forecasting in standard SAP R/3, and it is explicitly recommended by SAP for products with intermittent demand (SAP, 2013, p. 12). Considering that the Croston method is the most prevalent in practice, we selected it as the state-of-the-art baseline for testing against future predictive methods.

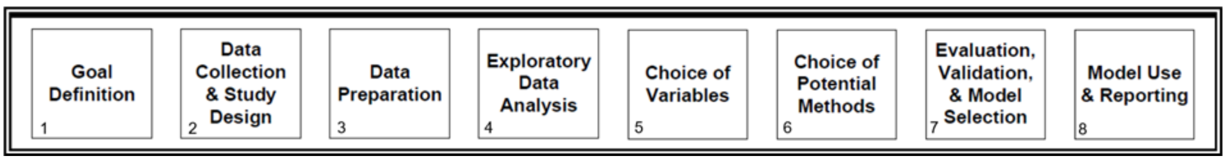

Figure 1. Steps for Building Predictive Empirical Models (Shmueli & Koppius, 2011)

In order to facilitate the use of predictive methods in IS, Shmueli & Koppius (2011) provide an eight-step process for building an empirical model including explicit guidelines on how to execute it for designing predictive models (see Figure 1). In Step 1, the prediction goal and the success benchmark are defined and outlined. In Step 2, key issues regarding data collection and study design need to be addressed, including using an experimental versus observational setting, choosing the data collection instrument(s), setting the sample size, and selecting candidates for observed variables. Step 3 deals with data preparation and outlines actions for data quality controls, including defining procedures for treating missing values and choosing a partitioning strategy. In Step 4, the data are to be evaluated for the purpose of defining variables for the analysis, and in Step 5, the variables are selected. In Step 6, the data transformation method is selected. In Step 7, the evaluation strategy, validation, and model selection are determined. Finally, in Step 8, the strategy for research dissemination is chosen and executed.

Although Shmueli and Koppius' guidelines on building empirical models constitute a significant step toward prescriptive instructions on how a class of predictive systems could be constructed, they are concerned only with the data model (Walls et al., 1992). While an empirical data model (predictive or explanatory) is a central part of any technological IS (Hevner et al., 2004; Orlikowski & Iacono, 2001), the design process of such systems also requires the consideration of several other aspects, including the knowledge base it uses, the organization within which the system is (to be) embedded, and the people who will use it. Therefore, the process suggested by Shmueli & Koppius (2011) provides a reasonably good starting point toward the development of predictive information systems but lacks adaptability to organizational challenges and contexts.

Shmueli & Koppius (2011) make a clear distinction between explanatory statistical models and predictive models. Typically, an explanatory statistical model is built for the purpose of testing causal hypothesis that specify how and why certain empirical phenomena occur. Predictive analytics include empirical predictive models (statistical models and other methods such as data mining algorithms) designed for predicting new/future observations or scenarios. Predictive power refers to a model's ability to generate accurate predictions of new observations. As seen in Fig. 1, their focus is on the data process. However, in practice, designers that instantiate predictive analytics systems with sensors need to respond to organizational challenges introduced by data complexity and veracity and by the contextual meaning of the data in the organization. When evaluating predictive analytics systems, this may introduce the need for both quantitative and qualitative evaluation, both to ensure the selected evaluation function performs well in the industrial setting and to enable a better contextual understanding of the results to improve the model in the next iteration. The context in industrial case settings is particularly relevant to the operational efficiency of the predictive analytics system.

Based on the above observations, we build on the predictive analytics process model proposed by Shmueli & Koppius (2011) and the three cycles model of Design Science proposed by Hevner et al. (2004) to develop a generic *yet enhanced* design process framework, which we call CAPS, that can be used to design and build industrial-scale predictive systems.

3 Research Approach

Context-aware predictive systems, like other IT artifacts, process information to enable or support predefined tasks. All context-aware systems collect context information using sensor technology to sense particular kinds of changes in the environment in which they are embedded. Predictive systems have a common modus operandi: they convert input data into a particular prediction about the future. Although an increasing number of industries embrace the utilization of sensor devices and IoT, there is no guidance for designers and developers, let alone for researchers, who are interested in predictive models that are based on data collected from sensor systems. Therefore, we follow the Design Science Research methodology to address this gap by developing the CAPS framework (Hevner & Chatterjee, 2010; Meth et al., 2015).

We started by exploring the literature on design as well as predictive analytics in order to gain a better perspective on the domain of interest, including specific challenges and problems facing context-aware sensor system development. Next, we explored and merged two relevant yet distinct approaches, namely DSR (Hevner et al., 2004) and Predictive Analytics (Shmueli & Koppius, 2011), into a new and enhanced CAPS framework that is better suited to support our design objective. The CAPS framework is developed and presented as one of the key contributions of our work. To evaluate its applicability and usefulness, we applied the CAPS framework to an ongoing problem at MAN Diesel & Turbo. In addition, we improved the framework based on expert feedback and qualitative field data. Finally, we derived a set of emerging propositions and reflected on how they could be applied in similar design problems. The propositions aim to assist organizations with the implementation of the CAPS framework and to realize its value in their context. The propositions constitute another contribution of this paper.

The research approach is explained with the help of a diagram in Fig. 2 below. Based on the literature on predictive analytics and DSR, we developed knowledge about the key features of such systems. Next, the author team developed a preliminary design of the CAPS framework. One of the authors was then embedded inside the MAN Diesel & Turbo Company for an extended period of time. Placement of the researcher in the organization facilitated data collection and contextual understanding. It is also important to note that although the CAPS framework was developed prior to the start of this project, we were able to refine its conceptualization and presentation as the engaged researcher witnessed firsthand the dynamic nature of the underlying issues and challenges. The researcher was able to obtain direct feedback and insights, which then led to the refinement of the CAPS framework. Finally, a set of propositions was established based on the researchers' direct experience with evaluating the CAPS framework inside MAN Diesel & Turbo.

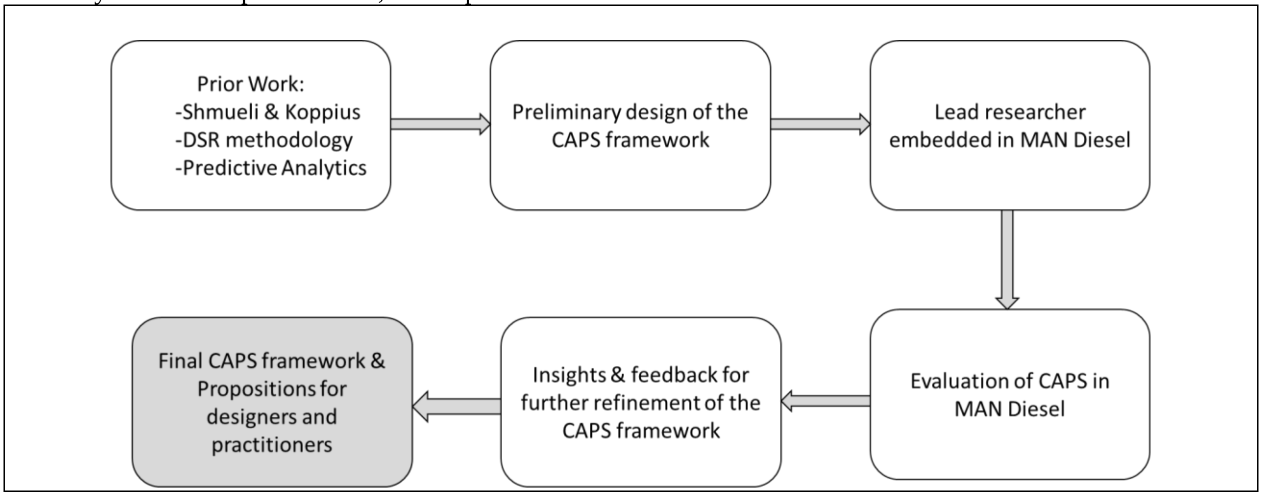

Figure 2. The Research Process Underlying the CAPS Framework's Development

It should also be noted that we developed the CAPS framework to address the problem of designing context-aware predictive systems. In that sense, the CAPS framework is also an artifact. However, in the context of this paper, we refer to the CAPS framework as a "framework" and to the resulting context-aware predictive systems as the design "artifacts."

4 Design Solution – Developing the CAPS Framework

In this section, we present how the Design Science Research methodology and Shmueli and Koppius' predictive empirical model were used to develop a framework for structuring a design process of Context-Aware Predictive Systems (CAPS), including predictive models that use contextual data (hereinafter, the CAPS framework).

The DSR methodology (Hevner et al., 2004) provides a *problem-driven* process to guide the design for IS artifacts, and Shmueli & Koppius' (2011) model provides a *data-driven* process to guide the development of predictive empirical models. We argue that building on both the DSR methodology and Shmueli & Koppius' (2011) model, we can develop a new comprehensive framework that is better suited to dealing with the design of predictive analytics systems that rely on using sensor data and similar digital traces (Figure 3). While the DSR methodology is geared towards general problem solving, the CAPS framework is designed to solve problems that are data-intensive and require predictive modeling.

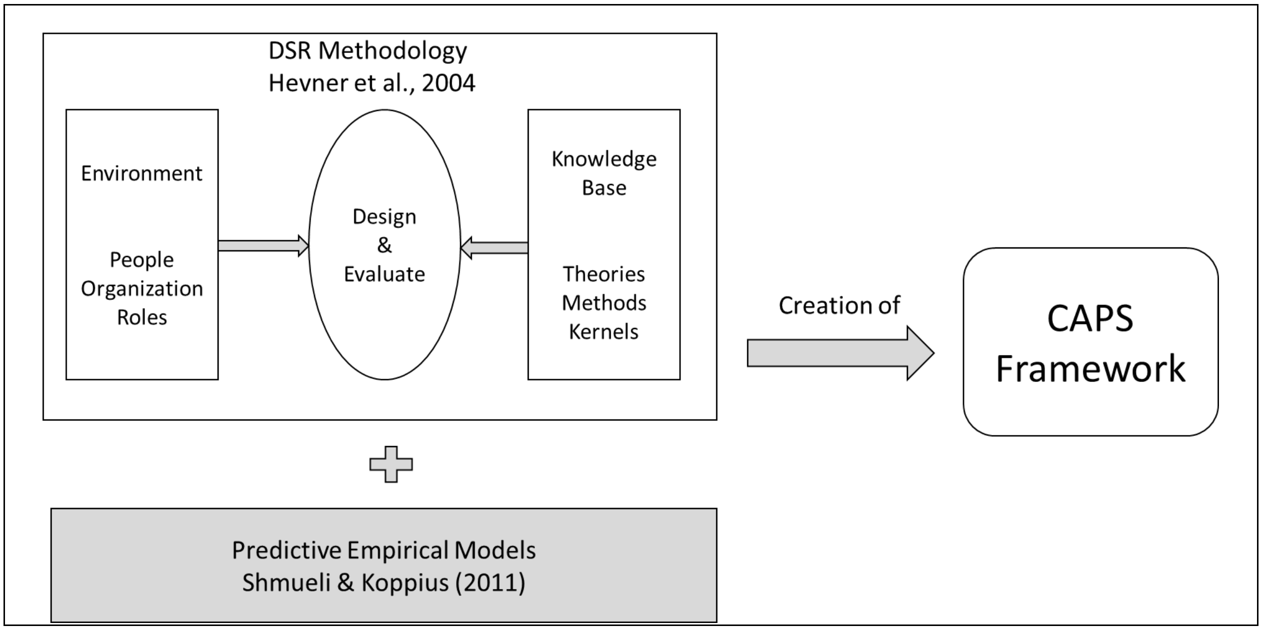

Figure 3. DSR Methodology and Shmueli & Koppius Models Provide Insight into the Creation of the CAPS Framework

As noted earlier, a comparative analysis of Hevner et al. (2004) and Shmueli & Koppius' (2011) models suggests a partial overlap of the two. These two conceptualizations demonstrate overlapping and complementary properties, as illustrated in Figure 4. The initial step in the Shmueli and Koppius model, goal definition, involves defining the purpose of the design process and properties constituting a good design for that purpose. Although the methodology of Hevner et al. discusses the evaluation concept thoroughly, it does not explicitly relate to the goal definition step. The following five steps are application-specific actions conceptually included in the generic develop/build step (marked in blue in Figure 4). Nevertheless, in the particular context of artifacts using predictive models, the output framework will benefit from a clear definition of activities related to the development/build steps. The Evaluation, Validation, & Model selection step in Shmueli & Koppius's process model corresponds to the Justify/Evaluate step of Hevner et al. (marked in green). Finally, the Model Use and Reporting step matches the Application in the Appropriate Environment and Addition to the Knowledge Base of Hevner et al. (marked in orange). Figure 4 provides a graphical representation that depicts the overlap between the DSR methodology and Shmueli & Koppius's (2011) models.

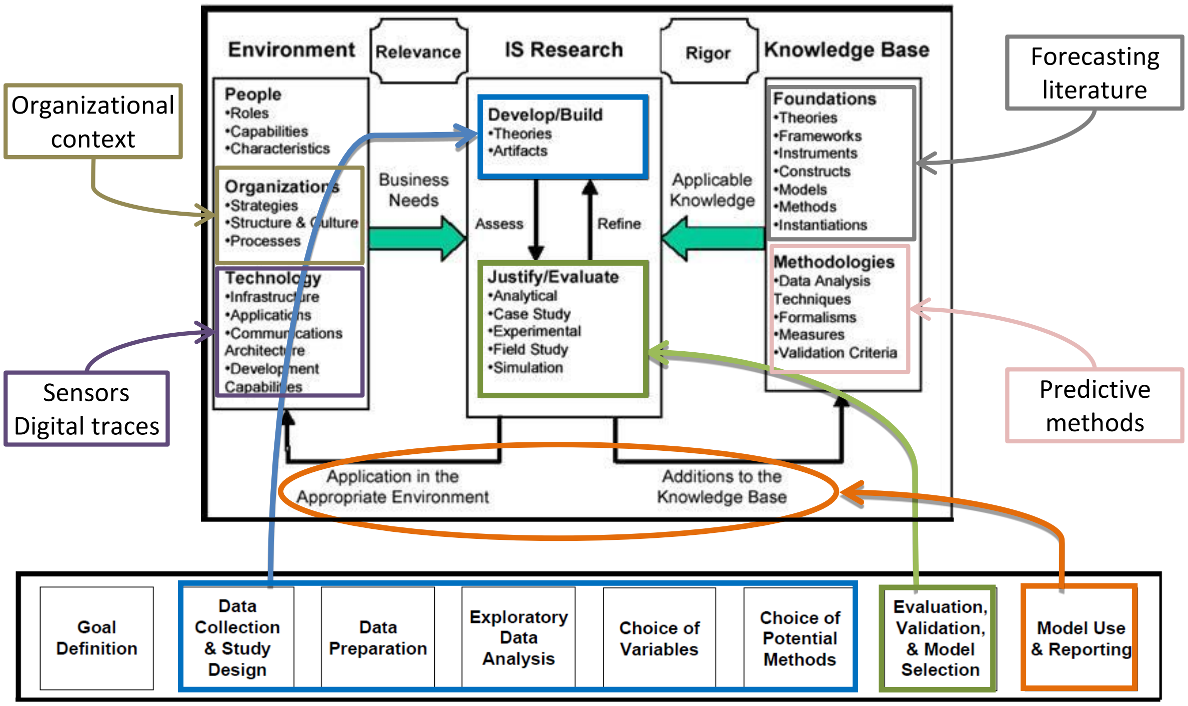

Figure 4. Alignment between DSR Methodology and Shmueli & Koppius Predictive Analytics Model

By leveraging these overlaps, we constructed a new framework for designing context-aware predictive systems – namely, the CAPS framework, which is shown in Figure 5. It is important to note that we used DSR methodology as a research method but also used the DSR methodology artifact to create a new artifact – that is, the CAPS framework.

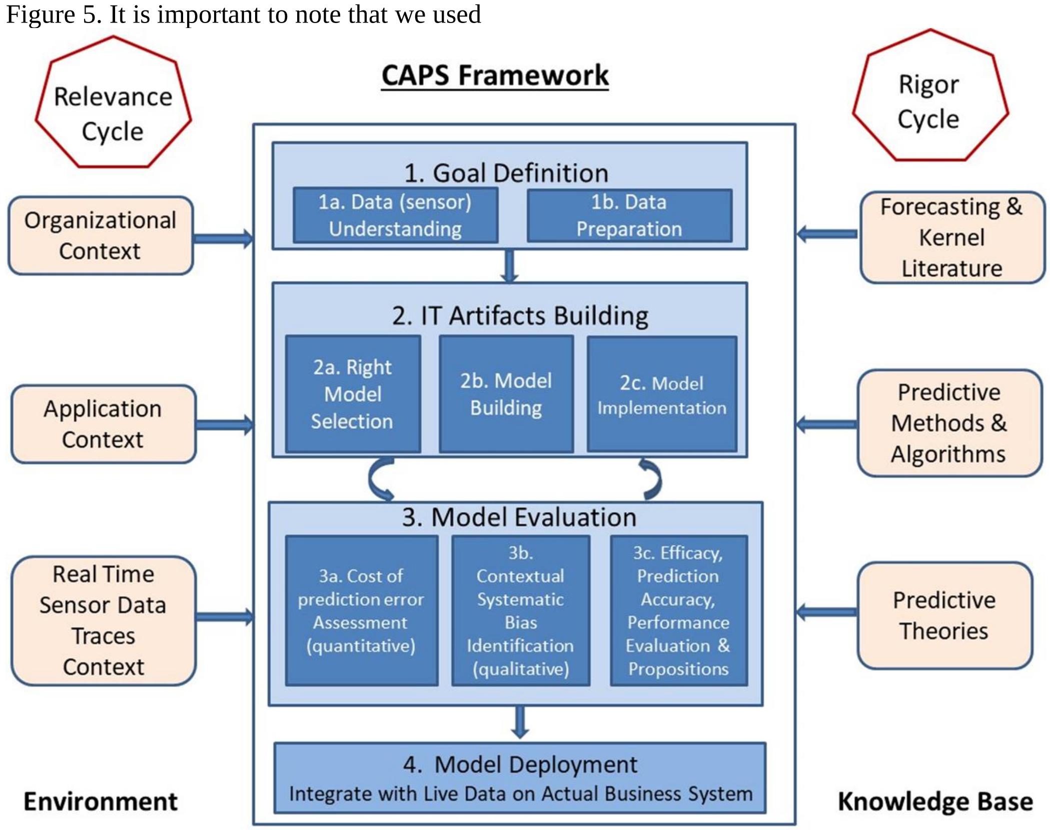

Figure 5. Framework for Designing Context-Aware Predictive Systems (CAPS)

As a starting point, in Step 1, ***Goal Definition***, we explicitly follow the Goal Definition step, which was previously only present in the DSR methodology in reference to Business Needs (Hevner et al., 2004). This is when the designer answers questions, such as what exactly needs to be designed, what needs to be predicted, and what could be a benchmark of good design that fulfills the design objectives in the underlying context. The Environment informs this step with explicit business requirements, and the Knowledge Base provides additional guiding reference points such as information about already developed measures and goal evaluation methods.

The *environment* on the left of the CAPS framework refers to real-time sensor data sourcing, the applications to be built, and the myriad of organizational challenges that a designer can face. A significant source of data at this stage is from sensors. These could be environmental sensors (e.g., pressure, temperature, on/off, wear/tear) that provide real-time context to our problem. Sensors help us augment our physical surroundings. However, sensor data may have weak signals, may be noisy, and may at times be difficult to make sense of when out of the overall context. However, once it is possible to process the data and make sense of it, then the organization has to decide which application best suits these data. A relevant question for the organization at this stage would be: Are we building predictive applications, building classification applications, or clustering data to visualize what is happening in our surroundings? The data analytics team often has to with the administration to address such organizational dilemmas in order to move forward with a particular approach.

Next, in Step 2, ***IT Artifact Building***, we follow DSR's Develop/Build logic, but with three sub-steps inspired by Shmueli and Koppius. We observed that four steps from Shmueli and Koppius's model (data collection and study design, exploratory data analysis, choice of variables, and choice of potential methods) are tightly coupled, lacking the required flexibility in step ordering. To avoid possible ordering confusion, we structure the Develop/Build step in three sub-steps: (2a) Right Model Selection for the task at hand (e.g., prediction, classification, clustering); (2b) Model Building, in which the model is constructed (using appropriate training and testing data sets); and (2c) Model Implementation, which details the actual implementation process. Specifically, in the case of sensor data and digital traces, Model Implementation, the third sub-step, requires a thorough investigation of the match between sensor-measured quantities and predicted values, including potential treatment of missing or incorrectly registered values, as well as specific data partitioning scenarios that address the threat of prediction over-fitting. The Environment informs this step with all the usable data and already existing business processes, and the Knowledge Base provides additional guiding information on a variety of machine learning algorithms for prediction purposes.

Although both models specify the evaluation/validation step (Justify/Evaluate in Hevner et al. and Evaluation, Validation & Model Selection in Shmueli and Koppius's model) as one of the keys to conducting a rigorous study, we concluded that further structuring of the evaluation process is required in the context of designing predictive analytics information systems. We defined the intended Evaluation process as an objective and quantified comparison of iterations. In addition, we also aimed to generate insight into why different methods produce better or worse quantitative results and use this insight to identify systematic biases that can be corrected.

In Step 3, ***Model Evaluation***, we follow DSR's Justify/Evaluate step in three sub-steps: Step(3a) Cost of Prediction Error Assessment, a quantitative evaluation in which we facilitate comparisons and a general understanding of the cost of prediction error developed in Step 1; Step(3b) Contextual Systematic Bias Identification, a qualitative analysis of the context of the study to identify any systematic/systemic bias using the Environment; and Step(3c) Efficacy, Prediction Accuracy, and Performance Evaluation. This step includes also Propositions generation to generalize the learning so far as a basis for the development of new iterations that yield better quantitative results with less systematic bias. Specifically in the case of sensor data and digital traces, in this step, an investigation clarifies whether there is a more direct way to monitor predicted values. There is a tradeoff between model precision and system complexity – in step 3c, the complete design process needs to be evaluated in relation to this tradeoff. Additionally, designers reflect on how well the current instantiation of the model is embedded in the Environment and the extent of its contribution to the Knowledge Base – typically through Organizational context, Sensors and Digital traces, as well as forecasting literature and predictive methods.

The framework concludes with Step 4, ***Model Deployment***. In this step, the designer integrates the context-aware predictive system with the rest of the production system. Specifically, the designer connects the sensor traces to enable the integration with the underlying operational process in the organization. The proposed CAPS framework is illustrated in Figure 5.

Overall, the CAPS framework offers guidelines for structured development of a context-aware predictive system, leading to shorter design and implementation times, as well as cost savings. The framework also has the potential to serve as an alternative approach to conventional forecasting by boosting explanatory forecasting method development and unwrapping the black-boxed, time-series-based forecasting routines.

5 Evaluation of the CAPS Framework

One approach to evaluating the CAPS framework is to present the framework to relevant industry experts and ask them to comment on the benefits and utility of the framework (Dey, 2001). A more robust approach is to utilize the framework to build context-aware predictive systems in an actual company. We took the latter approach, and we describe in this section how we utilized the CAPS framework in a development project at MAN Diesel & Turbo to design multiple artifacts. Figure 6 depicts the evaluation strategy and the design process. We built three artifacts with the goal of reducing forecasting errors by including sensor-based data. We also designed two additional artifacts that include more advanced use of sensors. The following section provides a short summary of the design iterations and describes the five artifacts.

Due to differences between candidate solutions in our case, multiple context-aware artifacts had to be built: some solutions could be completely implemented and evaluated within the period of the project, while others could only be evaluated, based on the current knowledge (see Figure 6).

Currently implemented solutions: Designing and deploying a sensor-based system to forecast spare part demand			**Anticipated later work: Evaluating advanced designs of a sensor-based system to forecast spare part demand**	
Artifact 1.1: Croston method (state-of-the-art)	Artifact 1.2: Croston with phase-out component	Artifact 1.3: Activity sensor	Artifact 2.1: Continuous multi-sensor monitoring	Artifact 2.2: Remote Monitoring Interface
Goal: Cost of forecasting error minimization			Goal: Evaluation of future designs feasibility	

Figure 6. Using the CAPS Framework to Design Multiple CAPS Artifacts

Following the CAPS framework, the initial step, Goal Definition (see Figure 5), calls for a thorough evaluation of the environment to determine what makes the predictive design suitable for a given context and how to quantitatively measure the cost associated with prediction error. The following two steps, IT Artifacts Building and Model Evaluation, are executed for designs under evaluation using the previously defined evaluation criteria (i.e., cost of prediction error, contextual systematic bias, and efficacy). We suggest starting with a state-of-the-art solution from the Knowledge Base to provide a baseline and ensure the necessary grounding in previous academic work. The model evaluation steps (Steps 3a, 3b, and 3c; see Figure 5) should then justify and evaluate the previously developed objective function as well as identify variables that are not monitored, perhaps introducing a systematic bias that can be removed in subsequent design iteration. When iteration cycles provide a satisfactory output, the environment and the installed base can be fed back with the newly designed predictive framework and the insights acquired during the design process.

5.1 Settings and context of the evaluation

MAN Diesel & Turbo is the world market leader for very large diesel engines for use in ships and power plants, and it 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. MAN Diesel & Turbo is a market leader – MAN has designed about 70% of the world's engines for active goods-carrying vessels, culminating in about 85% of seaborne share in world trade (IHS Global Services, 2009). Overall, MAN engines propel more than half of world trade (Maritime-Insight, 2013). Since the 1980s, MAN Diesel & Turbo has ceased building engines, and the manufacturing process has been outsourced mainly to Asian business partners. The company's strategy concentrates on (1) engineering-intensive engine design processes, (2) creating revenues from selling manufacturing licenses to third parties to build MAN engines, (3) generating revenue from engines in use, and (4) shifting the focus to the aftersales part of the business – namely, offering spare parts and services (Song & Zipkin, 1996).

The focus on aftersales introduces challenges to MAN's supply chain, especially in forecasting spare parts and service demand. Aftersales-based business models usually involve a higher level of heterogeneity and product variation than do initial sales environments, leading to higher levels of demand uncertainty and making demand predictions relatively more difficult (Teunter, Syntetos, & Zied Babai, 2011). Moreover, in the marine business design, changes introduced in the manufacturing process are widespread, typically due to local material availability or shipyard manufacturing limitations, causing alterations in instantiations of the same design and additional variation of the installed base (e.g., Dekker et al., 2010; Minner, 2011). Finally, the license-based business model implemented by MAN creates additional obstructions to aftersales activities because it introduces the engine builder (i.e., MAN licensee) as an intermediary between MAN and the end customer (i.e., ship owner). Thus, the engine builder both partners with MAN in the production phase and competes with MAN on the aftersales market of spare parts. This setup limits information flow between customers and MAN and also challenges the current process of spare parts demand forecasting. All the above is considered in the CAPS framework as organizational context.

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. In other words, if the demand is expected in advance, items or services can be ready at the time the customer requests them, thus 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, such as retrofit installations) enables stable production pipelines and lowers overall procurement costs by avoiding rush orders and expensive rush transportation, helping to keep the price at a level that is acceptable for customers. In an installed base environment characterized by high heterogeneity and a high number of offered products and services (and underpinned by potentially incomplete information on product build and use), an effective forecasting process can be considered non-trivial and very important.

5.2 CAPS Use Case 1 – Currently implemented solutions

The process of evaluation requires an explicit strategy, in which quantitative analysis is performed in an experimental setting; data are partitioned into learning and test-periods; and predictions are made for test periods, based on parameters fine-tuned with the learning sample and evaluated by an objective function. The qualitative evaluation follows, collecting insights related to the pros and cons of the chosen approach and also predicting possibilities to improve it. Based on these suggestions, which are verified in existing literature (Kourentzes, 2013; Tuenter & Duncan, 2009), refinements leading to new designs are made, which then not only are finalized and implemented but also undergo the same systematic evaluation process. Moreover, linking qualitative feedback to quantitative results should enable evaluation not only of holistic solutions but also of their systematic properties.

5.2.1 Three iterations of implementable artifact designs

Referring back to Figure 6, we conducted three iterations that resulted in slightly different artifacts, which aimed to minimize the cost of forecast error measured according to the formula developed in the Appendix. The initial state-of-the-art forecasting method, an implementation of the Croston method (considered the traditional approach), was used in order to provide a baseline for new solutions. We used a two-step approach intended for products with infrequent demand, calculating intervals that are separately expected between demand points and the magnitude of demand. After gaining a better understanding of the sales data (see the Appendix), a new artifact, called “Croston with a phase-out component” (see Artifact 1.2, Figure 6), was implemented and evaluated using the previously developed objective function, improving the baseline prediction by 4%.

To correct for the systematic bias in the dataset, we introduced another change to the baseline forecasting model. The basis for this approach is that if, in a given period, an engine is used 20% less than in an ordinary period, the calculation of the interval between spare parts replacements would extend the expected lifetime of spare parts in that engine by at least 20%. To achieve this goal, the unit of interval between spare part replacements would be measured in actual engine activity time, instead of in calendar time (months) as proposed by the Croston method. We had access to sensor measurements of engine activity (i.e., Figure 6, Artifact 1.3) – specifically, its load and a counter of its running hours. Based on these measures, we defined a new variable called engine activity, measured in running hours with maximum load. Intuitively, an engine can accomplish one running hour with maximum load by running either one hour at full speed or two hours at half of the maximum speed and so on. The implementation of this method improved the quality of prediction measure as the cost of forecast error by another 17% compared to the Croston with phase-out component solution, as well as by 20% compared to that of the baseline.

In summary, the implemented sensor-based design artifacts show prediction quality improvements when compared to the baseline Croston solution. Often, 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. Sensor-based solutions require specific additional information, and the quality improvement they provide is gradually coupled more tightly with the application; in turn, this tight coupling and specificity increase prediction quality. Furthermore, the additional information comes from sensor installation requiring some specific infrastructure, which introduces costs that are not present in the basic Croston scenario. Moreover, these observations suggest that sensor-enabled forecasting solutions can be financially feasible in environments characterized by, on the one hand, high level uncertainty and, on the other, a high cost associated with forecasting error.

The quantitative output of the three implemented designs in Figure 6, the first evaluation of CAPS, is presented in Figure 7. Croston with phase-out component improves the baseline prediction quality by 4%, while an activity sensor outperforms the baseline by 20% and the Phase-out solution by 17%. This activity sensor example shows that, together with the additional data, quality problems can occur, as the data might have been estimated or generated from a source that did not have that specific data usage in mind. In such a case, new data management routines should be implemented, leading to gradual quality improvement.

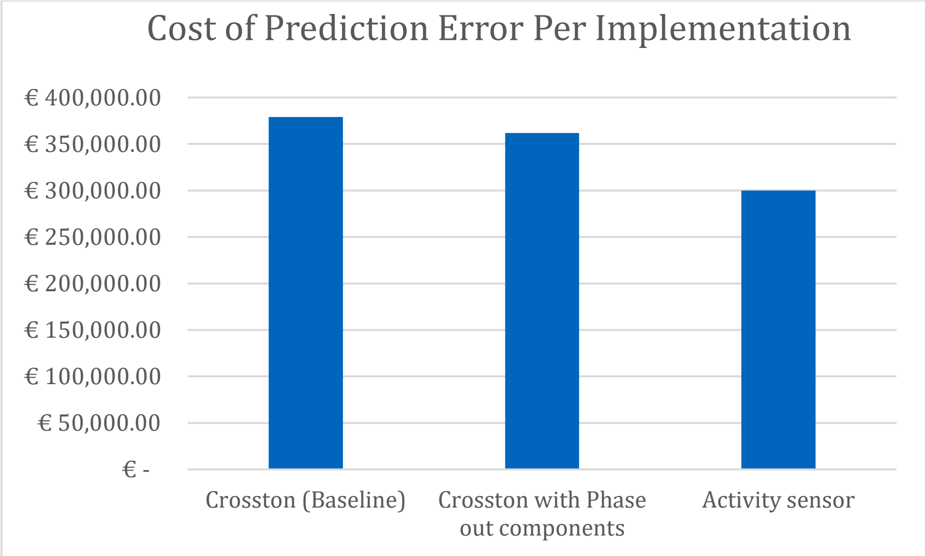

Figure 7. Result summary of Instantiation 1 of the CAPS framework

5.3 CAPS Use Case 2 – Anticipated later work

We used the CAPS framework a second time to review the feasibility of implementation of more advanced system designs. However it should be noted that these new designs could not be fully developed and evaluated within the time frame of the project. However, the overall goal of this exercise remains the same – designing, developing, and continuously improving a system using CAPS that can predict sales in the given context. To exemplify the benefit of the CAPS framework, we will use it to guide our evaluation in this section. Nonetheless, as new design artifacts will not be fully functional, there will be no means to quantitatively evaluate their performance as in the previously presented use of CAPS, but without a full deployment (step 4). As a result, a new quantitative evaluation function needs to be developed that, on the one hand, is able to provide insights into evaluated designs and, on the other, can be executed even without complete system implementation.

The idea behind the proposed objective function revolves around the concept that when it is not possible to calculate the cost reduction due to improved forecast, it is still possible to define what that reduction should be to make a given investment in a new system feasible: a requirement for developing a sensor-enabled forecasting system should be that the forecast error cost difference $\Delta COST_{FE}$ between the cost of forecast error before system implementation, $COST_{FE\_BEFORE}$, and the cost of forecast error after implementation, $COST_{FE\_AFTER}$, should out-weigh the system's implementation cost, $COST_{INV}$, within a reasonable time. Our enquiry in the case company about what exactly that reasonable time payback period should result in an interval from two to five years.

We decided to use a five-year threshold period because other innovative solutions toward better customer intelligence are currently within the strategic focus of the company. As a result, the new objective function becomes:

$$5 \cdot \Delta \mathrm{COST_{FE}} > \mathrm{COST_{INV}} \Leftrightarrow$$

$$\mathrm{COST_{FE\_AFTER}} < \mathrm{COST_{FE\_BEFORE}} - \frac{\mathrm{COST_{INV}}}{5}$$

5.3.1 Two iterations of advanced artifact designs

The main objective was to assess the feasibility of a solution based on the intended development environment. Two open projects from the case environment are evaluated: 1) continuous multi-sensor monitoring, a system collecting data from hundreds of simple sensors, and 2) a new interface for remote engine monitoring. The continuous multi-sensor monitoring system is based on a solution implemented in MAN that provides output from approximately 500 sensors installed on every engine design according to MAN specification, initially to facilitate the monitoring process for on-site engineers. The majority of sensors will monitor the temperature in various points of the engine, but there are some that monitor air pressure, angle of rotating elements, or other environmental settings. If the engine owner enables remote monitoring, MAN starts pooling the data into a central data warehouse that is used to detect abnormal sensor values, which trigger alarms.

Ideally, a forecasting model based on such a monitoring system could rely on measures observing the wear of spare parts. Subsequently, by means of general forecasting, the historical values could be extended into the future, estimating when a certain threshold value of wear is reached. As an example, a cylinder liner is a part of an engine installed as an inner wall of a cylinder, holding a piston in one plane and providing minimum friction when the piston is in motion. Currently, the decision to replace a cylinder liner is based on manual inspection and measurement of its thickness, for which an engineer needs to dismantle the engine and measure the exact thickness of the liner.

In relation to forecast level improvement justifying investment in the system, we can use the objective function from the advanced design of artifacts using CAPS and MAN historical data to evaluate it. Based on the current operational values, including stock-out percentage and surplus inventory value, the overall estimated cost of over-forecast is 1.8M€, and the cost of under-forecast is 4.3M€, totaling 6.1M€. With an investment cost estimate of 2.5M€, the requested payback period of five years is attainable with a forecast cost improvement of 8% (from 6.1 to 5.6 M€), assuming no additional maintenance cost of the new solution. However, such a significant lowering of forecast error cost might be difficult to gain when including additional maintenance costs of the new solution – assuming that 10% of the installed base will be monitored through the interface (about 2,000 engines) and that, with one remote diagnostic session costing 100€, including human cost and satellite communication cost, the actual forecast error cost after implementing the system needs to be 200k€ less, translating to 30% improvement (6.1M€ to 4.3M€).

The other system that was evaluated is the remote engine monitoring interface project. This project provides local engineers who need to perform system diagnosis of a MAN engine in a remote location (e.g., in a ship at sea) with a monitoring device built into a protective cap (helmet), equipped with a camera, microphone, and other necessary sensors, that is connected and networked to MAN headquarters, where experienced MAN engineers can “see” and “hear” the engine as if they were on site. Subsequently, the experienced MAN engineers are able to diagnose complex problems with the engine almost immediately and to guide the local engineers on how to solve them, thereby reducing expensive potential downtime costs of an engine on a running vessel. The project is currently in the prototype stage, and further development awaits market evaluation and input.

Overall, based on the cost assessment grounded in the case environment, we calculated a forecast improvement that would result in investment break-even, which amounts to 8% for the continuous sensor monitoring system and 30% for the remote engine monitoring project. Based on the improvement factors from the first CAPS instantiation, the first solution seems much more realistically feasible than the second.

Evaluating and reviewing two rounds of iteration of the Later Work evaluation (Figure 6) and the associated design cycles, we presented five predictive context-aware information systems using the following sensor technologies: a state-of-the-art Croston implementation, a Croston with phase-out component, an Activity sensor, the continuous multiple-sensor monitoring system, and the remote engine monitoring interface. The process of designing and analyzing the five artifacts (despite being specific to a single case and a particular, well-defined problem) not only led to gradual performance improvement but also uncovered some general aspects of designing IS predictive artifacts with sensor technologies that we discuss further in the following section.

6 Discussion

The application of the CAPS framework to the case provided an opportunity to evaluate the framework and engage in further reflection. Our experience with its application and evaluation in the MAN Diesel & Turbo case offers the opportunity to develop knowledge propositions that can be helpful in designing similar predictive systems in other companies and other problem domains. The propositions aim to help organizations to implement the CAPS framework and to realize its value in their context.

The first proposition stems from our design decision to utilize sensor data from already established sources and other data already collected for a different purpose.

This approach not only greatly lowers the cost and simplifies the implementation process but also ensures a certain level of data quality; if the data are already in use, they are being reviewed, and some potential quality issues might be spotted and corrected. The risk factor for the investment is also significantly lower; at the time of design the exact format and structure of the data are available, helping to avoid any surprises in terms of practical sensor limitations.

Generalizing these observations beyond the discussed case led us to formulate the first proposition for designing context-aware predictive systems:

Proposition 1 (Infrastructure Reusability): The readily available sensor data streams should be utilized to cut infrastructure costs, reduce system complexity, and ensure a certain level of data quality.

While developing an activity-sensor-based predictive system, our focus was on data quality and data estimation methods. Our experience shows that when sample size and data quality problems are predicted upfront, they can be mitigated, for example, by redesigning sensor infrastructure or data projection, capturing essential relationships from a small data sample, and encapsulating it into an estimation function (as in our implementation of activity sensor design). Generalizing these observations beyond the discussed case led us to formulate the second proposition for designing predictive systems:

Proposition 2 (Data Quality): Data quality measurement mechanisms should be provided as soon as possible to substantiate the predictive algorithms' assurance and mitigate prediction validity concerns.

The key takeaway from investigating a multi-sensor continuous monitoring system is that, when using sensor technologies, it is crucial to be clear about what is the most interesting variable data to sense and how this sensing activity should be implemented. Our example has shown that when a system measures a phenomenon of interest through a proxy measure, such as measuring the wear of piston rings through temperature around the rings, the process of analyzing the collected data might be both significantly and unnecessarily hindered. Generalizing these observations beyond the discussed case led us to formulate the third proposition for designing predictive systems:

Proposition 3 (Choice of Sensors): Sensors that measure metrics of interest should be selected to optimally balance between artifact predictive ability and complexity.

The insights related to implementing a remote monitoring interface evolve around the cost of forecast error. In the IoT world, it is common to dive into a sensor project merely because such projects are popular, without a well-defined project plan and benefits definition. For MAN, the design proved to be quite expensive, in terms of both initial implementation and running cost of the system, while some steps along the way were not completely focused on what factors added the most value. In general, the benefit of having the system in place must outweigh the cost, but this can only happen in an environment wherein the cost of prediction is very high and can be optimized. Generic solutions, such as our baseline approach or those presented in the literature review section, will sometimes be more affordable to implement, as they do not always require investment in new sensor infrastructure. This leads us to suggest that sensor-based predictive systems will usually be feasible for environments with high prediction error cost, typically characterized by high uncertainty, and we have formulated that observation as the fourth proposition for designing predictive systems:

Proposition 4 (Contextual Features): Context-aware predictive features should be developed to improve the prediction quality only when the estimated cost of prediction error is higher than the implementation cost.

All four propositions are presented in Table 2.

Table 2. Summary of Knowledge Propositions for Designing Context-Aware Predictive Systems with the CAPS Framework

Proposition 1: Infrastructure Reusability *The readily available sensor data streams should be utilized to cut infrastructure costs, reduce system complexity, and ensure a certain level of data quality.*
Proposition 2: Data Quality *Data quality measurement mechanisms should be provided as soon as possible to substantiate the predictive algorithms' assurance and mitigate prediction validity concerns.*
Proposition 3: Choice of Sensors *Sensors that measure metrics of interest should be selected to optimally balance between artifact predictive ability and complexity.*
Proposition 4: Contextual Features *Context-aware predictive features should be developed to improve the prediction quality only when the estimated cost of prediction error is higher than the implementation cost.*

Some additional *generalizations* regarding the financial feasibility of sensor-enabled predictive systems can be established based on the evaluated case. Building on Proposition 4, we refer to financial feasibility as the estimated difference between the cost and the benefit provided by the investment. Clearly, the latter is difficult to estimate reliably before the system is in place; thus, we select the cost of investment as the primary dimension in our feasibility analysis. Moreover, revisiting Proposition 3, the system can provide financial benefits when sensors provide added value to the prediction, which can be best achieved when the phenomenon of interest is measured directly (or, from the systems perspective, when the measure immediacy is high). The ability to obtain direct meaningful and unambiguous measures should be considered during the design phase, as also indirectly advocated in Proposition 2. Hence, the resulting two-by-two matrix for pre-assessing financial feasibility of predictive systems builds on two dimensions: investment cost and measure immediacy[2] (Figure 8).

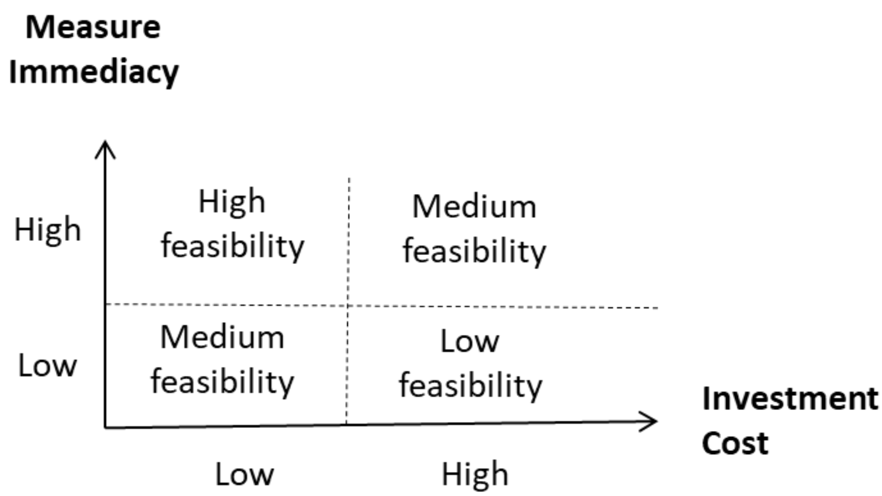

Figure 8. Matrix for Pre-Assessing Financial Feasibility of Context-Aware Predictive Systems

In evaluating the previously presented sensor-based solutions based on the matrix, we observe that artifact 1, Croston with phase-out component and activity sensor, will fall into the low investment cost half of the matrix (the left-hand side). The activity sensor, due to the need for estimation and the indirectness of data values, will fall into the medium feasibility quadrant, while the Croston with phase-out component will be classified into the high feasibility quadrant. The designs presented in the future work section will, on the other hand, fall into the right-hand side of the matrix. A multi-sensor continuous monitoring system, as described in the previous section (due to using a proxy measure of wear, i.e., temperature instead of the actual surface worn), will fall into the low feasibility quadrant, while the flexible monitoring system (monitoring exactly what makes the most sense in a given context) falls into the medium feasibility quadrant. The matrix in Figure 8 can also be useful to visualize possibilities for increasing the financial feasibility of a system, as the aim of the predictive IS designer is to move the system up and to the left.

The CAPS framework was evaluated based on a case that showed that it can provide useful guidelines to develop an environment-specific sensor-based predictive model that can outperform, in a given environment, the state-of-the-art generic predictive methods. Furthermore, in the context of Gregor & Hevner's (2013) knowledge contribution matrix, the CAPS framework falls within the "improvement" category. That is, the CAPS framework helps to develop new solutions for a known problem. Overall, the CAPS framework sheds light on an overlooked design blind spot pertaining to systems that deal with sensor data for predictive analytics problems within an organizational structure.

Another purpose of this research project is to theorize about the design of context-aware predictive systems. The systems that are in the scope of this research project are a class of IS artifacts characterized by the

[2] Immediacy is used to refer to how directly a phenomenon in question is measured. For instance, if engine operating temperature is readily available via dedicated sensors, then temperature has a high immediacy. If measuring the engine temperature is also not expensive, then the high immediacy and the low cost result in high feasibility (see Figure 8).

use of data describing the context of operation (e.g., from sensors) and using predictive methods. The utility of the CAPS framework rests on the fact that it can make the design process more complete and easier to manage, while mitigating most common errors and ensuring the quality of the outcome.

The presented contribution fulfills this predefined purpose. The CAPS framework for designing sensor-based predictive systems creates a starting point for academics and practitioners alike who are interested in designing such a system. The recommendations encompassed in the set of propositions can be valuable to mitigate some common errors, and the pre-assessment matrix helps to create a frame of reference when exploring new design opportunities.

It is likely that custom, context-specific predictive designs will continue to gain popularity in response to the inevitable growth of sensor dissemination and the digital traces they produce (Uckelmann, Harrison, & Michahelles, 2011). In this paper, we introduced a framework facilitating the process of designing context-aware predictive systems. The CAPS framework can provide both a structure and guidelines for developing context-aware systems that are environment-specific sensor-based predictive systems.

Naturally, the CAPS framework provides valuable guidance to practitioners. The knowledge propositions serve as guiding principles to design predictive IS artifacts that can result in shorter design and implementation times and lower the cost of implementation. As the popularity of context-aware predictive systems increases in the industry, the CAPS framework enables a better design of quality predictive systems. This potential improvement may also have significant implications for the forecasting field, especially spare-parts forecasting (Yin et. al., 2014), boosting explanatory forecasting method development (Sroginis, 2021) and unwrapping the black-boxed, time-series-based forecasting routines of the past.

7 Conclusion

Although digital sensing technologies have become ubiquitous and increasingly embedded in industrial environments, the ubiquity of smart objects that produce ever-growing streams of contextual sensing data still presents challenges for forecasting systems. We applied the Design Science Research methodology to develop a framework for predictive analytics on contextual data. Specifically, we developed a framework and a set of propositions for designing and evaluating context-aware predictive systems. The framework can be generalized to similar problems in related industrial domains.

Clearly, the usability and value of the proposed framework depend on its applicability to other case environments. Although a single case inherently poses a threat to external validity and generalization, in this study the case-anchored system development and evaluation served as a proof-of-concept and an exemplification (Zhang et al., 2002). Additionally, the scope of application of the framework may be perceived as somewhat limited, given that it focuses on predictive analytics of context-aware data. However, the growing availability of contextual information and fast-spreading sensor-based technologies are likely to help popularize predictive analytics with sensors.

Future work can focus on further replication and validation of the framework in new environments by instantiating it in different cases. MAN Diesel & Turbo is a large global company and a market leader in its target market, and thus the experience of the CAPS framework in the MAN case can possibly translate to new insights with regard to predictive maintenance and demand forecasting in industries such as airline maintenance, wind power generation, and off-shore drilling. More specifically, further contextual evaluation of systematic biases in other settings can provide insights to fine-tune the findings from the MAN case. Testing new sensor technologies can also provide a significant improvement in sensor-based predictions. Lastly, combining sensor data with other sources of information available in open data sources (Marton et al. 2013), such as pattern recognition algorithms, simulations, and even social data, may help to further fine-tune the performance of Context-Aware Predictive Systems and help managers to make better-informed decisions. In summary, with the growing ubiquity of sensor data, we believe that the utilization of predictive analytics on contextual data will become prevalent in organizational settings, and we hope that the CAPS framework and the propositions provided herein offer useful insights into the future development of Context-Aware Predictive Systems.

Acknowledgments

The authors are grateful to the senior editor Sandeep Purao and the anonymous reviewers for their support and valuable comments on earlier versions of this paper.

References

Aarts, E., & Marzano, S. (2003). *The new everyday: Views on ambient intelligence*. Rotterdam: 010 Publishers.

Abowd, G., Dey, A., Brown, P., Davies, N., Smith, M., & Steggles, P. (1999). *Towards a better understanding of context and context-awareness.* Proceedings of the First International Symposium on Handheld and Ubiquitous Computing, 304–307, Karlsruhe, Germany.

Aldahiri, A., Alrashed, B., & Hussain, W. (2021). Trends in using IoT with machine learning in health prediction system. *Forecasting*, *3*(1), 181-206.

Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega, 40*(6), 722-737.

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

Callegaro, A. (2010). *Forecasting Methods for Spare Parts.* (Doctoral dissertation) Universita'degli Studi di Padova. http://tesi.cab.unipd.it/25014/1/TesiCallegaro580457.pdf

Chatterjee, S. (2015). *Writing my next Design Science Research Master-piece: But How Do I Make a Theoretical Contribution to DSR?* Proceedings of the 22nd European Conference on Information Systems, Munster, Germany.

Chen, I. (2001). Planning for ERP systems: Analysis and future trend. *Business Process Management Journal, 7*(5), 374-386.

Cook, D. J., & Das, S. K. (2004). *Smart environments: Technology, protocols and applications*. Hoboken, New Jersey: John Wiley & Sons, Inc.

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. *International Journal of Production Economics*, 143(2), 536-545

Deming, W. E. (2000). *The new economics: For industry, government, education.* Cambridge MA: MIT press.

Dey, A. K. (2001). Understanding and using context. Personal and ubiquitous computing, 5(1), 4-7.

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS quarterly*, *37*(2), 337-355.

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

Gungor, V. C., & Hancke, G. P. (2009). Industrial wireless sensor networks: Challenges, design principles, and technical approaches. *IEEE Transactions on industrial electronics*, 56(10), 4258-4265.

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., & Chatterjee, S. (2010). *Design research in information systems: theory and practice.* New York: Springer Science & Business Media.

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

IHS Global Services (2009). *IHS Global Insight report.* Lexington, MA: IHS

Javaid, M., Haleem, A., Rab, S., Singh, R. P., & Suman, R. (2021). Sensors for daily life: A review. *Sensors International*, 2, 100121.

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

Kim, S., Lee, M., & Shin, C. (2018). IoT-based strawberry disease prediction system for smart farming. *Sensors*, *18*(11), 4051.

Kourentzes, N. (2013) Intermittent demand forecasts with neural networks. *International Journal of Production Economics,* 143(1), 198-206.

Kumar, B. A., & Sharma, B. (2020). Context aware mobile learning application development: A systematic literature review. *Education and Information Technologies*, 25(3), 2221-2239.

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

Marabelli, M., Hansen, S., Newell, S., & Frigerio, C. (2017). The Light and Dark Side of the Black Box: Sensor-based Technology in the Automotive Industry. *Communications of the*

Association for Information Systems, 40, doi:10.17705/1CAIS.04016

Maritime-Insight. (2013). Shipbuilding outlook report. London: Lloyd's List Intelligence

Marton, A., Avital, M. and Jensen, T. (2013), Reframing Open Big Data, *Proceedings of the 21th European Conference on Information Systems (ECIS)*, Utrecht, Netherlands.

Meth, H., Mueller, B., & Maedche, A. (2015). Designing a requirement mining system. *Journal of the Association for Information Systems, 16*(9), 799-837.

Minner, S. (2011). Forecasting and inventory management for spare parts: An installed base approach. *ts*In Altay, N. Litteral L. (Eds.) *Service Par Management* (pp. 157-169). London: Springer.

Müller-Wienbergen, F., Müller, O., Seidel, S., & Becker, J. (2011). Leaving the beaten tracks in creative work - A design theory for systems that support convergent and divergent thinking. *Journal of the Association for Information Systems, 12*(11), 714-740.

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

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

Nikkam, S. M. G. &. Pawar, V.R. (2016). Water parameter analysis for industrial application using IoT, *Proceedings of the International Conference on Applied and Theoretical Computing and Communication Technology (iCATccT)*, Bangalore, India, pp. 703-707, doi: 10.1109/ICATCCT.2016.7912090.

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.

Okano, M. T. (2017). IOT and industry 4.0: The industrial new revolution. *Proceedings of International Conference on Management and Information Systems, 25, 75-82*.

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

Pech, M., Vrchota, J., & Bednář, J. (2021). Predictive maintenance and intelligent sensors in smart factory. *Sensors*, *21*(4), 1470.

Perera, C., Zaslavsky, A., Christen, P., & Georgakopoulos, D. (2013). Context aware computing for the internet of things: A survey. IEEE communications surveys & tutorials, 16(1), 414-454.

Petropoulos, F., Nikolopoulos, K., Spithourakis, G. P., & Assimakopoulos, V. (2013). Empirical heuristics for improving intermittent demand forecasting. *Industrial Management & Data Systems*, 113(5), 683-696.

Pries-Heje, J. & Baskerville, R. (2008). The design theory nexus. *MIS Quarterly*, 32(4), 731-755.

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

Sampath, H. U., Erandi, I. A. S., Jinasena, T. M. K. K., & Panditharathna, J. H. (2019). Intelligent Traffic Light System Using Wireless Sensor Networks and IoT. *Proceedings of the 6th International Conference on Multidisciplinary Approaches (iCMA).* https://ssrn.com/abstract=3496594

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. *Naval Research Logistics, 43*(3), 381-396.

Shim, J.P., Avital, M., Dennis, A., Rossi, M., Sorensen, C., & French, A. (2019) The Transformative Effect of the Internet of Things on Business and Society, Communications of the AIS, 44(1), 129-140.

Sroginis, A. (2021). *The use of contextual information in demand forecasting*. (Doctoral dissertation, Lancaster University, UK).

Suma, V. (2021). Internet-of-Things (IoT) based smart agriculture in India-an overview. *Journal of ISMAC*, *3*(01), 1-15.

Storey, V. C., Burton-Jones, A., Sugumaran, V., & Purao, S. (2008). CONQUER: A methodology for context-aware query processing on the World Wide Web. *Information Systems Research*, (19)1, 3-25.

Teunter, R. H. & Duncan, L. (2009) Forecasting intermittent demand: a comparative study. *Journal of the Operational Research Society* 60.3: 321-329.

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.

Uckelmann, D., Harrison, M., & Michahelles, F. (2011). An architectural approach towards the future internet of things. In *Architecting the Internet of Things* (pp. 1-24). Springer Berlin Heidelberg.

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

Weiser, M. (1991). The computer for the 21st century, *Scientific American*, 265(3), 66–75.

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

Zhang, P., Benbasat, I., Carey, J., Davis, F., Galletta, D., & Strong, D. (2002). AMCIS 2002 panels and workshops I: Human-computer interaction research in the MIS discipline. *Communications of the Association for Information Systems, 9*(1), 20.

Appendix

Calculating the cost of forecast error

The key to a legitimate quantitative evaluation of a design is a meaningful objective function that can also work in a given organizational context. 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 asymmetry in the error cost. To address this diversity, two separate cost functions – over- and under-forecasting scenarios – are necessary. The actual cost associated with an under-forecasting situation occurs due to missed sales potential: because goods are not available when demand occurs, some customers will decide to drop the order rather than wait for the items. The percentage of customers dropping orders can be determined by the difference in the conversion ratio of quotes to orders (also referred to as hit rate) for in-stock quotes versus stock-out quotes.

For example, in the component group of piston ring sales at MAN, in a scenario where goods are in stock, the average hit rate is X, but in the case of a stock-out only Y of the issued quotes would convert to orders, where Y is 9 percent point lower than X, and, *ceteris paribus*, 9% of customers gave up the purchase probably due to lack of availability. In order to calculate lost profit, the average hit rate difference between in-stock quotes (**HR_{is}**) and the stock-out hit rate factor (**HR_{so}**) needs to be multiplied by under-forecasted volume **U_{vol}** (to compute sales volume missed due to stock-out) and, to convert sales turnover to earnings before interest and taxes EBIT, multiplied by the 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 the cost of potential depreciation and scrap, both proportional to over-forecast volume (**OF_{VOL}**). The opportunity cost is experienced because the investment in inventory is unnecessary and the capital can be invested differently, bringing certain profit to the company. Most of the firms have some baseline working capital ratio (**OC**) to be used for such calculations. In the context of MAN Turbo & Diesel, sales are expected every month, so over-forecasting in one month will lead to lower replenishment costs in the following month, and the frozen capital cost will be calculated for a single month. In the case of the cost of depreciation and scrap factor (**DF**), this reflects a possibility that unsold inventory will not move for a period of time. This leads the inventory to be written off by a certain depreciation factor, and, if parts are no longer sellable, to be written down completely and scrapped. Putting all the parameters together, the cost of forecast error, **$COST_{FE}$**, in the case context can be described as:

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

Developing three artifacts to minimize the cost of forecast error

The baseline "Croston" method had to be improved. This was done by exploring sales data as depicted in Figure A1. A product was sold four times in the baseline year – in January, March, June, and October – in quantities of 8, 7, 8, and 9 pieces, respectively. The forecast for the subsequent year was calculated in two steps: by calculating a demand magnitude when demand occurs and an interval between demand points. Then, the forecasted magnitude was computed based on historical quantities (the original method uses exponential smoothing; here, for clarity, we use an average), resulting in eight pieces (average of 8, 7, 8 and 9). The same was done for intervals between historical sales: the sale in March was two months after the one in January; the one in June was three months after the one in March; and the one in October was four months after the one in June. The mean of these intervals is three months, on average, between demand points. The forecast is the calculated magnitude spaced by the calculated interval (see Figure A1). This kind of calculation can be performed for every product.

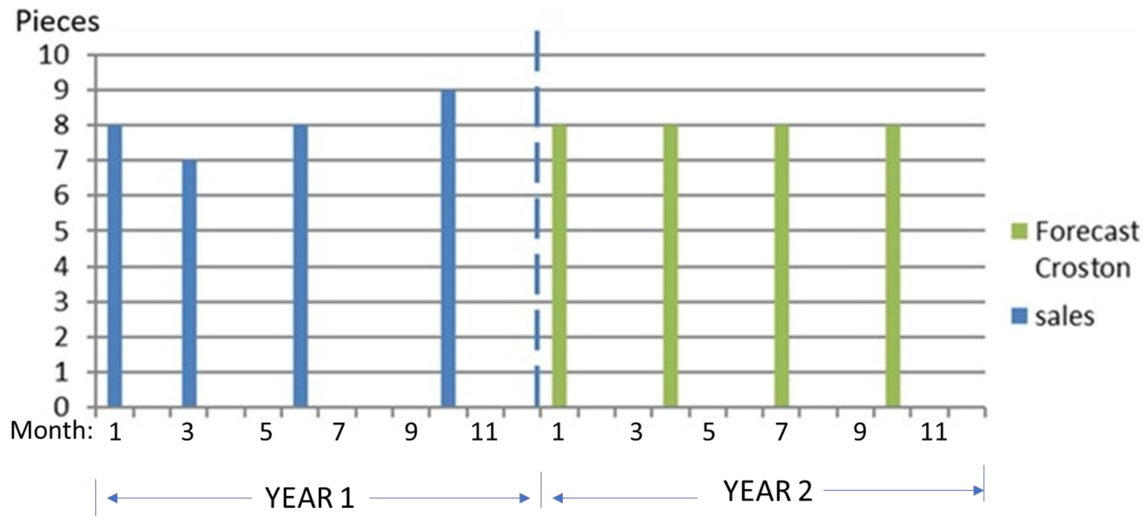

Figure A1. Croston Base-Line Method Forecast

During the qualitative evaluation of the Croston solution, the demand planning team from the organization pointed out that the test data set included years around the financial crisis, but the purchase patterns before and after the crisis were significantly different. According to their insights, as a consequence of financial crisis, the number of shipments dropped significantly, shrinking the margins of shipping companies. As a result, the shippers invested in new, bigger, and more fuel-efficient vessels, causing a massive turnover of ships. The difference in the engine product mix between the test and learning samples could lead to avoidable variance and lower quality of prediction in the test sample. To control for the changes in the product mix, we implemented an additional feature in the Croston model that removed from the test and the learning data sets the engine installations that were phased out. We could track the changes in the active product mix by tracing insurance registries data about the registration and removal of engine by ship owners.

The new model, called "Croston with a phase-out component" (see Artifact 1.2, Figure 6), was implemented and evaluated using a previously developed objective function, improving the baseline prediction by 4%. The quantitative evaluation was followed by the qualitative one. Demand planners, together with sales engineers, drew our attention to market changes occurring during the period of the study. As an effect of the lower demand for transportation services, which was the result of the global financial crisis, the sailing patterns of most of the customers had changed. Rather than cruising with maximum frequency and speed, the vessel management companies prioritized cost reductions and maximizing load per vessel at the expense of the transport time. Moreover, to optimize fuel consumption and vessel wear, ships would travel with the most efficient, rather than maximum, speed. After the global economy started to recover from the crisis, the situation started to slowly return to the previous status quo. All these changes could potentially lead to a significant change in wear and demand for spare parts, "polluting" the learning set with discontinued behaviors that were not present in the test sample (Notteboom & Cariou, 2013; Yin et al., 2014).

About the Authors

Michel Avital is a Professor of Digitalization at Copenhagen Business School. Michel is an advocate of openness and an avid proponent of cross-boundaries exchange and collaboration. His research focuses on the relationships between digital innovation ecosystems and organizational practices. He studies how emergent technologies are developed, applied, managed, and used to transform and shape organizations. He has published more than 100 articles on topics such as blockchain technology, the future of work, sharing economy, open data, open design, generative design, creativity, and innovation. He is a senior editor and editorial board member of leading IS journals and serves in various organizing capacities at major international conferences on digital technology and organization studies. Michel is a recipient of the AIS Fellow Award. Further information: http://avital.net

Samir Chatterjee is a Professor and Fletcher Jones Chair of Technology Management and Design at Claremont Graduate University. He has published more than 160 peer reviewed articles in the IS and Computer Science area. He serves as co-EIC of the journal *Health Systems*. He is currently on the editorial boards of ISR and JAIS and has served as AE on MISQ and several IEEE/ACM journals. He is the co-author of the best-selling book *Design Research in Information Systems* and is the founder of the DESRIST conference. He regularly serves as a chair or member of the technical program committees at numerous IEEE, ACM, and AIS conferences. He has received over $3.6 million in funding from federal (NSF, NIH) and private foundations. He is the Founder & CEO of DCL Health, a digital healthcare startup. When he is not busy with science and design, he is an avid musician and composer. In 2021, his debut new age album under the stage name Samir Bodhi made it to the Grammy ballots for nomination. He lives in Southern California with his wife and son.

Szymon Furtak earned his PhD from Copenhagen Business School. His research focuses on practical aspects of advanced data usage in organizations, especially regarding building advanced models for predictions. As a practitioner, he has been involved in innovative approaches to spare part forecasting and supply chain development, among others, in MAN Diesel & Turbo and MAN Energy Solutions. He currently works with green energy transformation at Sunnon Limited and is primarily interested in balancing smart electrical grids in photovoltaic power plants.