Julius-Maximilians-Universität Würzburg
Wirtschaftswissenschaftliche Fakultät

Design and Evaluation of Data-Driven Enterprise Process Monitoring Systems

Inauguraldissertation
zur Erlangung des akademischen Grades
doctor rerum politicarum (Dr. rer. pol.)

vorgelegt von
Felix Oberdorf, M.Eng.
geboren in Würzburg

Name und Anschrift:	Felix Oberdorf Keesburgstraße 4 97074 Würzburg
Erstgutachter:	Prof. Dr. Christoph M. Flath
Zweitgutachter:	Prof. Dr. Axel Winkelmann
Datum der Einreichung:	30. Mai 2022

Abstract

Increasing global competition forces organizations to improve their processes to gain a competitive advantage. In the manufacturing sector, this is facilitated through tremendous digital transformation. Fundamental components in such digitalized environments are process-aware information systems that record the execution of business processes, assist in process automation, and unlock the potential to analyze processes. However, most enterprise information systems focus on informational aspects, process automation, or data collection but do not tap into predictive or prescriptive analytics to foster data-driven decision-making. Therefore, this dissertation is set out to investigate the design of analytics-enabled information systems in five independent parts, which step-wise introduce analytics capabilities and assess potential opportunities for process improvement in real-world scenarios.

To set up and extend analytics-enabled information systems, an essential prerequisite is identifying success factors, which we identify in the context of process mining as a descriptive analytics technique. We combine an established process mining framework and a success model to provide a structured approach for assessing success factors and identifying challenges, motivations, and perceived business value of process mining from employees across organizations as well as process mining experts and consultants. We extend the existing success model and provide lessons for business value generation through process mining based on the derived findings. To assist the realization of process mining enabled business value, we design an artifact for context-aware process mining. The artifact combines standard process logs with additional context information to assist the automated identification of process realization paths associated with specific context events. Yet, realizing business value is a challenging task, as transforming processes based on informational insights is time-consuming.

To overcome this, we showcase the development of a predictive process monitoring system for disruption handling in a production environment. The system leverages state-of-the-art machine learning algorithms for disruption type classification and duration prediction. It combines the algorithms with additional organizational data sources and a simple assignment procedure to assist the disruption handling process. The design of such a system and analytics models is a challenging task, which we address by engineering a five-phase method for predictive end-to-end enterprise process network monitoring leveraging multi-headed deep neural networks. The method facilitates the integration of heterogeneous data sources through dedicated neural network input heads, which are concatenated for a prediction. An evaluation based on a real-world use-case highlights the superior performance of the resulting multi-headed network.

Even the improved model performance provides no perfect results, and thus decisions about assigning agents to solve disruptions have to be made under uncertainty. Mathematical models can assist here, but due to complex real-world conditions, the number of potential scenarios massively increases and limits the solution of assignment models. To overcome this and tap into the potential of prescriptive process monitoring systems, we set out a data-driven approximate dynamic stochastic programming approach, which incorporates multiple uncertainties for an assignment decision. The resulting model has significant performance improvement and ultimately highlights the particular importance of analytics-enabled information systems for organizational process improvement.

Kurzzusammenfassung

Der zunehmende globale Wettbewerb zwingt Unternehmen zur Verbesserung ihrer Prozesse, um sich dadurch einen Wettbewerbsvorteil zu verschaffen. In der Fertigungsindustrie wird das durch die die digitale Transformation unterstützt. Grundlegende Komponenten in den entstehenden digitalisierten Umgebungen sind prozessorientierte Informationssysteme, die die Ausführung von Geschäftsprozessen aufzeichnen, bei der Prozessautomatisierung unterstützen und wiederum Potenzial zur Prozessanalyse freisetzen. Die meisten Informationssysteme in Unternehmen konzentrieren sich jedoch auf die Anzeige von Informationen, Prozessautomatisierung oder Datenerfassung, nutzen aber keine “predictive analytics” oder “prescriptive analytics”, um datengetriebene Entscheidungen zu unterstützen. Daher wird in dieser Dissertation der Aufbau von “analytics-enabled” Informationssystemen in fünf unabhängigen Teilen untersucht, die schrittweise analytische Methoden einführen und potenzielle Möglichkeiten zur Prozessverbesserung in realen Szenarien bewerten.

Eine wesentliche Voraussetzung für den Auf- und Ausbau von “analytics-enabled” Informationssystemen ist die Identifikation von Erfolgsfaktoren, die wir im Kontext von Process Mining als deskriptive Methode untersuchen. Wir kombinieren einen etablierten Process Mining Framework und ein Process Mining Erfolgsmodell, um einen strukturierten Ansatz zur Bewertung von Erfolgsfaktoren zu ermöglichen, den wir aufbauend zur Identifizierung von Herausforderungen, Motivationen und des wahrgenommenen Mehrwerts (engl. “Business Value”) von Process Mining durch Mitarbeiter in Organisationen und Process Mining Experten nutzen. Auf Grundlage der gewonnenen Erkenntnisse erweitern wir das bestehende Erfolgsmodell und leiten Implikationen für die Generierung von “Business Value” durch Process Mining ab. Um die Realisierung des durch Process Mining ermöglichten “Business Value” zu unterstützen, entwickeln wir ein Artefakt für kontextbezogenes Process Mining. Das Artefakt kombiniert standard Prozessdaten mit zusätzlichen Kontextinformationen, um die automatische Identifizierung von Prozesspfaden, die mit den Kontextereignissen in Verbindung gebracht werden, zu unterstützen. Die entsprechende Realisierung ist jedoch eine herausfordernde Aufgabe, da die Transformation von Prozessen auf der Grundlage von Informationserkenntnissen zeitaufwendig ist.

Um dies zu überwinden, stellen wir die Entwicklung eines “predictive process monitoring” Systems zur Automatisierung des Störungsmanagements in einer Produktionsumgebung vor. Das System nutzt etablierte Algorithmen des maschinellen Lernens zur Klassifizierung von Störungsarten und zur Vorhersage der Störungsdauer. Es kombiniert die Algorithmen mit zusätzlichen Datenquellen und einem einfachen Zuweisungsverfahren, um den Prozess der Störungsbearbeitung zu unterstützen. Die Entwicklung eines solchen Systems und entsprechender Modelle ist eine anspruchsvolle Aufgabe, die wir durch die Entwicklung einer Fünf-Phasen-Methode für “predictive end-to-end process monitoring” von Unternehmensprozessen unter Verwendung von “multi-headed neural networks” adressieren. Die Methode erleichtert die Integration heterogener Datenquellen durch dedizierte Modelle, die für eine Vorhersage kombiniert werden. Die Evaluation eines realen Anwendungsfalls unterstreicht die Kompetitivität des eines aus der entwickelten Methode resultierenden Modells.

Allerdings sind auch die Ergebnisse des verbesserten Modells nicht perfekt. Somit muss die Entscheidung über die Zuweisung von Agenten zur Lösung von Störungen unter Unsicherheit getroffen werden. Dazu können zwar mathematische Modelle genutzt werden, allerdings steigt die Anzahl der möglichen Szenarien durch komplexe reale Bedingungen stark an und limitiert die Lösung mathematischer Modelle. Um dies zu überwinden und das Potenzial eines “prescriptive process monitoring” Systems zu beleuchten, haben wir einen datengetriebenen Ansatz zur Approximation eines dynamischen stochastischen Problems entwickelt, der mehrere Unsicherheiten bei der Zuweisung der Agenten berücksichtigt. Das resultierende Modell hat eine signifikant bessere Leistung und unterstreicht letztlich die besondere Bedeutung von “analytics-enabled” Informationssystemen für die Verbesserung von Organisationsprozessen.

Acknowledgements

> *"It always seems impossible until it's done."*
>
> (Nelson Mandela)

Proving the quote would not have been possible without many people's help, encouragement, and support throughout the last years. First and foremost, I would like to thank my doctoral advisor, Prof. Christoph M. Flath, for his excellent supervision and the chance to prove the quote. I am incredibly grateful for his continuous guidance, support, and patience in pushing me to improve my work. Whether in terms of scientific papers, cooperative projects, or visualizations and presentations, there was great feedback that helped me improve my skills. Therefore also, special thanks to Prof. Richard Pibernik for the insightful comments and constructive feedback. In addition, I would like to thank my second advisor, Prof. Axel Winkelmann, for the valuable feedback regarding research design.

Furthermore, I would like to thank all colleagues at the Chair of Information Systems and Business Analytics and the Chair of Logistics and Quantitative Methods in Business Administration for constructive discussions and valuable feedback during and beyond our OPIM seminars. I want to express special thanks to Myriam Schaschek and Nikolai Stein for the great collaboration on our research.

Finally, I would like to express my gratitude to my family and my wife, Elisa, for their endless and invaluable support over the last few years. I am deeply grateful for her endless patience and loving care throughout this journey.

Contents

1 Introduction

Increasing global competition necessitates enterprises to improve their processes to gain a competitive advantage. In the manufacturing sector, this is facilitated through tremendous digital transformation. Ubiquitous computing, connectivity combined, and continuous data collection have created a next-generation industrial infrastructure (Feng and Shanthikumar 2018; Reddy 2016) that facilitates organizational process improvement. Fundamental components in such digitalized industrial infrastructures are process-aware information systems that record the execution of business processes, assist in their automation, and unlock the potential of process analytics (Baines, Lightfoot, and Kay 2009) to provide valuable process insights or support decisions.

Figure 1.1 provides an overview of such analytics-enabled systems as the analytics and information system stack in the context of organizational process improvement. It includes different perspectives such as problem complexity, information system's focus, or analytics' maturity level. Further, it can be aligned to intersections with related domains, e.g., operations management or application fields such as process monitoring.

In terms of *analytics*, descriptive, predictive, and prescriptive analytics are distinguished (Lustig et al. 2010; Evans and Lindner 2012; Holsapple, Lee-Post, and Pakath 2014). Descriptive analytics focus on analyzing historic data to quantify and visualize past process performance and identify anomalies through techniques such as exploratory data analysis (Tukey 1977) or process mining (van der Aalst et al. 2011b). Predictive analytics uses historic data to estimate future situations, resulting in predictive services (Baines, Lightfoot, and Kay 2009) or predictive process monitoring systems (Mehdiyev, Evermann, and Fettke 2020). Beyond predictions, prescriptive analytics combine predictive analytics and the field of operations management to evaluate scenarios for identifying optimal decision policies.

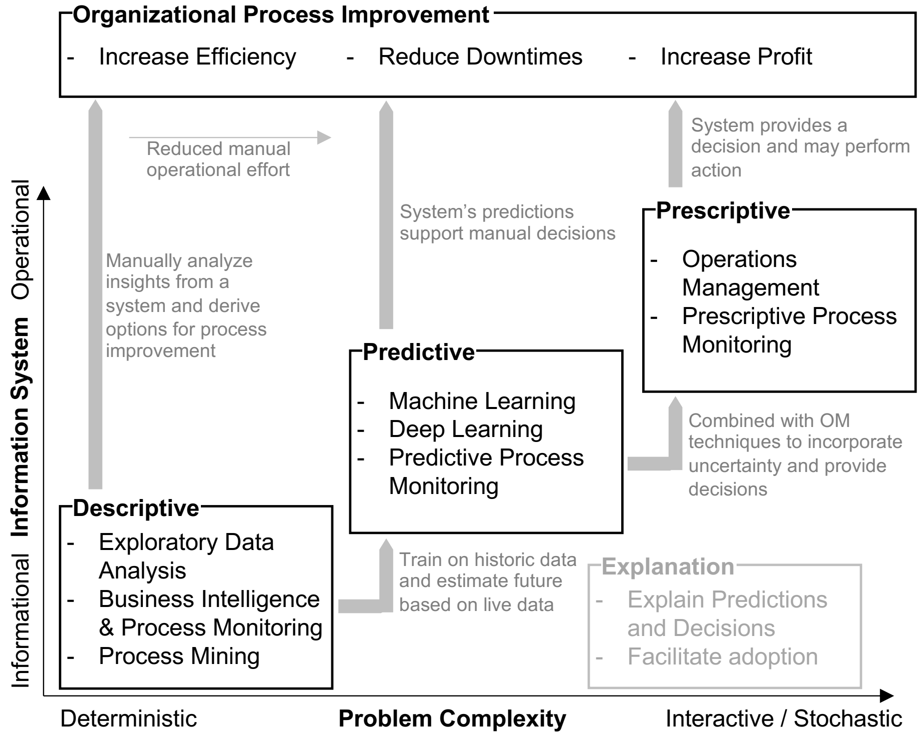

Figure 1.1: Analytics and information system stack.

From an *information systems* perspective, insights and decision support can be categorized whether a system has an informational or operational focus (Schwegmann, Matzner, and Janiesch 2013). Systems with an informational focus usually leverage descriptive analytics to gain process insights (Mehdiyev, Evermann, and Fettke 2020). Such insights can then be used to identify improvement potentials and trigger projects to improve processes. The informational perspective is also relevant beyond descriptive analytics. For instance, there is an increasing trend of explaining predictive or prescriptive models' predictions and decisions to facilitate a better understanding and user adoption (Arrieta et al. 2020; Senoner, Netland, and Feuerriegel 2021). Besides, there are operational-focused information systems for real-time support. Such systems leverage predictive and prescriptive analytics to assist processes or perform automated (prescriptive) actions, such as the “prescriptive control of business processes by using event-based process predictions” (Krumeich, Werth, and Loos 2016). In doing so, system complexity and interaction increase but, on the other side, lower the manual operational effort for organizational process improvement, such as increasing efficiency, reducing downtimes, or increasing the profit.

A case in point for using analytics-enabled information systems for process improvement is the efficient handling of disruptions during production processes. Given the complexity of today's manufacturing processes, problems and disruptions cannot be avoided entirely. For this reason, companies rely on disruption management systems to efficiently handle disruptions and, in turn, improve productivity (Lopez-Leyva et al. 2020). Such systems usually notify a responding agent to assist in solving a disruption (Macdonald and Corsi 2013). Typically, traditional disruption management systems only inform any responding agent to solve an initially unknown job caused by a disruption. If the type of a causing disruption was known, such a system could determine and assign the optimal agent, instead of any agent, and thus reduce downtimes. Yet, the system must operate within an organization's complex processes. Predictions in this context are not perfect, and the model's predictions include uncertainties that should be incorporated into the decision. To this end, an analytics-enabled disruption management system is envisioned that identifies the cause of a disruption and dispatches the optimal responding agent in the face of uncertainty. Going beyond predictions and providing prescriptive decision support could add value to disruption management but brings up the question of a prescriptive system's design.

1.1 Research Objectives

The importance of analytics-enabled systems, in general, is highlighted by Baines, Lightfoot, and Kay (2009) and Cheng and Johansen (2016) with the design of such systems as a resulting key question. Vater, Harscheidt, and Knoll (2019) extend this question and provide a "comprehensive review of key elements for prescriptive analytics in manufacturing" by pointing out two key aspects. First, the particular importance of prescriptive analytics and research—additionally emphasized by various calls for papers and special issues in leading journals (Giesecke et al. 2018; Hull et al. 2018; Sanders and Ganeshan 2015). Second and also underlined through calls for papers and special issues, a gap for research that presents the "current state-of-the-art in business analytics research and practice" (Gupta and Prakash 2001) and further with a particular process mining focus, how it is "... used and adopted at the enterprise level?" (vom Brocke et al. 2020). The combination of those findings motivates the guiding research objective of this thesis:

Guiding Research Objective *The design of a prescriptive process monitoring system for disruption management in production environments.*

To achieve the guiding research objective, I follow Gust et al. (2016) and "introduce analytics step-wise" by setting up and sequentially extending a system's analytics capabilities aligned with the analytics and information system stack. To do so, the initial focus is on a system for process improvement with descriptive analytics capabilities and factors that facilitate the system's design. With process mining as an appropriate descriptive analytics technique, the first subordinate research objective (**RO1**) is the *identification of success factors for process mining* in a real-world case study. The identified factors are subsequently leveraged to design and evaluate a process mining artifact.

The findings suggest that the analysis of historic data does provide insights into process improvement potentials. However, the transformation of processes is time-intensive and limits the system's use. I follow the analytics stack and try to overcome this limitation through predictive analytics and an operational focused system. Thereby, the second subordinate research objective (**RO2**) is the *design of a predictive process monitoring system for operational disruption handling*.

The designed system automates the disruption handling process and adds value through predictive analytics capabilities. Yet, the design of such a system and the analytics approaches is a challenging task, particularly considering practical needs such as the seamless combination of heterogeneous data sources or the model design through structured approaches. Recent advances in deep learning research, namely multi-headed neural networks, may bridge this gap from a technical perspective but require the design of such complex networks. To this end, the third subordinate research objective (**RO3**) is the definition of *guidelines for the design of multi-headed predictive end-to-end process monitoring models* in a real-world use case.

Based on the guidelines, predictive process monitoring models can be designed with a particular focus on an organization's needs. The system can combine heterogeneous data sources for predictions but should ultimately incorporate additional information, such as expected duration and availability, for the assignment decision of a responding agent. Yet, predicting such target variables is usually a challenging task entailing uncertainties. Incorporating such uncertainties in the decision may add value to the desired prescriptive process monitoring system but result in many potential scenarios, that must be evaluated—ultimately resulting in a prohibitively large state-space. To deal with such a state-space explosion, the fourth subordinate research objective (**RO4**) is the development of a *prescriptive analytics method to approximate the dynamic stochastic maintenance job assignment problem*.

1.2 Structure

This thesis consists of five independent articles that contribute to prescriptive process monitoring systems research.[1] Figure 1.2 provides an overview of this work's chapters aligned to the guiding and the four subordinate research objectives.

The first article, "Success Factors for Process Mining—A Multiple Case Study" (Chapter 2), addresses RO1 by identifying success factors for process mining as a descriptive analytics approach. In doing so, challenges and motivations for process mining are identified with employees of hierarchies from the production team up to the chief innovation officer and process mining experts and consultants.

The second article, "Context-Aware Process Mining — Disrupting Continuous Process Improvement" (Chapter 3), also addresses RO1 by applying the previously established findings to design a context-aware process mining artifact. We add to the understanding of process mining in practice through a quantitative and qualitative evaluation and point out the potential of leveraging context information for continuous process improvement.

The third article, "Analytics-Enabled Disruption Management: System Development and Business Value Assessment" (Chapter 4), addresses the design of a disruption management system with predictive analytics capabilities (RO2) to improve the operational disruption handling process. To explore the poten-

[1]See Appendix A for an exhaustive list of publications.

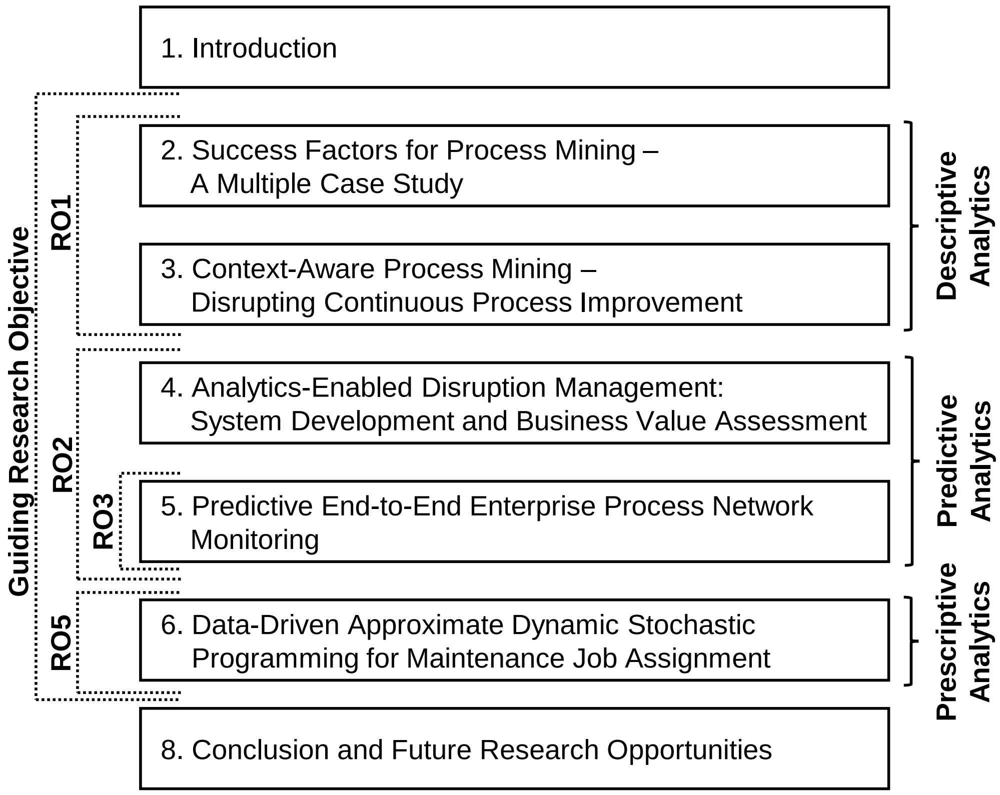

Figure 1.2: Chapters and structure of the thesis.

tial of such a predictive process monitoring approach, we illustrate the system's design through the action design research method (Sein et al. 2011) and evaluate the performance utilizing real-world manufacturing scenarios.

The fourth article, "Predictive End-To-End Enterprise Process Network Monitoring" (Chapter 5), extends RO2 and particularly focuses on RO3 by engineering an end-to-end method that guides the design of multi-headed neural network models for predictive process monitoring with heterogeneous data sources. The resulting method's five-phases include guidelines and suggestions for problem specification, data acquisition and preparation, model design, model evaluation, and model application in the organizational context. Subsequently, we demonstrate the usage of the method in a real-world use case with multiple context-aware event logs and additional disruption context information.

The fifth article, "Data-Driven Approximate Dynamic Stochastic Programming for Maintenance Job Assignment" (Chapter 6), presents the development of a prescriptive analytics approach to solve a complex dynamic stochastic assignment problem with multiple uncertainties (RO4). We combine weighted sample average approximation (wSAA) methods and leverage disruption context data to incorporate uncertainty about a maintenance job's duration and to account for a second uncertainty about upcoming disruption events through similarities in the context-aware event logs. We approximate the optimal solution and complement the prescriptive process monitoring system for disruption handling through the decision for the most appropriate responding agent to solve the maintenance job. As part of the working paper, an evaluation with synthetic data is presented that points out the competitiveness of the developed approach.

A summary of the results and insights is presented in Chapter 7. Besides, future research opportunities are outlined, and the work is concluded.

2 Success Factors for Process Mining—A Multiple Case Study

This working paper is currently under preparation for publication (Oberdorf et al. 2022a).

Process mining has great potential to generate business value in companies. However, there still seem to be challenges for practical deployment. As described in the literature, problems are often examined from a technical perspective, but ultimately there are challenges in implementation. To investigate these in more detail, a multiple case study was conducted in cooperation with manufacturing and production companies as well as process mining experts and consultants. Challenges and motivations in the context of business value generation through process mining are investigated. A comprehensive picture of process mining is provided by observing and conducting interviews with employees of different hierarchies (from production to chief innovation officer). From this, we derive four key lessons for generating business value through process mining and propose the process mining business value framework that guides the practical use of process mining.

2.1 Introduction

Companies need to navigate a constantly changing business environment. Their resilience depends on the ability to dynamically transform business processes. For this reason, there has been growing awareness towards business process improvement to alleviate the economic pressure of change. Process-aware information systems record the execution of an increasing number of business processes. As a result, the availability of data sources (Feng and Shanthikumar 2018; Reddy 2016) that can be leveraged to analyze business processes has increased. In turn, companies can extract process knowledge from the various data sources and apply business process management (BPM) (Dumas et al. 2013) and business process analytics (Muehlen and Shapiro 2015) techniques. Various information systems shape the IT infrastructure in companies, resulting in a big data environment with great potential for process analytics.

To tap into this potential, process mining is a powerful approach, and with the establishment of commercial process mining tools, its adoption in practice has increased rapidly. However, the academic debate on process mining is often technology-driven and largely neglects the practitioner perspective on process mining (van der Aalst 2019). Notably, process mining research is primarily concerned with the improvement and development of process mining techniques (Augusto et al. 2018; Tax, Sidorova, and van der Aalst 2019). Grisold et al. (2020a) observe the lack of academic investigation of organizational questions concerning process mining adoption in enterprises and focuses on the perception of process managers during the adoption of the technology. Consequently, management issues related to process mining have emerged as a new research direction. In this context, the relationship between process mining use in practice and the determinants of process mining success from the practitioners' perspective remains unclear.

The call for qualitative research for process mining (vom Brocke et al. 2020) was designed to answer such vague perspectives. Finally, it resulted in an editorial (vom Brocke et al. 2021b), which introduces a framework for research of process mining. The framework consists of five levels, namely ecosystem, organizational, group, individual, and technical level, each associated with potential fields for future research. Of these levels, they particularly point out that: "Information systems research, in addition, has a great opportunity to cover the many socio-technical aspects related to the use of process mining at the individual, group, organizational, and ecosystem level" (vom Brocke et al. 2021b). In a similar vein, Syed et al. (2020) shed light on "factors that influence process mining continuity in organizations" through an exploratory inductive case study. However, the results are limited to a single organization for a pension fund in the Netherlands. From a holistic perspective, Martin et al. (2021) investigate the general use of process mining in organizations by conducting a Delphi study with process mining experts. The aforementioned works form a starting point for a detailed understanding of the adoption of process mining in organizations. At a finer level of detail, a focused case study observing the specifics of implementing process mining in cyber-physical environments in the manufacturing industry is still missing. As Erasmus et al. (2020) highlight, the physical characteristics of manufacturing operations differ from highly digitized administrative processes. Most manufacturing-targeted case studies focus on business management activities and exclude shop floor functions. These differences between industries emphasize the necessity of more detailed research on the adoption of process mining in the manufacturing industry.

To address this research gap, we present an inductive case study approach dealing with process mining projects in real manufacturing contexts and shed light on motivations and perceived business value across organizational hierarchies (the research carried out for this case analysis is described in Section 2.3). In addition, we generalize our findings by incorporating input from process mining experts and consultants. We embed the practitioners' perspective by collaborating with production and manufacturing companies and include user experiences to identify challenges for process mining in practice and finally establish a set of success factors. The main contribution of our research is fourfold:

1. We present comprehensive insights about the adoption of process mining across different hierarchies in multiple manufacturing and production companies. Thereby, we highlight challenges, motivations, and the perceived business value of process mining across corporate hierarchies.

2. We propose a novel combination of the L*Lifecycle model (van der Aalst et al. 2012b) and the process mining success model by Mans et al. (2013). Thereby, we extend business value relevant success factors and measures to provide a structured guide for process mining projects in practice.

3. The results of our study establish a basis for future research to expand and develop new theories that generalize findings from the concrete experience of the case study (Sutton and Staw 1995). Beyond that, the case study promotes the academic debate on managerial questions on pro-

cess mining, supports practical use of process mining in the industry sector, and can help develop practical interventions, which we assist through formalized implications—each focused on the particular generation of business value through process mining.

4. We contribute to the theoretical understanding of socio-technical aspects in practical process mining and to the call for qualitative research for process mining (vom Brocke et al. 2020) as an organizational science (Grisold et al. 2020b) and its adoption.

2.2 Related Work

We first review the process mining literature. Subsequently, we highlight the scientific interest in the business value of process mining information systems.

2.2.1 Process Mining in Enterprise Systems

Process mining as a fledgling technology has mainly been concerned with the improvement and development of process mining techniques (Augusto et al. 2018; Tax, Sidorova, and van der Aalst 2019; Maita et al. 2018) and has scarcely examined related practical management questions (Turner et al. 2012; Emamjome, Andrews, and Hofstede 2019; van der Aalst 2019; Syed et al. 2020). Grisold et al. (2020a) address this research gap and shift the focus to the perception of process mining technology and software in organizations. A large share of the existing literature neglects the organizational impacts of the adoption of process mining and explores how to organize process mining projects (Aguirre, Parra, and Sepulveda 2017) or derive domain specific success factors and measures for applying such projects (Mans et al. 2013). Examples of methodologies to organize process mining projects are the L*Lifecycle model (van der Aalst et al. 2012b), the Process Diagnostic Method (PDM) (Bozkaya, Gabriels, and Werf 2009), the PDM specified for the healthcare domain (Rebuge and Ferreira 2012), and PM^2 (Eck et al. 2015). L*Lifecycle model covers the discovery of a single process model, process improvement, and operational support, whereas PDM only addresses a limited number of process mining techniques, thus inappropriate for complex projects (Suriadi et al. 2013). Although

PM^2 addresses some limitations and supports projects that specify the goal of process performance or compliance to rules and regulations, it lacks flexibility, a procedure described in detail, and a practical guideline (Diba 2019). However, the methodologies have not been thoroughly evaluated from a practical or qualitative view in research projects. The resulting question of what determines the success of process mining projects in practice is attempted to be answered by Mans et al. (2013). They developed a model that presents a set of initial factors and measures that influence process mining projects, derived from theoretical concepts in related fields such as process modeling and data mining. Following Nemati and Barko (2003), the impact on success is divided into measures and factors, the former being criteria for evaluating success and the latter being direct or indirect effects that support the success of the process mining project. It is interesting for researchers and practitioners to understand the needs of process mining users in terms of factors and measurements that influence the success of process mining projects. Thus, further qualitative research in process mining can advance process mining projects from an enterprise perspective (vom Brocke et al. 2020) for organizational science (Grisold et al. 2020b).

2.2.2 Process Mining Enabling Business Value

The business value of information systems and technology has been widely discussed within literature (Schryen 2010; Melville, Kraemer, and Gurbaxani 2004; Mooney, Gurbaxani, and Kraemer 1996; Barua, Kriebel, and Mukhopadhyay 1995). Information systems have been associated with positive effects on organizations, including lower operational costs (process automation), improved information dissemination, and process transformation (Daneshvar Kakhki and Gargeya 2019). Information systems or technologies are used mainly in areas that involve making evidence-based decisions to generate business value (Trieu 2017). Data-driven applications help enterprises understand their business and market to amplify timely business decisions (Chen, Chiang, and Storey 2012, p.1166). The creation and realization of value potentials depend on several factors, such as the development, communication, and provision of the information system (Kohli and Devaraj 2004). Furthermore, internal factors such as system integration, customer and supplier readiness might influence an organization's capability to create value potentials (Barua et al. 2004).

With the rise of big data analytics, the value of information systems subsumes the benefits generated by business analytics solutions (Grover et al. 2018; Wang et al. 2019b; Chen, Chiang, and Storey 2012). Two crucial success factors of analytics-based approaches are sophistication and data quality (Côrte-Real, Ruivo, and Oliveira 2020). Both factors are gaining importance given the increasing volume and scope of available data sources. Grover et al. (2018) examine the value creation of big data analytics, and point out the need for appropriate human business analytics capabilities to unlock relevant technologies and data-driven business opportunities. Process Mining, as a novel IT artifact (Eggers and Hein 2020), necessitates human knowledge on data science to initiate data-driven process insights (Abbasi, Sarker, and Chiang 2016). The application of process mining presents various challenges to non-experts (Grisold et al. 2020a), such as identifying and selecting suitable processes for process mining (Thiede, Fuerstenau, and Barquet 2018), or preparing data by cleaning and integrating event data (Andrews et al. 2018; Dumas et al. 2018b). These activities require human skills to analyze and interpret results for real-world process insights. Data-driven insights will then lead to business value in the form of opportunities to observe and analyze real-world process behavior and make evidence-based decisions (Grisold et al. 2020a). The derived business value often depends on the process mining scenario. The value might manifest itself in shorter production times through transparency on bottlenecks (Lee et al. 2014), or in improved customer satisfaction through an increase in service quality (Edgington, Raghu, and Vinze 2010). However, the value derived from data-driven process mining is only significant when the results are interpreted and translated into managerial actions (Dumas et al. 2018b). A major concern is that specifying the business value remains cumbersome (Eggers and Hein 2020), leading to a lack of foundation in communicating the benefits.

In this context, Syed et al. (2020) reconfirm the IT productivity paradox (Brynjolfsson 1993; Brynjolfsson and Hitt 1998) in the context of process mining adoption. The paradox asserts the common misconception in organizations that immediate benefits would accrue after implementing new technologies. Process mining is expected to provide various benefits, such as the key benefit of being able to trace processes realistically as opposed to opaque documentation of process participants (Bolt and van der Aalst 2015; Dumas et al. 2018b) or perceived increased efficiency and effectiveness (Grisold et al. 2021). However, Grisold et al. (2020a) mention that "... it remains unexplored how the effects of process mining can be translated into increased revenue and reduced costs".

Recent research sheds light on the generation of business value and guides its realization (vom Brocke et al. 2020, 2021b; Syed et al. 2020; Martin et al. 2021) through frameworks and success factors for process mining. However, a high-level perspective of process mining and differing characteristics, such as for manufacturing operations (Erasmus et al. 2020) are not reflected. We overcome this by going beyond business management activities and shed light on the shop floor process mining and associated success factors for business value generation.

2.3 Research Method

With respect to the lack of existing studies and frameworks that guide process mining projects in practice—particularly in manufacturing or production industry—and implicitly provide measures and factors to account for success, we set out to develop a framework that overcomes these limitations. To do so, we introduce our case study design and partners, present the research framework, and data collection and data analysis.

2.3.1 Case Study Design

For the case study, we follow a three-step procedure aligned to Bilgeri et al. (2019) consisting of a pilot study (1), practitioner interviews (2), and practitioner re-interviews with additional expert interviews (3).[2]

The **pilot study** phase had a duration of approximately half a year and had the particular focus of identifying relevant companies in the manufacturing and production industry. To separate technical process mining aspects, we focused on cases where open-source or commercial process mining software

[2] Due to the pandemic situation, we could not establish the planned cross-company workshop, but interviewed the employees in digital meetings with a duration of 52 minutes in average.

has been used to analyze and optimize processes. Further stable process models should have already been derived to omit the effect of the process models' quality on our research. Thus, we had an exclusive focus on the case study participants' (perceived) usefulness and challenges for process mining as well as non-technical aspects, which we identified as key issues of the *a priori* (without extension—Section 2.3.2) model established based on related research.

In the research procedure's **second phase**, we conducted practitioner interviews with employees of our partners to discuss success factors for process mining in the context of the extended *a priori* model. For open-ended interviews, we follow the interview checklist provided in the Appendix B. Based on the interviews, we refined the extended *a priori* model and provided it to the partners to establish it over a period of about nine months.[3]

In the research procedure's **third phase**, we re-interviewed employees of our partners to become familiar with their perception of the extended *a priori* model and in addition conducted interviews with process mining experts and consultants, to generalize our findings, finally resulting in the PM-BV framework.

2.3.2 A Priori Model

Guided through the three-step research process of Bilgeri et al. (2019), we first conducted a pilot study, subsequently conducted interviews with the involved employees, and finally generalized our findings through re-interviewing employees as well as expert and consultant interviews. Each of these phases adds to the development of a generalized framework for guiding business value generation in process mining projects. In the pilot study, we explore potentially interesting companies from the manufacturing and production industry. In addition to partner exploration, we were able to observe and discuss process mining frameworks and success models, to identify the *a priori* model that consists of useful constructs for our desired framework (Eisenhardt 1989).

There exist multiple process mining project methodologies that provide guidelines for the execution of process mining in practice. As there is no standard process mining methodology (Emamjome, Andrews, and Hofstede 2019), we choose the L*Lifecycle model (van der Aalst et al. 2012b) for two key reasons.

[3]Due to the pandemic situation, only P1 and P3 actually established the model on a daily basis.

First from a theoretical and research perspective, the L*Lifecycle model incorporates an extensive plan & justify stage that is research-driven. Second, the partners' process mining experts perceive the L*Lifecycle model as fundamental base and practical established. For this reason, we conduct our research on top of the L*Lifecycle model, nevertheless, the research could be extended to other frameworks in future work. The L*Lifecycle model establishes the key steps of the process mining projects and provides a structured implementation guide with five iterative stages: plan & justify, extract, control-flow model, integrated process model, and operational support. Yet, for practical use, the L*Lifecycle model lacks determinants of process mining success from the practitioners' perspective to derive practical interventions.

To overcome this, Figure 2.1 maps the process mining success model by Mans et al. (2013) to the L*Lifecycle model, resulting in the process mining business value (PM-BV) framework—yet without extension. The integrated success model comprises three success measures and six success factors. Success measures—process impact, process efficiency, and model quality—serve as metrics to assess the potential and effect of process mining projects. The success factors—management support, process miner expertise, structured process mining approach, data & event log quality, resource availability, and project management—guide the aspects that should be incorporated for a successful process mining project. From another perspective, success factors refer to challenges that (can) occur during a project.

This developed perspective allows us to simultaneously assess the "how?" and the "why?" of process mining projects. Initially, we do not consider the extension (marked with *), which we derive from the case study (section 2.4), but only the combination of L*Lifecycle and success model. We can map the success measures and factors to the stages and focus on operational process mining projects, where a stable process model is available. Thus, the process model relevant measure and factors are nearly constant and we notice a detachment of the success measure *model quality* as well as the success factors *process miner expertise*, *structured approach*, and *data & event log quality* from the operational process. As we explicitly relate our observed challenges to a process mining project, we depict that *project management* and *resource availability* are accompanying success factors as in any project in the enterprise. Therefore, we can particularly focus on the plan & justify as well as operational support & process transformation stages. There are only the success measures *process impact* and *project efficiency* as well as the success factor

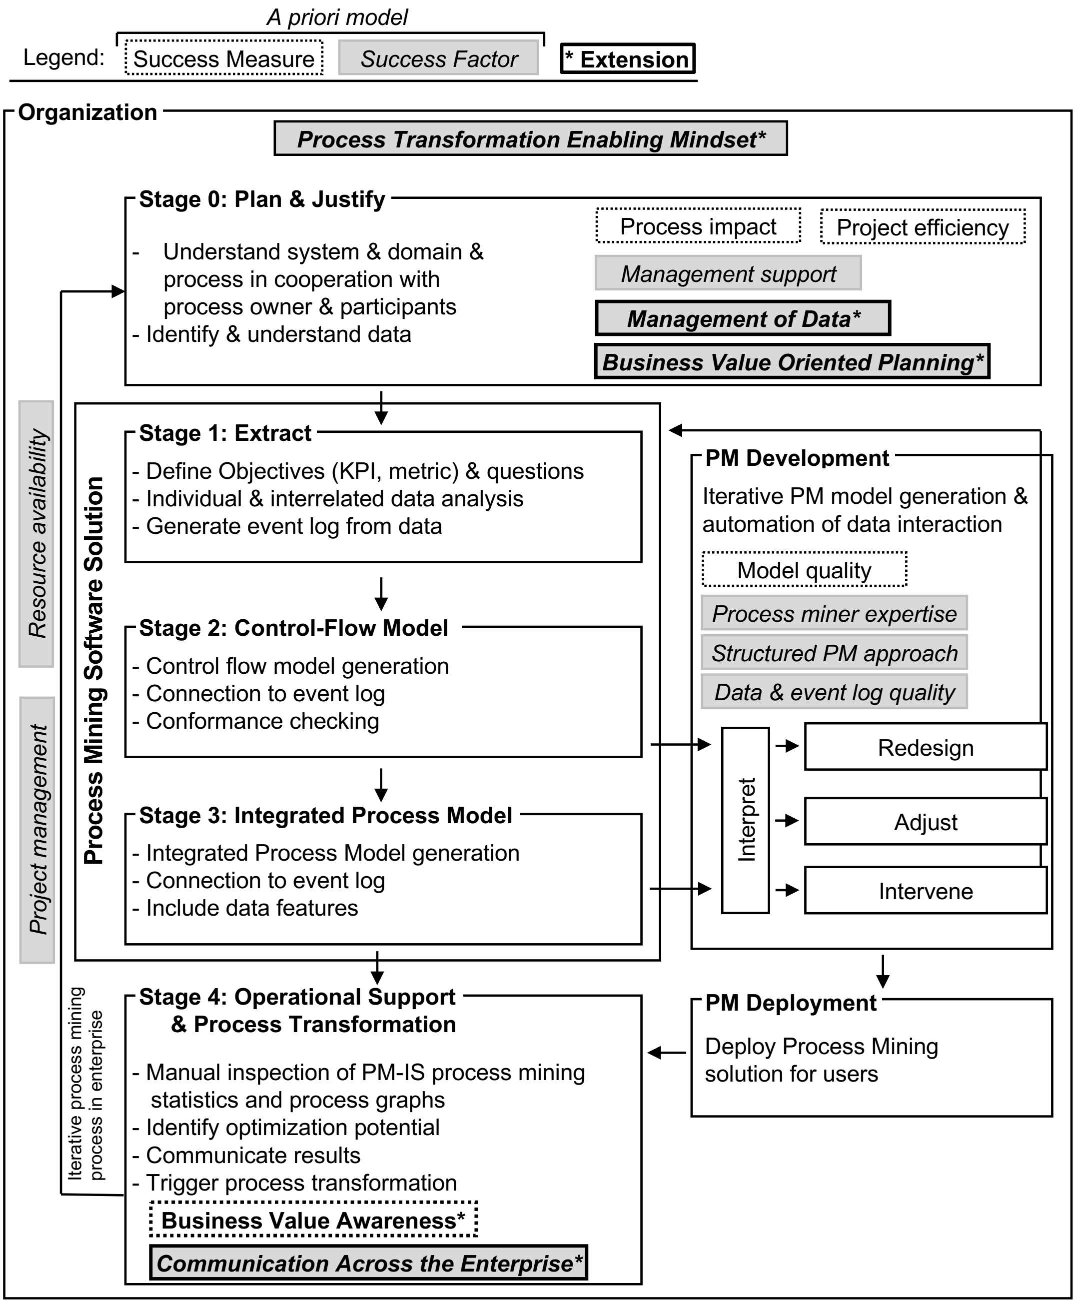

Figure 2.1: Process mining business value (PM-BV) framework based on the L*Lifecycle model (adopted from van der Aalst et al. 2012b) and the process mining success model (Mans et al. 2013) as well as the success model extension.

management support involved—mapped to the plan and justify stage. However, challenges arise not only during the planning and execution of projects. In particular, with regard to the practical implementation of process mining projects, companies face other challenges such as fragmented process knowledge or technical reasons (Dumas et al. 2018b; Andrews et al. 2018). To specify such challenges and derive an extension of the success measures and factors, we set out the case study. To this end, we interviewed and re-interviewed employees of the partner companies', as well as process mining experts and consultants, to finally derive the PM-BV framework and derive findings.

2.3.3 Case Study Partners

Our case study partners are four companies in the production and manufacturing industry, as well as process mining experts and consultants with associated expertise. Table 2.1 provides an overview of the case study partners categorized by the case study role (partner, expert, consultant) and provides information about the industry positions, as well as participating employees and the type of process mining solution.

All case study partners operate production lines to transform raw materials and parts into valuable goods. Each company runs some sort of manufacturing execution system (MES) which offers the opportunity to log production events. To analyze the logs, all partners use a process mining software solution. While P2 and P4 use commercial process mining software, P1 and P3 use open-source solutions. With respect to specialized solutions, the scope of analysis also differs. While P3 focuses on classic use-cases, P1 focuses on an individualized process mining solution, the process mining information system (PM-IS). This PM-IS offers (most) of the functions of commercial solutions and in addition enables the integration of heterogeneous data sources for process mining. For example, P1 has a disruption management system (DMS) that automates the notification of responding agents that help to solve a disruption during the production process, as a worker identifies and reports a disruption. The DMS provides additional production-related information that is then combined with the production log to identify disruption-related process paths. Clearly, this functionally (currently) exceeds the capabilities of commercial solutions, but

Table 2.1: Overview of the case study partners classified by their role as company partner (P), expert (E) and consultant (C).

Partner	Industry Position	Interviewees	Solution
P1	Produces high-precision mechatronic drive systems. *~3,000 employees*	15 employees in positions from production to CIO.	Open-source
P2	Major producer of high-class sun shading equipment. *~4,000 employees*	Five employees in business analyst and project manager positions.	Commercial
P3	Produces high-class windows. *~5,000 employees*	Two employees in process mining expert and process manager positions.	Open-source
P4	Major producer of construction materials. *~35,000 employees*	Two employees in management and process manager positions.	Commercial
E	Cross-industry process mining experts.	Three experts with experience from a wide range of process mining industry projects.	-
C	Consultants with special focus on process mining projects.	Current and former consultants for German based analytics and process mining consulting and vendor companies.	-

facilitates the analysis of the added value through customized solutions, which we discuss as part of the implications Section 2.5.

Based on the *a priori* model (without extension), one of the partners provided us with the opportunity for observations[4], where we could observe the

[4]Due to the pandemic situation, the possibility of live-observations was temporarily restricted and only limited to P1. Still, we could observe about 11 months live at P1, participate digitally during lock-downs, and discuss virtually with all partners over the course of the case study.

application of the *a priori* model and then identify challenges, motivations, and success factors under real world conditions. This provided the opportunity to conduct an in-depth case study that allows us to observe the application, impact, and acceptance of complex interrelated processes under real-world conditions (Eisenhardt 1989; Eisenhardt and Graebner 2007). Thereby, we are able to depict how users of different hierarchy levels were involved, even if "boundaries between the phenomenon and context are not clearly evident" (Glowalla, Rosenkranz, and Sunyaev 2014). Further, we describe the employees' corporate hierarchies as part of the data collection, which reveals the holistic approach for analyzing the adoption of process mining across organizations.

2.3.4 Data Collection

Our research collaboration includes more than a year of observations, passively participating in meetings and workshops, and training materials and documents as data sources across all partners (Table 2.2).

Additionally, we involve semi-structured and open-ended interviews as part of the case study. This format allows for a detailed understanding of the participants' motives, expectations, and challenges regarding business value generation through process mining.

At the partners, we interviewed employees from different hierarchy levels (Figure 2.2) to shed light on the challenges, motivations and business value of process mining across the enterprise. Depending on the position, we wanted to highlight different aspects. Through interviews with the board member (Chief Innovation Officer - CIO) and close to C-level management, we were able to gain insights into the strategic relevance of digitized processes in general and process mining in particular. Team leaders, project and process managers have experience with the implementation and roll-out of (digitalization) projects in the companies and are involved in process mining projects. The interviews with data scientists and BI analysts facilitated a better understanding of process mining and, in particular, for P1 the use of heterogeneous data for process analyses. We were also able to gain essential insights into the implementation and challenges of data science projects at the partners. In the production environment, we wanted to analyze the impact of process mining in planning (product designer) and process transformation results (production leader and

Table 2.2: Data collection overview with activities and data insights across all partners, experts, and consultants.

Data collection	**Activity**	**Data**
Interviews	Interviews with CIO, digitalization manager, project leader and consultants, business developers, process managers, process mining consultants and experts, as well as data scientists production employee and product designers.	Objectives, challenges, motivation, and business value of process mining with heterogeneous data sources.
Observation	Over 200 hrs. of observation.	Insights, results and acceptance of process mining with heterogeneous data. Data science project procedure and realization.
Workshops	Workshops and seminars.	Strategies, motivation and road-maps for digitalization and process mining. Communication, agile development and realization of process mining.
Documentation	Documentation and training material of data analysis, technical information system guidelines, process design, and product development processes.	Several hours of digital and physical training material—special focus were similarities between established data science projects and the process mining case study.

teams). This mix of employees across the partners allowed us to capture and analyze challenges, motivations, and success factors for applied process mining across the enterprise. For generalization purpose, we conducted the process mining consultant and expert interviews, to establish a holistic perspective for the PM-BV framework.

2.3.5 Data Analysis

For the data analysis, we followed an interpretive research approach (Clark et al. 2010; Maanen 1979) and leverage an iterative thematic analysis to identify patterns in large and complex data (Braun and Clarke 2006). To this end, we transcribed all interviews and analyzed them according to the process described in Braun and Clarke (2006). Initially, all data were analyzed, such as transcribed interviews, observation notices, and workshop protocols. With the objective of identifying phrases relating to our research questions, three researchers separately analyzed all transcripts by reading the full-text versions of the transcripts and descriptively coded them considering the guidelines of Saldaña (2015) to ensure the reliability of the results. For this purpose, we utilized the *MAXQDA* software following the guidelines put forward Kuckartz and Rädiker (2019).

In the second step, researchers derived conceptual themes from the coded data (Table 2.4), which we relate to as challenges for the use of process mining. Subsequently, the identified challenges were discussed and merged between the researchers. As part of an iterative process and after a preliminary catego-

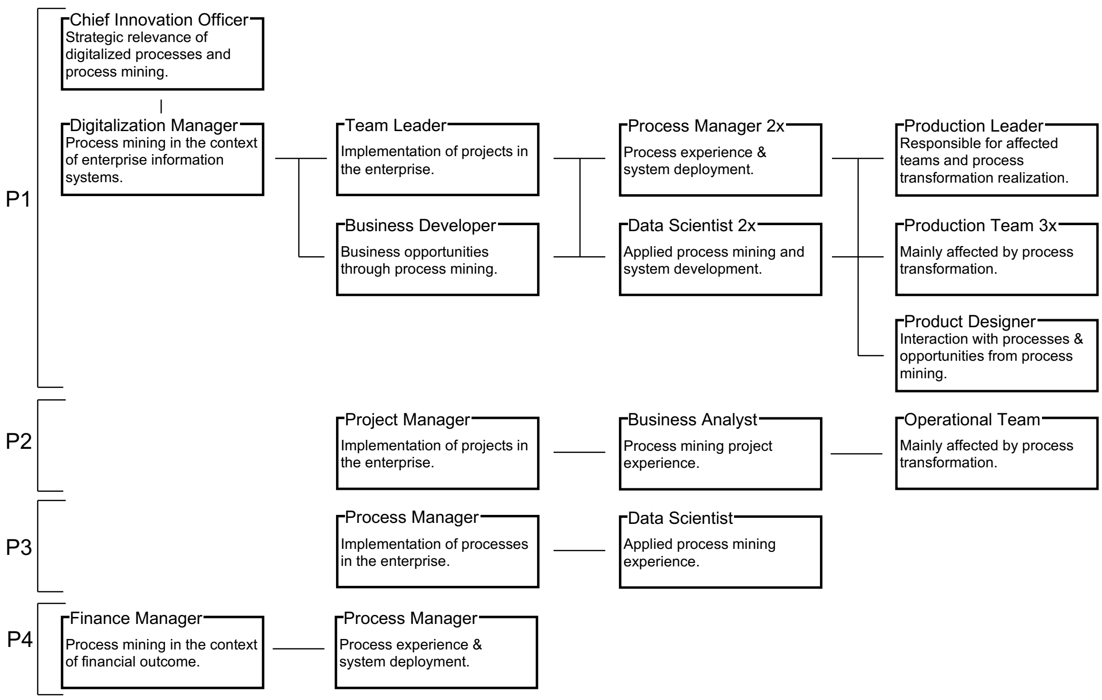

Figure 2.2: Organigram of interview participants.

rization, the challenge themes converged while facilitating a context-specific interpretation of the findings. The final step consists of finding the aggregation to provide abstracted findings and lessons. To this end, we put the coded and evaluated interviews in context with the observations and literature to derive findings. During this iterative process, we followed the approach of the inductive research method and analyzed the data at first. For the second stage and expert interviews, we followed similar approaches and compared the themes with those previously identified. Subsequently, we mapped our interview and observation results with the *a priori* model, enabling us to develop the PM-BV framework and derive general implications.

2.4 Findings

With respect to the design of the multiple case study, the findings are threefold. On the basis of extensive observations and interviews, we are able to elaborate on the motivation behind the adoption decision (Section 2.4.1). Therefore, we assess the challenges of process mining projects in practice (Section 2.4.2) and derive additional success factors and measures that complement the PM-BV framework (Section 2.4.3).

2.4.1 Motivation for Adoption

To shed light on the adoption of process mining across enterprises, we combined both the key motivations and perceived business value (Figure 2.3). We identified two key motivations—*financial* benefits and increased *transparency*. Furthermore, some interview partners (in particular data scientists and BI analysts) mentioned the objective of generating *informational insights*. Especially for interviews in production-related areas, we were able to identify a desire for **automation**. However, this motivation was universally applicable across all information systems and for this reason not specific to process mining solutions.

Financial

A general motivation to reduce costs by improving processes has been observed across all interview participants. However, it was particularly highlighted by the management that the investment in process mining must pay off to establish process mining in the long run.

> *"It is most important to trigger a process improvement, based on the [process mining] insights and then generate savings. These might be small and add over time, but ultimately the investment in process mining must be recouped."*
>
> (E - Manager)

Considering a desired return on investments for process mining projects, it is obvious that costs are a key inhibitor for adopting commercial solutions. Notably, this is also what triggered P2's decision to change to another commercial process mining software provider.

Monetary benefits from process mining projects can accrue through direct and indirect measures. Attributing savings to some measure or analysis can be a tricky task, because there usually are sizable time lags between the analysis, the implementation, and the realization of benefits. To this end, it was mentioned that it was important that effort and benefit must be in balance and that effectiveness has to be demonstrated in a timely manner.

> *"Therefore 'the faster, the better'. The faster I can demonstrate effectiveness, the better it is."*
>
> (P1 - Digitalization Manager)

However, there should be no rush throughout a process mining project, but the project must be set up on a solid planning. As mentioned by a leading consultant, it is fundamental to be aware of "the long-term effort and commitment" that a process mining project requires to unfold its full potential. A key step in doing so are discussions about the identified insights as such "discussions drive(s) process improvement."

Transparency

In all interviews, a frequent reason for process mining was the creation of transparency. However, there are different dimensions of this motivation depending on the hierarchy. Management seeks to obtain a transparent process overview using aggregated key performance indicators (KPIs), which are fundamental for assessing the results of a project. For process managers, data scientists, and business developers, one of the most important aspects of transparency was identifying optimization potentials. This highlights a key benefit of process mining.

> *"... to create visibility through data where things are perhaps not yet going well, where there is still a need for optimization. Being able to make objective decisions, not always relying subjectively on the perceptions of individuals [...], but making data-based decisions to design an optimized production process."*
>
> (P1 - Business Developer)

Value of Informational Insights

We highlight the perceived value of insights for the individual interviewees (Figure 2.3). In the management areas, value tends to be perceived as context information for KPIs. Regarding process mining, process managers, business analysts, and data scientists perceive the value of insights most—and are in some cases motivated by these. This value relates to the informational insights which immediately result from process mining and thus are most aware, even in the short term. Yet, there are more dimensions of process mining, which are relevant for a holistic understanding of process mining's perceived adoption across the enterprise.

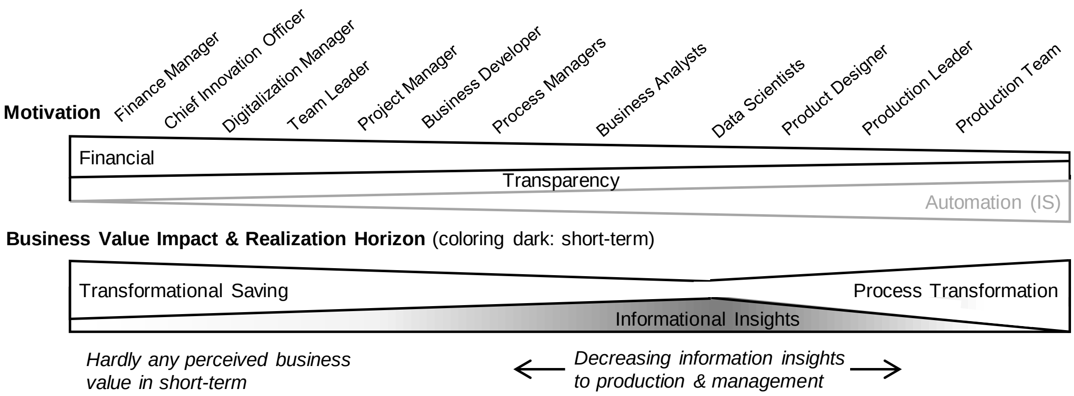

Figure 2.3: Employees motivations combined to desired business value with realization horizon.

Business Value

We adopt the framework proposed by Mooney, Gurbaxani, and Kraemer (1996) to categorize business value potentials. This framework distinguishes automational, informational, and transformational business value. Table 2.3 provides an overview of the motivations and benefits of process mining in practice.

Table 2.3: Business value opportunities facilitated by process mining.

Business Value	Motivation	Horizon	Benefit
Automational	-	-	-
Informational	-	short	Business value quantification
	Transparency	short	Process discovery
		short	Informational insights
Transformational	Objectivity	medium	Data-driven decision support
	Financial	long	Process enhancement

Automational business value was not mentioned as direct result of process mining. While process mining automates the process identification itself, the result of the identification is knowledge of about a process, which the interviewees referred to as informational value. Informational business value is an obvious benefit of process mining. Organizations can benefit from process discovery which creates informational insights. Transformational value emerges in the medium and long term through the adoption of data-driven decision support and the implementation of process improvements.

Concerning informational insights, these were generated by the users—in our case, data scientists and process managers. They were discussed with direct colleagues. However, across different hierarchies, the perception of insights decreased significantly. These insights were less critical for management and production—for both the results of the transformation matter. However, the processes themselves must first be transformed in order to tap into saving potentials. The gap between process mining theory and practice becomes apparent here. Process mining primarily generates process insights. In practice, however, it is not the insights, but the ultimate results that are relevant. Therefore, it is clear that process mining will likely struggle to generate *immediate added value* as desired by management. Although this was known to some extent, it highlights the challenge of *result communication*. Being aware of the motivations and context of perceived business value across the enterprise allows us to better understand the challenges that must be overcome for process mining in practice.

2.4.2 Challenges

There are several challenges for establishing process mining in companies. The findings are summarized in Table 2.4. It classifies the key insights into four main categories: technical, business, managerial, and organizational challenges. The multifaceted challenges result in associated lessons learned. In essence, the challenge categories follow the dimensions of the ecosystem established by Vidgen, Shaw, and Grant (2017) to analyze how process analytics create business value.

2.4.3 Lessons Learned

Regarding the identified challenges, we were also able to depict additional success factors for process mining in practice. The resulting lessons arising from the case study lead to the PM-BV framework (Figure 2.1 with extensions). The lessons are formalized implications for CIOs, IT managers, and researchers facing related problems and are subsequently discussed in the context of the literature.

Lesson 1: Leverage Heterogeneous Data

Heterogeneous data can enhance the quantification of business value in the project planning phase, as all partners pointed out. In addition to existing approaches to extend the analysis perspectives (van der Aalst 2013), it offers possibilities to render process mining more attractive for companies. For example, all partners (P1-P4) leveraged additional data to overcome the challenge of *business value quantification*. By relating the process under consideration and (potential) savings to a monetary scale, the savings potential could be quantified for the process mining project. Furthermore, P1 also uses hetero-

[5] Please note that by means of conceptual themes, the *Observations & Participant Statements* relate to first-order concepts and the *Challenges* second-order themes.

Table 2.4: Perceived challenges of the process mining project from participants' statements and related lessons learned aligned to the categories technical (T), business (B), organizational (O) and managerial (M).[5]

Categories (T: Technical, B: Business, O: Organization, M: Managerial)	Challenges	Observations & Participant Statements	Lessons
T, M	Data availability and management	• Process mining with heterogeneous data has huge potential but necessitates extensive coordination and authorization of data access • Knowledge of data structure is needed to reveal data potential	**Leverage Heterogeneous Data**
T	Customized production	• Process diversity due to highly customized production	
B	Business value quantification	• Quantification of process improvement in the context of business value • Relating a process improvement to monetary scale	**Foster Immediate Business Value through Process Analytics**
B	Immediate added value	• Delivery of transferable and transformational results is desired • Proof of effectiveness must be provided quickly in the light of volatile economic situation	
B	Delayed improvement	• Delayed improvements impede direct trace back of improvement origins such as process mining	
O	Process Improvement awareness	• Awareness for enterprise process network and interrelated processes • Resistance to change paired with sceptics about process mining results	**Establish Process Transformation Enabling Mindset**
O	Fear of transparence	• Process mining must not be interpreted as monitoring or fault detection mechanism	
M	Result communication	• Just the words big data and process mining result in misleading and-sometimes too-high expectations • Even generated business value does not reach top level management	**Facilitate Stakeholder Management**

geneous data from a more operative perspective to identify disruption-related processes. There, additional data enable a more targeted analysis of existing production event logs, e.g., by serving as an explanatory variable for context events, as we could observe for the combination of process mining with the DMS. This facilitates the identification of disruption-related processes. This is despite the presence of branched process graphs, e.g., due to the challenge of *customized production*. As products are customized, the identification of the main process path, without additional data, is difficult. The data scientists implemented a relational logic[6] between the disruptions and the process log. By only selecting a certain disruption, a related process graph can be generated and even compared with production graphs, excluding the disruption-related events automatically. Comparing both process graphs facilitated a more targeted identification of disruption-related process anomalies.

Of course, the quality of heterogeneous data is becoming increasingly important. Data not only provides added value on its own, but also combined. In addition to data and event log quality, which Mans et al. (2013) have already defined as a success factor, the management aspect of data procurement—**Data Management**—is becoming an increasingly important success factor (Dong and Srivastava 2013; Doan, Halevy, and Ives 2012).

> *"In particular if you don't know what data you need or what is in [the data], it is difficult to find arguments to get the data. This is not only time consuming at the beginning, but also critical for the successful completion of projects."*
>
> (E - Data Scientist)

For this purpose, it is crucial to identify relevant data sources early on and take care of authorization and data access in a timely manner. In the case of collaboration between different departments and areas, obstacles should be expected. Nevertheless, the effort to obtain and combine data for process mining projects offers tangible results.

Lesson 2: Foster Immediate Business Value through Process Analytics

Process mining generates business value in the short term, but this is often limited to informational insights. The insights depict options, but do not yet offer the managements' desired transformational business value by means of financial savings. However, to establish process mining in companies, it is precisely this presentation of *immediate value* that is crucial to obtain the necessary support in projects. Eggers and Hein (2020) express this with the question "how organizations develop the antecedents necessary to implement process

[6] The logic mainly consists of limiting the number of process events for a given time window and (if available and intended) additional filters for production lines.

mining and how lasting business value can be created". One crucial antecedent is **Business Value Oriented Planning**, which we have identified as a success factor for process mining projects. Already during the planning of projects, how immediate business value can be generated must be considered. If immediate business value can be achieved, the *delayed improvement* challenge's relevance becomes less critical. It is usually challenging to demonstrate financial savings after a short period of time. However, the presentation of qualitative feedback can be a target-oriented approach:

> *"At the beginning [of the process transformation], it is difficult to demonstrate savings. We have learned that it is often good to get feedback from people about measures–perhaps we may even obtain an estimate from those affected."*
>
> (P1 - Data Scientist)

Lesson 3: Establish a Process Transformation Enabling Mindset

Process mining has great potential to optimize processes and in turn generate business value. However, the measures developed must be implemented and accepted by those involved in the process. To this end, a **Process Transformation Enabling Mindset** in the organization is a decisive success factor. Problems should not be viewed negatively, but should be understood and leveraged as opportunities for improvement. Digital methods must support this and should not trigger a feeling of monitoring. Process mining creates enormous transparency right from the start, especially in combination with additional data sources. From an opportunity point of view—it offers enormous optimization potential. Ideally, these opportunities are recognized, accepted, and leveraged, so that the processes can be examined and improved from within the respective areas themselves, without external pressure from superiors. How such factors can be incentivized remains an exciting question for future research.

Lesson 4: Facilitate Stakeholder Management

Stakeholder management is a crucial component of many projects (Karlsen 2002), and process mining projects are no exception. Of particular importance is the success factor **Communication Across the Enterprise**. It is necessary to manage result expectations to avoid compromising the project implementation and its success with false expectations. It applies to the expectations of production and either management.

> *“This expectation when it comes to just these words big data, data analytics, process mining—most people don’t really even know what it means. I think it raises such high expectations because of that. Or maybe hopes. [...] So you have to be careful.”*
>
> (P1 - Project Consultant)

Just as crucial is *results’ communication*. The perceived benefits of developers and users may differ from the final business value. Although metrics such as accuracy or root mean squared error are often used in data analysis, they are often inappropriate for management. To overcome such misinterpretations, it is crucial to include **Business Value Awareness** as a success measure. For all parties involved, the goal must be clear: to generate business value. Furthermore, everyone must be aware of how this is achieved, typically through cost and time savings realized through process transformation. Finally, the business value must be communicated to all stakeholders and establishes the foundation for a follow-up project.

2.5 Implications and Opportunities for Research

Throughout the cooperation, we pinpointed insights into the implementation phase of process mining projects in practice. Since the implementation phase often exceeds the observation horizon of process mining case studies, we can formalize and discuss implications based on the insights set out by the conduct of the case study. Table 2.5 provides an overview of the implications discussed.

Table 2.5: Overview of the implications discussed.

Business and Managerial	Technical and Organizational
• Increase practical aspects of process mining frameworks • Consider informational insights in (gated) project frameworks • Extension of success factors to more general frameworks • Management reporting is a decisive factor for perceived value • Facilitate top management support	• Additional data sources should only be integrated with specialized information systems • Facilitate end-to-end insights through process mining • Combine information system and process mining adoption research • Integrate digital, physical and social perspectives • Increase framework's focus on deployment phase

2.5.1 Business and Managerial Implications

In general, there are multiple process mining methodologies (van der Aalst et al. 2012b; Eck et al. 2015; Bozkaya, Gabriels, and Werf 2009), which offer structured frameworks for the use of process mining technology. However, the scope is limited mainly to technical aspects of process mining, and managerial questions are left out. Emamjome, Andrews, and Hofstede (2019) define criteria for the maturity of process mining case studies in general and synthesize the phases of existing process mining project methodologies to provide a basis for process mining maturity models. They reveal that most studies investigate process mining as a technology, and thus process mining methodologies are mainly technology-driven. Practical guidelines for process mining deployment within an organization are neglected, leading to a gap between process mining theory and process mining in practice. Within the case study, we observe that detailed guidance is needed to know, recognize, and address challenges during the implementation process from the practitioner's perspective. Consequently, such frameworks should also be practically motivated if the practical transfer of process mining is desired.

A distinct factor of our research was the combination of the L*Lifecycle model and the process mining success model. The combination of success measures and factors showed in which phases factors are already present and where our particular research focus should be. Furthermore, in discussions about the applicability of process mining frameworks, we depicted the question for a gated framework.[7] From a practical perspective, there is a desire, especially of management, for frameworks that define the process of an analytics project in general or process mining in particular. At the same time, however, it would also be desirable to have clear *gates* that must be achieved to continue a project. For example, if valuable insights cannot be committed to a certain extent, a continuation of the process mining project should be critically discussed with decision makers. According to our findings, in particular, in the short term, these are primarily informational insights. Assessing these must be considered in the definition of project gates' acceptance criteria. Their accurate description may be challenging, but at the same time, it offers opportunities for future research.

These findings might not only be limited to process mining projects and frameworks. More general data-science frameworks (e.g., Flath and Stein 2018) can also benefit from the extension of success factors and an extensive deployment phase. Often, the technical implementation of machine learning or data science projects is described extensively and in detail in such frameworks. However, the practical transferability is limited. There is a lack of suggestions on how the trained models and insights gained can be adopted in practice in a target-oriented manner. Future research should investigate the use of (data science) frameworks in practice in more detail, following Neff et al. (2017), and add outstanding components to the frameworks. It would also be interesting to consider the connections between existing frameworks, incorporating different domains such as machine learning, process mining, or operations management. Furthermore, we look forward to more research on the combination of success factors with such frameworks and how they affect the application, particularly in practice.

[7] The *gates* imply some kind of milestones. However, the wording was distinct to highlight that a gate must be open to passing.

From a business perspective, practical frameworks enriched with success factors might even guide how to address obstacles in realizing perceived benefits in business value. The case study covers observations not just from a management perspective, but across the enterprise. Ultimately, this unlocks insights on the dissemination of perceived business value over time and hierarchical organization. These findings provide the basis for establishing the generation of business value in general as a success factor for process mining projects. Similarly, the general need for perceived value in the face of technology adoption is found in previous studies on the implementation of information system technology (Joshi 1991; Kim and Kankanhalli 2009). However, studies differ in explaining perceived value as a factor for technology acceptance as they measure perceived usefulness, perceived ease of use for users associated with the new information system. However, the observations in the case study draw different conclusions about value generation. The perceived value is not related to the user itself, but to evidence of business value in quantifying effectiveness and process improvement for top-management reporting.

This type of reporting is crucial as the communication of project outcomes is strongly related to top management support for the project under consideration. We observed top-management support as a significant success factor within process mining projects, which is influenced by the business perspectives of the project, such as business value generation. In this context, Markus and Robey (1983) investigate the purpose of understanding the resistance of organizations to change in the implementation of information systems. An essential factor they observe is that organizations cannot successfully adapt to changing environments without top management support because users' resistance stems from a loss of power due to the new technology.

2.5.2 Technical and Organizational Implications

Process mining technology refers to an emerging approach that amalgamates business process management and data analytics. Its focus is data-driven; thus from a technical and organizational perspective, process mining implementation projects face similar obstacles and can be considered in the light of related domains such as data mining or analytics. Developing technical capabilities within enterprises remains a well-known challenge in the analytics domain. Gust et al. (2017) grasp the challenges of new analytics implementation in traditional enterprises and identify four key lessons to help practitioners execute future analytics seed projects successfully. They observed difficulties with respect to the business, organizational, and technical dimensions. In the technical dimension, especially in data management, they noticed in a similar vein to our case study silo thinking about inter-departmental data management and sharing due to hierarchical organizational structures. The findings of Gust et al. (2017) differ from the findings of the process mining projects in the point of extensive requirements on pre-processing data for process mining, which is needed after extracting data from the various sources of information systems. The data structure extracted from information systems must be transformed to appropriate formats (e.g., XES[8]) and requires domain knowledge on the processes considered. This means process mining project managers have to deal with heterogeneous data sources and their use for process mining, creating a struggle already. They also require process knowledge for the area under consideration to prepare the data accordingly for the application of process mining algorithms. If process managers are not involved in the technical details, they must organize the availability of process and data experts. On the one hand, the amount of data increased the potential to generate value, both through insights and then transformations. However, the complexity for the process manager increases accordingly. Here we could observe that the development and application of a specialized and simplified process mining information system offers substantial benefits (Tiwari, Turner, and Majeed 2008). The effort to acquire and process the data is integrated into the information system. It can be leveraged subsequently, significantly simplifying the use of process mining in practice.

Our case study is situated in the manufacturing and production context. Some of the project members were involved in process analysis according to the lean production concept, among others, with the help of value stream mapping. We could observe that this process knowledge was an advantage for the execution of the projects. This indicates that general process knowledge can promote process mining in the short term, as it lowers the barriers to entry. However, process managers are often highly specialized in their processes, and knowledge of the interrelated processes is lacking. The phenomenon of

[8] http://www.xes-standard.org/

silo thinking is not uncommon in traditional enterprises with hierarchical organizational structures where collaboration across the organization is weak. In this regard, digitalization and emerging technologies, such as process mining, could enable end-to-end insights into processes and a holistic view on the heterogeneous IT landscape of enterprises (Armengaud et al. 2020).

Companies that adopt new technologies often face resistance to change from an organizational perspective (Markus 1983; Ram and Sheth 1989; Joshi 1991). For example, despite the potentials of data analytics in general, organizations remain skeptical about its adoption and application to unlock its potential benefits (Dubey et al. 2020). As process mining tools are part of the information system landscape, implementation strategies and tactics may resemble each other in their adoption process. Thus, the adoption strategies and tactics of the novel research field of process mining could benefit from existing phenomena in the information system domain. A wide variety of studies have examined technology acceptance models for the adoption of data analytics services (Gursoy et al. 2019; Klumpp et al. 2019; Ostrom, Fotheringham, and Bitner 2019; Mohd Salleh, Rohde, and Green 2017).

In summary, the studies underpin the findings of the presented case study in the following points: operational performance can be achieved through the adoption of technology, but people should be the focus and supported. For good change management to be possible, it is crucial to generate insights through operational support using predictive or prescriptive insights that benefit from time or cost savings. The observed downside of technology adoption is often the perceived loss of control or privacy due to greater transparency. Müller et al. (2016) analyzed the relationship between an enterprise's performance and data analytics. They show a positive correlation between data analytics and business performance improvement of 3-7%. These results might serve as hypotheses for future research investigating the relationship between process mining and increased operational performance (Grisold et al. 2020a). Especially in the production environment, these findings would be of great importance.

In the manufacturing industry, especially in the production sector, efficiency improvements are a key objective. Often, there are some strategic directions according to which the whole way of working within the organization is aligned. In P1, the predominant strategic concept was lean production, which puts the customer at the center. The customer-centric perspective leads to a customized production with a high process diversity. Process managers were using process mining as a new technology to enable digitized customer-centric services through data analytics. Therefore, practitioners might benefit from drawing implications from customer experience research. Partners face similar challenges as Bolton et al. (2018), who examine the specific challenges of enterprises in integrating digital, physical, and social perspectives to create a holistic customer experience. With the help of designed realm cubes, they classify projects or project topics according to digitalization density, physical complexity, or social presence to address practical and social implications of customer experience development for managers. Applying such a framework to the management strategy of process mining managers in practice could generate perceived value for the customer.

It has emerged from the interdisciplinary perspectives across the enterprise that the most critical factor for successful deployment of process mining concerns how to implement new technology and integrate the human factors that should interact with the latest technology (Ahmad 2015). For process mining, we need additional research to have a more precise understanding of what human factors are and how practitioners are using process mining in their daily working routines. As process mining covers the interplay between business process management and data analytics, its roots in BPM research are traditionally based on the combination of knowledge from information technology and management science to support operational business processes in organizations (Dumas et al. 2018b). The approach of technology support in management decision-making is reflected in research on process mining, where the approaches are mainly technology-driven (Garcia et al. 2019). The effects between the adoption of information systems and management are known (Mohd Salleh, Rohde, and Green 2017; Thong, Yap, and Raman 1996) and thus could be further transferred to the process mining domain, including process mining information systems. While we draw our results on a single case-study with comparative analysis against other projects, pre- and post-process mining deployment, and literature, we agree with Thiede, Fuerstenau, and Barquet (2018) and look forward to extending the case study and including multiple companies and adoption across multiple (production) teams within an organization. Following Ongena and Ravesteyn (2019), who mention benefit differences between product and service organizations, might transfer their research to departments within an organization. Generally, future case studies may also take a stronger focus on the deployment phase of process mining initiatives, as those often go beyond the scope of inducted case studies.

2.6 Conclusion

We present a multiple case study on investigating the success factors of process mining projects in practice, with a particular focus on manufacturing companies. We shed light on the generation of business value given the context of challenges and motivations for process mining's practical use, and do so by observations and conducting interviews with employees of different hierarchies. Given the context of our observations, we were able to develop the PM-BV framework and additionally point out motivations, challenges, and additional success factors that are relevant for process mining in practice, which we generalize through multiple partners and additional interviews with process mining experts and consultants. Following Grisold et al. (2020a), we present how "process managers perceive process mining" or according to Eggers and Hein (2020): "What collaboration practices influence the implementation and usage of process mining artifacts?". In answering these questions, we focus on the business value of such process mining projects. In particular, we depict challenges that have to be overcome to establish process mining in an enterprise process network, which allows us to provide insights into business, organizational, and managerial aspects. In doing so, we derive lessons for research and practice and provide practical insight into the use of process mining in businesses. Finally, we formalize our findings, implications, and opportunities for research and discuss the findings in the context of the literature to propose upcoming research streams. In doing so, we look forward to providing a fundamental part that facilities the successful implementation of process mining projects and that leads the way for the generation of process mining enabled business value.

3 Context-Aware Process Mining — Disrupting Continuous Process Improvement

This paper is in review at the *Thirtieth European Conference on Information Systems* (Schaschek et al. 2022).

Established methods for Continuous Process Improvement include value stream mapping and Six Sigma to identify optimization potentials and derive recommendations for future process cycles. Contemporary manufacturing systems with high complexity and variety increase the effort required to perform such analyses. By leveraging recent advances in Process Mining we address the shortcomings of traditional methods by leveraging the potentials of big data analysis in manufacturing environments. In an Action Design Research project, we develop a context-aware Process Mining Information System artifact, which emerged in interaction with practitioners in a sociotechnical context. In doing so, we leverage contextual information and process information in a combined manner to provide automated visual and statistical support to support process improvement. The artifact is deployed in a manufacturing environment and evaluated both quantitatively and qualitatively, enabling the formalization of learnings.

3.1 Introduction

An established means for achieving productivity gains in manufacturing is Continuous Process Improvement (CPI) to. The associated benefits are typically quantified by the speed with which parts progress through a manufacturing process (Schmenner and Swink 1998). Such system flow improvements mainly comprise of the reduction of bottlenecks, process variation, and non-value-adding activities (Schmenner 2012). However, exactly these problems stem from increasing manufacturing process complexity due to internal and external influences (Schuh et al. 2019). An increasing customer demand for specialized products results in ever greater customization (ElMaraghy et al. 2013). As a result, internal influences impact manufacturing complexity by implementing product diversity in the value stream (Schuh et al. 2019).

Traditional methods to manage product complexity and CPI in the manufacturing environment, such as Value Stream Mapping (VSM) or 5S analyses, are paper-based and can only provide a snapshot of de facto process flows (Sunk et al. 2017). The manual recording of processes is time consuming and tedious, especially in a dynamic production environment where efficiency gains must be quickly identified (Lorenz et al. 2021). Six Sigma is another established method in the CPI context. Here, the focus lies on reducing variation to achieve a defect-free process (Valles et al. 2009). However, the method does not consider the entire value stream of a process. In summary, CPI encompasses numerous methods which, taken on their own, cannot unlock all the improvement potential. Therefore, some companies take advantage of the complementary nature of Lean and Six Sigma and implement hybrid methodologies that combine fundamental characteristics, such as Lean Six Sigma (LSS) (Bhuiyan and Baghel 2005). For example, Drohomeretski et al. (2014) address this by analyzing the scope of each method individually and introducing a combined LSS method in the context of a multi-company case study. They point out the common objectives of the methods, notably the reduction of waste manifested by nonvalue-added work, cycle time, or instability. These improvements ultimately allow the value of LSS to be quantified. However, LSS is still a manual paper and pen-based analysis with associated shortcomings.

In parallel, continuous digitalization has favored the collection of extensive data related to the production process. This promotes synergies between traditional CPI methods and the analysis of digital traces. For example, data-driven analysis of production processes can reduce the value stream complexity (Schuh et al. 2019). Recent studies emphasize the fusion of traditional CPI with advanced techniques from big data analytics (Kregel et al. 2021; Knoll, Reinhart, and Prüglmeier 2019; Lorenz et al. 2021). In particular, Process Mining

(PM) has found great favor in addressing the pitfalls of traditional CPI methods. Using data recorded with process-aware information systems, its central objective is to derive as-is process models to provide postmortem insights on process executions to derive process improvement recommendations (van der Aalst 2016b). Thus, making use of these insights enables diagnostics of process behavior and promotes efficient CPI.

The focus of most PM studies is on discovering control flows of processes from event logs (Garcia et al. 2019). However, PM should go beyond the analysis of historical event log data to include all available data sources and techniques (Tiwari, Turner, and Majeed 2008; Zerari and Boufaida 2011). This results in a big data environment with heterogeneous data sources (Zhang et al. 2017). Analyzing these in a combined manner can provide valuable context information and add to the CPI objective. Especially in the manufacturing environment, the big data environment adds useful contextual information, such as sensor data, time data, location data, frequency or information from related devices (Becker and Intoyoad 2017).

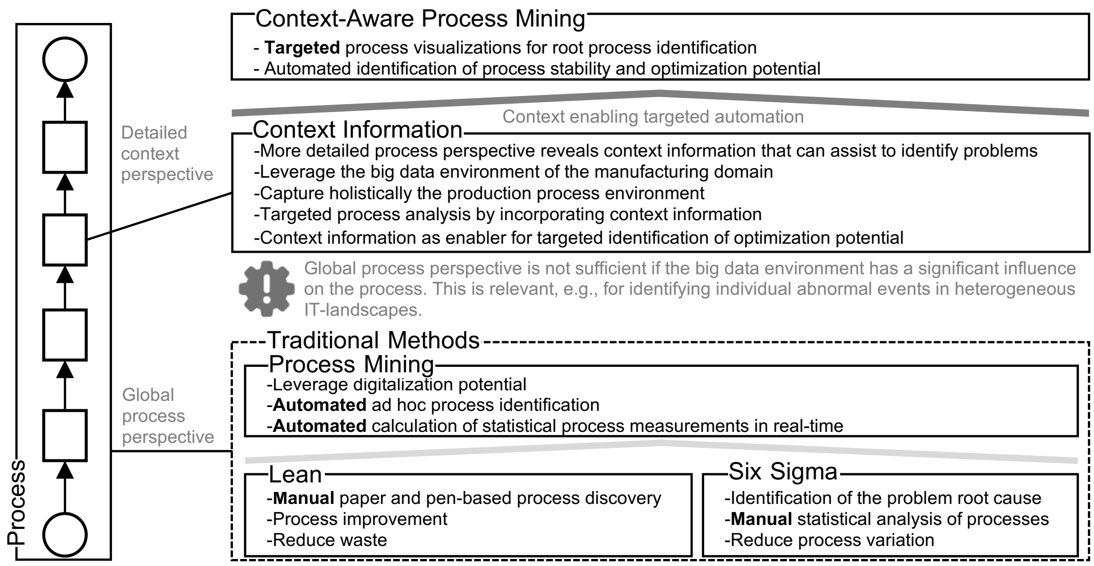

Figure 3.1: Synergies between Lean, Six Sigma, Process Mining and the benefit of context-awareness.

We address the mentioned limitations by developing a context-aware Process Mining Information System (cPM-IS), which tackles the manufacturing-specific task of optimally describing as-is processes in the context of production disruptions. Consider a production line where associated production information (e.g., a production log with events for starting and finishing certain production steps) is continuously collected during the production process. Due to complex processes for customized products, disruptions, such as incorrectly recorded events, maintenance problems, or product damages, occur over the course of the production process. Analyzing the entire production log for such events usually does not provide satisfactory results, as information on disruptions is not included. This complicates the mapping between disruptions and disruption-related processes. However, there is also context information on the disruptions collected (with disruption management systems similar to Lopez-Leyva et al. (2020) and Oberdorf, Stein, and Flath (2021)), that can be leveraged to identify disruption related processes in an automated and efficient manner.

To do so, we base our research approach on the synergies between big data analytics and traditional CPI methods to derive recommendations for process improvement through context-aware information. In doing so, we build on a hybrid concept inspired by context awareness as well as Lean and Six Sigma (Figure 3.1). Contrary to the manually executed LSS method, we remedy its shortcomings with dynamic analysis capabilities through automated and visual process analysis. In detail, by describing the production process statistically and visually in the context of disruptions, we enable *preventive countermeasures* for future process flows and aim to reduce the *risk of critical situations and disruptions* as well as *stabilizing processes*. To this end, we design a valuable IT artifact—the cPM-IS—by leveraging the Action Design Research (ADR) methodology in cooperation with a German medium-sized manufacturing company. To deploy the artifact in the partner's IT landscape, we follow the ADR process and iteratively design the artifact considering real-world requirements. The practical relevance enables us to quantitatively and qualitatively evaluate the artifact in real-world scenarios, to highlight the value of the cPM-IS, and to formalize our learnings.

3.2 Related Work

In the context of the current state of the literature, we first present PM as a central process data analysis tool and show the connection of PM to process improvement initiatives. Subsequently, we highlight the scientific interest in combining the concept of context awareness in PM and pinpoint out research contributions.

3.2.1 Process Mining and Process Improvement

PM is an emerging method for data-driven processes analysis and works as a nexus between traditional Business Process Management (BPM) and Data Science (van der Aalst et al. 2012a). By visualizing end-to-end value streams in organizations, PM provides process managers with valuable insights into real-world process behavior. The central pillar of PM contributes to the field of descriptive data analytics with its process discovery functionality (van der Aalst 2016b). Beyond process discovery, it provides diagnostic process analysis (i.e., conformance checking) to identify root causes of problems and can enable operational support (Munoz-Gama et al. 2016). Real-time process monitoring paves the way towards predictive and prescriptive data analysis approaches (e.g., Rama-Maneiro, Vidal, and Lama 2020b; Oberdorf et al. 2021a; Weinzierl et al. 2020). Given its many capabilities, it has been adopted in various domains, including healthcare, logistics, or manufacturing (Garcia et al. 2019).

In general, CPI enhances operational performance by reducing waste, process variations, cycle time, and improving overall quality (Bhuiyan and Baghel 2005), in addition to creating value for stakeholders (Näslund 2008). The systematic approaches of CPI guide companies in integrating improvement initiatives to achieve and sustain alignment of the main objectives (Snee 2010). The common denominator of process improvement allows the methods to be well suited to data-driven process analysis procedures. This is precisely what PM research is concerned with. PM uses data-driven insights for process mapping and enhancement. For example, Knoll, Reinhart, and Prüglmeier (2019), Schuh et al. (2020), and Lorenz et al. (2021) investigate the potential of PM for end-to-end processes in organizations. The results demonstrate the applicability of PM to the proven principles of Lean and VSM. They highlight PM's superior effectiveness and operational value compared to manual VSM and point to real-time improvement potential. On the contrary, Kregel et al. (2021) integrate Six Sigma and PM and present a proof-of-concept for incorporating PM into Six Sigma's improvement procedure. We present a novel *automated* and *LSS-aligned* PM artifact to leverage historic production-related data and evaluate its potential in a real-world manufacturing use case.

3.2.2 Context-Aware Process Analytics

In today's manufacturing environments, context awareness and its ability to improve process performance is not a new perspective. The concept of context awareness goes back to adaptive control systems, and its application areas are diverse (e.g., healthcare, Big Data, IoT) (Oprea, Moisescu, and Caramihai 2021). It is seen as a key component of future manufacturing solutions, as it can handle changes in the environment and improve the performance of intelligent systems (Lenz et al. 2020). For example, Dhuieb, Laroche, and Bernard (2016) presents a context-aware system to support daily activities in managing dynamic and complex manufacturing systems. Another example is Bertram et al. (2020), which focuses on the links between workers' activity recognition and the information about disruptions and their causes in production. They visualize the manual processes of workers with Petri nets and model worker support with hidden Markov Models derived from contextual aspects of activity recognition.

Context awareness is also considered an essential component of a successful business process analysis (vom Brocke et al. 2021a). The notion of context in BPM can be traced back over several years (Hallerbach, Bauer, and Reichert 2008) and Rosemann et al. (2006) define context as "the relevant subset of the entire situation of a business process that requires a business process to adapt to potential changes in the context variables." As a result, context can encompass all implicit and explicit impacts that effects the inherent situation of a process (Janiesch and Kuhlenkamp 2018). Kerpedzihev et al. (2021) consider aligning BPM methods with context-sensitive tools. Additionally, they provide guidelines, which point towards leveraging non-process-related data. However, research on data and context-aware process analysis is still in a formative stage. Zerari and Boufaida (2011) and Mounira and Mahmoud (2010) use PM to discover context information in the process environment and specify business rules for process flexibility. Another direction in research concerns the development of context-aware process discovery algorithms that link context data to relevant events (e.g., Shraga et al. 2019, 2020). Most research related to context-sensitive processes is focused on the medical domain.

The manufacturing domain is predestined to incorporate contextual information into decision making, as it is a heterogeneous environment with a wide variety of accessible information associated with the production processes (Becker and Intoyoad 2017). Process context information is the characteristics that describe the environment in which the process is performed (Abowd et al. 1999; Cunha Mattos et al. 2014; Rosemann, Recker, and Flender 2008), such as information on time, resources involved or location. Becker, Lütjen, and Porzel (2017) propose a novel framework for PM in heterogeneous logistic processes tested with simulated data, and Becker and Intoyoad (2017) examine its validity in regular practice. Unlike the approaches previously presented, Wang et al. (2021) investigate the context awareness of process recommendation methods to provide process modeling assistance and improve processes. They propose a process visualization tool to automatically visualize and annotate process nodes with contextual information to help process modelers model a production process. With regard to the process outcomes in the manufacturing area, Ehrendorfer, Mangler, and Rinderle-Ma (2021) investigate to what extent the context data streams collected during the process have an influence.

A comprehensive understanding of the use of contextual information in the environment of automated manufacturing processes is still pending. We contribute to this research gap by designing the cPM-IS artifact consistent with the fundamental LSS principles to improve processes and productivity in manufacturing. Based on the reference model for context engine integration, developed by Janiesch and Kuhlenkamp (2018), we focus on implementing a data-driven and automated cPM-IS artifact. As detecting root processes of disruption for process improvement in the manufacturing environment requires special knowledge due to the complexity of production processes, a handy tool is needed that automatically supports process specialists with visualizations. PM still requires manual work, therefore, we include contextual information to derive automated and targeted process analysis results. We define context data as information about events that occur in the background of a business process, in our case, in production and production disruption processes.

In the case of the cPM-IS we use the term context ambiguously, as we refer to it to both the augmentation of process event logs with process environment variables and to the grouping of event log observations based on information originating from the heterogeneous production data landscape. The developed cPM-IS enables *automated* process and context analysis in *real-time* by leveraging *contextual information* and *process data* in combination.

3.3 Research Methodology

The development of a viable cPM-IS artifact depends on the socio-technical environment of the artifact, which is why it involves many different stakeholders. As a result, it is fundamentally shaped by its organizational context during development, deployment, and operation. Based on this diagnosis, we decide to follow the ADR methodology (Sein et al. 2011). Unlike other design science research methods, ADR aims at designing a problem solving artifact while iteratively evaluating and learning from the continuous interventions (Peffers, Tuunanen, and Niehaves 2018). The ADR process consists of four iterative phases, *problem formulation*, *building, intervention, and evaluation*, *reflection and learning*, and *formalization of learning*. Taking into account the desired context-aware PM artifact, in addition to the ADR framework, PM frameworks become relevant for artifact development.

The well-established L*Lifecycle framework introduced by van der Aalst et al. (2012a) guides PM projects in practice through the five stages *plan & justify*, *extract*, *control flow model*, *integrated process model*, and *operational support*. As we interact in both domains, PM as well as organizational artifact development, we align both frameworks. Figure 3.2 shows the aligned ADR and L*Lifecycle frameworks, which we subsequently present according to the four iterative ADR phases.

3.3.1 Problem Formulation

Within the first ADR phase *Problem Formulation* and *stage zero* of the L*Lifecycle model, we place a particular focus on identifying potentials in the process analytics domain, determining potentials in existing data, and gaining extensive

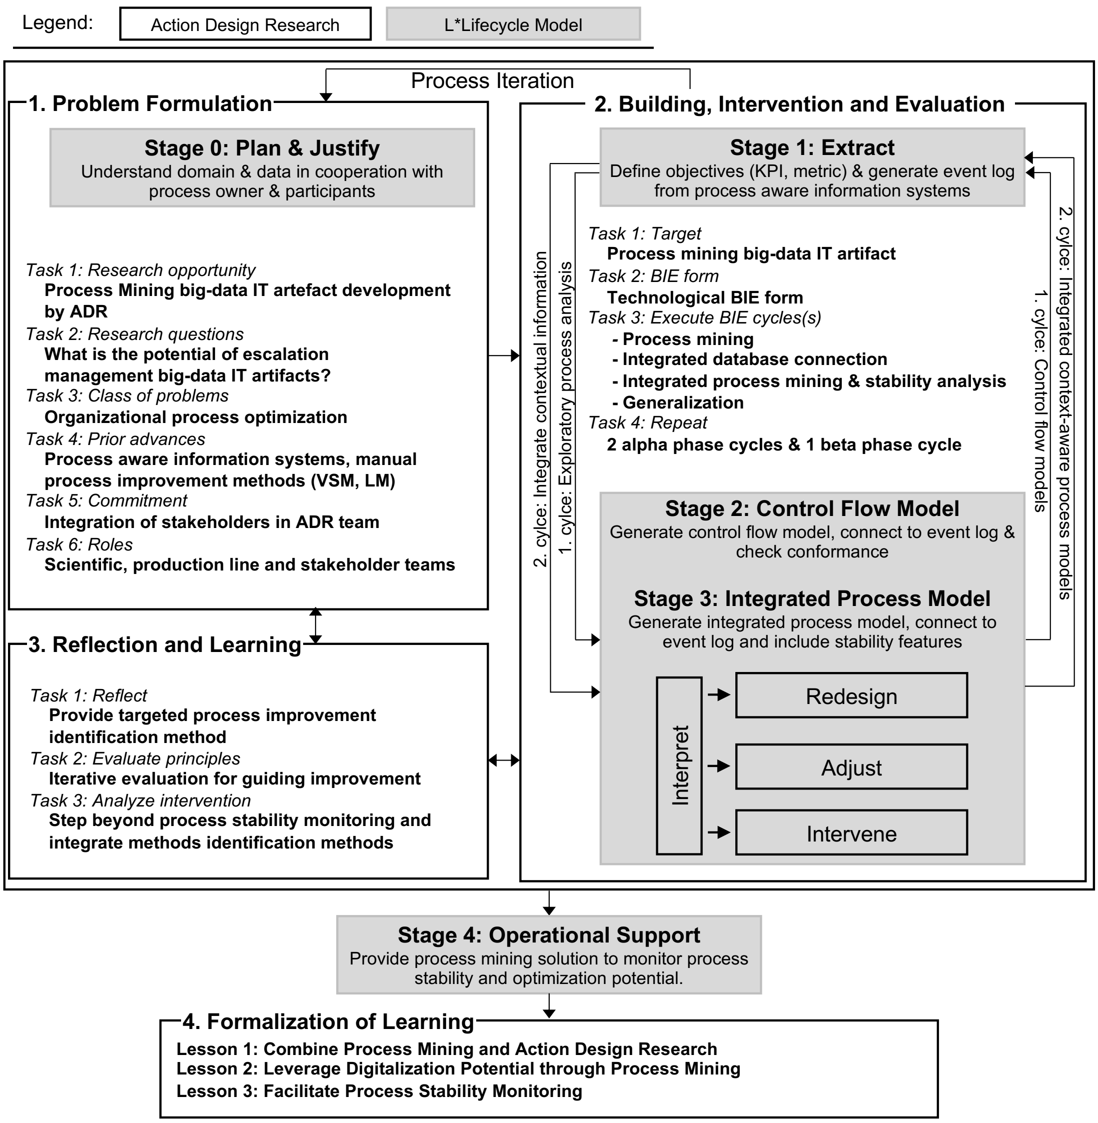

Figure 3.2: Stages of ADR with a task overview aligned with the L*Lifecycle model components. Adapted from Sein et al. (2011) and van der Aalst et al. (2012a)

understanding of stakeholders. Furthermore, we explore previous advances in the areas under consideration to build on research questions and justify the project plan. To identify relevant related work for our phenomenon of interest, we conduct a literature search. We compile a summary of previous advances in organizational and context-sensitive process analysis with a focus on the manufacturing environment. We detect reusable concepts for new ways of creating value in production processes with context information for CPI. The body of literature exposes data-driven process analysis methods, namely PM, which seeks to derive valuable insights into as-is processes in a descriptive manner.

On the basis of this foundation, we present the research opportunity to stakeholders responsible for managing and developing the production environment in our cooperation company. Subsequently, we establish an ADR team of three scientists, two project managers, and two data scientists for the entire artifact development process. To ensure long-term commitment, we extend the ADR team by stakeholders with positions at and close to C-level. Additionally, we assign roles and responsibilities according to the expertise of the ADR team member.

3.3.2 Building, Intervention and Evaluation

In the second stage of ADR *Building, Intervention, and Evaluation (BIE)*, we design the cPM-IS artifact and iteratively refine as well as evaluate it. Meanwhile, *stages one to three* of the L*Lifecycle model are performed sequentially to develop the part of the cPM-IS associated with PM. These stages represent the core of a PM project according to L*Lifecycle (van der Aalst et al. 2012a) and are repeated several times during the BIE cycles. Several iterations of stages one to three are performed to complete a BIE cycle, as the production line teams continuously evaluate and report to the scientific team. Additionally, we incorporate the reflected feedback from the ADR evaluation workshops at the end of the BIE cycles.

To start the iterative BIE cycles, we compile the main requirements of the ADR team, e.g., *automated context data analysis* as well as *integrated database interaction* for the artifact users. Additionally, we conducted interviews with stakeholders associated with the upcoming project to improve the perspective of the collected requirements. The stakeholders confirmed the requirements and added the requirements *enhanced visualization* and *analytics capabilities*. Subsequently, we abstract the specific requirements into more generic design principles (Table 3.1) and include design principles, such as automation and simplicity, for artifact design to address minimum user interaction and (automated) identification of context-event-related paths. Taking into account these principles, we draft our initial artifact consisting of four engines, each addressing a previously defined design principle. Beyond that, we have considered stakeholders' concerns about usability, functionality, and financial issues and built the artifact based on open-source software. As a side effect, the decision to use open source software proved valuable in terms of customization, ease of use, and simple integration into the heterogeneous IT landscape.

Table 3.1: Main requirements and design principles for the context-aware PM artifact.

Requirement	**Design principle**	**Addressed by**
Minimum user interaction	Automation	IT artifact
Connection of context data sources	Integration	Context data engine
Connection of process log sources	Integration	Process engine
Identification of context event related paths	Simplicity	Analytics engine
	Transparency	Decision support engine
Fit artifact to organization	Step-wise implementation	ADR process

Within the first BIE cycle, we address the stakeholder requirements by performing stages 1 and 2 of the L*Lifecycle model (Figure 3.2). In doing so, we identify candidate data sources and initiate an exploratory process analysis to discover the underlying processes. Data is extracted from process-aware information systems, and event logs are created that contained information about the production process of production lines. We use the event log data as input for an initial descriptive process analysis using publicly available open source PM tools (e.g., PM4PY). A descriptive data analysis reveals valuable insights into the production processes at hand. Production orders are selected as process instance identifiers and the process steps performed are defined as activities. As activities are one of the most essential elements of a process, special attention is paid to them. For example, the recurrence of activities is often an unwanted behavior, as it can be an indicator of inefficiency, interruptions, or disruptions in a production process. The findings of the initial process analysis are essential for creating a first process-driven prototypical design of the intended context-aware PM information system that meets the previously identified requirements.

Of particular importance in the artifact's design phase was the development of methods for *automated* PM. Beyond automated event log generation for general and context-aware processes, we focused on the control flow model and the development of the integrated process model. Thereby, we created methods for automated model generation and connected them to the event log data to provide possibilities, e.g., for conformance checking. Based on the initial models, we iteratively refined and evaluated the PM procedure. Subsequently, the design is explained to the participants in a focus group workshop, and we conduct a workshop to evaluate the first cPM-IS design based on interview and observation methods (Sonnenberg and Brocke 2012).

Initially, we only provided general graph visualizations and primary statistical analyses. Based on feedback, we refined the prototypical design of the cPM-IS artifact, for example *analytics capabilities* and *context event related process visualization*. During the second cycle, the artifact functionality was enhanced with context event-related process identification, and a user interface is created to organize the results of the process analysis. Regarding the feedback, we develop a method to leverage the context event data to identify context event-related process graphs and measurements. These refinements allow the evaluation of its practical use (Sonnenberg and Brocke 2012) and the feedback in the ADR workshops improved increasingly. Taking into account the feedback, the context-aware PM artifact is eventually integrated into the production environment, enabling descriptive ad hoc answers to *what is* and *why did this* happen on the production line by providing advanced diagnostic functionality.

As we finally established an automated real-time standard analysis and visualization[9], the ADR team was satisfied and we completed the alpha phase. In the ongoing beta phase the focus shifts from the main improvement to a generalization of the artifact application to cope with the selection and analysis of multiple chosen context events.

[9] E.g., context-aware PM analysis, including process graphs and statistics.

3.4 Context-Aware Process Mining Artifact Design

In the context of organizational CPI, we present a context-aware PM approach which extends state-of-the-art methods to to create a novel information system artifact. We integrate the various organizational data sources into this information system. Since the main goal of PM is to extract sequences of activities from an event log, our context-aware approach aims at the expansion of the process analysis dimensions in two aspects. First, we enrich the process event logs with context features, such as frequency of process paths, cycle time, or process stability. Second, we categorize process flows regarding the context in which they are in and visualize context-event-related process paths. Depending on the particular task, context-awareness is added to enrich the global process perspective or used to zoom into a detailed context perspective. For the initial design of the artifact, we use multiple process event logs from process-aware information systems, such as production and logistics event logs, to implement PM algorithms in the first cycle. In the second cycle we incorporate context information from a disruption event log [10] to facilitate targeted process visualizations. These various data sources are the input to the artifact, ultimately enabling automated decision support for process engineers and specialists.

To do so, our artifact (Figure 3.3) interacts with the organization—employees and databases—to improve processes through context-aware PM with a particular focus on the manufacturing environment. It consists of four components (gray boxes) with their respective features (light blue) and data (blue).

3.4.1 Context-Aware Process Mining Artifact Engines

The four artifact's engines have distinctive purposes. While the process and context data engines focus on data collection, the analytics and decision support engines aggregate the collected data and transform it to provide valuable decision support. To this end, the user is provided with a front-end (Figure 3.4) that assists in specifying certain processes (e.g., time information, order, or disruption IDs), which are input parameters for the process engine.

[10] This log data is provided by the disruption management system which automates the handling of disruptions during the production. In addition, it enables the collection of data, by means of context information, as described by Oberdorf et al. (2020) for a similar system.

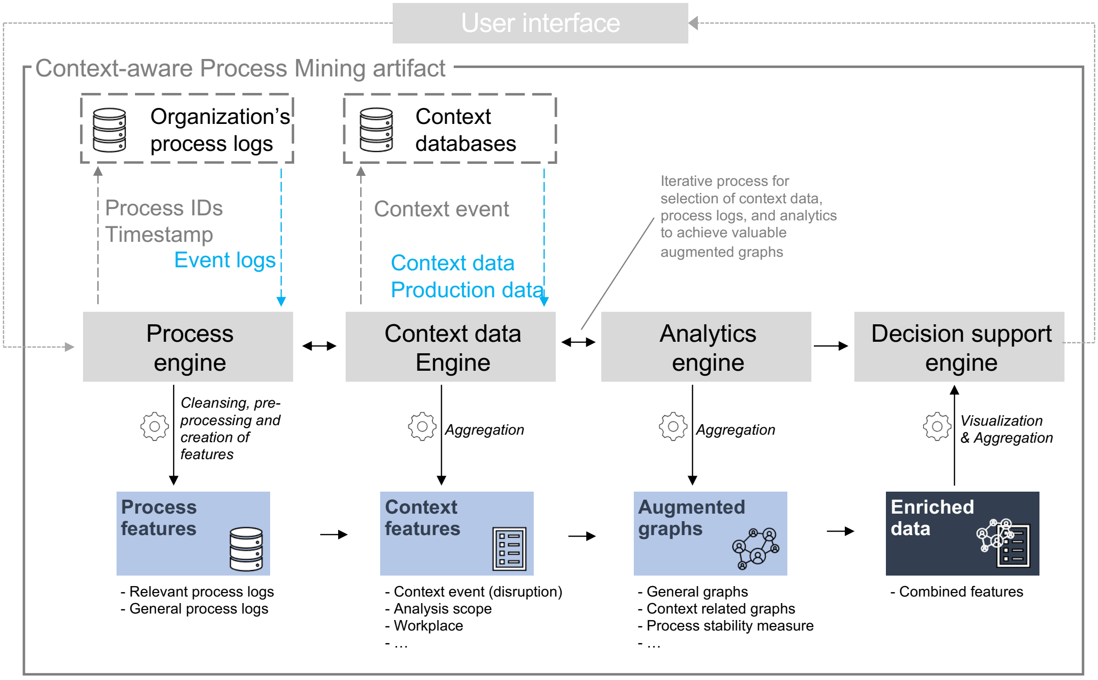

Figure 3.3: Context-aware PM artifact with organizational integration and data-source interaction.

The **process engine** is responsible for the automatic mapping of events from distributed process-aware information systems to a holistic event log. If processes originate from separately managed data sources documented in separate log files, the process engine combines the various sublogs for an integrative end-to-end process flow. To do so, the process engine's input parameters do not limit the scope of process events under analysis. Based on the input parameters, relevant process events are selected, however, some events might not be included due to time input parameters. Instead of neglecting such events, the process engine collects all associated events that match the events chosen due to input parameters, finally resulting in end-to-end processes. In addition, the process event log is cleaned and preprocessed for the use of PM algorithms. Process instances and associated features are forwarded to the context data engine.

The **context data engine** retrieves context information from an additional data base in the heterogeneous and scattered IT landscape.

There is additional information collected that assists in identifying processes that are (more likely) associated with a context event. In our case, the context event log is recorded with a disruption management system and comprises disruption events (context events) and associated information, such as work place and time. Having such context (disruption) events available enables us to combine the (complete) process log and disruption events. This combination with the complete log is of particular importance because the workplace (or event) where a disruption is noticed usually differs from the workplace where the disruption is caused. For example, small coating damage may originate from transportation to the final assembly and does not necessarily originate from the coating workplace itself. To this end, the extracted context events and related features of the specialized system are then used by the analytics engine to enrich the global process perspective processed by the process engine.

The **analytics engine** automatically combines process-related and contextual information and enables users to receive targeted and augmented graph visualizations. It also allows users to specify the scope of analysis through the artifact's interface by means of choosing context events and the processes to be analyzed. The user triggers this process by selecting single context events to identify context-related processes following the approach visualized in Figure 3.4. For example, based on the work order IDs[11], the process activities of the currently produced items are selected. This allows accounting for process events in the context of events' root process analysis. However, some of the selected processes could be unrelated to a context event of interest. For this reason, the sub-logs of several identical and comparable context events are additionally grouped and then processed by the analytics engine. Based on the user's selection, the engine simultaneously computes context-related process graphs and augments the generated graphs with process context information. Thus, context events serve as reference points for the automated identification of context-related processes. Process context information includes discriminative features (e.g., process stability, process duration, processing time, or re-occurrence of activities) that further describe process characteristics at the event log and activity level to enrich the global process perspective. This helps users identify process anomalies that cause context events. Statistical measures, such as the average processing time, are computed using the pro-

[11]Each production order receives a work order ID, which facilitates a unique assignment of process instances.

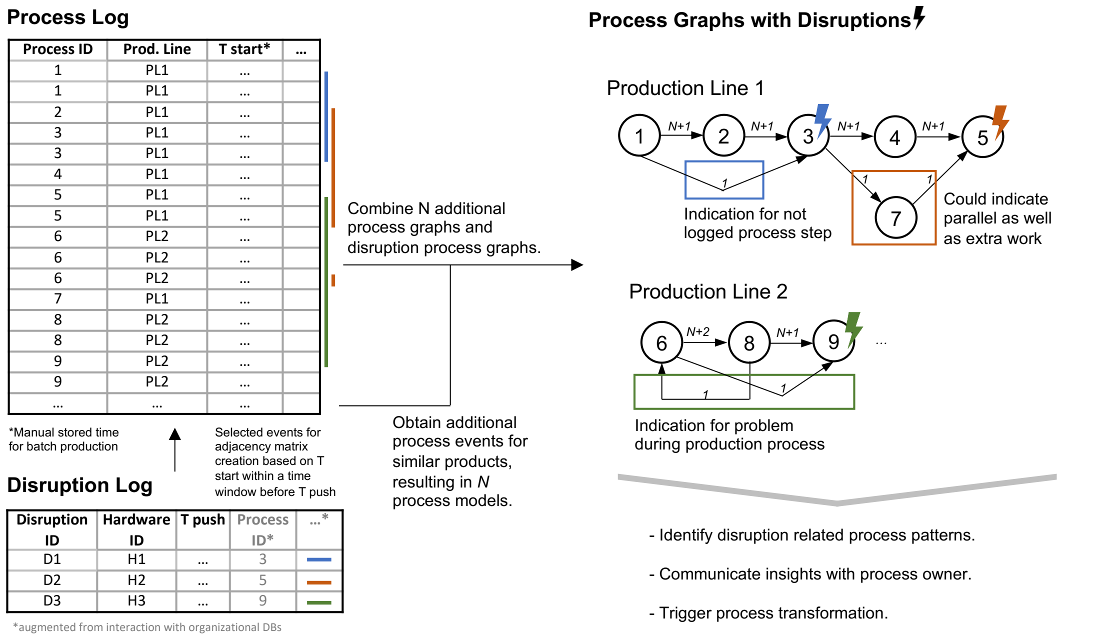

Figure 3.4: Process graph generation with visualization of context event causes.

cess features provided by the process engine. Additionally, the analytics engine uses the available context data from the disruption management system and process data to derive the measure of process stability, which measures the similarity of context-related process paths to main process paths. The process stability measure is calculated by creating trace profiles (Song, Günther, and Aalst 2008) and compute the Euclidean distance as a similarity-based measure between process instances. With this approach, we follow the methods of PM deviation detection for complex event logs (e.g., Li and Aalst 2017). Beyond context-related analysis, the analytics engine allows users to select process activities of similar products and exclude context-event-related production events, for a detailed product-related perspective.

To finally provide valuable insights, the **decision support engine** aggregates the resulting enriched data with combined features. To this end, the process graphs are visualized, and general and context event-related information (e.g., statistics) are provided in a user interface. Graph visualization is particularly important, as it enables the user to compare the general and context event related graphs. The visualization of process flows facilitates the comparison of processes considering the correlations with context events, and thus a targeted evaluation to finally depict differences and optimization potential.

3.4.2 Artifact Deployment

Following the iterative ADR process, we design the cPM-IS and deploy the artifact as part of the iterative BIE cycles. Moving from a prototype to a deployed solution, we operationalize the process described in Figure 3.3. Regarding the analytics capabilities, we automate data collection (process and context data) and feature engineering and implement the context-aware process mining approach (Figure 3.4) as well as the decision support engine. For implementation and automation, we rely on PM4PY (Berti, Zelst, and van der Aalst 2019) and Python, combined with a customized front-end (Figure 3.6). The Python-based backend is deployed on a standard virtual server and connected to organizational systems such as the SAP, MES, and Disruption Management System, thus enabling automated data processing.

Users access the cPM-IS front-end through standard web browsers a self-service analytics fashion. Thereby they individually select (context) events of interest to analyze associated processes. Having deployed the system with the collaboration partner, we can rely on a two-fold evaluation that covers a quantitative evaluation of the cPM-IS method and a qualitative evaluation from a user perspective.

3.5 Evaluating the Artifact

To evaluate the proposed cPM-IS artifact, we demonstrate that the instantiation of the artifact operates satisfactorily in helping process engineers and workers perform contextual process root analysis for process improvement. Therefore, our evaluation aims to justify the artifact, focusing on whether the artifact works well in the problem setting under consideration (Venable, Pries-Heje, and Baskerville 2012; Peffers, Tuunanen, and Niehaves 2018).

Generally, the analysis of production processes requires access to special knowledge because of its vast complexity and variety. Identification of the root processes of production disruption was a manual process performed by production process specialists prior to instantiation of the artifact. This was time consuming, as root-process events are not obvious from the direct occurrence and recording of a production disruption with the disruption management system. Therefore, the objective of cPM-IS is to facilitate process and technical specialists in identifying process events associated with disruptions. In parallel, it supports the achievement of the basic principles of LSS through automated and disruption-aware process visualization in real time, as well as statistical process metrics. To achieve this goal, state-of-the-art process analysis technologies were used to provide decision support. We first evaluate this decision support quantitatively and subsequently provide a qualitative evaluation based on interviews with cPM-IS users.

3.5.1 Quantitative Evaluation

The cPM-IS provides process specialists with statistical process measures, such as average process duration (process instance level) or processing time (activity level). These measurements reveal performance-related process variations and anomalies. In addition, the measurements enable a dynamic process flow analysis from different perspectives, namely activity, process instance, or resource level based on historical data. The as-realized process flows reveal that not all executed process traces are compliant. To draw conclusions about rare process variants or deviant process flows, cPM-IS establishes a process stability measure. This is needed to detect the connection of processes to context events. Standard PM methods enable conformance verification by comparing the as-designed with the as-realized process model and deriving the measure of fitness (Munoz-Gama et al. 2016; Aalst, Adriansyah, and Dongen 2012). In the companies' dynamic process environment, manually created process models quickly become obsolete due to customized production. For this reason, the artifact uses the process stability measure to compare the similarity of main process paths to context event-related process paths. The process stability is calculated for each process instance and can help to detect abnormal process variants. When comparing the process stability measure for all process instances in the process log, those related to production disruptions occur to have lower similarity measures than general process paths (see Figure 3.5).

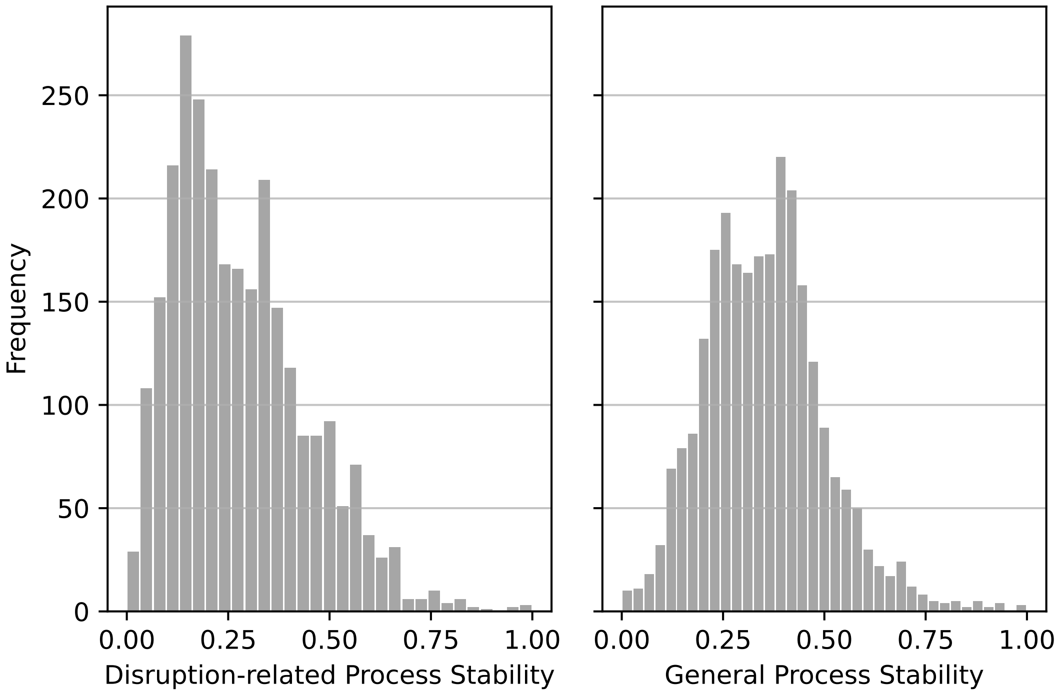

Figure 3.5: Histogram demonstrating the distribution of process stability for general and disruption-related processes.[12]

The identification of context-related paths is additionally assisted by automated reporting and visualization of process performance metrics that facilitate process analysis. Using process discovery algorithms to visualize process flows completes the artifact design to capture a holistic and dynamic mapping of production process flows. Comparing as-realized general and context-sensitive process graphs and associated measurements allows process specialists to draw conclusions on process flows related to the context event (Figure 3.6). Beyond that, the overall conformance of the process for each context-aware process graph can be evaluated by comparing them to the general process model.

[12] Note that lower process stability values denote dissimilarity and vice versa high values denote similarity of processes to the main process.

[13] We apply the heuristic miner to derive heuristic nets provided by the open-source PM platform PM4PY.

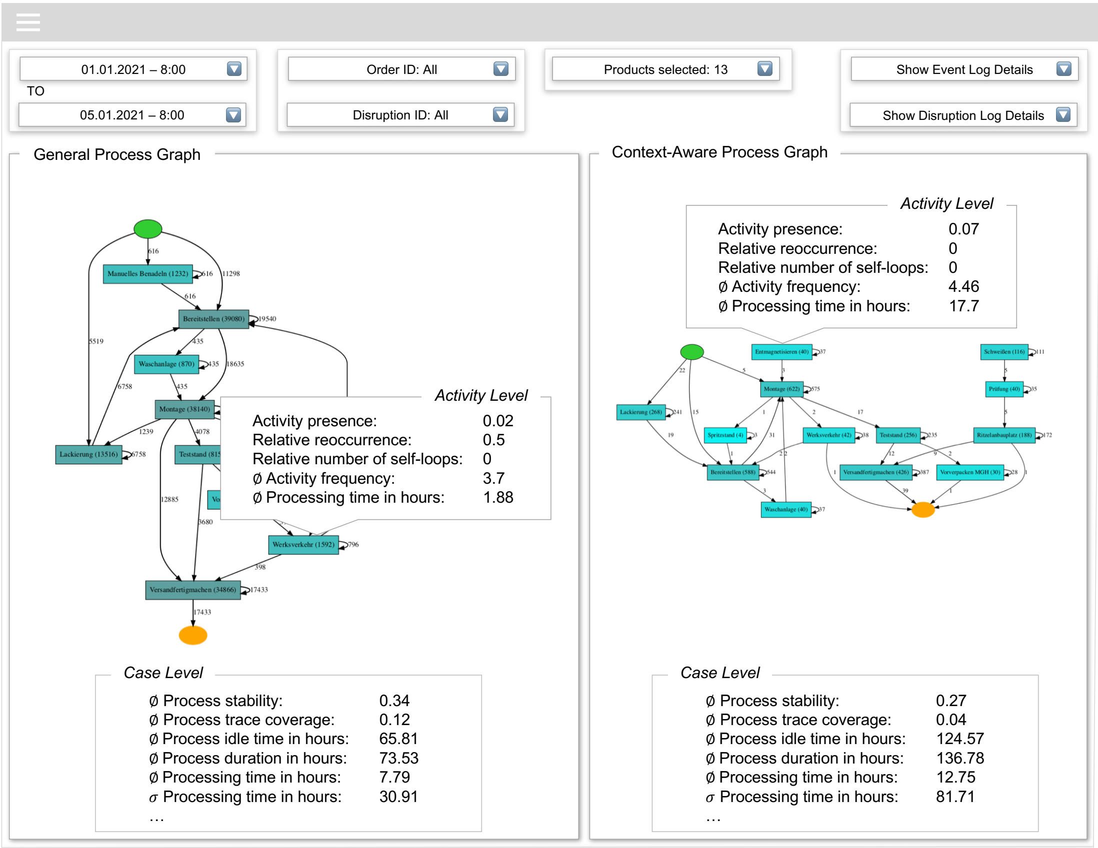

Figure 3.6: An exemplary comparison of an as-realized main and a context-aware process graphs[13], which are annotated with statistical measures on process instance and activity level, as well as the process stability measure.

3.5.2 Qualitative Evaluation

Beyond the application examples of the cPM-IS, we can provide qualitative feedback from the ADR scientific team members. From a technical perspective, the particular importance of automated data processing was mentioned. Besides, the data scientists pointed out the practical importance:

> *"In particular the developed method [for context event related process path identification] is highly valuable. It enables us to perform more targeted process analysis and thus even reduces the required amount of time."*
>
> (Data Scientist)

The method is integrated into the artifact, which was designed concerning the initially collected requirements, mainly minimum user interaction and identification of context event-related paths. Thus, we expanded the scope of employees, which can perform process analysis massively. Previously, data scientists were necessary to obtain and analyze the data, which is now automated through the artifact.

> *"The tool is quite helpful to get process insights, even without help. This facilitates the work and coordination."*
>
> (Process Engineer)

As we evaluated the artifact instantiation, we discussed our findings within the extended ADR team. Finally, we were able to formalize our learning as the last stage of the ADR process.

3.6 Formalization of Learning

In the fourth ADR stage, we formalize our learning in terms of general organizational implications based on the experiences of the ADR process. The two resulting lessons have formalized implications for CIOs, IT managers, and researchers facing related problems.

Lesson 1: Open-source facilitates context-aware Process Mining

Given the context of our research and cooperation, we frequently asked and discussed the following question: *Why not commercial Process Mining software?*. Usually, companies prefer established software because they provide support and more sophisticated specifications or interfaces to the company's IT infrastructure. It is convenient but also associated with high costs. Weighing the costs against the efficiency of the project and the business value created by using commercial software solutions yields insufficient results in light of limited experimental capabilities. Especially in research projects or research collaborations, this is often a decisive factor in pursuing open-source options. Exactly this point has made the decision in our case.

Open source alternatives offer a wide range of functions and also facilitate specific adaptations. These turned out to be essential in the course of our cooperation. To pursue automated and context-aware PM, it is necessary to extend standard PM methods and integrate them into existing systems. Our cooperation partner is already using Python across different projects, and this facilitated the development of the cPM-IS with individualized functions and its usage. This confirms observations on open source analytics solutions put forward by Gust et al. (2017). Our custom solution, unlike standardized procedures (Zelst et al. 2020), allows to combine contextual data and process data from heterogeneous data sources for PM.

Lesson 2: Leverage and extend context data usage

In the context of the cooperation, we could access a wide range of resources within the company. In particular, the newly established disruption management system provided a rich environment for our research. The access to context data facilitated the implementation of cPM-IS in the company. Although we evaluated a single use case, following Lorenz et al. (2021), we consider our practical findings to be generalizable to similar production environments. In our case, contextual information on production processes is required on top of process-aware production data. For some companies, access to context data will not be available, although there are many promising alternative data sources (e.g., software failure data (Gruszczyński 2019)). These should be examined for their suitability, e.g., for the targeted identification of optimization potential, and combined with PM methods. In summary, more studies are needed to reach a more general conclusion. In particular, it would be desirable to adopt this method to other domains.

3.7 Conclusion

Recent advances in business process analytics enable real-time and automated support of value stream analysis in the manufacturing environment (Lorenz et al. 2021). Building upon the technology of PM, we develop a context-aware PM artifact in an action design research project. We contribute to traditional CPI approaches such as Lean Management and Six Sigma using state-of-the-art process analysis technologies with our IT artifact. In particular, we address a lack of research in the PM literature that deals with the available data in the heterogeneous big data environment. Using and analyzing these in a combined manner can provide additional insights and create business value. Given the vital role of process improvement in the manufacturing environment, we show with the instantiation of our designed cPM-IS artifact that context awareness guides the way towards better and more automated approaches for CPI.

However, quantifying the success of the instantiation and the resulting reduced downtime and process improvement would require a numerical evaluation. Under the current pandemic situation, reliable numerical evidence of process improvement from post-instantiation of the cPM-IS was impossible. The decrease in production volumes during this period distorted process improvement actions in terms of manufactured products. We plan to evaluate the cPM-IS in future research initiatives quantitatively.

However, the results of the cPM-IS artifact reinforce the importance of integrating contextual information in PM analytics initiatives. In doing so, they underline the business value created by analyzing the past of business processes. However, process analysis can go beyond automated descriptive data analysis by exploring the future of business processes. Taking into account the classification of the information systems of Schwegmann, Matzner, and Janiesch (2013), our system has a primarily informational character. Beyond plain information, it could be extended with operational functions, such as real-time forecasting and decision support, as indicated by Oberdorf, Stein, and Flath (2021) in the context of disruption management. We also look forward to integrating the associated aspects in future work. In doing so, we follow Wang et al. (2021) and contribute to the overall goal of process improvement, which is a crucial driver of business success. Adding and extending the fundamental aspects of our cPM-IS will be an essential building block for automated organizational process improvement.

4 Analytics-Enabled Disruption Management: System Development and Business Value Assessment

This paper is published in *Computers in Industry* (Oberdorf, Stein, and Flath 2021).

Industry 4.0 initiatives can help traditional manufacturing industry cope with increasing global competition. Such solutions facilitate transparency, automation as well as business process transformation. This paper elaborates on a collaboration with a medium-sized manufacturing company. We highlight the design, evaluation and roll-out of an disruption management system with integrated data-driven decision support. We do so by applying an action design research process. Thereby, our study focuses on the system design concerning the creation of business value.

The system leverages state-of-the-art machine learning algorithms for disruption type classification and disruption handling duration prediction. These predictions can be embedded in an integrated planning procedure leveraging diverse organizational data sources (e.g., personnel availability, production plans) to instantiate a prescriptive analytics solution. Combined with informative analytics insights, this allows the proposed system to generate significant business value by reducing disruption durations. In the long run, the transformational business value enabled by the system is likely to exceed the automational business value. This highlights the special importance of tight integration of industrial analytics applications within business processes.

4.1 Introduction

In recent years increasing global competition (Matschewsky, Kambanou, and Sakao 2018) as well as disruptive supply and demand shocks (e.g., COVID-19 pandemic, US-China trade war) have increased cost pressure across all industries. At the same time, there is a trend towards highly customized products running against the traditional efficiency lever of scaling up production. To better cope with these challenging situations, manufacturing companies seek to increase productivity by adopting lean manufacturing practices or process design and facility layout improvements (Kovacs 2020). More recently, firms started to push forward industrial internet initiatives (Rüßmann 2015; Gilchrist 2016; Kagermann et al. 2013) in order to support or automate labor-intensive processes and increase their productivity (Ghobakhloo and Fathi 2020). To achieve this, manufacturing systems are sensorized and connected to IT systems (Monostori et al. 2016; Müller et al. 2017) allowing the automated and continuous collection of information. Wang, Törngren, and Onori (2015) notes, that such Industry 4.0 initiatives “emphasize[s] the extension of traditional manufacturing systems to full integration of physical, embedded and IT systems including the Internet.". We adopt the framework proposed by Mooney, Gurbaxani, and Kraemer (1996), to assess the business value created by Industry 4.0 (I4.0) from a theoretical lens. To account for the recent developments in business analytics, we map the respective success value tiers to the corresponding levels of analytics sophistication (Camm et al. 2020) and database interaction to establish potentials and requirements (Figure 4.1).

The automation of processes, such as data collection, creates *automational value*. This is the basis for *informational value*, which emerges from information collection and subsequent dissemination, e.g., through process monitoring and dashboards. Leveraging automational and informational business value is supported by existing commercial software solutions for automated (disruption) processing, information visualization or descriptive analytics.[14] To go beyond automational and informational value companies have to transform business processes to become data-driven and thereby create *transformational business value*. This necessitates integration of live production

[14]Specialized firms offering such solutions for manufacturing firms include `www.tulip.co`, `www.peakboard.com` and `www.l-mobile.com`

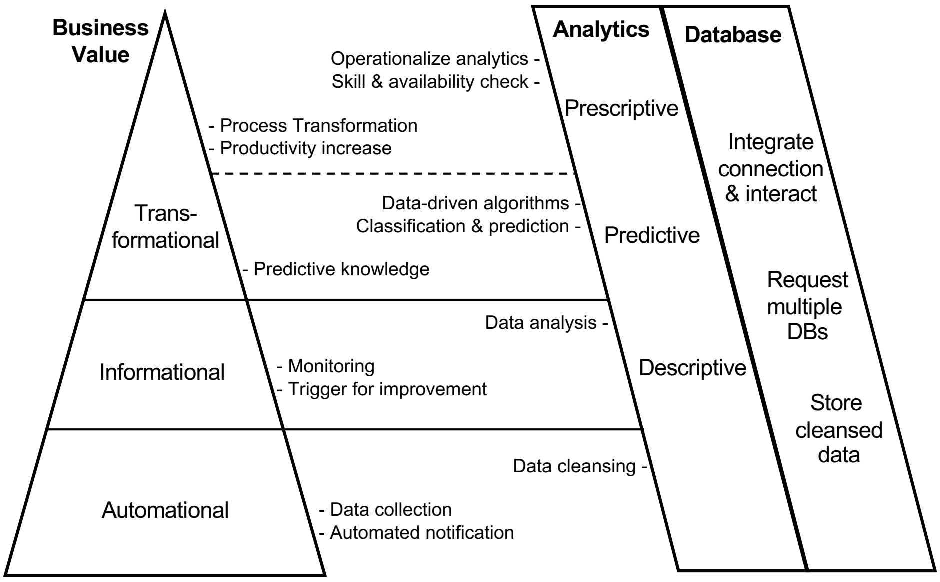

Figure 4.1: Business value in relation to the adoption of analytics. (adopted from Mooney, Gurbaxani, and Kraemer 1996)

data with other organizational information systems such as manufacturing execution systems (MES) or production planning systems (PPS). I4.0 applications can better justify the required investments by creating business value along multiple dimensions. Furthermore, information processing and system integration for a wide range of applications is necessary.

This paper is concerned with efficient handling of production process disruptions via an analytics-enabled disruption management system. In manufacturing settings, disruptions[15] result in a situation where a worker cannot continue the current task. Hence, production at single workstations or even across the complete line is interrupted resulting in sizeable disruption costs. However, due to today's manufacturing processes' complexity, such disruptions cannot entirely be avoided. To reduce the cost of downtime, companies rely on disruption management systems to efficiently handle disruptions and in turn improve productivity (Lopez-Leyva et al. 2020). Such systems usually prompt a responder[16] in order to assist in solving a disruption (Macdonald and Corsi 2013).

[15]Typical reasons include, e.g., missing materials, damaged parts, or non-functional machines.
[16]e.g., production team leaders, or logistics or maintenance specialists

Existing disruption management systems automate the notification process as well as the collection of disruption data and increase shop-floor transparency by means of monitoring tools such as dashboards. These systems typically inform any available response person (comparable to pushing the service button in an airplane). This approach is well-suited for simple settings where tasks can be handled by any responder. However, disruptions in more complex production environments often require specific skills and the notification of any available responder is inadequate. In such settings, predefined responders with a broad skill-set are deployed to analyze the cause of a disruption and subsequently notify a suited expert. This approach avoids sending wrong responders at the cost of an expensive two-step approach.

In contrast, we envision a data-driven disruption management system that automatically identifies the underlying disruption cause and dispatches the best-suited expert. Clearly, the performance of such a system is driven by the ability to correctly predict the disruption type. While good predictions significantly reduce down-times and costs by sending the correct expert directly, bad predictions lead to wrong dispatches—and therefore unresolved disruptions requiring re-dispatches of other specialists—and incur additional costs.

Such a system is currently not offered by commercial available systems. We collaborated with WITTENSTEIN SE, a German medium-sized manufacturing company with multiple distributed production and assembly lines for highly customized mechatronic products. As part of a large scale company-wide digitalization strategy, a disruption management system has been developed. In the short run the objective was to create automational and informational business value. Beyond these initial benefits the system should generate transformational business value in the future.

This paper discusses how to design integrated disruption management systems facilitating transformational business value creation. In particular, we illustrate how to integrate analytics-enabled decision support and highlight corresponding use-cases.

4.1.1 Status Quo Process

Our research starts with a review of the current disruption process. We consider a production process spanning the process steps supply of parts, component production, assembly, testing and shipment, each with respective work stations. At each workplace disruptions (e.g., component damage or missing materials) occur frequently. Responders are dispatched to resolve these disruptions. However, dispatched responders are oftentimes not close by, may lack the necessary skills to resolve a given disruption or may be unavailable. This means that the current disruption handling process (Figure 4.2) requires workers to search for an appropriate and available responder. After a disruption is communicated and additional information, e.g., on a certain machine or product is obtained, the responder assists in solving the disruption. This evolved process has some obvious shortcomings:

- *Searching for an appropriate and available response person* is time consuming. Most time is taken by the search itself, especially if the first contacted response person cannot assist in solving the disruption.
- *Interruptions of colleagues* are common during the status quo process, as possible response persons are interrupted throughout the search process[17].
- The *disruption is only communicated once a response person is found*, which adds additional time in which the response person has to think about a possible solution before being able to assist the worker in problem solving.
- The *lack of disruption information* results in the necessity to interact with organizational databases to find additional information on a machine, workplace or product.

[17] e.g., if they are occupied—due to phone calls or meetings—or do not have the skills to solve the disruption that caused the disruption

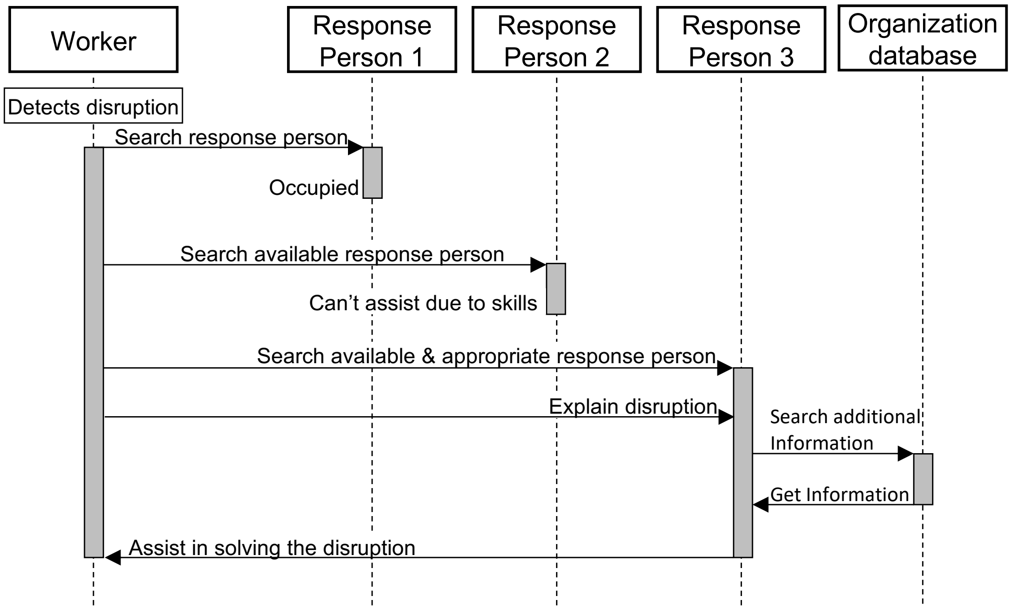

Figure 4.2: Traditional disruption management process.

4.1.2 Research approach

The shortcomings highlight significant improvement potential. By adopting a resource-based view to identify supply chain productivity potentials (Chae, Olson, and Sheu 2014), we seek to improve the employees (e.g., worker or response person) utilization through an IT system. To achieve this, the new system must automate communication between a worker and the appropriate response person. In order to notify an appropriate and available person, we need to know *what is happening* and if the worker has the skills to assist in solving (appropriate) as well as *how long the disruption's solution will likely take*, to ensure availability and avoid overbooking of responders. To this end, the triggered disruption type has to be classified, the duration predicted, and a response person dispatched. Using advanced machine learning models we establish a data-driven decision support system which assists disruption handling during the production of highly customized products. Data-driven decision support facilitates business value creation through advanced analytics (Davenport and Harris 2017; Brynjolfsson, Hitt, and Kim 2011). Yet, in operational processes such systems are underrepresented in research. Our research sheds light on how analytics-enabled I4.0 applications generate business value along three dimensions—automational, informational, and transformational. In the context of disruption management, we contribute methodologically by integrating suitable analytics approaches.

4.2 Related Work

We first provide an overview of current information systems research with focus on analytics-enabled business value facilitation. Taking into account the production environment, its digitalization, and possible use-cases of analytics-enabled IT systems, we review recent advances in industrial internet applications and advanced analytics. Subsequently, we highlight recent advances in disruption handling with a special focus on disruption management systems.

4.2.1 Value of Information System and Operational Information Systems

Companies strive to generate business value (Mooney, Gurbaxani, and Kraemer 1996). Typical benefits of information systems include lower operational costs (process automation), improved information dissemination as well as process transformation (Daneshvar Kakhki and Gargeya 2019).

However, the decisive factor here is not only the support but the generation of business value. Therefore, several factors are essential. Firstly, the development, communication, and provision of the system (Kohli and Devaraj 2004). On the other hand, internal factors shape the possible business value. Barua et al. (2004) analyze such factors based on a structural model. They identify system integration, customer and supplier readiness as important factors. Comparable to the results in Kohli and Devaraj (2004), a decisive factor is the cooperative development of such systems. To ensure this for our study, we apply the action design research (ADR) approach (Sein et al. 2011).

Another factor that is gaining importance for business value generation are business analytics solutions (Wang et al. 2019b; Chen, Chiang, and Storey 2012). Crucial success factors of analytics-based approaches is, among other things, the competence of the analysis as well as the data quality (Côrte-Real, Ruivo, and Oliveira 2020). Both become increasingly important as we consider the amount and extent of data sources, available due to the current digitalization of production environments.

4.2.2 Industry 4.0 and Advanced Analytics

The global manufacturing industry is facing a fourth industrial revolution fueled by the internet of things and servitization. To realize I4.0 capabilities, traditional manufacturing systems are sensorized and connected to IT systems (Monostori et al. 2016; Müller et al. 2017), creating cyber-physical systems (CPS) (Wang, Törngren, and Onori 2015). These developments are geared towards the vision of smart factories facilitating highly customized production orders (Thoben, Wiesner, and Wuest 2017).

In order to understand the differences between traditional manufacturing systems and CPS, Penas et al. (2017) note that "...the whole traditional industrial setting of methods, processes and tools in manufacturing system design and analysis then have to be deeply reconsidered." This reconsideration leads to completely new solutions. Special importance is placed on sensors, networks, services and interfaces with respective layers of representation (Boyes et al. 2018; Li, Da Xu, and Zhao 2015). "Automation, big data, analytics, and the Internet of Things (IoT) [...] create opportunities for substantial gains along the entire industry value chain" (Gürdür, El-khoury, and Törngren 2019). Ultimately this will lead to enhanced productivity, especially in combination with traditional lean management practices (Ghobakhloo and Fathi 2020).

Predictive and prescriptive analytics applications play a central role in this context. Instead of manually classifying defects or disruptions (Lopez et al. 2010; Rubin et al. 2003), analytics approaches enable automated classification or the prediction of future defects. They facilitate decision support in areas such as quality management (Fahey, Jeffers, and Carroll 2020; Lyu, Liang, and Chen 2020; Ma and Chu 2019; Sanchez-Marquez et al. 2020; Stein, Meller, and Flath 2018), reliability analysis, and predictive maintenance (Al-Dulaimi et al. 2019; Zhang, Zhang, and Li 2019; Zschech 2018).

4.2.3 Disruption Management

Given the complexity of manufacturing processes, interruptions and disruptions cannot be entirely avoided. Disruption management systems are deployed to improve the disruption handling processes and in turn productivity (Lopez-Leyva et al. 2020). Disruption management systems assist in organizations' disruption handling, e.g., in manufacturing (Müller et al. 2017). A disruption triggers a disruption process, which is solved within the organization. Thereby, different dimensions of disruption management—hierarchical and functional—have to be considered (Malega 2014). Disruptions on a functional level address a colleague on the same hierarchy level, while hierarchical disruptions are forwarded to a supervisor. Kassner et al. (2017) describe these different types of human interaction, depending on the hierarchy level, as a social factor. As shown by Romero et al. (2017), the social factor is a crucial design principle during the development of a disruption management system or in general for CPS (Cardin 2019).

Jaech et al. (2018) present a promising approach as part of availability checking. They predict the duration of energy system outages, based on historical data. Facing the problem of unknown disruption duration, the adoption of such an approach seems beneficial. Yet, the transfer to a production system, the extension to live data as well as the integration in the system are open points in the context of disruption management.

Lopez-Leyva et al. (2020) show the development and advantages of a disruption management system by the means of a case study. The practical main contribution is an automated notification process—yet without operational analytics capabilities. They follow Dombrowski, Richter, and Krenkel (2017) and Da Silva and Baranauskas (2000) by creating an automated communication process leveraging Andon and I4.0 design principles.

A more operationally focused approach is taken by Yang, Qi, and Yu (2005). They re-plan the entire production whenever a disruption occurs. However, this approach mainly relates to general production planning. This is less suitable for disruption management since the disruptions should be solved as fast as possible, without massively affecting the planned production. To minimize the impact on the production plan, we focus on downtime minimization—through automation and notification of appropriate and available responder.

This is achieved through a combination of automation, communication and analytics.

In terms of I4.0, advances in automation and communication technology are beneficial to improve the disruption handling and thus increase productivity. In particular, more options for integration and connection to enterprise systems arise (Mohamad et al. 2019)—either through an integration to the manufacturing execution system (MES) or integrated database connections (Răileanu et al. 2018). Additional organizational data sources shift the disruption management towards a big-data problem (Chen, Chiang, and Storey 2012; Russom 2011), including the regarding potential of system intelligence (Qin, Liu, and Grosvenor 2016), e.g., decision-making (Alcacer and Cruz-Machado 2019).

However, current research, as well as commercial systems, are primarily focused on process automation. Disruptions are recorded, automatically processed, and a responder is notified. While such systems can generate automational and informational business value they do offer an avenue towards process transformation. To achieve such a comprehensive scope of value generation one needs to incorporate analytics capabilities. Consequently, analytics-enabled disruption management systems and their relation to transformational business value is an open research gap.

We address this gap as we depict how a classical system—primarily focused on automation—can be extended by integrated analytics, in order to generate additional benefits. In particular we go beyond existing approaches (Lopez-Leyva et al. 2020; Dombrowski, Richter, and Krenkel 2017; Da Silva and Baranauskas 2000) and shed light on analytics-enabled dimensions for business value generation. This contribution is of particular importance in view of the current I4.0 transformation. Manufacturing companies already have (information) systems in place, which generate business value from an automational perspective. In addition, the digitalization of production units expands the number of available data sources. However, the potential of the resulting data often remains unused. While the systems automate activities, they necessitate additional analytics efforts to gain insights into the data—and thus the underlying processes. Often, the added value of such insights is not conscious or tangible in a timely manner, which might limit the adoption of analytics-enabled systems (Gust et al. 2017). Through our research, we aim to depict how insights can be generated and practical value can be derived by means of established analytics methods. Ultimately, the approach has the potential pave the way towards further process transformations.

4.3 Advanced Disruption Management

To address the shortcomings of the existing disruption process, we design, implement and evaluate an advanced disruption management system (DMS 4.0). Doing so we adopt an action design research process in collaboration with the industrial partner.

4.3.1 Action Design Research Process

The disruption management system serves many different users across different functions and hierarchy levels. Accordingly, it is fundamentally shaped by the organizational context and has to meet different expectations. To take these expectations into account during development and roll-out as well as during operation, we use ADR methodology (Sein et al. 2011). ADR's advantage is the fact that design and evaluation are not separated. Instead, ADR incorporates an iterative process of design, evaluation, and learning from intervention (Peffers, Tuunanen, and Niehaves 2018).

We have adopted this process to our problem definition and emphasize the generation of (long-term) business value (see Figure 4.3). Our key points for the generation of business value in the iterative phases of the ADR process are supplemented respectively. We have taken these key points into account both in the composition of our ADR team[18] as well as during the project implementation. For example, during the problem formulation phase, we focused on aligning with corporate goals. During the building, intervention, and evaluation phase, we increasingly ensured regular direct communication at short intervals. For more ADR process details, we refer to Oberdorf et al. (2020) and limit this paper to a critical discussion of business value enabling factors in Section 4.5.3.

[18] The ADR team consists of three scientists, two production line teams, and stakeholders close to C-level to ensure long-term commitment

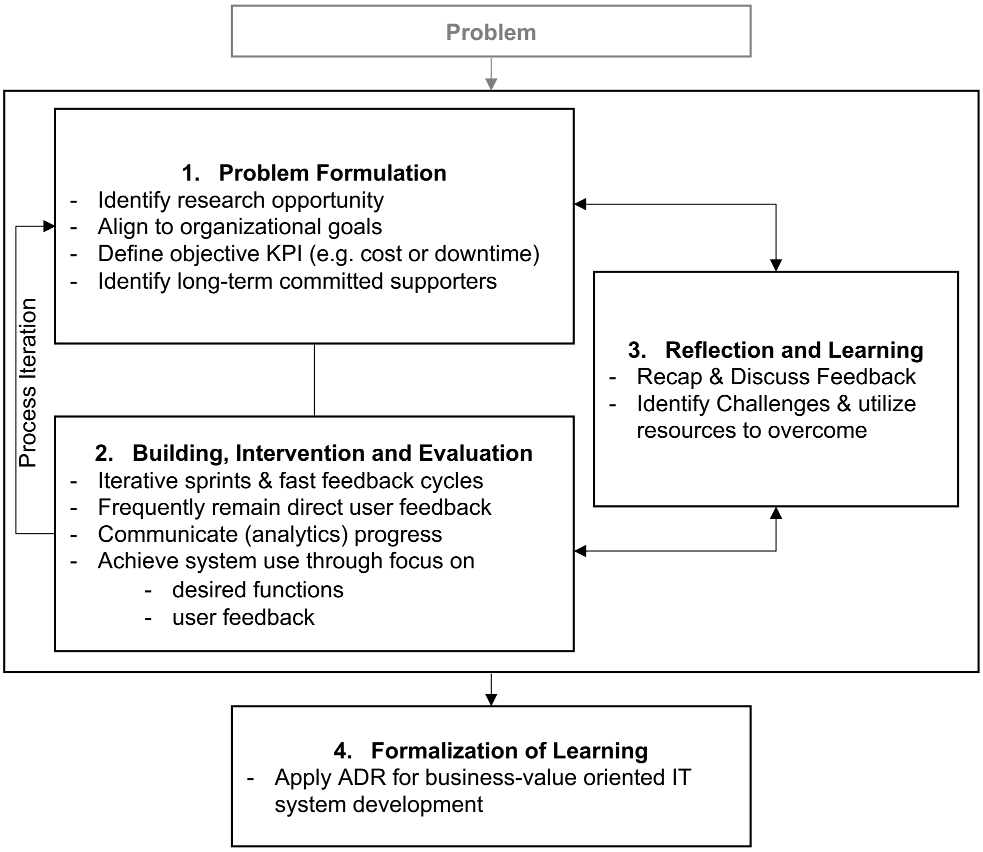

Figure 4.3: Stages of ADR with key-points for business value oriented development. Adapted from Sein et al. (2011)

4.3.2 The Disruption Management System 4.0

The outcome of the ADR process is DMS 4.0. The system consists of a combination of hardware components, database systems and an analytics back-end. On first glance, it establishes a new *disruption handling* process (see Figure 4.4).

As soon as a disruption occurs, a worker pushes the hardware button[19] to trigger the disruption process (digital andon). Immediately, the generated disruption timestamp and the hardware ID is transmitted to the DMS 4.0. By means of an integrated database connection, relevant context data from MES and PPS is automatically retrieved. In particular, we augment workplace, prod-

[19]The BeagleBone based hardware consists of a push button and lights indicating the current disruption status as well as a USB-C port for initial software flashing and power supply. The hardware connects to the DMS 4.0 via a W-LAN connection and communicates through MQTT.

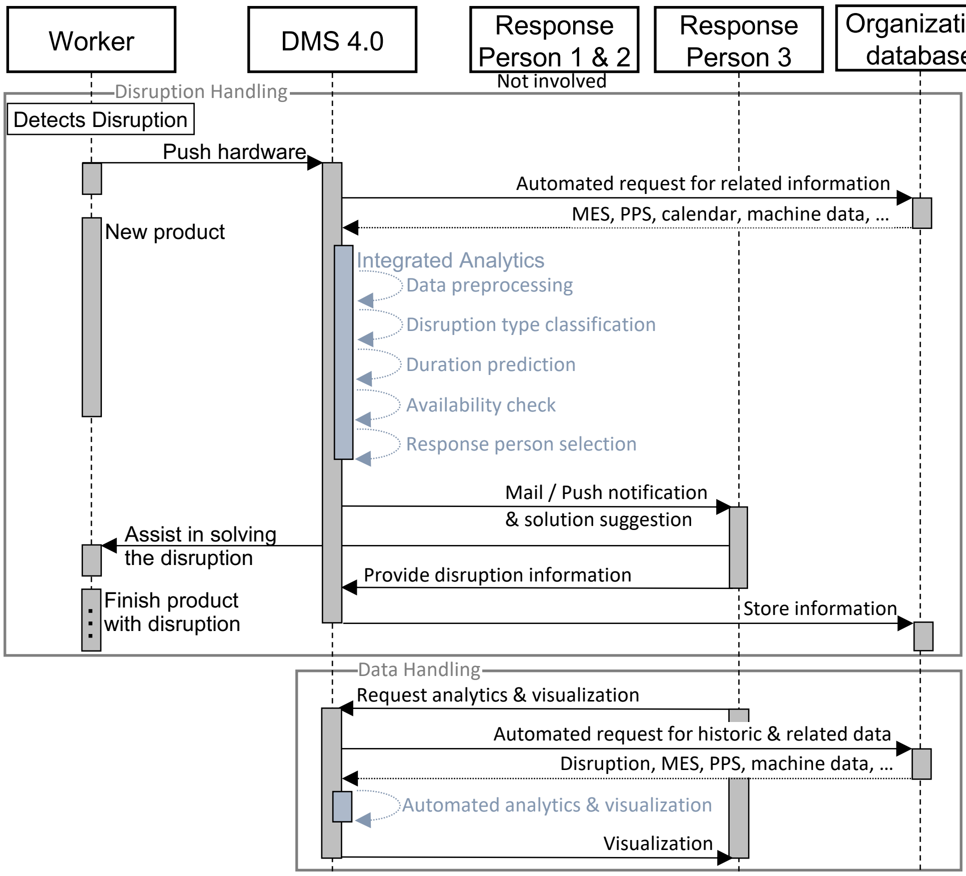

Figure 4.4: Digital disruption management process.

uct as well as production type information, which we subsequently process to identify a suitable and available responder. This automated search no longer requires the worker to set aside the current task for a prolonged time but rather allows him to continue working on the next job. The responder is automatically dispatched and notified. After resolving the disruption, the responder enters disruption related information in the DMS 4.0. For example the occurred disruption type or necessary material resources as well as an initially created timestamp and the device id are stored. In addition, the DMS 4.0 automatically creates and stores a finish timestamp, which can further be utilized for duration evaluation.

Regarding analysis tasks, the new disruption management system incorporates an integrated *data handling* process. On request, the responder is provided with a solution suggestion or visualizations with related analyses to assist root cause analysis. To this end, the system interacts with the organizational databases and obtains both current and historical disruption, MES, and production data[20]. The communication is implemented based on Python and entails the management of database addresses, tables and credentials. The various system functions can be accessed via REST-API endpoints and a web interface. Analyses are also offered with respective insights, which simplifies the access to the information significantly.

The system facilitates efficient disruption processing and facilitates business value across all three categories put forward by Mooney, Gurbaxani, and Kraemer (1996). Table 4.1 provides an overview of these benefits and assesses the corresponding realization time following Scheepers and Scheepers (2008).

Table 4.1: Business value opportunities facilitated by the DMS 4.0

Business Value	**Horizon**	**Benefits**
Automational	short	Fewer interruptions
	short	Automatic messaging
Informational	short	Continuous disruption tracking
	medium	Analytics insights about disruptions
Transformational	medium	Responder scheduling
	long	Process improvement

The automational and most of the informational business value emerge initially from the roll-out of the system due to process automation, simplification, and end-to-end data collection. In contrast, transformational business value emerges in the medium and long-term due to process improvements. Such improvements are triggered by insights from disruptions and related processes allowing an automated data-driven decision support for the responders.

4.4 Integrated Analytics

To tap into the system's transformational benefits, an analytics foundation for responder scheduling is required. We follow Wuest et al. (2016) by relying on supervised machine learning algorithms to deploy the proposed data-driven

[20] They consist of various types such as MS SQL, MongoDB or MySQL.

decision support system. Instead of just collecting data and providing minimal information, we expand the system to include integrated analytics. In particular, we train classification models to predict the disruption type which caused a given disruption. Additionally, we train regression models to predict the time required to resolve this disruption.

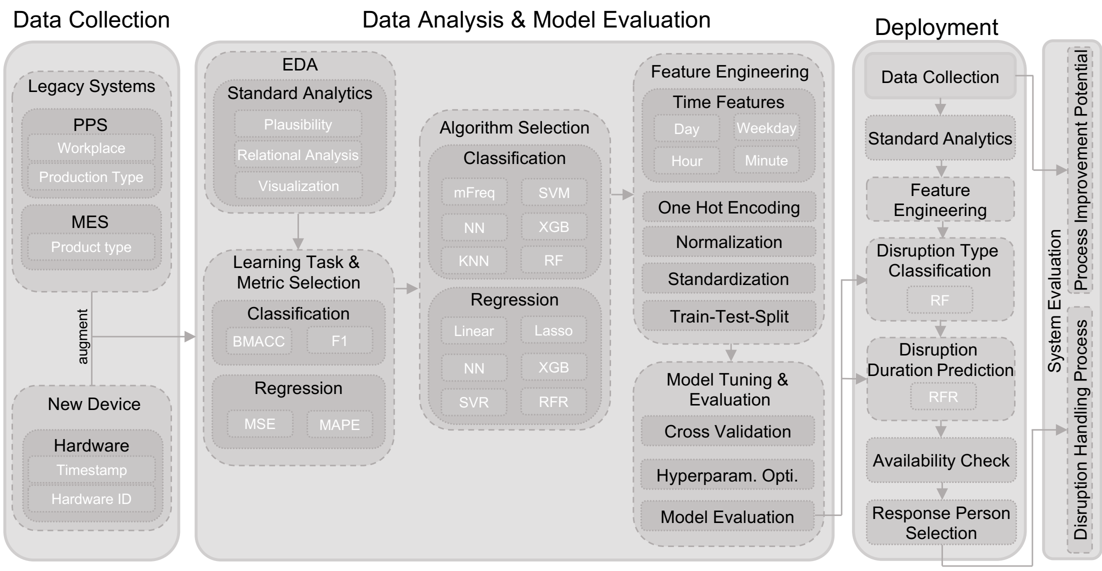

Figure 4.5: Integrated Analytics Workflow for machine learning model evaluation (adapted from Flath and Stein 2018).

We adapt the Flath and Stein's 2018 data science toolbox. In particular, we include additional steps regarding deployment and system evaluation (Figure 4.5). These extensions are crucial for transferring theoretical results into practice. While the original workflow focuses on data processing and model training, the deployment extension depicts how the extended workflow's components can be combined for practice. By means of the new disruption handling process (Figure 4.4) the integrated analytics functionalities correspond to the components of the deployment[21]. As soon as a disruption occurs, the DMS 4.0 prepares data and selects a responder based on the trained models and methods, as described next.

[21] For the system figure we combine data collection, standard analytics, and feature engineering by their function, data preprocessing

4.4.1 Data Analysis and Model Evaluation

During data collection, we first augment the data collected by the digital andon (timestamp and lot ID) with organizational process data to identify workplace, process and product information. Furthermore, we include the labels for the underlying disruption classes (out of 32 possible) and the respective durations. Subsequently, we perform an exploratory data analysis (EDA) to develop a better understanding of the available data and check for plausibility. As shown in Figure 4.6, we find that there is a high class imbalance across the different disruption types.

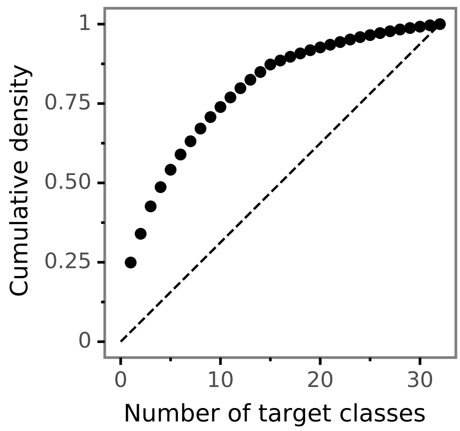

Figure 4.6: Imbalance of target classes

Learning Task and Metric Selection

Based on the data properties, suitable machine learning metrics have to be selected for the tasks at hand. In our problem at hand we account for the high class imbalance and use the balanced multi-class accuracy (BMACC) (Brodersen et al. 2010) and F1-score for the classification task. In contrast to the normal accuracy score, both metrics allow sample weighting to cope with the class imbalances in the data. We use the BMACC for model training and selection as it exhibits robust behavior against imbalances while focusing on the detection of as many errors as possible.

For the regression task we choose the root mean squared error (RMSE) and mean absolute percentage error (MAPE). The training is basically performed on RMSE due to its property to overweight higher deviations. Discussions with company stakeholders revealed that small deviations are acceptable as they can be compensated by small time buffers. In contrast, large deviations, mean that upcoming deadlines cannot be met or work has to be interrupted. Accordingly, higher deviations have to be weighted more heavily rendering the RMSE a well suited metric. For the final comparison of regression algorithm performance, we also report MAPE, as the percentage scale is more intuitive for representation.

Feature Engineering

Having essential suitable metrics for the two ML tasks, we have to perform feature engineering to transform the raw data into valuable features. The timestamp is converted into further features, such as day, weekday, hour and minute separated, to incorporate temporal similarities between disruption events. Further, we have information about the workplace, where a disruption was detected and DMS hardware pushed. Subsequently, we obtain additional features, on products produced at the certain workplace, from the organizational databases. Despite of production quantities, we remain mainly categorical features.

To this end, we compare different methods to encode categorical variables (such as workplace or product category). Performing a detailed comparison of one-hot encoding and ordinal encoding shows that one-hot encodings yield the best results in our setting. Additionally, we normalize numerical variables to ensure stable model training therefore stable results.

Algorithm Selection

We consider white box models (linear regression, lasso regression) as well as black box algorithms (random forest, support vector machines) for the two different learning tasks. While white box models yield highly interpretable results, black box models are able to also capture non-linear relationships in the data. Our evaluation shows that highly connected and interrelated production processes (that are captured by a large number of possible feature combinations) lead to non-linear relationships in the data. Given our metrics and features, we tune the models and evaluate their performance using historic training data. We follow the workflow proposed by Brownlee (2018) to perform hyperparameter tuning for each model. Based on a grid search, common algorithm-specific parameter search spaces are processed. After training a model with the parameters of the current tuning run, the metrics for the test data are calculated and stored. We repeat this process for each model with 15 different random initializations to ensure robustness. Finally, we evaluate the results of the hyperparameter-tuning and select the best-performing models and parameters for the classification as well as the regression task.

Model tuning and Evaluation

A particular focus of our work is to ensure that the selected models generalize well and perform good across the different production lines. This procedure is motivated by our observation that some configurations perform well on some lines and poorly on others. To achieve robust and generalizable results, the minimum samples per leaf are specified to avoid overfitting. Satisfactory results were obtained with a minimum sample size of three samples (without limiting the decision tree depth).

4.4.2 Disruption Type Classification

To identify a suitable responder (necessary qualifications, permission, availability) we need to know the type of the underlying disruption. We address this classification problem by training five different machine learning models (support vector machine, k-nearest neighbors, neural network classifier, XGBosst, and random forest). Following standard protocols we perform a train-test split retaining the last month of data for model testing/evaluation and the remaining data for model training. We use one hot encoding to encode categorical variables and generate additional features such as hours, day and weekday of the disruption based on the timestamp in both data sets. We compare the different models to a naïve benchmark always predicting the most frequent type of disruption.

As summarized in Table 4.2, all evaluated machine learning models outperform the naïve benchmark in terms of BMACC as well as the F1-score.[22] As the

[22]Note that the resulting multi-class accuracies relate to the 32 class classification problem. Accordingly the 78 % random forest accuracy is a sufficient result, allowing more reliable responder selection.

RF algorithm significantly outperforms the other algorithms it was selected for predicting the disruption type of a new disruption. The map the predicted error type to a list of response persons and their respective skill sets. Filtering the list based on the required skills results in a first shortlist of potential responders.

Table 4.2: Comparison of algorithms for disruption type classification.

Algorithm	**BMACC** (%)	**F1**
Most Frequent Baseline	3.1	0.12
Logistic Regression	6.7	0.21
k-Nearest Neighbor	21.2	0.41
Multi Layer Perceptron	14.7	0.33
XG Boost	59.7	0.56
Support Vector Machine	42.7	0.70
Random Forest	**78.0**	**0.83**

4.4.3 Disruption Duration Prediction

To ensure efficient and feasible responder schedules, we need reliable estimates of the disruption duration. This disruption duration prediction is formalized as a regression machine learning task. Focusing on the mean absolute percentage error (MAPE), we compare a standard linear regression model, a lasso-regularized linear regression model, a support vector regression, and a random forest regression model. We perform hyperparameter tuning on all models and find that the random forest performs best (Table 4.3). To account

Table 4.3: Comparison of algorithms for disruption duration regression.

Algorithm	**RMSE** ($1e^{-3}$)	**MAPE** (%)
Linear Regression	3.60	40.8
Lasso Regression	3.97	42.1
Multi Layer Perceptron	3.78	41.0
XG Boost	3.42	39.3
Support Vector Regression	1.30	24.4
Random Forest Regression	0.78	18.6

for the prediction MAPE of approximately 18%, we add small buffers to the estimated disruption handling duration to prevent responder overbooking.

4.4.4 Responder Availability Check

The objective of the availability check is to avoid notifying responders that have upcoming appointments during the predicted solution time. To this end, the disruption type as well as the required duration are estimated and compared against upcoming appointments of suitable responders. If there is an overlap between the schedule of a responder and the estimated solution time, the response person is removed from the list of suitable responders. According to this logic we obtain a list of suitable and available responders. Currently, the responder is randomly selected from this list and subsequently notified.

4.4.5 System Deployment

Going from a prototype to a deployed solution, we operationalize the process described in Figure 4.4. Concerning the analytics capabilities, we automated data collection, feature engineering and implemented the random forest classifier and regressor as the default methods for classification and duration prediction. Combining these functionalities with availability validation using responder schedules we instantiated the disruption management system. The DMS 4.0 was then rolled out as a pilot project within a gear production unit. This unit consists of manual assembly, machine-assisted part manufacturing workplaces as well as logistics with about 30 shop-floor workers in total. The workers are split across two production lines with seven workplaces each.

Establishing novel data-driven processes in traditional manufacturing environments is challenging as reservations and unfamiliarity with the system have to be overcome (Yiu, Yeung, and Cheng 2020). To overcome these difficulties we follow the framework proposed by Almeida Marodin and Saurin (2015). As part of the ADR process, we organized workshops with the system deploy team as well as all involved employees to gather concerns about the system. For example, one of the concerns was that the system or process for reporting disruptions was too complicated. As a result, we deploy to the Andon-inspired hardware design with only a single button for disruption reporting. This simple design eliminated further concerns regarding the integration of the system into the workflows and workstations.

We received positive feedback regarding the pilot system as well as requests from other areas. At the same time, there were some objections regarding the predictive functions. One primary concern was the handling of mis-classified disruptions. This is still an open issue that we will consider in more detail in future research through extended analyses and surveys. Nevertheless, we address this point as we schedule regular meetings to collect qualitative feedback from the participants. The feedback is evaluated and implemented establishing a continuous improvement process. One of the measures was the extension of a back-up option. If no responder remains on the selection list or if the certainty of the algorithm is too low[23], a predefined responder is notified. These extensions have satisfied all participants and we were able to put the system into operation as well as evaluate its functionality.

4.5 Evaluation

To evaluate the system, we analyze the collected disruption data of the first eight months after roll-out. We evaluate the disruption handling process as well as the process improvement potential, where we highlight the special importance of additional data sources.

4.5.1 Disruption Handling Process

For the evaluation of the disruption handling process, we analyze the development of the time required to handle disruptions. Due to automational benefits, the average disruption duration has been reduced by about 15 % since the initial system roll-out (Figure 4.8). In contrast, the informational and transformational value of the new system materializes in the medium and long term as workers and responders get used to the new system. Simultaneously, the disruption type classification improves as more training data becomes available over time (Figure 4.7). While the BMACC of the RF model is only 10% af-

[23]We compare the prediction probability with a threshold.

ter 1 month, it reaches 78% after 8 months. The increased accuracy implies a downtime reduction, as an appropriate responder is going to be notified more frequently. In combination, both effects—informational and transformational—enable us to further reduce the average disruption handling duration by 70 % compared to the initial disruption duration over a period of 8 months after the roll-out. The reduced downtime directly affects the organization by means of business value as direct and indirect downtime costs are significantly reduced.

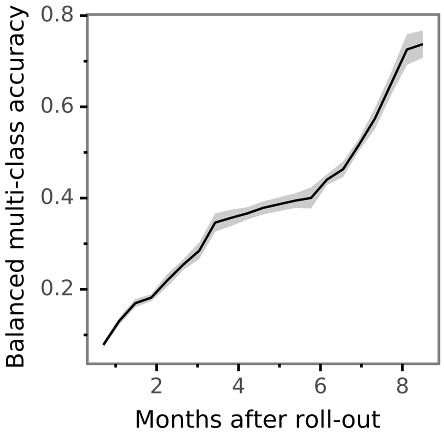

Figure 4.7: Random forest multi-class accuracy depending on training size with standard error.

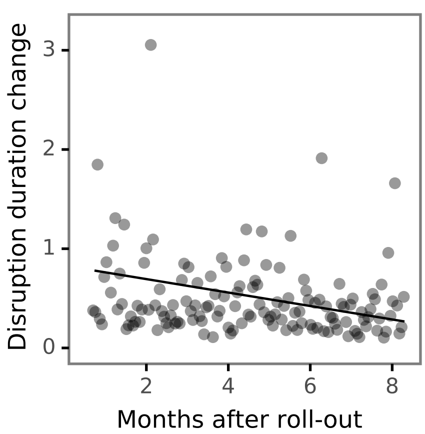

Figure 4.8: Disruption duration trend (line) analysis—related to the previous duration.

As part of the ADR process and based on our objective to further improve the system, we collected feedback on the system and concerns in follow-up workshops with the ADR team. The feedback on the system and the implementation was consistently positive. Our partners especially highlighted the easy handling of the system and the reduced downtime. In addition, they also called out the business value of simplified and automated data handling. Through the DMS 4.0 web-interface, (as part of the data handling process) disruption information and analysis can be obtained easily.[24] As a result, the analyses are used more frequently for meetings and discussions. The resulting dissemination of the information has an additional added value. By discussing the data, correlations are critically questioned, and e.g., outliers (over 150 % of previous duration) are identified, at which point an improvement process was initiated.

[24]Prior to the roll-out of the new system an expert, e.g., a data analyst, had to provide analysis results on demand.

To illustrate such a process, we evaluate the process improvement potential and highlight the special importance of additional data sources.

4.5.2 Process Improvement Potential

As mentioned in the follow-up workshop, the data handling process creates additional value for the organization through advanced process insights. These insights should result in long-term business value through process improvement. Indeed, the available data suggests that such improvement processes have been triggered by employees. Figure 4.9 shows the number of disruptions in relation to the production quantity. Figure 4.10 reports the number of disruptions per output (production) unit as a function of the time since roll-out.

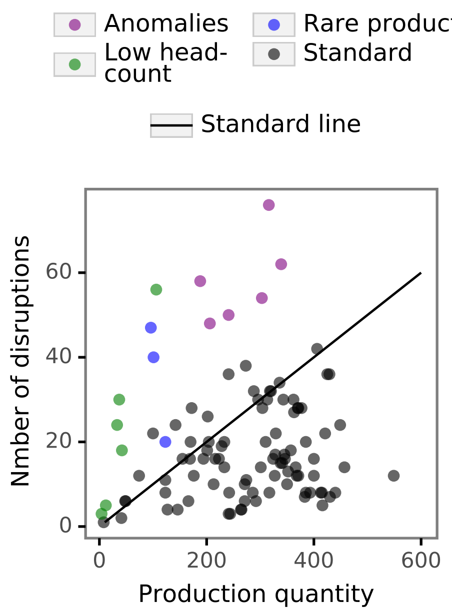

Figure 4.9: Number of disruptions related to production amount with highlighted outliers.

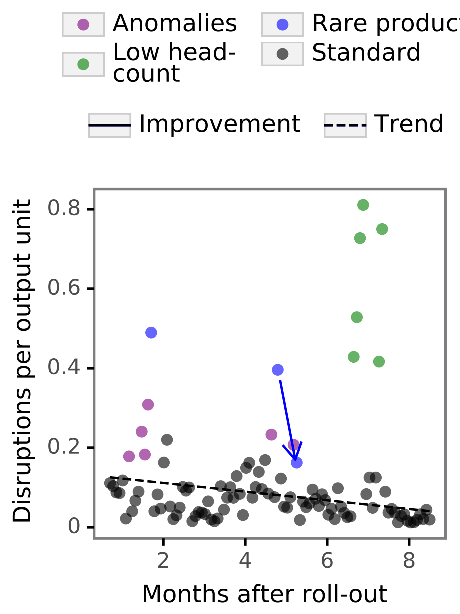

Figure 4.10: Disruptions per output unit with highlighted outliers and process improvement result.

Focusing on Figure 4.9 we find that the number of disruptions tends to increase in production quantity, but usually remains below 10 % (diagonal line). While there are some exceptions to this pattern, we could better understand the context by incorporating product type (standard vs. rare products) and worker availability. These variables explain some of the dispersion present in the number of disruptions. However, some anomalies have remained unexplained so far but may of course be nonsystematic.

By taking a dynamic perspective, we want to shed light on organizational learning (Figure 4.10). Note that the prevalence of worker shortages in months six to eight corresponds to the exogenous shock created by lockdown measures during the COVID-19 pandemic. As this is a unique situation we decided to omit these values in the subsequent analysis. The data suggests that the production-normalized number of disruptions has continuously been falling since the introduction of the system. In particular we can see learning effects with respect to rare products. These findings confirm the viability of continuous improvement processes and in turn generate transformational business value.

4.5.3 Discussion

For a reflection of our results, we want to refer back to Table 4.1 which highlighted the envisioned business value opportunities facilitated by the DMS 4.0. We were able to tap into all of these opportunities. Automated messaging and dispatching reduce the disruption duration which directly translates into automational business value. From an informational view, data collection and continuous monitoring as well as the improved data handling generate initial value for descriptive analytics applications such as reporting or dashboards. Pursuing data collection over a longer period leads to the availability of sufficiently large training data to create predictive analytics models which generate insights on disruption occurrences and the underlying processes. These informational insights form the base for the execution of the new responder scheduling process. This improves both worker and responder productivity and hence manifests transformational business value. Even though we do not obtain perfect predictions we can provide business value through an analytics-enabled disruption management system.

We identified some concerns of relevant stakeholders during the integration of the DMS 4.0. Especially at the beginning of the system roll-out, the predictive power is limited due to the low initial data availability leading to reduced user acceptance. However, the performance of the algorithms improves as more data becomes available. We found that communicating such details in advance as well as during the project helped to mitigate the concerns. This necessitates a profound cooperation between the development team and the system users (Barua et al. 2004; Kohli and Devaraj 2004). We can emphasize that the ADR process is well suited for implementation. In particular, the iterative building, intervention and evaluation cycles as well as the reflection and learning stage with the employees helped to increase the acceptance of the system. Acceptance is key for the successful creation of business value. A change—particularly related to processes—can only be successful if accepted and executed (Holtzblatt and Beyer 1997, 1993).

Even though we show the implantation of DMS 4.0 in a single case study, we believe that the proposed methodology generalizes well across other Industry 4.0 settings. While new operating conditions and data sources will require different machine learning models, the potentials for the mentioned business value improvements are fundamentlally generic. A methodological challenge may arise from companies collecting an increasing amount of unstructured data from sensorized production equipment. We currently analyze the use of multi-headed neural networks to integrate various unstructured and structured data sources into the analysis and will report on the results in upcoming research papers (Oberdorf et al. 2021a)

Our research contributes to literature by extending established research on technology enabled business value (Schryen 2010; Melville, Kraemer, and Gurbaxani 2004; Mooney, Gurbaxani, and Kraemer 1996; Barua, Kriebel, and Mukhopadhyay 1995). Against the backdrop of growing big data and analytics-enabled opportunities, current research enhances the existing benefits of IT systems by additional analytics-enabled benefits (Grover et al. 2018; Wang et al. 2019b; Chen, Chiang, and Storey 2012). We contribute here by pinpointing how classical systems can be enriched with analytics capabilities, to finally provide additional business value. Thereby, we extend existing disruption management approaches (Lopez-Leyva et al. 2020; Dombrowski, Richter, and Krenkel 2017) and integrate analytics capabilities.

4.6 Conclusions and Implications

Our research explores the potentials of an I4.0 enabled disruption management system in the traditional manufacturing industry. Our study sheds light on how analytics-enabled industrial applications help create business value. To explore the interplay between analytics and IT business value we relate descriptive, predictive and prescriptive analytics to automational, informational and transformational IT business value. We posit the central importance of analytics for the generation of transformational business value. As process automation and data collection are prerequisites for such a business process transformation, the corresponding business value will in a sense emerge as a byproduct.

We collaborated with WITTENSTEIN SE, which is faced with challenges such as highly customized products, large shares of manual tasks as well as a (yet) limited degree of digitalization across the manufacturing processes. To enable prescriptive analytics and in turn generate transformational business value, we structure the underlying process as machine learning tasks for disruption classification and duration prediction. Direct integration of such analytics applications is of special importance to establish an automated process, without additional interaction.

In addition to facilitating improved disruption handling, the DMS 4.0 simplifies the interaction with organizational databases as interactions with multiple systems are bundled in one system. The interaction with databases is automated and provides stakeholders with timely reports and visualizations. This in turn may prompt a more frequent use and in turn better dissemination of information. Ultimately, such systems can and should "…provide insights for many of the traditional manufacturing operational issues …" (Babiceanu and Seker 2016).

As with any data-driven analytics application, sufficient training data is a perquisite. This training data is the foundation for learning reliable prediction models (classification and regression). The models in turn are the key input for identifying and scheduling a response person in the disruption handling process. Creating transformational business value is by no means a quick win but necessitates a certain level of patience on behalf of the organization. Trust in organizational IT systems and processes is an essential success factor. Considering the experiences from our collaboration, we want to highlight the special importance of the continuous improvement process meetings which shaped the system.

While the disruption management system shows promising first results, we want to highlight some limitations and potential for further research. In the current implementation, a response person is chosen from the short-list according to a random selection. Future research should model the deployment of the available response persons in the sense of a generalized assignment problem. However, integrating uncertainty on upcoming disruptions (i.e., prediction errors) results in a stochastic optimization problem leading to computational expensive models.[25] Unstructured data such as error reports could be leveraged to further increase the quality of the predictive models. To this end, deep learning models also be considered to better handle unstructured data.

[25]The RF's 78 % accuracy will obviously result in some miss-assignments. Integrating prediction accuracy in the stochastic optimization algorithm will lead to decision that take the uncertainty into account result in more reliable responder assignments.

5 Predictive End-to-End Enterprise Process Network Monitoring

This paper is published in *Business & Information System Engineering* (Oberdorf et al. 2023).

Ever-growing data availability combined with rapid progress in analytics has laid the foundation for the emergence of business process analytics. Organizations strive to leverage predictive process analytics to obtain insights. However, current implementations are designed to deal with homogeneous data. Consequently, there is limited practical use in an organization with heterogeneous data sources. This paper proposes a method for predictive end-to-end enterprise process network monitoring leveraging multi-headed deep neural networks to overcome this limitation. A case study performed with a medium-sized German manufacturing company highlights the method's utility for organizations.

5.1 Introduction

Business processes are the backbone of organizational value creation (Dumas et al. 2018a). The progressing digitalization of business processes results in massive amounts of historical process data (van der Aalst 2016a). In parallel, analytics capabilities facilitate the use of this data (Vera-Baquero, Colomo-Palacios, and Molloy 2013; Beheshti, Benatallah, and Motahari-Nezhad 2018). Business process analytics refers to a set of approaches, methods, and tools for analyzing process data to provide process participants, decision-makers, and other stakeholders with insights into the efficiency and effectiveness of operational processes (Zur Muehlen and Shapiro 2015; Polyvyanyy et al. 2017; Benatallah et al. 2016).

Among others, business process analytics aims to reduce a decision-maker's distance to observing a business event (Zur Muehlen and Shapiro 2015). Two classes of information systems serve this purpose, which promise to assist decision-makers but have been discussed independently (Schwegmann, Matzner, and Janiesch 2013). First, business intelligence systems query historical event log data to address descriptive problems (Mehdiyev, Evermann, and Fettke 2020) or to prepare predictions of future process behavior (Schwegmann, Matzner, and Janiesch 2013). Second, monitoring systems provide real-time decision support, e.g., through predictions, based on historical event log data (Janiesch, Matzner, and Müller 2011).

As a methodological basis for predictive monitoring systems, predictive process monitoring (PPM) is gaining momentum in business process management. PPM provides a set of methods that allow predicting measures of interest based on event log data (Maggi et al. 2014). By gaining insights into the uncertain future of a process, PPM methods enable decision-makers to prevent undesirable outcomes (Marquez-Chamorro, Resinas, and Ruiz-Cortes 2017).

Recent research in PPM proposes various methods, which can be arranged into two general groups according to the prediction task (Mehdiyev, Evermann, and Fettke 2020). The first group of methods addresses *regression* tasks and refers to the prediction of continuous target variables, such as the completion time of a process instance (e.g., van der Aalst, Schonenberg, and Song 2011; Wahid et al. 2019). In contrast, the second group tackles *classification* tasks and refers to the prediction of discrete target variables, such as the next activity (e.g., Mehdiyev, Evermann, and Fettke 2017; Breuker et al. 2016), process violations (e.g., Di Francescomarino et al. 2016), or process-related outcomes (e.g., Flath and Stein 2018; Kratsch et al. 2020).

PPM typically predict measures of interest based on a single event log documenting a specific process or multiple sub-processes (Cuzzocrea et al. 2019; Senderovich, Di Francescomarino, and Maggi 2019). Oftentimes, the (process) control flow information is feature-encoded yielding one target variable per process instance or prefix (part of the process instance) (e.g., Tax et al. 2017). More sophisticated approaches append (process) context information to control flow information of a single event log to increase the explainability of input variables concerning the target variable (e.g., Yeshchenko et al. 2018; Brunk et al. 2020).

In organizations with a process-oriented design (Eversheim 2013), business processes typically flow through multiple departments. Departments are not only connected via the *organization and department layer*, but also via the *enterprise process network* layer, connecting departments, processes, and information systems (Figure 5.1).[26] More specifically, it establishes inter-department

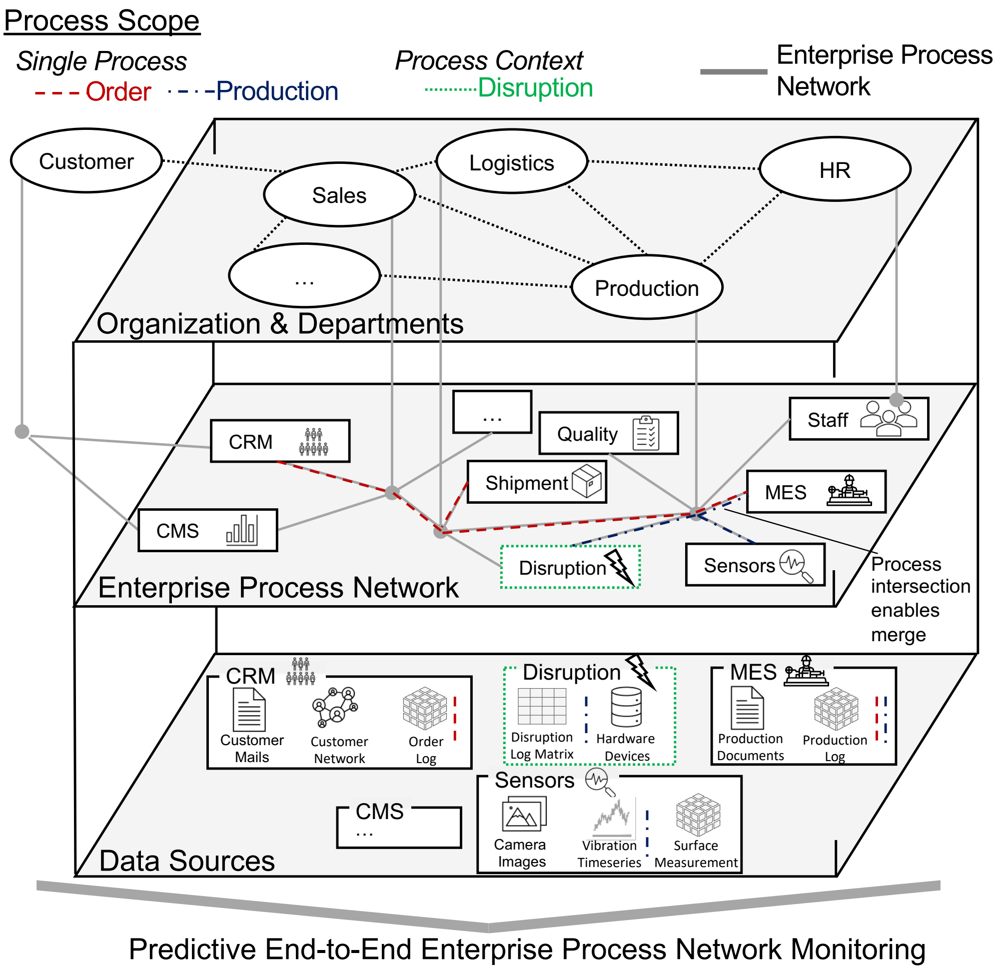

Figure 5.1: Overview of process scope in the organizational context.

and inter-process dependencies, as departments will usually be involved in

[26]We consider the enterprise process network as the intra-organizational process network, based on the control view concept of the ARIS framework (Scheer 2013) for the architecture of integrated information systems. For the representation, we adapt vom Brocke and Rosemann (2014, p. 54) "business process trends pyramid" with its distinct layers for enterprise (organization & departments), business processes, and implementation (both enterprise process network).

a multitude of processes (e.g., a disruption in the process of the production department affecting the shipment process in the logistics department or may influence the sales processes) and a process will often involve multiple departments (e.g., an order process (red) that spans sales, logistics, and production department).

Consequently, the enterprise process network extends the process scope. Processes are then manifested in information systems and interact with various *data sources* on the respective layer. In addition to the data stock of the information systems, which mainly consists of master data and operational transactions, other original data sources can be extracted. These additional data sources are often used to obtain data for analyses or predictive tasks. In general, they consist of a combination of one or multiple event logs with control flow information and additional event-log-related context information. Other data sources not directly related to the process are possible, such as process context information. For example, additional internet of things (IoT) data (e.g., temperature, humidity, vibration, sound time-series measurements, or surface roughness, and roundness protocols, or even (machine) acceptance logs) are collected in digitalized environments. Besides recording structured IoT data, unstructured data, such as images and videos can help detect quality defects in a manufacturing environment. Moreover, text documents can also add predictive value. Given this data scope definition, Figure 5.1 distinguishes data sources such as an order event log (red-dashed), a production event log (blue-dash-dotted), both with control flow and process-related context information, as well as disruption context information (green-dotted).[27] A product's dispatch time prediction may benefit from additional information from the disruption and logistics process. Such an interplay between the different processes of an enterprise process network may increase the predictive power, as more data potentially results in additional relevant features, which can contribute to the description of a prediction target variable. However, existing PPM approaches do not consider this (Borkowski et al. 2019), limiting their practical use as the seamless combination of heterogeneous data is not applicable.

[27] Note: The combination of multiple data sources requires a common denominator, such as a process intersection or timestamp, to synchronize and merge, e.g., individual process event logs and context information.

We address this limitation with a method for predictive end-to-end enterprise process network monitoring. The main contribution of our research is threefold:

1. This paper presents a method for predictive enterprise process network monitoring in the business process management (BPM) domain. The method establishes an end-to-end perspective on predictive process network monitoring in an organizational context. In doing so, it facilitates the combination of heterogeneous data sources for predictive tasks and guides the problem specification as well as the design and application of a multi-headed neural network (MH-NN) model.
2. This paper proposes a multi-headed deep neural networks (DNN) model that integrates multiple data sources of an enterprise process network, such as the color-highlighted process logs or context information in Figure 5.1. With this deep learning (DL) architecture, the heterogeneous data are processed in dedicated neural network (NN) input heads and concatenated for prediction, based on cross-department information.
3. The results from a case study conducted with a medium-sized German manufacturing company shed light on the practical relevance. We evaluate our method against traditional machine learning (ML) and state-of-the-art DL approaches in terms of predictive power and runtime performance based on real-world data. While the DL model constructed with our method exhibits somewhat higher computational costs, its predictive power is significantly higher than the considered baselines.

5.2 Background and Related Work

We first review recent advances in PPM with a special focus on prediction methods. In doing so, we highlight the research gap and position our method for predictive end-to-end enterprise process network monitoring.

5.2.1 Prediction Methods in Predictive Process Monitoring

Process Mining (PM) is an established process analysis method in BPM that involves data-driven (process model) discovery, conformance checking, and enhancement of processes (van der Aalst et al. 2011a). PM's general idea is to gain process transparency from event log data. It is thus an approach for process analytics, particularly focusing on ex-post process diagnostics. With the advent of predictive analytics, new potentials of gaining insights from event log data have been unlocked (Breuker et al. 2016). Using these methods, PPM has emerged as a new subfield of PM (Marquez-Chamorro, Resinas, and Ruiz-Cortes 2017). PPM provides a set of techniques to predict the properties of operational processes, such as future process behavior (e.g., next process activities) or process outcomes (e.g., a process performance indicator). A branch of early PPM approaches augment discovered process models with predictive capabilities but require certain model structures to support prediction tasks. In doing that, the process model is transformed into a predictive model. For example, van der Aalst, Schonenberg, and Song (2011) introduce a technique that uses an annotated transition system with the capability to predict process completion time based on historical event log data. Another example is Rogge-Solti, van der Aalst, and Weske (2013), who mine a stochastic Petri net with arbitrary delay distribution from event log data. These approaches can be described as *process-aware* because they utilize "(…) an explicit representation of the process model to make predictions" (Marquez-Chamorro, Resinas, and Ruiz-Cortes 2017, p. 4).

However, real-world processes are usually more complex than the discovered process models (van der Aalst 2011). The process-model-dependence limits the predictive power (Senderovich, Di Francescomarino, and Maggi 2019). To overcome this restriction, another, more recent branch of PPM approaches proposes to encode sequences of process steps as features vectors for the straightforward use of ML models. This transforms the event log's sequential process information into a predictive model without discovering a process model. Leveraging the generalization power of ML models, sequence-encoding approaches often outperform predictive models built on top of discovered process models (Senderovich et al. 2017).

The multi-layer perceptron[28] is another NN architecture that has been leveraged for PPM. The MLP does not explicitly model temporality, and therefore, received sequential data has a two-dimensional data structure. For example, Theis and Darabi (2019) used MLPs to predict the next activities. DNNs have been applied to PPM, due to the conceptual similarities between next event prediction and natural language processing tasks (Evermann, Rehse, and Fettke 2016). DNNs can outperform statistical (e.g., Verenich et al. 2019) and traditional ML approaches (e.g., Kratsch et al. 2020; Mehdiyev, Evermann, and Fettke 2020; Evermann, Rehse, and Fettke 2016). DNNs perform multirepresentation learning, which "(…) focuses on extracting the multiple representations from the single view of data" (Zhu et al. 2019, p. 3) and are good at unveiling intricate structures in data (LeCun, Bengio, and Hinton 2015). A popular subclass of DNNs are recurrent neural network (RNN) approaches (Rama-Maneiro, Vidal, and Lama 2020a), including long short-term memory (LSTM) and gated recurrent unit (GRU) neural networks, providing the capability to capture temporal dependencies within sequences (Rumelhart, Hinton, and Williams 1985). Another DNN architecture, which allows the processing of temporal patterns, is the convolutional neural network (CNN) (Zhao et al. 2017). This DNN type requires grid-like input, such as image data. To leverage the potential of CNN for PPM, a preprocessing of sequences from temporal to spatial structure is needed. For example, Pasquadibisceglie et al. (2019) show the validity of such a sequence preprocessing for predicting the next process activity using the *helpdesk event log* and *BPI challenge 2012* data. Graph neural networks (GNN) are recently used in PPM because the process control flow follows a graph structure (e.g., Stierle et al. 2021) and can directly be processed through GNNs. Beyond the four general architectural types MLPs, RNNs, CNNs, and GNNs, extensions (e.g., transformer networks with dense layers like MLPs; Moon, Park, and Jeong 2021) or combinations (e.g., long-term recurrent convolutional networks; Park and Song 2020) were proposed for PPM.

[28] Note: As a MLP is a mathematical function composing of many simpler functions, it can be considered as a feed-forward DNN (Goodfellow, Bengio, and Courville 2016).

5.2.2 Data Scope vs. Prediction Methods in Predictive Process Monitoring

Statistical approaches in PPM (e.g., van der Aalst, Schonenberg, and Song 2011; Rogge-Solti, van der Aalst, and Weske 2013) start with the control flow information of event log data (Figure 5.2). This type of information is key for process predictions, as the control flow of processes describes their structure.

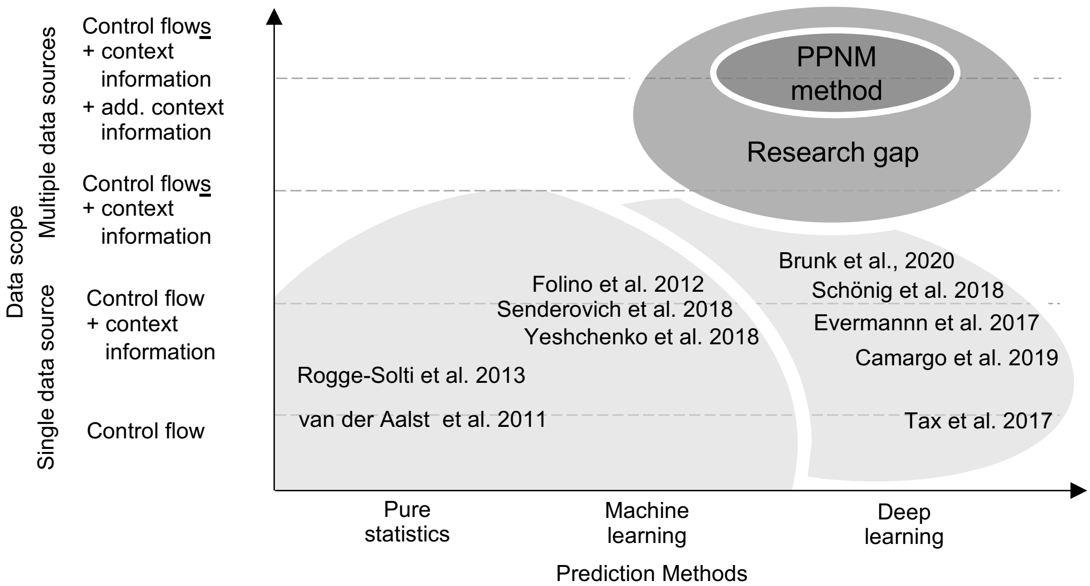

Figure 5.2: Classification of exemplary PPM techniques by *data scope* and *prediction method* with highlighted research gap and our proposed method.

By using ML, the scope of data is extended and PPM techniques can encode further event log information in feature vectors (e.g., Folino, Guarascio, and Pontieri 2012). This additional information is called *process context information*. It characterizes the environment in which the process is performed (Cunha Mattos et al. 2014; Rosemann, Recker, and Flender 2008), and represents, for example, information about the resource that performs an activity.

In recent years, PPM research has suggested DL architectures that integrate context information to improve prediction results (Rama-Maneiro, Vidal, and Lama 2020a). However, such architectures ignore that in practice control flow and context information stem from multiple data sources. Current PPM approaches receive single event logs as input and do not leverage information from multiple data sources. Thereby, an event log can also contain several subprocesses, such as in the event log shared at the BPI Challenge 2012.[29]

Currently, there are no PPM techniques using multiple data sources to perform end-to-end enterprise process network predictions. However, new time series forecasting techniques (e.g., Canizo et al. 2019; Mo et al. 2020; Wan et al. 2019) offer a promising way to realize such predictions through multi-headed NN. These networks process data from each input head (e.g., from a machine sensor) individually and merge the heads' outcomes subsequently. Motivated by this idea, we set out to adapt this method for end-to-end enterprise process networks.

5.3 Method Engineering Process

We develop our end-to-end enterprise process network monitoring (PPNM) method based on the method engineering research framework for information systems development methods and tools proposed by Brinkkemper (1996). Methods describe systematic procedures "to perform a systems development project, based on a specific way of thinking, consisting of directions and rules, structured in a systematic way in development activities" (Brinkkemper 1996). The method engineering process consists of three phases (Gupta and Prakash 2001): requirements engineering, method design, and method implementation. First, we define requirements for the construction of the PPNM method (Table 5.1). Second, we present the design, evaluation, and implementation of the PPNM method (Section 5.4) and describe the method's phases in detail in the context of a case study of a medium-sized German manufacturing company. Finally, we discuss the PPNM method critically and provide implications (Section 5.4.4).

We arranged our related work insights for the requirements engineering and set up a workshop with organizational partners to collect and discuss key requirements for the PPNM method. In this workshop, we developed *business* and *performance requirements*. From the business perspective, it is of particular importance to enable end-to-end analyses (R1). Thus, the engineered method requires a problem definition phase on the organizational layer, and

[29] https://www.win.tue.nl/bpi/doku.php?id=2012:challenge&redirect=1id=2012/challenge

the method's results must also apply to the organizational layer. Moreover, requirement R2 requires the integration of diverse organizational data sources. Additionally, the method must scale well so that new data sources can seamlessly be included to increase the predictive power. The prediction target should be designed (R3), such that the engineered method addresses classification as well as regression tasks.

From the business requirements, we can derive some performance requirements. The predictive model resulting from the method should outperform the predictive power of traditional ML and DL approaches (R4). The combination of data sources must therefore add predictive value. In addition to greater predictive power, the model must also provide predictions in a sufficiently quick manner (R5).

5.4 Predictive End-To-End Enterprise Process Network Monitoring

We propose PPNM, a novel five-phase method for predictive end-to-end enterprise process network monitoring (Figure 5.3). First, the underlying problem is specified. This includes (business) problem identification, (business) pro-

Business Requirements	
R1	End-to-end approach for predictive enterprise process network monitoring
R2	Multiple data sources can be processed through specialized input heads individually and the predictive power combined
R3	The prediction target is variable, such as either classification and regressions tasks can be performed
Performance Requirements	
R4	Outperform traditional ML and state-of-the-art DL approaches regarding predictive power
R5	Runtime performance suited for timely predictions in real-world use case

Table 5.1: Overview of business and performance requirements for the PPNM method.

cess understanding, and predictive task specification. Second, the method prescribes to acquire and prepare the input data for the MH-NN model. Third, the MH-NN model is designed and subsequently evaluated in the fourth phase. Lastly, PPNM describes aspects of the model application.

5.4.1 Problem Specification

The first phase specifies the problem by adapting the approach of Benscoter (2012), beginning with the *problem identification* at the business department or enterprise process network layer. Their approach to "identify and analyze problems in your organization" (Benscoter 2012) has a particular focus on identifying a situation's impact on processes, workers, as well as problem-relevant metrics. Subsequently, the establishment of an *understanding of the interdependent processes* and data sources is crucial. Within an organization's layers, all relevant processes and data sources, which can add value to the predictive analysis task, should be identified. Subsequently, their dependencies should be understood to identify common denominators for synchronizing heterogeneous data sources and how they relate to the organizational problem or situation. Based on this process and data understanding, the method prescribes to *define the organizational objective and the type of predictive task* (regression or classification).[30]

5.4.2 Data Acquisition and Preparation

Given the identified relevant processes and data sources, we then acquire and prepare the input data for the desired MH-NN. *Data acquisition* relates to activities seeking to obtain the heterogeneous data. This data is analyzed to gain insights about the data source and subsequently prepare it for the MH-NN. The network processes each data source individually, without the need for prior aggregation and combination. In doing so, it leverages standard preparation techniques (Han, Pei, and Kamber 2011) for the individual data sources. In addition, it follows the general stream of DL methods (LeCun, Bengio, and

[30] A regression relates to estimating a numerical output, such as the forecast of financial, sales, downtime information, or organizational key performance indicators. In contrast, a classification's output incorporates the estimation of categorical types, such as if an event may happen (binary) or if an event has a particular type (multi-class).

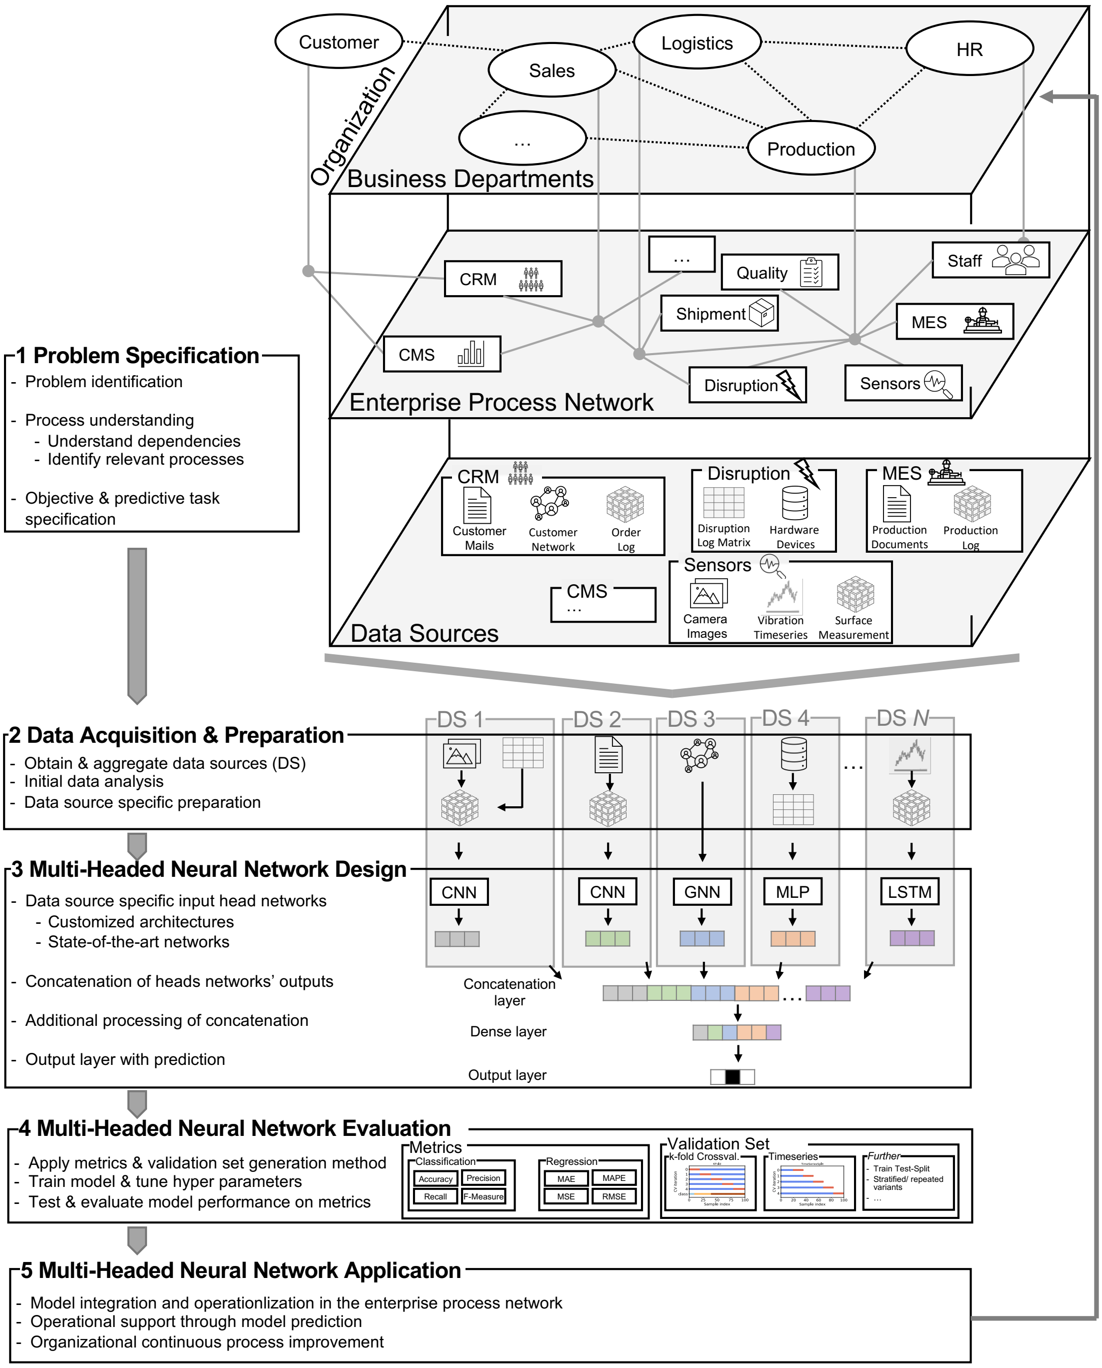

Figure 5.3: Five-phase method for predictive end-to-end enterprise process network monitoring.

Hinton 2015), which limit the extensive preparation by focusing on generalized DL architectures for feature extraction.

As part of preparation, PPM requires appropriately encoding events and sequences. *Events* can be encoded based on the attributes' type. *Sequences* of events can be encoded as feature-outcome pairs (van Dongen, Crooy, and van der Aalst 2008), n-grams of sub-sequences (Mehdiyev, Evermann, and Fettke 2020), feature vectors derived from Petri nets (Theis and Darabi 2019), or weighted adjacency matrices (Oberdorf et al. 2021a).

5.4.3 Multi-Headed Neural Network Design

We now design the multi-headed NN. Thereby we follow recent work on PPM methods, which move from explicit process models and traditional ML approaches to NN-based approaches (Mehdiyev, Evermann, and Fettke 2020). Yet, for some scenarios, the sequential structure of these NNs is not sufficiently flexible, such as, if data from different sources with different dimensions are required to explain the output variable. Following Chollet (2018, p. 301), the proposed architecture for these cases is a multi-head NN. Architectures with multiple heads use *independent single-channel input heads* to process each input individually. With this approach, each *data source can be processed*, according to its data type and structure. Head outputs are then *concatenated* and *further processed* to ultimately yield a *prediction in the output layer*.

For the design of the multi-headed NN, the method facilitates the use of a multitude of architectures (Figure 5.4). In general, it distinguishes customized and state-of-the-art architectures.

For *customized architectures*, a combination of NN layers can be selected (compare 5.2.1). Following Goodfellow, Bengio, and Courville (2016), combining various layers in a task-specific manner enables the implicit extraction of valuable features. To this end, distinct properties of architectures can be leveraged, such as the particular suitability of LSTM layers to process time-series or CNN layers for matrix data. These properties can even be combined to process time-series, such as a combination of LSTM and CNN layers (Brownlee 2017).

[31] The network symbols are from https://colah.github.io/posts/2015-08-Understanding-LSTMs/, https://theaisummer.com/Graph_Neural_Networks/, https://en.wikipedia.org/wiki/Convolutional_neural_network, and https://www.quora.com/What-is-max-pooling-in-convolutional-neural-networks.

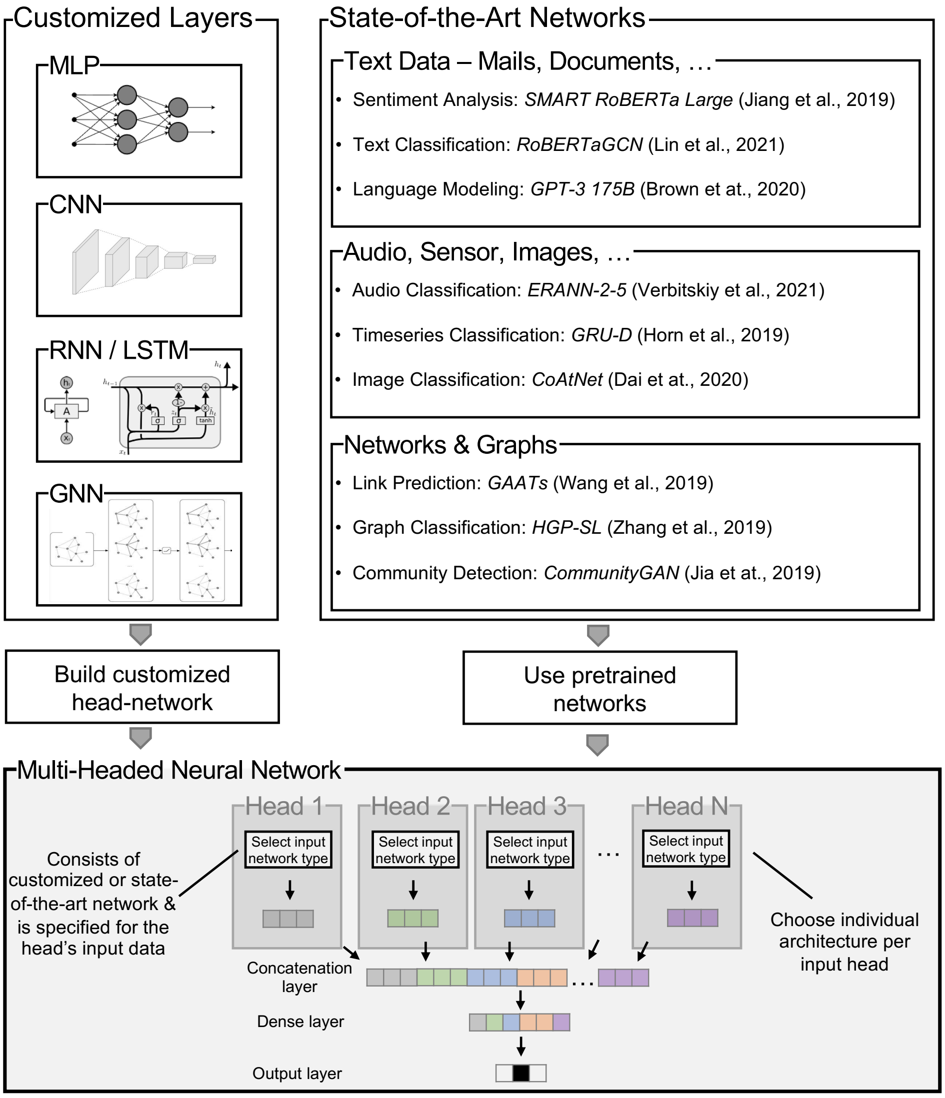

Figure 5.4: Overview of potential NN layers and state-of-the-art networks (Papers with Code 2021) for the multi-headed NN's input heads.[31]

In addition to the customized architectures, the method taps into recent advances in the DL domain by incorporating established architectures. There are *state-of-the-art* architectures for the various domains such as image, text, or signal processing. As the numbers of available architectures are constantly changing, we suggest checking for currently available state-of-the-art networks during a model's design phase to build on recent research advances.[32] Figure 5.4 provides an overview of currently established state-of-the-art methods for various tasks. Depending on the data type, we show current DL solutions for problems, such as sentiment analysis (Jiang et al. 2019), language modeling (Brown et al. 2020), text, time-series, audio, image, or graph classification (Lin et al. 2021; Horn et al. 2020; Verbitskiy and Vyshegorodtsev 2021; Dai et al. 2021; Zhang et al. 2019), as well as link prediction (Wang et al. 2019a), or community detection (Jia et al. 2019) in networks.

The common denominator for such models is that they consist of complex DL architectures with many hidden layers and trainable parameters. Because the training of such models is computationally demanding, they are usually provided with pretrained weights, which can then be leveraged for the prediction task at hand or even fine-tuned based on the task's specific data.

5.4.4 Multi-Headed Neural Network Evaluation

The method next requires to consider aspects of model evaluation. For this purpose, we follow Brownlee (2020)'s approach, including the generation of a validation set and the use of performance metrics to assess a model's performance. The evaluation of the resulting model is crucial for the selection of a proper configuration. It reveals whether the model is suitable to estimate the desired target variables. To this end, test and validation sets are artificially generated through validation methods. In particular, in the field of PPM, selecting an appropriate validation set method is challenging. A time-series method should be chosen if the data has time-dependent features such as time-series or process logs and graphs with relevant time information. Commonly, it can be chosen between three *validation set generation* methods (Figure 5.3). In addition to the validation set generation, it is common to keep a holdout set containing exclusive data for a final model evaluation.

The most common method used is a straightforward strategy, referred to as a train-test split procedure (James et al. 2017, p.176-178). An alternative evaluation procedure is k-fold cross-validation for estimating the prediction error

[32] Besides recent publications, more practical related sources for recent advances are `https://paperswithcode.com/`, `https://github.com/sebastianruder/NLP-progress`, or `https://github.com/rwightman/pytorch-image-models`.

(James et al. 2017, p.181-186). It splits the data set into k folds, uses $k-1$ of folds for training and the other fold for validation.

In some settings, regular k-fold cross-validation is not directly applicable. This is the case for time-series data, where observations are samples with fixed time intervals. The constraint is the temporal components inherent in the problem. Here, a time-series split is an appropriate method, where in the k^{th} split, the first k folds are used as a train set, and the $(k+1)^{th}$ fold is used as a test set. Time-series splits have the drawback that there is overlap between the training and testing data. This limitation is solved by forward testing techniques where the model is automatically retrained at each time step when new data is added (Kohzadi et al. 1996).

After selecting an appropriate validation technique, the next step is choosing a *performance metric* for the predictive problem. For classification tasks, accuracy is a very commonly applied metric. It measures the ratio between the number of correctly predicted target labels and the total number of predictions. The accuracy metric is only designed for tasks considering all classes as equally important, and its usefulness suffers if the samples within the classes are not equally distributed. For imbalanced data sets, the preferable metrics are balanced accuracy or the weighted f-score. The most common metrics for evaluating predictive regression tasks are mean absolute error (MAE), or the mean squared error (MSE). To provide relational insights, in particular in an organizational context, the mean absolute percentage error (MAPE) is useful. One of the metrics is then chosen for model training, yet it is common to provide an overview of multiple metrics for the evaluation.

Based on the chosen validation set generation and performance metrics, the model is *trained and tuned*. For the hyper-parameter tuning, search algorithms such as grid, random, or Bayesian search for the optimization parameters (Bergstra and Bengio 2012; Snoek, Larochelle, and Adams 2012) should be leveraged. To do so, well established packages, such as Hyperopt (Komer, Bergstra, and Eliasmith 2019), keras-tuner (O'Malley et al. 2019), or auto-sklearn (Feurer et al. 2019), with associated guidelines can be used. Finally, the tuned models are *tested* and the learning curves *evaluated*, to ensure a robust model for the prediction task.

5.4.5 Multi-Headed Neural Network Application

In the last phase, the method describes aspects for MH-NN application. This includes the *operationalization* of data acquisition and preparation as well as the deployment of an evaluated MH-NN. Of particular importance is the live connection to the enterprise process network and the data sources. Instead of training on historical data, the MH-NN must handle live data to provide real-time predictions. Thus, besides model performance, runtime performance becomes particularly relevant during model deployment.

If the model is integrated into the enterprise process network and connected to (live) data sources, it facilitates the prediction of the desired variable. Such a prediction then affects an organizational process, for example, through the prediction of upcoming events or the classification of an event's type, which can be used to provide better solutions in organizations. As the processes are improved due to the prediction, the designed model then assists in the organizational goal of *process improvement*.

5.5 Method Evaluation

To evaluate the PPNM method, we use a real-world use case. Figure 5.5 summarizes the application, evaluation, and discussion scenarios.

5.5.1 Problem Specification and Industry Background

For our research, we collaborated with a medium-sized German manufacturing company. The firm has multiple distributed production and assembly lines for highly customized mechatronics products. Competitive pressure necessitates the firm to offer high-quality products with (mass) customization options. This combination can lead to fairly complex production processes. Here, disruptions[33] where a worker has to interrupt work, are not uncommon.

To efficiently handle such disruptions (Lopez-Leyva et al. 2020), our cooperation partner has deployed a disruption management system. The system

[33]Typical reasons include, e.g., missing materials, damaged parts, or non-functional machines.

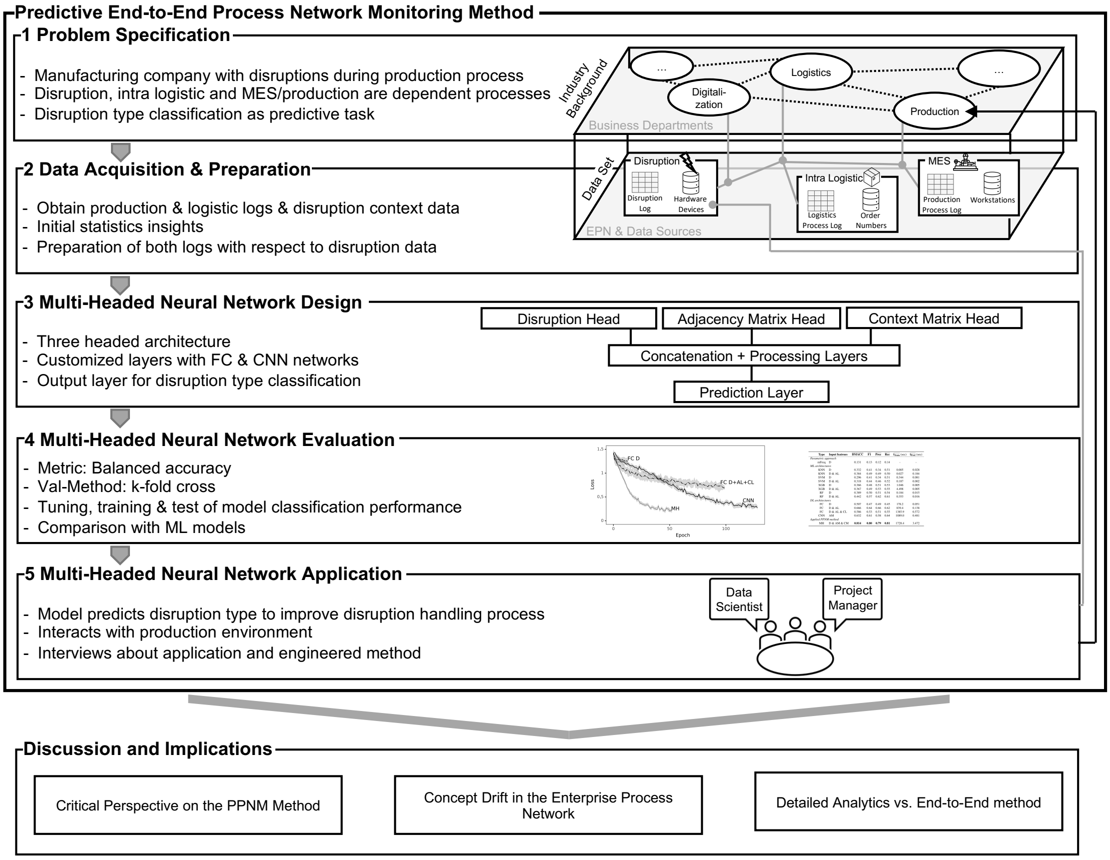

Figure 5.5: Overview of PPNM method evaluation.

automates responder notification for solving a disruption.[34] As a disruption is solved through the responding agent, the agent provides the system additional information, such as one of 32 disruption reasons (types). We identified the disruption's type as a central component of the problem specification. If the type was already known, an agent could already prepare the solution process (e.g., bringing relevant tools or documentation), which reduces the disruption associated downtime.

In parallel, the production processes have been analyzed with PM techniques to identify optimization potentials. However, due to the enterprise process network's complexity, interrelations, and dependencies, the respective analyses are very time-consuming. Consequently, the realization horizon of

[34]As an employee detects a disruption during the production or logistics process, the employee presses one of the system's hardware devices. In doing so, the system automatically notifies a responding agent (employee with specialized skills for disruption solving), who assists in solving the disruption.

possible benefits is long. Striving for immediate benefit with minimal analysis effort, we adopt the PPNM method and provide an end-to-end PPNM solution. Thereby, the MH-NN is integrated into the organizational enterprise process network. The organizational objective is to improve the production process through better disruption handling, resulting in reduced downtime. We do so by predicting the disruption type and providing a solution suggestion to a notified agent based on the prediction. Thus, accurate predictions are essential to ensure target-oriented suggestions which speed up the solution process.

We cooperate with various departments (digitalization, logistics, and production) to evaluate the PPNM method in practice. Thereby, we need to deal with each department's process event log and related databases.[35]

5.5.2 Data Acquisition and Preparation

We compute basic statistics and advanced event log characteristics such as sparsity, variation, or repetitiveness (Heinrich et al. 2021; Di Francescomarino et al. 2017) to better understand the *production* and *logistics* event log data used (Table 5.2) as well as the disruption context information (Table 5.3). The descriptives demonstrate the high complexity of the semi-structured event logs with many unique process variants and activity types. Furthermore, we combine both event logs and obtain the combined production event log, which contains information about the logistics and production process, its control flow, and context information.

The disruption log is closely related to the intra-logistics and production departments and processes, as disruptions occur in both departments. It contains information about historical disruptions with features such as the disruption hardware id and timestamp. This way disruptions can be mapped to a workplace through the hardware device database. This enables us to retrieve product information from the respective data sources, which we can also leverage as features for the predictive task.

With the data at hand, we follow the PPNM method by preparing data and designing a multi-head NN. We start with the data preparation for the disrup-

[35] Production and logistics processes span across the departments, such as logistics events are performed in the production department. However, the respective logs mainly originate from one of the departments.

	Data sources	Production	Logistics
Number of	process instances	24581	24581
	process variants	859	240
	activity types	156	69
Events per instance	minimum	4	2
	average	5	4
	maximum	34	20
Process	sparsity	0.006	0.002
	variation	0.034	0.010
	repetitiveness	0.425	0.434

Table 5.2: Overview of the *production* and *logistic* event log with a summary of descriptive statistics.

	Data source	Disruption
Number of	events	4739
	numerical features	4
	categorical features	20

Table 5.3: Overview of the disruption context information features.

tion log. Concerning the hardware id, we include additional workstation and product information, which we one-hot encode. Besides, we can extract time features, such as days, weekdays, hours, and minutes, from the disruption-associated timestamp, which we subsequently normalize.

By aggregating the logistics and production log, we obtain a process event log with context information. To transform the event log into valuable features, we follow the approach of Oberdorf et al. (2021a) and select process instances within a time window, which we subsequently transform into a matrix representation. By doing so, rows and columns relate to specific workstations and the value of a distinct cell to the production quantity within the time window. For NN preparation, we scale each matrix by the maximum production quantity of all matrices. This process is used for the control-flow data (*process matrices*) as well as for the context data (*context matrices*).

5.5.3 Multi-Headed Neural Network Design

We choose a three-headed DNN architecture (Appendix C). The *disruption vector* is the first input for the multi-head NN and is processed with an MLP (head), including a batch normalization. For both input matrices (weighted adjacency and context matrices), we use CNN architectures, consisting of stacked CNN and fully connected (FC) layers. For the context information, we apply a CNN-FC architecture to perform best in combination with the other heads. It consists of three CNN-layers and a subsequent FC layer. The third head's design—the process event head—posts a more challenging task. We tested the architecture from the context information and appended the adjacency matrices to the context matrices in the fourth dimension.[36] However, none of these approaches delivered satisfactory results. For this reason, we use knowledge about the fundamental production processes in the definition of the CNN kernel sizes. Basically, multiple sequential CNN layers extract features with distinct kernels.[37] After feature extraction, both matrix head outputs have a 4D shape. To combine both with the disruption head's output vector, we flatten the matrix head outputs. The flattened features are subsequently processed by a dense layer and the final output dense layer for the multi-class classification task.

5.5.4 Multi-Headed Neural Network Evaluation

To numerically evaluate the proposed method, we classify the type of each disruption event with the constructed MH-NN. In addition, we compare traditional aggregation-based approaches, where we append the disruption input vector with engineered (process) adjacency list features and, in addition, a vector of context information. Instead of 24 disruption vector features, we use 291 input features for adjacency list combination. In combination with the 267 additional adjacency list features, we use a total of 558 features.

To ensure reliable results, we perform a five-time repeated five-fold cross-validation with random initialization. To prevent the DNN models from overfitting, we integrate an early stopping rule for validation accuracy. We store

[36]The first dimension relates to the batch size, dimensions two and three to the matrix, and the fourth dimension to the heads of a CNN. In image processing, it represents multiple color channels.

[37]A small kernel is leveraged to extract information within a production line, up to large kernels, which extract information across multiple production lines.

the best-performing models during each training cycle and used a Bayesian optimization algorithm (O'Malley et al. 2019) for hyperparameter tuning. Our tuning objective is the validation accuracy with a maximum retrial of 50 configurations.

For the tuned FC, CNN, and multi-headed (MH) models, we at first compare the validation loss (Figure 5.6) at the stopping time. The multi-headed approach's loss clearly outperforms the other DNN architectures. In addition, it reaches a solid model with fewer epochs compared to the CNN or FC architecture with flattened feature inputs.

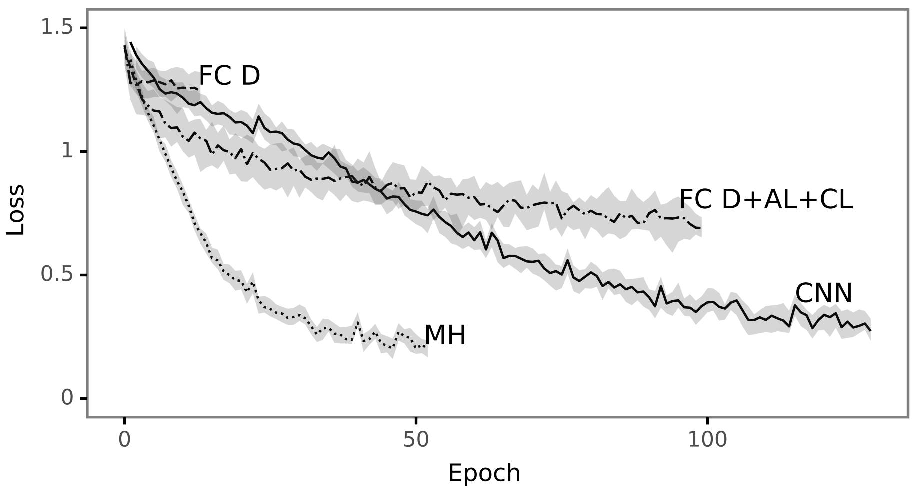

Figure 5.6: Comparison of validation loss of FC, CNN, and MH algorithms for disruption classification with input scenarios for disruption vector (D), the combination with adjacency list (AL) as well as context list (CL) vector.

The final models are subsequently evaluated on the hold-out set, resulting in the metrics summarized in Table 5.4. All evaluated algorithms, ML, and DNN models outperform the naive benchmark in terms of BMACC as well as the (weighted) F1-score, Precision, and Recall-score. We observe that the FC architecture benefits from the additional adjacency list features. However, we also see that the additional context list features lead to a predictive power decrease, suggesting that the FC architecture cannot prevent overfitting completely.

A comparison of CNN with only adjacency matrix features shows that they contain some basic information. However, this performance does not match the FC architecture with disruption and adjacency list features. The proposed multi-headed NN approach outperforms all benchmark architectures. Besides the better training behavior of the multi-headed NN approach, the higher aggregation of the data seems to result in this information loss. Due to the matrix properties, the CNN can identify patterns in the data that lead to improved results. Note that the resulting multi-class accuracy refers to a 32-class classification problem. Accordingly, the 81 % multi-class accuracy is a good result, allowing a reliable solution suggestion.

Type	Features	$\overline{\mathbf{BMACC}}$	$\overline{\mathbf{F1}}$	$\overline{\mathbf{Prec}}$	$\overline{\mathbf{Rec}}$	$\overline{\mathbf{t_{Train}}}^{*}$	$\overline{\mathbf{t_{Pred}}}^{*}$
Basic benchmark approaches							
mFreq	D	0.131	0.13	0.12	0.14	-	-
KNN	D	0.332	0.41	0.34	0.51	0.005	0.028
KNN	D & AL	0.384	0.49	0.49	0.50	0.027	0.184
ML architectures							
SVM	D	0.296	0.41	0.34	0.51	0.344	0.081
SVM	D & AL	0.318	0.44	0.46	0.52	0.187	0.002
XGB	D	0.366	0.48	0.51	0.53	1.046	0.005
XGB	D & AL	0.367	0.49	0.53	0.55	4.498	0.005
RF	D	0.389	0.50	0.51	0.54	0.184	0.015
RF	D & AL	0.442	0.57	0.62	0.61	0.353	0.016
DL architectures							
FC	D	0.507	0.47	0.49	0.45	178.2	0.051
FC	D & AL	0.666	0.64	0.66	0.62	839.4	0.138
FC	D & AL & CL	0.586	0.53	0.51	0.55	1,385.9	0.572
CNN	AM	0.632	0.61	0.58	0.64	1,089.0	0.481
Applied PPNM method							
MH	D & AM & CM	**0.814**	**0.80**	**0.79**	**0.81**	1,728.4	3.472

Table 5.4: Comparison of algorithms for disruption classification with input scenarios for disruption vector (D) and combinations of adjacency and context list or matrix. *Time in seconds.

The experimental results of the multi-headed architecture are in line with recent research in computer vision (He et al. 2016) in general and predictive process monitoring (Rama-Maneiro, Vidal, and Lama 2020a) in particular. The DL algorithms show superior performance for the specific use case of multi-class classification. However, the superiority of the MH-NN architecture in terms of predictive power is tied to some drawbacks regarding implementation and training time. Compared to the standard ML models, that are readily implemented using libraries such as Scikit-learn (Pedregosa et al. 2011), finding and implementing optimal NN architectures for each network head is a complex and time-consuming task. Additionally, the training of the multi-headed NN takes significantly more time.[38] Clearly, this is a limitation of the MH-NN model. For our use-case, however, the prediction duration is more relevant, which is acceptable and facilitates the application of the model.

5.5.5 Multi-Headed Neural Network Application

In the last phase of the PPNM method, we deploy data acquisition and preparation as well as the best-evaluated model. The method's resources are deployed on a cooperation partner's standard commercial virtual machine with Linux OS. It is connected to the organizational enterprise process network through an MQTT connection, which enables the live interaction with the disruption management system. Whenever a disruption occurs and the worker triggers the notification process, the disruption data is transmitted through the MQTT connection and triggers the prediction process. Recent production and intra-logistic event log data are automatically obtained, and all data are prepared as well as forwarded to the MH-NN. The prediction result is then transmitted to the disruption management system and improves the information, which a responding agent receives as part of the disruption notification. Therefore, better preparation for the disruption task at hand is possible, which ultimately reduces disruption downtimes.

Under the current pandemic situation, a reliable live evaluation of the method was not applicable, as the production amount decreased. We could deploy and test the method in such circumstances, but the available amount of practical evaluation data is too scarce for a numerical evaluation. Instead, we interviewed a data scientist and a project manager.

The collaboration facilitated the awareness for the great interdependence of the processes. Clearly, processes affect each other, even across organizational borders, which the employees were aware of. However, combining these

[38]We trained all models on a NVIDIA GeForce GTX 1080 TI with 11 GB GDDR5X RAM.

heterogeneous data sources meant great efforts. The proposed method provides a valuable tool for structured data combination across departments.

> *"Of course, we are aware of interdependent processes, but leveraging the data was usually not practical. The multi-headed NN approaches bridge this gap, as we can further combine data without the downside of extensive aggregation. And due to the deployment, even without first searching and collecting the data."*
>
> (Data Scientist)

We presented the initial results to data scientists, project managers, and managers of the cooperation partner and discussed the practical implications. Aligned with the data scientist's perspective, the project manager depicts the potential on an organizational scale. Beyond the digitalization, production, and logistics departments, applications to financial and controlling are of particular focus. Connections to the customer relationship management (CRM) system or website user statistics may enable a better prediction of incoming orders, leading to improved production planning. In addition to better predictions, the deployment is then of special importance.

> *"We do not just want to have the [multi-headed NN] approach, but really looked forward to deployment of services. Without deployment, we can not generate the desired value."*
>
> (Project Manager)

5.6 Discussion and Implications

The presented method enables predictive end-to-end enterprise process network monitoring by leveraging a multi-headed NN architecture. Through the cross-organizational end-to-end view, interrelationships and dependencies between different departments, processes, and information systems can be jointly analyzed.

5.6.1 Critical Perspective on the PPNM method

Through the first and last phase with particular focus on the organizational layers, we enable *end-to-end analyses* and fulfill R1. Leveraging the multi-headed DNN architecture provides a scaleable solution to combine *multiple data sources* (R2) from across the organization and processes, each with specialized input heads. Even if we only presented a classification case study, both regression and classification tasks are applicable as *prediction target* (R3). For the case study, we applied the PPNM method to a real-world use case and designed a three-headed DNN architecture with multi-log and context data input heads. This architecture fulfills the performance requirements R4 and R5 for *predictive power* and *runtime*. Combined with the employees' feedback, we can summarize that the PPNM method helps guiding the development of predictive end-to-end enterprise process network monitoring.

Considering the MH-NN, architecture alternatives may enhance predictive power. Thus, it may be worth comparing multiple architectures for the same input. We did so during the MH-NN design, resulting in the design with three customized heads. However, with ongoing advances in NN development, new layers or even (pre-trained) state-of-the-art methods may emerge. Thus the chosen MH-NN should be regularly reviewed.

5.6.2 Concept Drift in the Enterprise Process Network

The fifth phase consists of the final step of model integration and operationalization in the enterprise process network. It comprises the final online deployment, where (live) data sources are fed into the trained model for real-time predictions. Once the predictive model has been put into production, it draws on the knowledge from the historical data used for training. Deployed models inevitably face the phenomenon of structural changes in data over time, which is referred to as *concept drift* and usually leads to a deterioration of the prediction performance. Maisenbacher and Weidlich (2017), Denisov, Belkina, and Fahland (2018), and Spenrath and Hassani (2020) mention respective observations in various organizational PPM contexts. Yet, the concept drift problem is neither limited to PPM, but also known in the more general fields of PM (Adams et al. 2021; Sousa et al. 2021) and ML (Widmer and Kubat 1996).

For valid process predictions and analyses, the phenomenon of concept drift has to be detected and counteracted at an early stage. Currently the PPNM method, does not account for concept drift. To detect a concept drift, multiple methods are known (Seidl 2021; Kahani, Behkamal, et al. 2021), such as local outlier detection, which can initiate retraining of the model with updated data to avoid wrong predictions and achieve temporal stability (Teinemaa et al. 2018).

5.6.3 Detailed Analytics vs. End-to-End Method

A common phenomenon of traditional enterprises with hierarchical organizational structures is silo thinking. The symptoms of it are weak collaboration throughout the organization. As a result, isolated process analysis within departmental boundaries is often observed, as there is little responsibility for end-to-end processes (Eggers et al. 2021). Nevertheless, a holistic view of the organization is necessary as processes often span several departments. Connected through information systems, inter-departmental information about processes is available. In this regard, digitalization and emerging technologies, such as PM or PPM, enable end-to-end insights into processes and a holistic view on the heterogeneous IT-landscape of enterprises (Armengaud et al. 2020). Both PM and PPM provide tools for generating insights on processes on an organizational scale, as they can process large amounts of data. For example, Lorenz et al. (2021) provide an end-to-end perspective for PM to improve the productivity in make to stock manufacturing processes, and Eggers et al. (2021) show how management decisions can drive an end-to-end perspective on process data by creating new process owner positions. However, the capability of end-to-end process analysis is hardly considered in research as well as in practice.

Our proposed PPNM method contributes to this field of research by integrating the enterprise process network with all its interrelations and dependencies. In addition, for PPM as a subcategory of PM, our research has shown the benefits of taking an end-to-end view of processes for predictive tasks. The PPNM method and the fusion of inter-departmental data sources significantly increase the predictive power. This is already a first contribution, but it should not be the end of the research. Our approach for end-to-end PPNM is only an avenue towards general approaches for end-to-end PM. Therefore, future research should focus on leveraging the resources of the enterprise process network for PM and derive end-to-end insights.

5.7 Conclusion and Outlook

We present the PPNM method, for end-to-end enterprise process network monitoring, leveraging a MH-NN approach. In doing so, we overcome the phenomenon of silo-thinking and separated analysis of in data sources, as we enable the seamless combination of multiple data sources, combined with specialized processing and NN computation for each input. The resulting MH-NN outperforms classical ML and DL models and was applied and evaluated in an organizational context.

From a more general perspective, the method is an essential piece of research, enabling end-to-end PPNM on an organizational scale. Further, it guides the path towards a more general end-to-end PM, which then overcomes silo-thinking and enables an organization's enterprise process network's potential (Jokonowo et al. 2018). However, the approach is not limited to single organizations. Due to the approaches' extend-ability, additional data sources, even across multiple organizations, could be combined and leveraged each best. Thus, we further contribute to research towards holistic supply chain analytics. Respective inter-organizational PM analyses are proposed by Hernandez-Resendiz et al. (2021) for descriptive supply chain analytics, yet predictive insights are neglected. Our research extends the scope and enables the inter-organizational combination of data, even for predictive tasks. With larger data integrated, additional analytics research streams such as federated learning or aspects such as data ownership become more relevant and should be investigated in future research. The transfer of improved process predictions within and across organizations is not only relevant for research, but especially for enterprises by means of scaling the respective solutions. Thus, our method not only enables new research but could be a fundamental component for scaleable enterprise-ready PPNM solutions with heterogeneous intra- and inter-organizational data sources.

6 Data-Driven Approximate Dynamic Stochastic Programming for Maintenance Job Assignment

This working paper is currently under preparation for publication (Oberdorf et al. 2022b).

Across various industries, companies face the challenge of assigning employees to jobs in a cost-minimizing way. Due to heterogeneous abilities of the employees (agents), the time to complete a job differs, and the decision for an agent thus directly affects the costs. In addition, the decision affects the future state through agents' availability due to previous and current assignments. With randomly occurring jobs during a day and uncertainty about the time to complete a job, this results in a dynamic stochastic assignment problem with a prohibitively large state-space.

Motivated by a real-world maintenance assignment problem of a manufacturing company, we propose and study a new data-driven approximate dynamic stochastic programming approach, which addresses both uncertainties. To this end, we leverage local machine learning methods to approximate the conditional distribution of the uncertainties regarding a set of features. Based on these distributions, we solve the dynamic stochastic optimization problem and benchmark it with a set of existing state-of-the-art approaches that do not account for both uncertainties. The proposed approach provides superior performance, and we can additionally shed light on the practical limitations and future research directions.

6.1 Introduction

In the manufacturing and service industry, companies face the challenge of assigning employees to jobs in an efficient manner. For instance, call center operators need to assign service agents with heterogeneous abilities to randomly arriving service requests. The same applies to medical units with incoming patients and the assignment of specialist doctors. Another example are production companies, where the requests relate to maintenance jobs due to disruptions during production. When assigning appropriate agents to solve a job, a decision-maker considers the assignment cost of agents and attempts to minimize them. Given different agent specializations, the costs vary across the available agents. Thus, the agents chosen to solve a job directly affect its cost. The case of assigning maintenance jobs illustrates the direct relationship between agents and costs. The downtime of a machine depends on the specialization of an assigned agent. It is assumed that specialists for specific jobs perform them faster, while generalists take longer but can cover a broader range of maintenance jobs.

Making such decisions in real-world use-cases is a challenging task due to the complexity and dynamic nature of the problem and because real-time decision support is increasingly needed. Therefore, companies that face assignment decisions (e.g., medical units, manufacturers, or service providers) can benefit from decision support systems instead of trusting in gut feeling about an assignment. However, determining the right agent for a problem is difficult, particularly because of the dynamic problem, where the current decision impacts the system's (unknown) future state by means of the availability of agents for (potentially) upcoming jobs. In addition to the unknown future, the necessary time to solve a current job (concerning the chosen agent) is usually uncertain when the decision has to be made. More formally, a company has to solve a complex dynamic stochastic assignment problem with recourse, where the vector-valued decisions are interrelated, and a current assignment's costs and future jobs are unknown. Variants of the dynamic stochastic assignment problem are well studied; however, we emphasize that most variants of solution systems proposed for dynamic stochastic assignment problems involve only one kind of uncertainty (Powell 1996, 2007).

Research in operations management (OM) has focused on building models and solving these through (manual designed) approximations to deal with such complex problems. However, recently the abundance of data, combined with advances in computer science and statistics, has led to a shift in operations management research (Mišic and Perakis 2020). Using historical data and related features to support operational decisions is becoming increasingly important to improve decisions and address natural uncertainties (Gambella, Ghaddar, and Naoum-Sawaya 2020; Kraus, Feuerriegel, and Oztekin 2020). In general, the accurate solution of a dynamic stochastic assignment problem is not tractable and results in a prohibitively large state space due to the problem's dimension and complexity. Therefore, the recent body of literature on dynamic decision-making under uncertainty follows data-driven optimization approaches and proposes local machine learning methods to take advantage of historical data and related features (Bertsimas and Kallus 2020a; Ban and Rudin 2019; Bertsimas and Koduri 2021). Typically, an unknown joint distribution of uncertain parameters and available data is assumed to model future uncertainty and approximate the optimization problem (Bertsimas, McCord, and Sturt 2019). Bertsimas and Kallus (2020a) prove the asymptotic optimality of data-driven optimization problems under mild conditions. In the fundamental literature on data-driven problems in optimization, on which we rely, there are two different paradigms. Prescriptive analytics approximate the objective using machine learning to predict the optimizer (Bertsimas and Kallus 2020a; Bertsimas and Koduri 2021) and predict-then-optimize, where the prediction models aim to minimize the decision error instead of the prediction error (Elmachtoub and Grigas 2021). A common denominator for any of these approaches is to take a point or vector-based predictions or prescriptions into account. However, facing the dynamic stochastic assignment problem, the sequence of jobs during a day is particularly important and must be considered.

Motivated by a manufacturing company's real-world maintenance assignment problem, this paper proposes and studies a new data-driven approximate dynamic stochastic programming approach, which addresses this issue in a time-continuous and event-discrete formulation with finite horizon. We estimate the current assignment costs (duration) and the sequence of unknown upcoming jobs (outlook) as part of the optimization to provide a data-driven agent assignment. We build on recent prescriptive analytics approaches (Bertsimas and Kallus 2020a) and leverage local machine learning methods to estimate weights, which can be understood as an approximation of the conditional distribution of the uncertainties with respect to a set of features. In addition, we extend the approach and expand it to provide a novel approach to account for sequences of (potentially) upcoming jobs. Using real-world data, we evaluate our approach relative to relevant benchmark policies. From a more general perspective, our approach establishes a crucial link to incorporating historic information to overcome myopic decisions and deal with upcoming jobs. In doing so, we shed light on *how valuable is the outlook information* and discuss it in the context of real-world problems.

Figure 6.1 provides a high-level overview of our approach: At first, we formalize the optimization problem incorporating the occurrence of jobs and the assignment of agents with respect to real-world constraints (Section 6.3). Crucial inputs for solving this problem are the estimated current job duration as well as the estimation about upcoming jobs. To incorporate uncertain future jobs, we utilize the similarity between disruption-job and production-log data and present a kernel method-based approach. Combining these three components results in the data-driven maintenance assignment to minimize the total disruption downtime considering the interdependence on (potentially) upcoming jobs.

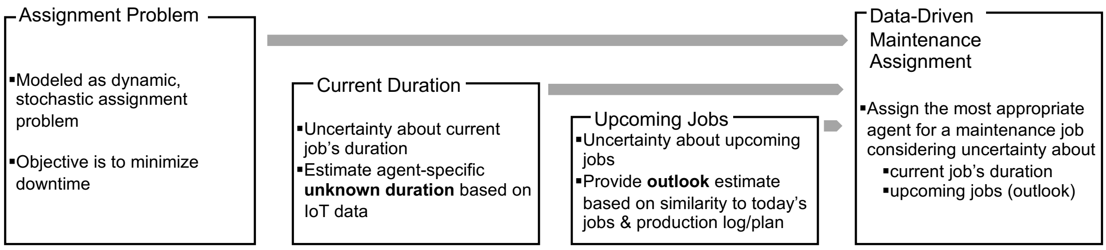

Figure 6.1: Overview of the proposed data-driven approach.

6.2 Related Work

Our research in the context of data-driven maintenance job assignment builds on two literature streams. On the one hand, we build on traditional operations management problems, such as the assignment problem. On the other hand, our research is based on research in machine learning applications. As we aim to combine machine learning and mathematical optimization, we seek to build on prescriptive analytics approaches.

6.2.1 Combinatorial Optimization Problems

Combinatorial problems arise in operations research and typically comprise tasks such as assignments of a discrete, finite set of objects that satisfy certain constraints (Ausiello et al. 2012). The generalized assignment problem (GAP) is a classical operations research model defined as a deterministic minimization problem that seeks the minimum cost assignment of tasks (jobs) to agents, with each task (job) assigned to one agent subject to capacity restrictions of the agent (Srinivasan and Thompson 1973). Generally, the GAP problem differentiates from the classical assignment problem such that an agent may be assigned to multiple jobs ensuring that each job is performed exactly once while being limited by the resource availability of the agents (Ross and Soland 1975). More precisely, the formulation of the assignment problem is a special instance of the GAP. The application scenarios of the GAP problem are manifold, which is why extensions have been proposed to describe the specific problems more accurately. The various extensions relate to the types of resource constraints, which describe the circumstances of jobs and relationships between agents and jobs.

For instance, in the manufacturing environment, an application scenario is loading in flexible manufacturing systems (Kuhn 1995) or production scheduling (Farias Jr, Johnson, and Nemhauser 2000). For the production scheduling case, the GAP can be further generalized to include the problem of allocating jobs to time periods, which is referred to as GAP with a special ordered set (Farias Jr, Johnson, and Nemhauser 2000). However, the production scheduling problem can also be formulated as a bi-objective GAP (Zhang and Ong 2007), where, for example, a job assignment's time and cost are taken into account. The example shows that the various scenarios to be solved may require different extensions of GAP, of which we present the relevant extensions for the problem described in the paper. See Kundakcioglu and Alizamir (2009, p. 1153-1162) for an extensive overview of existing applications, extensions, and solution methods of the GAP.

In a scenario where each job is executed with the objective of fulfilling a specific demand profile, the sequence of executing jobs gains in importance. This can be readily explained with the example of production scheduling: If a production job is completed before the demand, inventory costs incur, and if the job deadline is exceeded, shortage costs incur. However, the GAP does not incorporate the sequence of jobs to be performed into the problem definition. For this reason, the dynamic extension of GAP adds a time dimension, assuming a due date for each job, and allocates costs if a job is finished before or after the due date (Kogan and Shtub 1997). With the continuous-time optimal control model Kogan and Shtub (1997) derive analytical properties of optimal behavior of a dynamic system. Similar jobs are grouped in the dynamic GAP model, agents work at a fixed rate, and due dates are interpreted as occurring demand. However, the original formulation of the dynamic GAP only considers deterministic demand and agent-related constraints. Both restrictions can be addressed by extending the model to deal with stochastic demands and multiple agent-job relationships, such as the time-dependent capacity of agents (Kogan, Khmelnitsky, and Ibaraki 2005). The dynamic GAP can also be examined in an online setting where the future is revealed step-wise and decision making follows after each step or decisions are made only once. Feldman et al. (2009) proposed the GAP for display ads allocation with the objective to maximize the value of all assigned impressions. The online GAP is enhanced by Alaei, Hajiaghayi, and Liaghat (2013), adding two sources of uncertainty to the online stochastic GAP for subscription-based advertising.

Other studies focus on dealing with incomplete information about the future. The stochastic environment of GAP forms an entire research stream that focuses on the mapping of uncertainty in GAP. Uncertainty transforms the basic GAP into a stochastic model that considers the uncertainty of agents' resource consumption (Dyer and Frieze 1992) or capacities (Toktas, Yen, and Zabinsky 2006), penalties (Spoerl and Wood 2004), assignment costs (Dyer and Frieze 1992), or job emergence (Albareda-Sambola and Fernandez 2000; Albareda-Sambola, Der Vlerk, and Fernandez 2006). The stochastic models incorporate the stochasticity with random variables and represent the uncertainty in different scenarios. Research on stochastic GAP proposes different ways to model the stochastic environment. For instance, while Mine et al. (1983) develop a heuristic for the stochastic side constraints, Dyer and Frieze (1992) infer a probabilistic analysis for cost and resource parameters that are drawn from a uniform distribution on the unit interval, and Blower and Dowlatabadi (1994) and Chalabi et al. (2008) investigate Monte Carlo Sampling to sample uncertain variables. Spoerl and Wood (2004) were the first to propose a two-stage stochastic GAP with an exact algorithm with normally distributed resource consumption coefficients with known means and variances. The stochastic extension of the GAP is also studied by Errarhout, Kharraja, and Corbier (2016), which investigate the problem in the home healthcare environment and formulate a two-stage stochastic GAP with constraints related to the agent's skill, travel load, and capacity. The uncertainty in their model is the time required for care, which is related to the learning care factor. They use Monte Carlo sampling to generate different scenarios (Verweij et al. 2003). The model description shows similarities with the stochastic decision approach we describe in this paper. However, in the approach of Errarhout, Kharraja, and Corbier (2016), the patients' demands and required skills are already known, whereas, in our model in the production environment, neither the required skill set nor upcoming jobs are known.

With the desire to cope with uncertainties in dynamic optimization problems, Powell (1996) proposed the first model for a stochastic, dynamic assignment problem in a continuous-time setting by combining the assignment model with an approximate recourse function. The proposed model handles uncertainties with forecasts and integrates actual as well as forecasted demands. It outperforms the standard myopic models and is seen as a milestone for newer models with integrated approximated expected recourse functions. Subsequently, Powell (2007) developed the approximate dynamic programming framework following the basic idea of traditional dynamic programming. This framework for stochastic optimization solves large multiperiod optimization problems by decomposing the temporal dependency into small subproblems and incorporating mathematical programming.

6.2.2 Data-driven Optimization

The GAP's complexity is shown to be NP-hard, (Sahni and Gonzalez 1976) and determining a feasible assignment of an instance of a GAP is NP-Complete (Fisher and Jaikumar 1981). Consequently, exact solution approaches face limitations for large-sized instances because they become computationally intractable. The limitations of exact algorithms are overcome with the help of heuristics, metaheuristics, or relaxations (Öncan 2007). One of the earliest approximation algorithms for the GAP is a polynomial-time approximation considered by Shmoys and Tardos (1993). The problem class of decision-making under uncertainty requires special, simplifying assumptions and alternative solution techniques.

Traditional models for stochastic optimization problems assume perfect information on probability distributions of random variables in parametric form (Liyanage and Shanthikumar 2005), whereas non-parametric approaches replace the true distribution with an empirical distribution (Shapiro 2003; Kleywegt, Shapiro, and Homem-de-Mello 2002). In reality, information asymmetries exist with regard to the uncertain future, and strategies derived from optimization solutions with incorrect assumptions might lead to poor performance in practice (Bertsimas and Thiele 2006). In the light of this problem, several studies focused on data-driven non-parametric optimization approaches to provide a near-optimal approximation of the problem solution (Bertsimas and Kallus 2020a). The body of literature on data-driven optimization centers on uncertainty in dynamic and stochastic optimization problems and divides into two paradigms: prescriptive analytics (Bertsimas and Kallus 2020a; Bertsimas and Koduri 2021) and predict-then-optimize approaches (Elmachtoub and Grigas 2021).

Solution systems with distribution-free data-driven approaches comprise operational statistics (Liyanage and Shanthikumar 2005; Chu, Shanthikumar, and Shen 2008), sample average approximation (SAA) (Bertsimas, Gupta, and Kallus 2018b), and robust optimization (Bertsimas and Thiele 2006; Bertsimas, Gupta, and Kallus 2018a). However, these approaches do not directly prescribe the optimal solution and instead follow a two-step approach. In contrast, Ban and Rudin (2019) consider training machine learning models on data and directly predict the outcomes of the newsvendor problem and show its effectiveness in practice. They proposed linear empirical risk minimization (ERM) and kernel optimization methods for finding a near-optimal solution for the newsvendor problem. Ban and Rudin (2019) also show that the ERM approach might work for non-linear function spaces by applying kernels and Notz and Pibernik (2021) demonstrate how out-of-sample guarantees for various kernels can be derived for non-linear solutions. Another data-driven optimization approach by Bertsimas and Kallus (2020a) compare a set of different prescriptive analytics approaches and assign weights to historical data based on additional information available. They propose a weighted SAA and ERM to provide a near-optimal approximation of the problem solution.

Building on the results of Bertsimas and Kallus (2020a), Bertsimas and Koduri (2021) compare two methods that apply the framework of regression in reproducing kernel Hilbert space to solve stochastic optimization problems with historical data. The methods distinguish in the approximation objective: The first approximates the objective function, whereas the second approximates the optimizer. Instead of estimating the conditional distribution of the uncertain parameter given the covariate vector, the first method estimates the conditional expectation. In the second method, the optimal decision is predicted directly without determining and estimating the expected realizations of the uncertain parameters under the covariate vector condition. The experimental results of Bertsimas and Koduri (2021) show superiority for the second method over the first in overcoming the curse of dimensionality, performing better on unseen data, and showing performance for high-dimension data. They additionally exploit global machine learning methods to overcome the obstacles of local machine learning methods (Hastie, Tibshirani, and Friedman 2009) in prescriptive problems.

We draw on the contributions of Bertsimas and Kallus (2020a) to demonstrate that the data-driven stochastic optimization approach by incorporating additional historical information can be particularly effective in real-world maintenance job scheduling. Bertsimas and Kallus (2020a) apply several machine learning methods, such as k-nearest neighbor, random forests, Nadaraya-Watson kernel regression, and local linear regression and exploit the distance between an observation of auxiliary data and existing data to predict the uncertain parameters in the objective function. The choice of machine learning methods can be attributed to the fact that the proximity between the observations are used for prediction. As a result, the weighted observations can be used as an approximation of the conditional distribution of the uncertain parameters. We build on and extend this approach to approximate the solution of a dynamic stochastic program (section 6.3). The solution approach is subsequently presented (6.4), where we propose a novel method to incorporate outlook information and highlight in the evaluation that the value of outlook information is discriminative.

6.3 Problem Description

Consider the following dynamic stochastic assignment problem faced by a manufacturing company that operates multiple manufacturing lines: On any given day, the company faces multiple disruptions at its production lines. Each disruption event triggers a maintenance job. We use index j to represent individual maintenance jobs associated with the disruptions and denote the set of all disruptions as $\mathcal{J}$. Disruptions occur during a day according to some unknown stochastic process. This process results in J disruptions for a single day; J is a random variable whose realization is unknown before the end of the day. The company employs an online dispatching system—that is, whenever a disruption occurs, the dispatcher immediately assigns an agent a from a set A of agents to the corresponding maintenance job and determines the time at which an agent starts the maintenance job. Let q_{aj} denote a binary variable that indicates whether agent a was assigned to job j, and let the vector $\boldsymbol{q}_j = (q_{1j}, ..., q_{|A|j})$ represent the assignment decision for maintenance job j. Since exactly one agent is assigned to job j, we have $\sum_{a \in A} q_{aj} = 1$. We denote by s_j the time at which one of the agents starts job j. The time at which a maintenance job j can be started depends on the availability of the agents at time t_j, the time at which disruption j occurs. We denote by $\hat{s}_{aj}$ the earliest time at which agent a can start job j, and by $\hat{\boldsymbol{s}}_j$ the vector of earliest starting times of all agents for job j. The dispatcher assigns an agent and the starting time of the maintenance job, s_j, is given by $\sum_{a \in A} \hat{s}_{aj} q_{aj}$.

The time an agent requires to complete maintenance job j is a random variable, and we denote it by $R(q_{aj})$. From a practical perspective, the distribution of $R(q_{aj})$ depends on the skills required to solve a job and the skills of an agent. Shortly after an agent starts job j, the agent can provide a sufficiently accurate estimate of the time required to complete the job so that $r(q_{aj})$, the realization of $R(q_{aj})$, is known for $q_{aj} = 1$. Given s_j, $r(q_{aj})$ and all previous assignment decisions, the earliest time at which agent a can start job j $(j > 0)$ can be computed as follows: $\hat{s}_{aj}$ = $\max\{t_j; (s_{a0} + r(q_{a0}))q_{a0}; ...; (s_{a,j-1} + r(q_{a,j-1}))q_{a,j-1}\}$. We assume that all agents are available at t_0, the time at which the first disruption ($j = 0$) occurs, i.e., $\hat{s}_{a0} = t_0$ (for all $a \in A$). Realize that in this system, the dispatcher is free to assign an agent to job j that is available immediately (i.e., $\hat{s}_{aj} = t_j$) or an agent that becomes available at a later point in time $\hat{s}_{aj} > t_j$.

The dispatcher's objective is to determine a sequence of assignment decisions and starting times that minimizes the total downtime during a day. At time t_j ($j > 0$) the dispatcher knows the downtimes $d_0(\boldsymbol{q}_0), ..., d_{j-1}(\boldsymbol{q}_{j-1})$ associated with the assignments $\boldsymbol{q}_0, ..., \boldsymbol{q}_{j-1}$ made for previous disruptions. For example, the downtime $d_{j-1}(\boldsymbol{q}_{j-1})$ is given as $s_{j-1} + r_{j-1}(\boldsymbol{q}_{j-1}) - t_{j-1}$. In contrast, the downtimes induced by the current disruption j and all future disruptions ($j + 1, ..., J$) are uncertain. They are only known after an agent is assigned and has started the corresponding maintenance job. We denote by $D_j(\boldsymbol{q}_j), ..., D_J(\boldsymbol{q}_J)$ the random downtimes associated with the assignment decisions made for the current and all future disruptions. Clearly, the subsequent assignment decisions are linked: The assignment decision $\boldsymbol{q}_j$ influences the availability of agents for subsequent maintenance jobs ($j + 1, j + 2, ...$), which is captured by the vectors of earliest starting times. Therefore, when assigning an agent, the dispatcher should not only consider the random downtime $D_j(\boldsymbol{q}_j)$, induced by the current assignment decisions, but also its impact on (all) future decisions for disruptions that are yet unknown. For instance, the vector of earliest starting times $\hat{\boldsymbol{s}}_j$ is known for job j, but we do not know the vectors for future jobs $j + 1, j + 2....$ We denote by $\hat{\boldsymbol{S}}_j = (S_{aj}, ..., S_{|A|j})$ the vectors of uncertain starting times of future jobs ($j + 1, j + 2, ..., J$). We can then express the random start time of future jobs as $S_k = \sum_{a \in A} \hat{S}_{ak} q_{ak} (k = j+1, ..., J)$ and the random downtimes associated with future jobs by $D_k(\boldsymbol{q}_k) = S_k + \big(\sum_{a \in A} R(q_{ak}) q_{ak}\big) - T_k (k = j+1, ..., J)$, where T_k denotes the random time at which the k-th disruption occurs.

The dispatcher wants to solve the following dynamic stochastic optimization problem for every disruption j :

$$\min_{\boldsymbol{q}_j,\dots,\boldsymbol{q}_J} \mathbb{E}\big[\sum_{k=j}^{J} D_k(\boldsymbol{q}_k)\big] \tag{6.1}$$

s.t.

$$D_j(\boldsymbol{q}_j) = s_j + \big(\sum_{a\in A} R(q_{aj})q_{aj}\big) - t_j \tag{6.2}$$

$$D_k(\boldsymbol{q}_k) = S_k + \big(\sum_{a\in A} R(q_{ak})q_{ak}\big) - T_k \qquad \forall\, k = j+1,\dots,J \tag{6.3}$$

$$s_j = \sum_{a\in A} \hat{s}_{aj} q_{aj} \tag{6.4}$$

$$S_k = \sum_{a\in A} \hat{S}_{ak} q_{ak} \qquad \forall\, k = j+1,\dots,J \tag{6.5}$$

$$\hat{s}_{aj} = \max\{t_j; (s_h + r(q_{ah}))q_{ah} | h = 0,\dots,j-1\} \quad \forall\, a \in A \tag{6.6}$$

$$\hat{S}_{ak} = \max\{T_k; (s_j + R(q_{aj}))q_{aj}; (s_h + r(q_{ah}))q_{ah} | h = 0,\dots,j\} \tag{6.7}$$

$$\forall\, a \in A; k = j+1,\dots,J$$

$$\sum_{a\in A} q_{ak} = 1 \qquad \forall\, k = j+1,\dots,J \tag{6.8}$$

$$q_{ak} \in \{0,1\} \qquad \forall\, a \in A; k = j+1,\dots,J \tag{6.9}$$

The objective of the model, stated in (6.1), is to minimize the expected sum of total downtimes depending on the J assignment decisions. Constraints (6.2) and (6.3) define the current and the future downtimes depending on the assignments. Constraints (6.4) and (6.5) specify the current start time for maintenance job j and the uncertain future start times for jobs $k = j+1,..,J$ depending on assignment decisions q_{aj} and q_{ak}. Constraints (6.6) and (6.7) express the earliest current start time and the earliest future start times for each agent $a \in A$. (6.8) ensures that exactly one agent is assigned to each maintenance job. (6.9) is a binary constraint for the decision variables.

Clearly, this optimization problem cannot be solved to optimality: The future states of the system (S_k) are uncertain and not only depend on previous and future assignment decisions q_{aj}, but also on the number of disruptions J, the times T_j at which they occur, and the time $R(q_{aj})$ required to complete the maintenance jobs. Even if we were able to estimate a distribution for each of these random variables, we would face a prohibitively large state space. Therefore, we can neither derive nor characterize an optimal assignment policy. The conventional way to deal with this problem is the application of approximate dynamic programming techniques relying on an approximation of the value function and a state-space reduction.

We propose a different—data-driven—approach to solving this dynamic assignment problem. Our approach leverages extensive historical data (observations of past disruptions and numerous co-variates that may be predictive of the random variables of interest) and is computationally tractable. In the next section, we develop this approach.

6.4 Solution Approaches

We develop our data-driven approach to solve the stochastic dynamic assignment problem in two steps. First, we address the myopic single-stage problem, in which the dispatcher only aims at minimizing the downtime associated with the current disruption at hand. Thereafter, we explain how we account for the dynamic nature of the problem—that is, we incorporate the impact of the assignment decision at time t_j on future decisions at times $t_{j+1}, ..., t_J$ and their expected outcomes.

The myopic single stage problem can be expressed as follows:

$$\min_{\boldsymbol{q}_j} \mathbb{E}\big[D_j(\boldsymbol{q}_j)\big] = s_j + \Big(\sum_{a \in A} \mathbb{E}\big[R(q_{aj})\big] q_{aj}\Big) - t_j \tag{6.10}$$

$$\text{s.t.}$$

$$(6.4), (6.6), (6.8), (6.9)$$

The distributions of $R(q_{aj})$ are unknown, but the decision-maker has access to a set of data $S_N = \{(r_0(\boldsymbol{q}_0), \boldsymbol{x}_0), ..., (r_N(\boldsymbol{q}_N), \boldsymbol{x}_N\}$ that contains historical observations of times of completion $r_n(\boldsymbol{q}_n)$ and corresponding observations of features represented by p-dimensional vectors $\boldsymbol{x}_n \in \mathcal{X} \subseteq \mathbb{R}^p$. The p features describe, for example, a disruption's type, the assigned agent, the production line at which the disruption occurred, the time of occurrence, and various features derived from the production log. Such production features include, for example, the utilization of each of the production lines, the types of products produced, and time features capturing the duration of individual production steps per production line.

We assume that at least some of these features are predictive of the time $r(q_{aj})$ required to complete the job associated with a particular disruption and that there exists an (unknown) joint distribution $X \times R$ of features and times for resolving the disruption. Therefore, we would like to solve the following problem:

$$\min_{\boldsymbol{q}_j} \mathbb{E}_{R|X}\big[D_j(\boldsymbol{q}_j)\Big|\boldsymbol{X} = \boldsymbol{x}_j\big]. \tag{6.11}$$

However, we do not know the conditional distribution $R|X$; instead of trying to estimate feature dependent distributions of $R(q_{aj})$ and optimizing (6.11), we follow the prescriptive analytics approach proposed by (Bertsimas and Kallus 2020a) and solve:

$$\boldsymbol{q}_j^* = \operatorname*{argmin}_{\boldsymbol{q}_j} \sum_{n \in S_N} w_n(\boldsymbol{x}_j) d_j(\boldsymbol{q}_j, \boldsymbol{r}_n) \tag{6.12}$$

$$\text{s.t.}$$

$$(6.4), (6.6), (6.8), (6.9)$$

In the most general terms, this approach approximates (6.11) and can be viewed as a weighted form of the well-known sample-average-approximation technique (SAA). The function $w_n(\cdot)$ can be considered as a weight function that captures the similarity between the features pertaining to disruption j ($\boldsymbol{x}_j$) and the features pertaining to the N previous disruptions. We impose $w_n \in [0,1]$ and $\sum_{n \in S_n} w_n = 1$. In essence, (6.12) intends to optimize over the empirical density distributions of the downtime, conditional on $\boldsymbol{x}_j$. Bertsimas and Kallus (2020a) proposed multiple specifications of the weight function $w_n(\boldsymbol{x}_j)$, based, for example, on k-nearest-neighbors (KNN), kernel methods, or decision trees and random forests (RF). We explain our choice of the weight function for solving (6.12) in Section 6.5.4.

Solving (6.12) yields a myopic data-driven solution to the stochastic dynamic problem (6.1)-(6.9). To account for the impact of the current decision $\boldsymbol{q}_j$ on future assignments and downtimes, we follow a similar data-driven approach as for the myopic problem but now consider the downtimes $d_k(\boldsymbol{q}_k)$ associated with potential future decisions $\boldsymbol{q}_k$ ($k = j+1, ..., J$) dependent on the current decision $\boldsymbol{q}_j$. We again leverage historical observations of the times required to complete maintenance jobs and associated features $(r_n(\boldsymbol{q}_n), \boldsymbol{x}_n)$ from the set of data S_N. Now, however, we are interested in sequences of historical disruptions on individual days, because we want to account for the impact of decision $\boldsymbol{q}_j$ on the downtimes associated with subsequent disruptions $j+1,..,J$. For a given day $o \in O$ we define such a sequence by $P_o = \{(r_0(\boldsymbol{q}_0), \boldsymbol{x}_0), ..., (r_J(\boldsymbol{q}_J), \boldsymbol{x}_J(o))\}$. P_o is an ordered set that comprises all tuples $(r_n(\boldsymbol{q}_n), \boldsymbol{x}_n)$ for a given day o, sorted in ascending order according to the times t_n at which disruption n occurred. O denotes the set of historical observations—that is, the days in the past, for which we observed a sequence P_o. Thus, our entire data set is now $S_O = \{P_o | o \in O\}$. For each new disruption j, we split the data set S_O into subsets $S_O(t_j)^+$ and $S_O(t_j)^-$. $S_O(t_j)^+$ contains all sequences $P_o^+(t_j) = \{\, P_o \mid t > t_j \,\}$, associated with disruptions that occurred at a time $t > t_j$. Likewise, $S_O(t_j)^-$ contains all sequences $P_o^-(t_j) = \{\, P_o \mid t < t_j \,\}$, associated with disruptions that occurred at a time $t < t_j$. The rationale for splitting the data set into $S_O(t_j)^+$ and $S_O(t_j)^-$ can be explained as follows: Assume a new disruption j that occurs at t_j = 11:30 am. To evaluate the decision $\boldsymbol{q}_j$ for this disruption, we only utilize information for past sequences of disruptions that occurred after 11:30 am. We therefore assume that any earlier disruptions do not have informational value for the decision taken at t_j.

Based on these considerations, we can extend the data-driven myopic approach (reflected by (6.12) and the associated constraints) to account for the effect of the current decision $\boldsymbol{q_j}$ on future downtimes. Solving the following problem yields a data-driven solution to the stochastic dynamic assignment problem stated in (6.1)-(6.9).

$$\boldsymbol{q_j}^* = \underset{\boldsymbol{q}_j}{\operatorname{argmin}} \left[\sum_{n \in N} w_n(\boldsymbol{x}_j) d_j(\boldsymbol{q}_j, \boldsymbol{r}^m) + \sum_{o \in O} v_o(\cdot) \min_{\boldsymbol{q}_k, \dots, \boldsymbol{q}_{J(o)}} \sum_{k \in P_o^+} d_k(\boldsymbol{q}_k, r_k | \boldsymbol{q}_j) \right] \tag{6.13}$$

s.t.

$$s_k = \sum_{a \in A} \hat{s}_{ak} q_{ak} \qquad \forall k \in \{j \cup P_o^+\}; o \in O \tag{6.14}$$

$$\hat{s}_{ak} = \max\{t_k; (s_h + r(q_{ah}))q_{ah} | h = 0, ..., k-1\} \quad \forall\, a \in A; k \in \{j \cup P_o^+\}; o \in O \tag{6.15}$$

$$\sum_{a \in A} q_{ak} = 1 \quad \forall k \in \{j \cup P_o^+\}; o \in O \tag{6.16}$$

$$q_{ak} \in \{0, 1\} \quad \forall\, a \in A; k \in \{j \cup P_o^+\}; o \in O \tag{6.17}$$

The first term in the objective function (6.13) is the same as in (6.12)—it captures the (weighted) downtime associated with the current disruption j, dependent on the assignment decision $\boldsymbol{q}_j$. The second term approximates the effect that $\boldsymbol{q}_j$ has on future downtimes. More specifically, it reflects the weighted sum of minimum downtimes associated with subsequent disruptions $k \in P_o^+$ on past days $o \in O$, given $\boldsymbol{q}_j$. Here, $v_o(\cdot)$ is a weight function that measures the similarity between features of the current day and each past day $o \in O$ with features $x_0, .., x_{J(o)}$. Again, we impose $v_o \in [0, 1]$ and $\sum_{o \in O} v_o = 1$. Constraints (6.14)-(6.17) are the deterministic, data-driven counterparts of constraints (6.4)-(6.9) of the original stochastic dynamic problem. In Section 6.5.4 we explain how we derive features and how we employ machine learning techniques to determine the weight function $v_o(\cdot)$.

For doing so, we must consider a limitation of our solution approach. From a practical perspective, the historic observations only consist of realized manufacturing jobs, where an agent was assigned to a job. However, the effect of assigning another agent to a job j is unknown, what refers to the fundamental problem of causal inference (Bertsimas, McCord, and Sturt 2019). To deal with this limitation, we present a practical solution approach (Section 6.5.3) that leverages the available features to even out the unknown realizations.

6.5 Evaluation

This section presents numerical analyses to evaluate the out-of-sample performance of the data-driven solution approach for the stochastic dynamic assignment problem. We follow the process summarized in Figure 6.2 and at first describe a real-world data set from organizational databases (Section 6.5.1) as well as the preparation of the data (Section 6.5.2), obtain features that deal with the limitation of unknown realizations for not assigned agents (Section 6.5.3), and evaluate the data-driven weight estimation (Section 6.5.4). In addition to the proposed myopic and outlook prescriptive approaches with respective policies, we additionally present benchmark policies such as classical SAA or predictive approaches. By comparing the multiple policies, we can assess the properties of our approach and introduce the value of features as well as the value of outlook (Section 6.5.5).

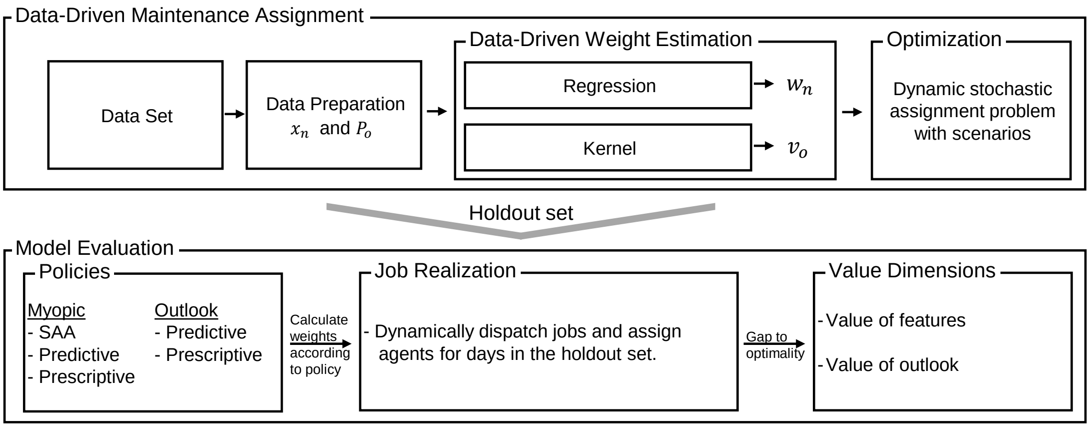

Figure 6.2: Overview of the numerical evaluation.

6.5.1 Data Set

For our research, we use databases of our cooperation partner that include disruption and production data over twelve months. Figure 6.3 provides an overview of the underlying relational data structure and annotates the connection to the organization's disruption management system.

The disruption management system handles the process of notifying a responding agent, as a worker detects a disruption and presses the **disruption hardware**. Each hardware device is associated with a device id and a workplace id that enables the allocation of disruptions in relation with the **workplace** and the **disruption log** tables. As the hardware is pressed, an event in the disruption log table is created with relational information to the device and a timestamp of disruption occurrence. Further, an agent from the **agent pool** is notified to assist in solving the job. As an agent solves the job caused by a disruption, the agent provides additional information to the system, and

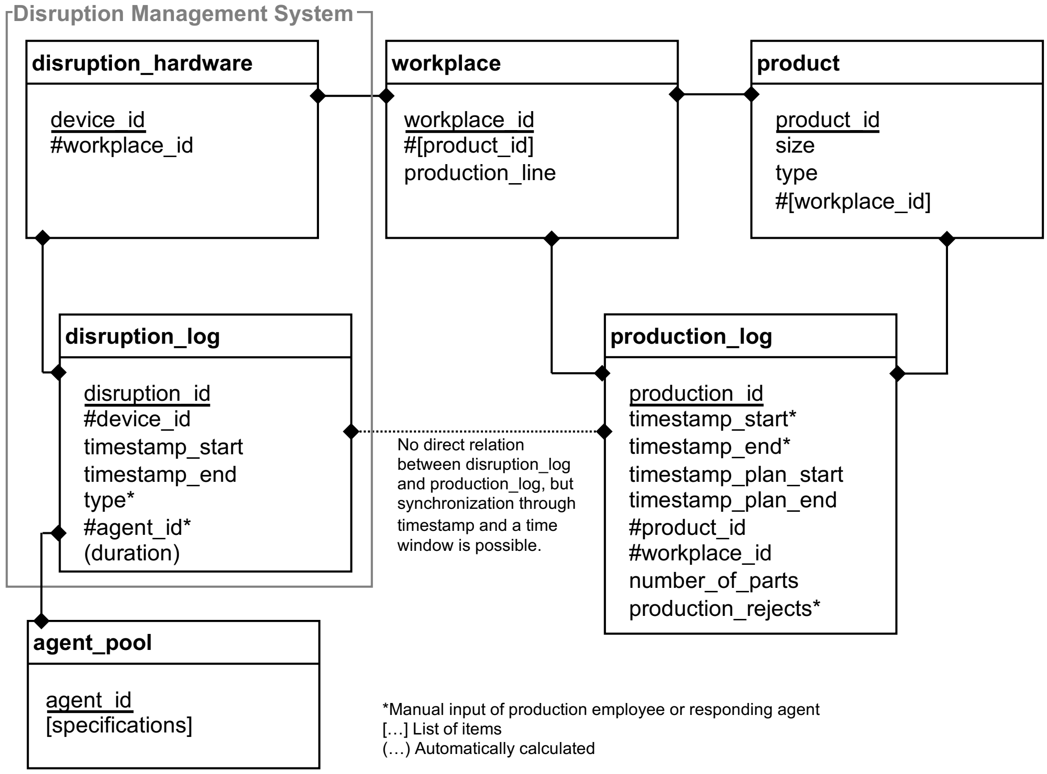

Figure 6.3: Overview of the existing relevant data structure.

a timestamp is added. We can calculate the solution duration and prepare features based on this information. Yet, the number of available features that may add to the assignment decision of an agent to the job caused by a disruption is limited.

To augment the data with additional data, we use the relation between the workplace and the **product** table. Specific products can be produced at each workplace, and thus additional information about the product size and type can be obtained. Combining the workplace, product, and also the time information from the disruption log enables a relation to the **production log** table, which is part of the central production system and entails most production-relevant information. It is used to plan production and store production start and end times for each product, including the workplace and number of parts. During the production itself, the actual start and end times as well as production rejects[39] are recorded at each workplace. Naturally, this is a manual process and thus results in inaccuracies in recording the respective timestamps. For instance, instead of logging each production step in the system, it is common to manufacture a product and afterward provide the production information to the system or only provide information about the first and last production step. This implies that the start and end timestamps may not reflect the current production situation. Due to this, we subsequently present the implementation of a time window to extract valuable features from the relational data structure.

6.5.2 Data Preparation

We aim to transform the raw data into features $\boldsymbol{x}_n$ and sequences P_o that we can leverage for the subsequent data-driven weight estimation. To this end, we combine and prepare the available data to subsequently calculate duration w_n and outlook weights v_o (Figure 6.4).

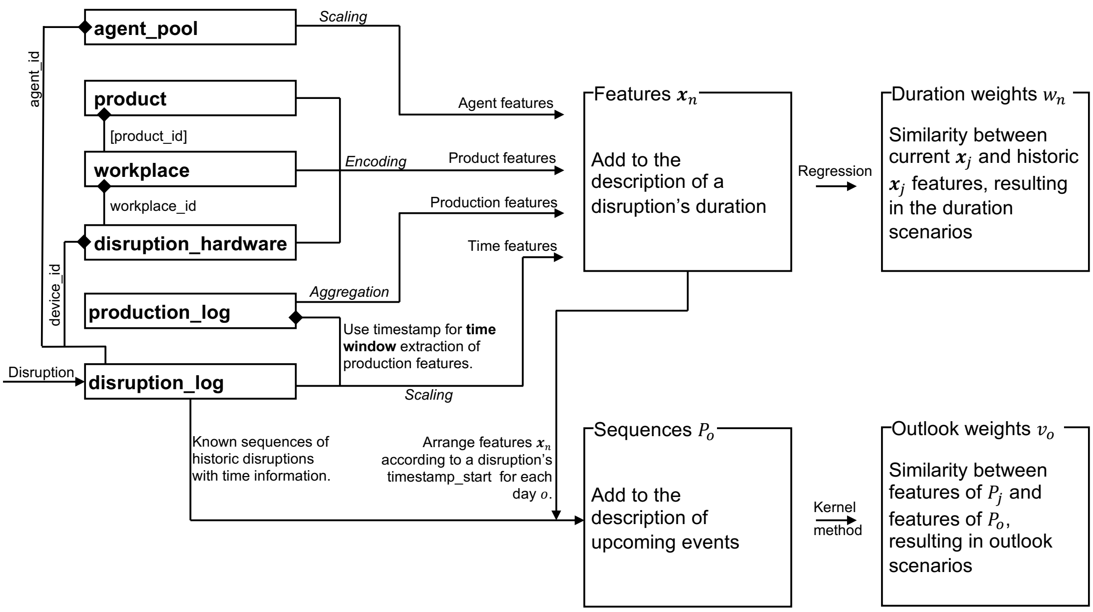

Figure 6.4: Workflow for the combination of data tables (rhombus) and the data preparation (*italics*) for the calculation of duration weights w_n and outlook weights v_o.

[39] Result when a product has a damage and there is no use of continuing the production or manufacturing process.

For each disruption, we have a corresponding event (job) in the disruption log with information about the time, disruption hardware device, and responding agent. We use these to combine the available tables to prepare $\boldsymbol{x}_n$, which consists of

$$\boldsymbol{x}_n = \begin{cases} x_{n0}, ..., x_{n4}, & \text{— Product features (size, type, workplace, ...)} \\ x_{n5}, ..., x_{n28}, & \text{— Production features (quantity per workplace, ...)} \\ x_{n29}, ..., x_{n32}, & \text{— Time features (day, weekday, hour, ...)} \\ x_{n33}, ..., x_{n38} & \text{— Agent features (specifications).} \end{cases} \tag{6.18}$$

Based on the disruption hardware's device id, we can use the relational data structure to augment the **product features**. For instance, we map the device id to a workplace and then to the products associated with a workplace. From these, we obtain features, such as the type or size of a product and the workplace where the disruption was triggered (hardware was pressed)[40] or the corresponding production line. We ordinal encode these based on internal information about dependencies between the distinct product sizes, types, and associated workplaces.

Because of the absence of a direct relation between the disruption log and the production log, we leverage both logs' time information to obtain relevant events that we aggregate to the **production features**. To do so, we use a disruption event's start timestamp combined with a time window[41] to select included production events. We then aggregate the resulting sub-log per workplace and sum the production quantities, reflecting the workplaces' utilization within the time window. We use the disruption log data itself and extract **time features** from a disruption's start timestamp. For instance, we extract and scale the day (of the month), weekday, or hour for the resulting features $\boldsymbol{x}_n$. The features described so far are available when a disruption occurs and can be leveraged for the weight calculation.

[40]Note that the workplace where a disruption is recognized usually is not the workplace where the cause of a disruption results from. For instance, a damage can be recognized some production steps (and workplaces) after the actual damage.

[41]With respect to the mean production duration, the time windows was chosen with five hours.

There are additional features that the responding agent provides after completing a maintenance job. Thus, these features cannot be considered for weight calculation, but we leverage them to overcome the limitation of unknown realization times for the assignment of other agents. We obtain the assigned agent and associated specifications from the agent pool based on the agent id. For the generation of the **agent features**, we particularly focus on the specifications, which reflect how specialized an agent is for specific job types.

For the preparation of the sequences P_o, we leverage the prepared features in combination with the time information of the disruption log, resulting in the sequences

$$P_o = \begin{cases} \boldsymbol{x}_0, & - \text{ Features from first disruption} \\ \boldsymbol{x}_1, & - \text{ Features from second disruption} \\ \dots, & \\ \boldsymbol{x}_{J(o)}, & - \text{ Features from last disruption.} \end{cases} \tag{6.19}$$

The feature vectors $\boldsymbol{x}_n$ are arranged with respect to the occurrence time (timestamp start) of the disruptions during a day. This results in O sequences that reflect the historic occurrences of disruptions over the course of each individual day o.

6.5.3 Synthetic Data Set for Generalists and Specialists

We draw on a data set enriched with synthetic data for a detailed evaluation of the solution approach. There are two reasons for this. The fundamental problem of causal inference arises in problems with predictive models where decisions affect the uncertainty of realizations and result in the absence of counterfactual information. However, this statistical bias has been shown to have little effect on predictive analytics and still leads to good decision-making (Bertsimas and Kallus 2020b). Further, data collection as ground truth for the trained model affects decision-making. In the considered data, the historical process of disruption management is represented. Agents are assigned to jobs without considering information about the causing disruption. Because no information is taken into account for the (historic) assignment decision, usually generalists are assigned that can handle a wide range of tasks. Therefore, it is not possible to directly observe in the data how the problem-solving of agents of different specialization shapes up for a disruption. As a consequence, we enrich the data with synthetic data on solution durations for agents of different specializations. This results in a data set with solution duration information for each observation and agent.

We show that this approach is compelling in practice, as it can derive direct recommendations for decision-making. A direct benefit arises from the fact that it would be beneficial to directly assign a specialist because specialists usually solve jobs of their subject faster. In contrast, the solution duration increases massively if a specialist is assigned to a wrong job (type). To preclude this risk, primarily generalists are assigned historically what led to the reasons for using a synthetic data set. Only taking generalists into account has been reasonable for the historical process; however, the proposed approach enables an organization to consider specialists and deal with uncertainties and consider the up- and downsides of a specialists' assignment. To this end, we envision a data set that includes information about the solution durations for specialists (correct and misassignment) and generalists, where features are available for the realized assignments and the counterfactuals.

We set up an empirical study where agents with different specificity levels (according to their type, e.g., maintenance or logistics specialist) solve jobs of their subject and jobs from other fields. Based on these empirical duration measurements, we calculate mean scaling factors[42] For each agent type (e.g., logistics, maintenance), we duplicate the historical observations and scale the solution duration distributions according to the empirical results. The resulting synthetic data set enables a data-driven weight estimation with the product, production, time, and an agent's type as input features.

6.5.4 Data-Driven Weight Estimation

For the data-driven weight estimation, we follow the myopic and outlook solution approaches from Section 6.4 and choose weight specifications with respect to Bertsimas and Kallus (2020a) for the prescriptive models. In addition

[42] A correct assigned specialist needs a fourth solution duration, compared to a generalist. A misassigned specialist needs the double solution duration compared to a generalist.

to the prescriptive policies, we use SAA and predictive approaches as benchmarks and present how to express the respective policies through weights.

We use a 70 % subsample of the data to train the models for predictive and prescriptive tasks.[43] The remaining 30 % are split into a test set to evaluate train-test performance (20 %) and the holdout set $P_{O^{\mathrm{h}}}$ (10 %). For the model's basic training, we follow the workflow proposed by Brownlee (2018) to perform hyperparameter tuning for each model. Based on a grid search, common algorithm-specific parameter search spaces are processed. After training a model with the parameters of the current tuning run, the metrics and estimated weights for the test sample are calculated and stored. We repeat this process for each model with 15 different random initializations to ensure robustness. Finally, we evaluate the results of the hyperparameter-tuning and select the best-performing models and parameters for the predictive models and the prescriptive weight estimation.

Duration Weight Estimation

The duration weights approximate the current job's (unknown) solution duration. A common form for the approximation is to use a SAA and weight the historic samples with respect to their occurrence:

$$w_n^{\mathrm{mSAA}} = \frac{1}{|N|} \quad \forall n \in N \tag{6.20}$$

We use the weights w_n^{mSAA} to solve the myopic objective (6.12) and refer to it as the myopic SAA (mSAA) policy.

Going further and leveraging the available features more, we can make use of *predictive* models. To use the best-performing model, we compare the performance of regression trees (Breiman et al. 1984), k-nearest-neighbors (Trevor and Robert 2001, chap. 13), and random forest (Breiman 2001) algorithms and report the mean absolute error (MAE), mean squared error (MSE), and root mean squared error (RMSE; Chai and Draxler 2014) as metrics (Table 6.1) while using MSE as training metric.

[43] The subsamples are based on the original data set to omit overfitting for distinct types on the synthetic data set.

Table 6.1: Comparison of regression algorithm performance.

Model	MAE	MSE	RMSE
Regression Tree	0.089	0.071	0.226
k-Nearest Neighbor	0.236	0.105	0.324
Random Forest	**0.081**	**0.024**	**0.154**

As summarized in Table 6.1, the random forest (RF) algorithm outperforms the other models. To this end, we choose the RF algorithm as the myopic predictive policy model. To reflect the RF's prediction through a weight, we use:

$$w^{\mathrm{mPred}} = \{\mathcal{R}^{\mathcal{T}}(\boldsymbol{x}_j) : 1\}. \tag{6.21}$$

with $\mathcal{R}^{\mathcal{T}}(\cdot)$ as the prediction of the $\mathcal{T}$ estimators.

Aligned to the model evaluation and the empirical results of Bertsimas and Kallus (2020a) and Notz and Pibernik (2021), we also leverage the RF for the myopic prescriptive weight function

$$w_n = \frac{1}{\mathcal{T}} \sum_{t \in \mathcal{T}} \frac{\mathbb{1}\big[\mathcal{R}^t(\boldsymbol{x}_j) = \mathcal{R}^t(\boldsymbol{x}_n)\big]}{|\{i : R^t(\boldsymbol{x}_i) = R^t(\boldsymbol{x}_j)\}|} \quad \forall n \in N \tag{6.22}$$

with the partition rules R^t and predictions $\mathcal{R}^t$ for the $t^{\mathrm{th}} \in \mathcal{T}$ decision tree.

Outlook Weight Estimation

The outlook weights relate to the simulations of potential upcoming disruptions. In contrast to the random forest approach (6.22), we can not account for the similarity based on supervised local learning techniques because for the sequences P_o no loss function can be established. Providing labels or quantities for the individual days o is challenging and, particularly from a practical perspective, not a useful task. Having no quantities available, we focus on kernel functions $K(\cdot)$ that compute the similarity of vectors and use it to calculate weights for a wSAA (Bertsimas and Kallus 2020a). This results in a outlook prescriptive policy (oPres) with weights v_o^{oPres}:

$$v_o^{\mathrm{oPres}} = \frac{K\big((P_o - P_{o_j})/h\big)}{\sum_{i \in O} K\big((P_i - P_{o_j})/h\big)} \quad \forall o \in O. \tag{6.23}$$

The kernel function's input is the calculated distance between the current disruption's sequence P_{o_j} and the historic sequences P_o, which we can visualize as a kind of similarity matrix between the distinct observation days (Figure 6.5). For the definition of P_{o_j}, we combine the current day's (o_j) known features $P^-_{o_j}(t_j)$ and $P^+_{o_j}(t_j)$ as an estimate about the upcoming production, based on the day's production plan. We compute the distance between all the observation days O aligned to the kernel distance calculation based on the euclidean norm. Note that we currently calculate the distance for each feature individual and average the resulting distances, whereas this could be extended in future research to account for spatial and temporal dependencies. With respect to the current implementation, a small distance implies more similar vectors and a lighter color between two observations.

To only incorporate the most similar samples, the kernel function considers a bandwidth $h > 0$, which restricts the observations and thus affects the number of resulting scenarios. As the features in the sequences are scaled

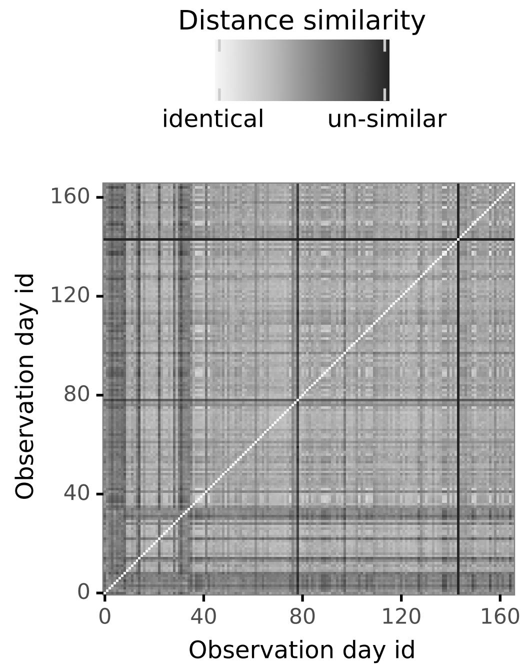

Figure 6.5: Similarity matrix based on distance calculation between outlook scenario samples.

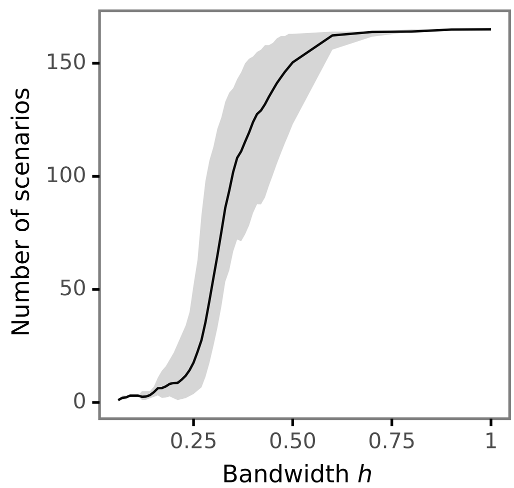

Figure 6.6: Number of outlook samples for the kernel-based weight calculation affected by the bandwidth.

($\boldsymbol{x} \in [0,1]$), $h = 1$ implies that all days are considered, whereas for $h < 1$ the number of scenarios is reduced. To assess the impact of the bandwidth on the resulting outlook scenarios, Figure 6.6 provides a sensitivity analysis of the number of scenarios with respect to h. Clearly, the number of scenarios increases for increasing h, whereas we depict a certain non-linearity due to the distance similarities from Figure 6.5.

Bertsimas, McCord, and Sturt (2019) present a number of kernel functions and additionally proof the asymptotic optimality for those: The naïve kernel $K(\boldsymbol{x}_j) = \mathbb{1}_{||\boldsymbol{x}_j|| \leqslant 1}$, the epanechnikov kernel $K(\boldsymbol{x}_j) = (1 - ||\boldsymbol{x}_j||^2)\mathbb{1}_{[||\boldsymbol{x}_j|| \leqslant 1]}$, and the tri-cubic kernel $K(\boldsymbol{x}_j) = (1 - ||\boldsymbol{x}_j||^3)^3\mathbb{1}_{[||\boldsymbol{x}_j|| \leqslant 1]}$. While the naïve kernel weights all samples within the bandwidth h equally, the epanechnikov or tri-cubic kernels account for the between feature vectors. Yet, the polynomial consideration of feature vectors' distances in combination with increasing bandwidth parameter h might result in discontinuous weights, which is relevant for searching for an optimal prescriptive bandwidth parameter. The weights are scaled to the sum of kernel distances, which implies discontinuous behavior as additional features are included for higher bandwidths h across the ranks O^r (Figure 6.7). For the sake of traceable results, this limits the clarity. Although the naïve

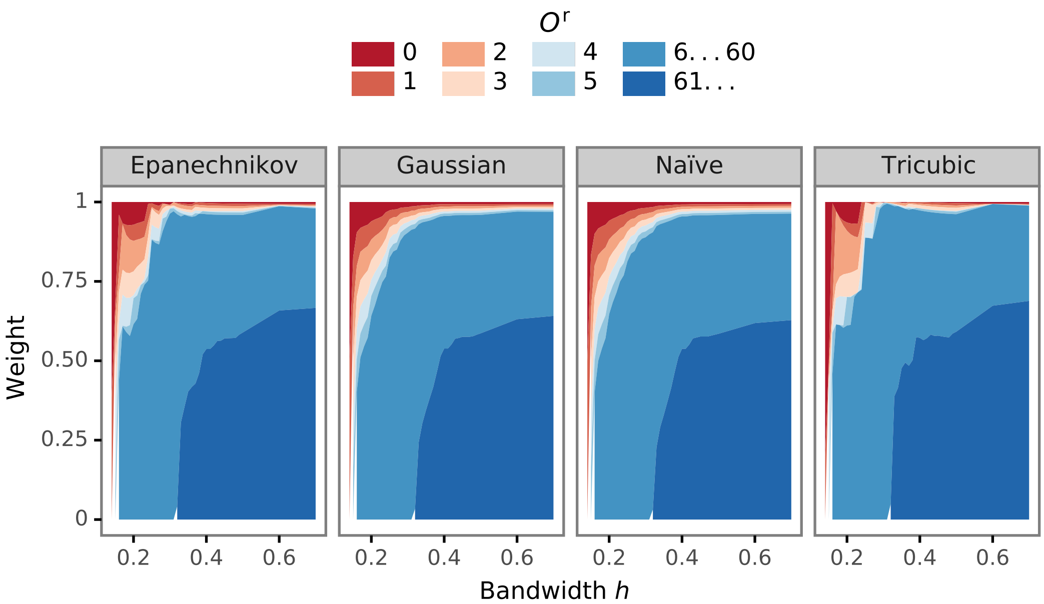

Figure 6.7: outlook weights for bandwidth and kernel scenarios per rank.

kernel does not have this shortcoming, the distance information within the bandwidth h is largely unused due to the equal weighting of observations.

To overcome these shortcomings, we propose the weight estimation with a Gaussian kernel $K(\boldsymbol{x}_j) = e^{(-\frac{1}{2}||\boldsymbol{x}_j||^2)}\mathbb{1}_{[||\boldsymbol{x}_j||\leqslant 1]}$. Following the kernel definition, the features' distance is considered for the weight calculation and the kernel should result in continuous weights. To prove the applicability of the Gaussian kernel: For the Gaussian kernel, theorem 16 of Bertsimas, McCord, and Sturt (2019) holds in combination with theorem 3 of Walk (2010), where they explicitly propose the Gaussian kernel as an example. To estimate the optimal bandwidth for the prescriptive approach, we numerically evaluated the data-driven maintenance assignment model and chose a bandwidth of $h = 0.28$ for the prescriptive policies.

To additionally assess the value of considering and weighting the potential realization paths, we additionally applied a kind of prediction weighting for upcoming jobs:

$$v_o^{\mathrm{oPred}} = \mathbb{1}[\min(P_o - P_{o_j})] \quad \forall o \in O \tag{6.24}$$

and limit the number of the outlook scenarios only to the most similar observation with the smallest distance to the current observation sequence.

6.5.5 Numerical Evaluation

For the numerical evaluation, we leverage the synthetic data set and evaluate myopic and policies with outlook information on the holdout set. For each job $j \in P_o$ an agent is assigned, with previous assignments having an impact on the system's future state by means of the agents' availability. We calculate a days' total downtime with respect to the ex-post optimal policy $\boldsymbol{q}_j^{\mathrm{opt}}$, resulting in the gap to optimality λ^i:

$$\lambda^i = \frac{1}{|P_{O^{\mathrm{h}}}|} \sum_{o \in O^{\mathrm{h}}} \sum_{j \in P_o} \frac{d_j(\boldsymbol{q}_j^i)}{d_j(\boldsymbol{q}_j^{\mathrm{opt}})} \quad \forall i \in I. \tag{6.25}$$

Value of Features

Our initial analysis considers myopic solution approaches with the corresponding gaps to optimality λ^i to highlight *how the integration of features* contributes to the assignment decision and yields the *value of features* (Figure 6.8). The results suggest that the gap to optimality can be reduced by about $\Delta\lambda = 10\,\%$ by using myopic solution approaches (prescriptive) instead of a common SAA. The myopic prescriptive approach leads (in median) to the lowest gap to optimality and has a low spread, particularly compared to the myopic predictive approach. Conversely, the myopic predictive approach leads to few competitive results that even outperform the myopic prescriptive approach in the tail. On the other hand, there is a large spread, and also some results perform worse than the SAA approach.

Considering the initial results from a real-world perspective implies that the myopic predictive approach is, in some cases (and over the day often) right with the predictions, resulting in competitive results. However, there are also cases where the predictions tend to be wrong and lead to miss-assignments, such as the results with $\lambda^{\mathrm{mPred}} > 6$. Such miss-assignments increase the downtime for a current job and affect upcoming jobs (the system's future state) through the availability (or non-availability) of agents due to previous assignments. In contrast, both the SAA and the myopic prescriptive policies lead to less of such outliers. This is due to a more frequent assignment of generalists that can solve each job, but with a slightly longer downtime. Yet, compared with the downtime due to a miss-assignment, assigning generalists can provide more stable results. In contrast, it then also limits the potential for low gaps to optimality, as depicted from the SAA policy.

The myopic prescriptive approach combines the properties of the myopic SAA and predictive approaches through accounting not only for the features but also associated uncertainty, which is considered through the prescriptive weight calculation. Instead of trusting the random forest's prediction, the myopic prescriptive approach weights historic features $\boldsymbol{x}$ and their known downtimes based on the similarity with the current features $\boldsymbol{x}_j$. With more similar features in a leaf node, the weights facilitate the assignment decision of specialists, whereas less similar features tend to a generalist's assignment. Beyond the performance of the approaches, this may also impact the practical adoption of the approaches, as discussed in Section 6.6.

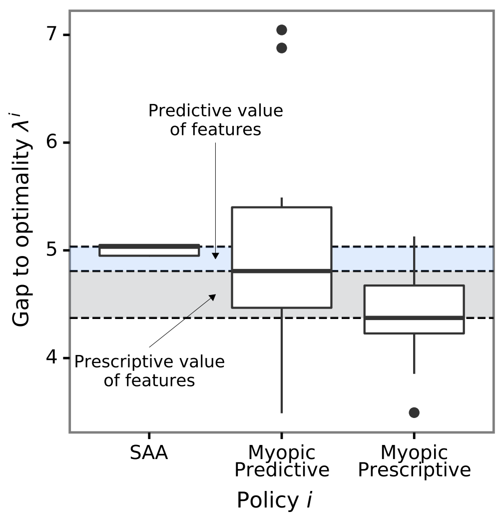

Figure 6.8: Gap to optimality of myopic solution approaches with only information about the current job, resulting in the predictive and prescriptive value of features.

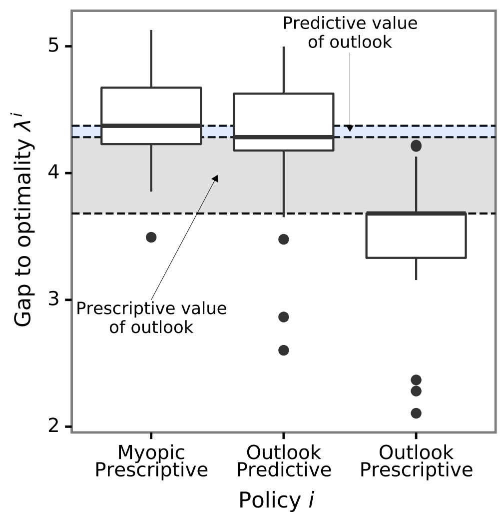

Figure 6.9: Gap to optimality between myopic prescriptive and outlook prescriptive approaches reveals the prescriptive value of outlook.

Value of Outlook

Only information about the current disruption is taken into account for any of the myopic approaches, and the impact on the system's future state is neglected. However, with disruptions occurring over a day, the system's future state should be considered, as pointed out in Sections 6.3 and 6.4. To this end, we leverage the solution approach with outlook information and shed light on the *value of outlook*. As we previously evaluated the myopic prescriptive policy to be competitive, we subsequently leverage it as the benchmark for the solution approaches with outlook (Figure 6.9). The outlook information leads to an additional improvement of the assignments and reduces the gap to optimality for the outlook predictive and prescriptive approaches. Yet, taking the outlook information predictively into account leads to results close to the myopic prescriptive approach. Thus, the associated predictive value of outlook is limited. In contrast, the prescriptive outlook approach significantly adds to the assignment decision and the resulting prescriptive value of outlook. In addition to the improved median performance of the outlook approaches, we depict some (positive) outliers with low gaps to optimality. These result from correct assignment decisions, combined with limited miss-assignments over the course of a day.

To account for the differences between the predictive and prescriptive outlook approaches, we consider the nature of the problem. In the real-world problem, the historic observations sequences o have a certain similarity but are usually not (perfectly) similar, as we could already depict from the distance matrix (Figure 6.5). Thus, only considering the closest vector (predictive) can still result in future states where a (potentially) upcoming disruption is not considered because the most similar day has no information about such. Instead, the most similar (upcoming) realization paths are considered and weighted according to their similarity with the prescriptive outlook approach. Taking this uncertainty into account enables the model to estimate the uncertain future. For instance, if there is a high probability that a specialist is promptly needed for an upcoming job, the model would not assign the respective agent. This could then prevent the miss-assignment of another agent for the upcoming job. However, a single day is usually not that similar, that it adds much to the estimation of future jobs. But, including (potentially upcoming) realization paths and the associated uncertainty in the current decision facilitates assignments in the face of how the system's future state may be impacted and clearly adds to the assignment decision.

6.6 Discussion and Implications

In the course of our research and cooperation, interesting discussions have arisen around the topics of algorithm performance, deployment, and practical analytics adoption. We want to examine these in the context of a critical analysis and additionally provide a critical assessment from a generalization perspective.

6.6.1 Upcoming Event Forecasting

We investigated data-driven approaches for a maintenance assignment problem and estimated the consequences of decisions on current and future downtimes. The proposed method's key contribution is to overcome point-wise forecasts or complex sequence-based forecasts and instead leverage a wSAA approach to consider historic observation sequences. In doing so, we show the discriminative value of features and outlook information as well as prescriptiveness from a data-driven approximate dynamic stochastic programming approach for maintenance job assignment. Basically, this involves the challenge of assigning agents to jobs where both a current job's completion duration and upcoming jobs (outlook) are unknown. Yet, the current decision affects the system's future state and should therefore be considered. Foresight is a challenging task here, as either a point forecast or a sequence of (potentially) upcoming jobs is required. In parallel, no loss function can be provided for the similarity between days. Conventional point or sequence forecasting methods such as ARIMA or exponential smoothing become increasingly complex for complex forecasting tasks with trends, many cycles, non-stationarity, or random inference. The number of parameters to be estimated grows with the number of outputs. Pattern similarity-based methods overcome this obstacle and predict the output vector at once (Dudek 2015). Models based on similarity extract regularities and patterns and extrapolate detected relationships in data to simplify the forecasting problem (Dudek and Pełka 2021). In particular, similarity-based methods are advantageous for repetitive, similar-shape cycles in time series (Dudek and Pełka 2021), such as in the production environment.

We exploit similarity-based local machine learning methods for the duration and outlook weight estimation. The assumption behind the algorithm selection relates to Bertsimas and Kallus (2020a), who demonstrate that weight functions of machine learning methods can be considered to approximate the conditional distribution of uncertain parameters in an objective function. This approach differs from the approaches proposed in Ban, Gallien, and Mersereau (2019) and Sen and Deng (2018), where the predictive model generates a point forecast of the uncertainty parameter with the help of additional data. Sen and Deng (2018) use stochastic learning (e.g., ARIMA) to model uncertainty in stochastic programming, whereas Ban, Gallien, and Mersereau (2019) propose the residual tree method to approximate uncertainty in multi-stage stochastic programs with the help of contextual data and prove asymptotic optimality of the approach. Kannan, Bayraksan, and Luedtke (2020) name these solution approaches empirical residuals-based SAA because the different scenarios are obtained by adding residuals observed during training to the point forecast. They propose leave-one-out residuals and prove a better approximation of conditional distribution for small sample sizes. The various SAA frameworks proposed in the literature are appropriate for applying parametric, nonparametric, or semi-parametric regression methods. Bertsimas and Kallus (2020a) demonstrate nonparametric regression techniques in a re-weighted SAA problem and show convergent approximations. Kannan, Bayraksan, and Luedtke (2020) note that parametric or semiparametric might outperform nonparametric approaches if the functional dependence of the covariate vector and the uncertain model parameters is a good approximation of the true dependence.

In the dynamic procurement problem, of Ban, Gallien, and Mersereau (2019) or the inventory management problem, (Sen and Deng 2018; Bertsimas and Kallus 2018) the uncertain parameter is the demand, described as a random vector, and a linear relationship between the demand and contextual data is assumed. In contrast to predicting demands, time-series sequence forecasts of upcoming maintenance jobs are hardly feasible. Thus, we rely on the similarity of days in the production environment to be prepared for the occurrence of similar situations, similarly to Bertsimas and Kallus (2018). Additionally, the occurrence of a maintenance job occurs in a random manner, and the agent assignment decision of current downtimes affects the realizations of upcoming assignments. Bertsimas and Kallus (2018) incorporate the consequences of decisions on the realization of random variables and show that even for non iid additional data, the asymptotic optimality is given under mild conditions.

We extend the respective framework to a real-world multistage problem in a production context, which includes two types of dependent uncertainties. Due to our data-driven maintenance assignment approach, we can improve the assignment policy and highlight the predictive and prescriptive values of features and outlook information.

6.6.2 Generalization of Data-Driven Approach

In our research, we considered applying the data-driven maintenance assignment approach for handling disruptions in production. However, the presented method is transferable to other domains, e.g., service hotlines, ticket systems, or fulfillment problems. Existing data-driven approaches mainly focus on three application areas (Mišic and Perakis 2020): *supply chain management* in the context of location (Glaeser, Fisher, and Su 2019), omnichannel (Acimovic and Graves 2015), and inventory decisions (Notz and Pibernik 2021; Ban and Rudin 2019), *revenue management* covering choice modeling and assortment optimization (Feldman, Paul, and Topaloglu 2019), pricing, and promotion planning (Ferreira, Lee, and Simchi-Levi 2016), and personalized revenue management (Baardman et al. 2020) and *healthcare operations* (Bertsimas et al. 2016; Bertsimas et al. 2017). To follow a data-driven similarity-based approach for optimization problems under uncertainty, there must be a certain similarity of events in the data (in our case, occurrence of disruptions), which may be leveraged for the data-driven weight estimation. If this is given, the presented method for taking upcoming events into account can be transferred to further problems by weighting historical samples.

Considering the nature of real-world problems, such a similarity is often at hand. For instance, in healthcare operations such as patient scheduling in hospitals, we face a similar problem with uncertainty about the duration of medical treatment and which treatments will occur during a day. With the objective of minimizing the patients' waiting time, the treatments should then be scheduled with respect to constraints for medical staff and probably also treatment rooms. Yet, the decision must not be made without any information. Beyond process information about internal processes, even medical patient records can be leveraged to estimate upcoming situations. Particularly for such a scheduling problem, the usefulness of point-wise predictions is clearly limited, as interdependencies between (potentially) upcoming patients and available staff and rooms must be taken into account. Leveraging the additional information to re-weight historic observations and using these re-weighted sequences can then add value to the scheduling.

For problems like the fulfillment problem, the approach may even be simplified. Basically, the problem describes the question from which stock keeping units (SKU) of a logistics network an ordered item should be shipped to achieve the maximum revenue (Acimovic and Graves 2015). The current uncertainty is reduced or eliminated because the ordered product is already known. The data-driven outlook weight estimation can reduce uncertainty about upcoming orders (jobs) and assign a warehouse (agent) from which to ship an item. This is only possible if information on the ordered products is available in the data. Although not yet evaluated, our method contributes to answering the question of Acimovic and Farias (2019): “Can existing methods from network RM [Revenue Management] be adapted, or can new tractable methods be developed to explicitly incorporate forecast error and other data for slow- and medium-moving SKUs when maximizing expected reward?". We build our methods on the body of relevant literature that develops general frameworks for problems in optimization under uncertainty (Bertsimas and Kallus 2020a; Bertsimas, McCord, and Sturt 2019; Ban and Rudin 2019) and show that these can be transferred to the problem of maintenance job assignment.

The aforementioned data-driven optimization methods under uncertainty have in common that they process similarities in data in a supervised way and face the challenge of deriving uncertainty sets from historical data. Future research might consider the perspectives of sophisticated deep learning techniques showing superior performance due to their ability of representing complex phenomena (Gambella, Ghaddar, and Naoum-Sawaya 2020). In many real-world manufacturing scenarios, most of the observed data is unbalanced due to the low occurrence of anomalies in production. Using a limited set of past observations to predict uncertainties could be inefficient, as the model is more prone to anomalies. As a result, supervised machine learning methods are hardly ever applied for anomaly detection in the production environment. Unsupervised models can deal with the high imbalance by comparing the similarity of inputs and calculating anomaly scores. Unsupervised machine learning methods, similar to the kernel methods applied for estimating the outlook weights, measure the closeness between data points, cluster centroids, and then set a threshold value to flag anomalies in the data. Future research could consider unsupervised machine learning methods like clustering or unsupervised deep classification for data-driven scenario-based optimization to construct uncertainty sets (Goerigk and Kurtz 2020). Beyond that, future research should leverage decision-making insights within the optimization problem. Ning and You (2019) propose closed-loop data-driven optimization frameworks, which allow feedback from the decision-making optimization problem to machine learning. Future avenues of research can build on existing general data-driven optimization frameworks, which can be transferred to different real-world settings.

Yet, a limitation of the approach is the evaluation based on a synthetic data set. Due to the fundamental problem of causal inference caused by the absence of counterfactual information, we generate a synthetic data set for evaluation. The results suggest that the proposed approach is promising, and thus, we look forward to improving the evaluation process. For instance, we suggest using nested conditional weights, as proposed by Bertsimas, McCord, and Sturt (2019) and shed light on alternative approaches for dealing with causal inference.

6.6.3 Prescriptive Analytics Facilitates Adoption

With the objective of deploying the proposed approach, we frequently discussed: *What facilitates the adoption of analytics methods in production?* One central point has emerged from the discussions: It is elementary that as few wrong decisions as possible are made, and thus not appropriate agents are assigned. This may still be overlooked in the case of individual incorrect assignments; however, if the frequency increases, it was assumed that the system's acceptance decreases, which was also confirmed in a conversation with production employees. This aligns with current research about analytics and information system adoption in industry. Jacobs et al. (2019) state that for device adoption, it must be perceived as useful and the expectations be confirmed. By means of the assignment decision, this relates to a few inappropriate assignments. Given the context of our results, this depicts a more distinct advantage of the data-driven maintenance assignment approach. For this purpose, we consider the evaluation results (Section 6.5) in the context of the distribution spread of the different policies. Basically, the prescriptive policies were necessary to outperform an always generalist policy. At the same time, the lower distribution spread of the prescriptive policies is striking. Accordingly, the number of non-optimally assigned agents must be small; otherwise, the distribution spread would be more similar to the predictive policies.

As more often a generalist is assigned by the prescriptive approach, and this decision is closer to the current decision heuristic; this should facilitate trust in the system, which is a key point for adoption (Yang, Lee, and Zo 2017; Kim and Song 2020). Another decisive factor is stakeholder management by means of clarification about potentials and risks (Brougham and Haar 2018;

Alkawsi et al. 2021). Yet, the risks are currently limited to a miss-assignment but could, in turn, influence trust in the system. Accordingly, it must be clear to the users that there may be not directly obvious assignments. This opens up a broad field of research opportunities that can be incorporated into future work.

6.7 Conclusion and Outlook

Advancements in machine learning and operations research revolutionize optimal decision-making in management science by using data predictions. We leverage those methodological advances of data analytics in optimization for a data-driven approximate dynamic stochastic programming approach for complex decision-making in the production environment. Our multistage approach involves the challenge of a maintenance job assignment problem and incorporates stochasticity of current and upcoming production disruptions. We focus on predicting the uncertain parameters of a dynamic stochastic optimization problem and demonstrate the gap to optimality between the different approaches. We formulate a data-driven weight estimation to estimate the uncertain parameters and take advantage of historic observations of disruptions and related features. Thereby, we assume an unknown joint distribution of historic observations and unknown joint distribution of the feature-dependent upcoming production disruptions. We apply three local machine learning algorithms for the current weight estimation, namely regression trees, k-nearest-neighbors, and random forests, and compare their performance on the maintenance job duration regression. The outlook weights are estimated with various kernel methods, namely naïve, gaussian, epanechnikov, and tricubic. Their performance for the prediction of upcoming maintenance jobs is evaluated respectively. We evaluate the performance of our data-driven maintenance assignment approach by demonstrating the discriminative value of features and outlook information on a data set that is prepared with respect to the results of an empirical study.

The future avenues of our work could also include different important directions. For instance, the temporal weighting of the disruptions for the derived uncertainty scenarios, giving more weight to near-time disruptions at the decision point, could improve predictive capabilities. Furthermore, the data-driven solution system could be expressed as a surrogate optimization model constructed from the available data (Kim and Boukouvala 2019). Another possibility is to enhance the data-drive approximate dynamic stochastic programming approach with feedback between machine learning and mathematical programming (Ning and You 2019). Such extensions could add to assignment decisions and, combined with the incorporation of outlook information, could enable improved data-driven decision-making.

7 Conclusion and Future Research Opportunities

Analytics-enabled information systems are a key driver for organizations to gain a competitive advantage. Such systems facilitate process-aware data collection, which results in big-data environments that enable gaining process insights and providing operational support. A respective example is the efficient handling of disruptions during production through an analytics-enabled information system, which is the subject of this thesis.

7.1 Summary

This dissertation contributes to the guiding research objective by *designing a prescriptive process monitoring system for disruption management in a production environment* and sheds light on process improvement potentials, considering all the analytics and information system stack levels.

7.1.1 Descriptive System

The first article (Chapter 2) focuses on the identification of success factors for process mining as a descriptive analytics approach for process improvement (**RO1**). To this end, we set up a multiple case study including observations, employee interviews, as well as expert and consultant interviews. Based on an a priori model that maps the L*Lifecycle model for process mining projects (van der Aalst et al. 2012b) and a process mining success model (Mans et al. 2013), we identify additional success factors for project mining projects in practice. We observe the application of process mining systems and interview hierarchies from the production team to C-level employees to point out the perceived value of process mining across the enterprise. Thereby, we consider the business value in terms of automational, informational, and transformational value to point out the spread of information in process mining projects. In the short term, primarily informational insights realize and directly involved employees are aware of the informational value. Because such insights do not monetize in turn, projects may be stopped due to the absence of returns. To overcome this, we extend the a priori model through additional success factors that we formalize in four key lessons that guide the business value generation in process mining projects.

Based on these findings, we leverage the action design research method in Chapter 3 to design a context-aware process mining information system that incorporates heterogeneous data to identify relevant processes, which contributes to **RO2**. The fundamental problem in the manufacturing case company is the high degree of customization of products and thus a high level of diversity in production processes. Facing challenges such as limited personal resources due to illness and work restrictions resulting from the pandemic situation, the processes must be optimized to cope with the requested demand despite pandemic restrictions. However, the identification with state-of-the-art solutions (without heterogeneous data) is a challenging task. The designed system leverages context data from a disruption management system to identify disruption-related process paths to overcome this. Thus, relevant processes, which often result in disruptions, can be identified through heterogeneous data and subsequently improved. Yet, there is only an user interface for comparing process models available, but no direct suggestion provided, which process should be addressed.

7.1.2 Predictive System

Building on the identified success factors, Chapter 4 contributes to **RO2**, as it showcases the development of a predictive process monitoring system for disruption handling, with a particular focus on the system design, evaluation, and roll-out for business value generation. The system leverages state-of-the-art machine learning algorithms for disruption type classification and duration prediction. Combined with additional organizational data sources (e.g., personnel availability, production plans) and a simple assignment procedure, the system generates the desired business value—initially through automation. In the long run, the transformational business value enabled by the system is likely to exceed the automational business value. This highlights the importance of tight integration of industrial analytics applications within business processes.

Business processes are a key component of organizations (Becker, Mathas, and Winkelmann 2009; Fischer et al. 2019) and provide additional information that can be leveraged for analytics tasks. However, current implementations are designed to deal with homogeneous or highly aggregated heterogeneous data, limiting its practical use in organizations. To overcome this, Chapter 5 proposes a novel five-phase method for predictive end-to-end enterprise process network monitoring leveraging multi-headed deep neural networks. The method facilitates the integration of heterogeneous data sources through dedicated neural network input heads, which are concatenated for a prediction. As the method guides the design of a multi-headed deep neural network from an end-to-end perspective, this work addresses **RO2** and **RO3** to enable predictive process monitoring on an organizational scale. For the evaluation, the engineered method is applied to a real-world use case with multiple context-aware event logs from different departments and additional disruption context information. The resulting multi-headed neural network shows superior performance for the specific use-case of multi-class disruption classification and thus highlights the applicability of the engineered method.

7.1.3 Prescriptive System

Leveraging improved predictions already results in business value in some situations. However, systems operate in complex organizational processes that limit predictions in practice. Such predictions may reduce the uncertainty about upcoming situations, but decisions still have to be made under uncertainty. In the presented real-world use case, where disruptions occur over the whole day, the decisions also affect the availability of responding agents for upcoming jobs. Thus, the number of potential scenarios massively increases and results in a prohibitively large state-space, limiting the problem's solution and thus the use of a respective system. Therefore, Chapter 6 sets out to address **RO4** focusing on a prescriptive solution approach for the complex assignment problem.

In particular, the prescriptive assignment of responding agents in the disruption handling process is analyzed. The decision to be made is whether a generalist or a specialist is assigned to a maintenance job caused by a disruption. While generalists can address a multitude of jobs, specialists usually solve an appropriate job faster but lack the solution of jobs out of their scope, resulting in more extended downtimes. The objective is to assign the agent that minimizes the total downtime of a day, considering previous and current disruptions as well as upcoming disruptions in an online setting. Clearly, in a complex manufacturing environment, the occurrence of disruptions is uncertain. In addition, a disruption's duration is uncertain, too, until the assignment decision is made and the responding agent knows the caused maintenance job. A data-driven approximate dynamic stochastic programming approach is developed and applied to a real-world use case to incorporate uncertainties in the assignment decision. It combines a random forest wSAA (current disruption's duration) and a kernel method wSAA (upcoming disruptions) to account for the uncertainties and approximate the optimal solution. Compared to predictive or single-uncertainty-prescriptive approaches, the presented approach significantly reduces downtime and facilitates the practical application of prescriptive process monitoring.

7.2 Future Research Opportunities

These findings point out a prescriptive process monitoring system's potential and assess how the distinct levels of the analytics and information system stack can add value. A general objective should be to enhance this value and realize it in practice. With a particular focus on the practical use of analytics-enabled systems, fields for future research arise.

7.2.1 Heterogeneous Data Sources

Facilitated through increasing digitalization, organizations increasingly operate in a big-data environment. As such environments evolve, organizations often lack an effective and unified data collection infrastructure, resulting in a set of dispersed and isolated databases (Naedele et al. 2015). Increasingly in manufacturing companies, there can additionally exist unstructured data, such as images and videos, which may help to detect quality defects in a manufacturing environment, or text documents, e.g., emails or production manuals, that can also add predictive value. Considering all these data and the potential value, combined with an organizational desire for improved models, research should focus on heterogeneous data sources.

The findings of this thesis emphasize this for each level of the analytics stack. Yet, the results of this thesis indeed rely on heterogeneous data, but mainly in a structured or semi-structured form. Future research should also include unstructured data and point out the resulting (additional) value (Wang et al. 2018) and shed light on challenges and limitations. This is facilitated through recent developments in sentiment analysis (Jiang et al. 2019) or text classification (Lin et al. 2021), which simplify the information extraction from textual data. Combining such models with structured and semi-structured organizational data sources might add even more value to analytics tasks and is another interesting field of research.

7.2.2 Federated Learning

Navigating in an organization's big-data environment can add value to analytics tasks but, in parallel, entails risks, such as aspects of data privacy and ownership (Li et al. 2020). Federated learning can be a crucial pillar to deal with such aspects (Yang et al. 2019) as it "involves training statistical models over remote devices or siloed data centers, such as mobile phones or hospitals while keeping data localized" (Li et al. 2020). The application is not limited to mobile phones and hospitals but includes the entire industry, supply chains, and retail. Suppliers or stores may not want to share their data but are interested in predictions with combined data. Or a machine tool manufacturer may want to add data from digitalized machine components, which are not publicly available. Federated learning can ensure data privacy and ownership for analytics tasks with shared knowledge in such scenarios.

The general field of federated learning provides many topics for future research, such as the design of federated learning systems with complex interactions between multiple actors and the shared system (Bonawitz et al. 2019). With increasing focus on operational application, efficient communication between the actors becomes more critical (Konecny et al. 2016) and should be analyzed. Future analyses can then focus on a value of collaboration and a value of information privacy, whether actors share data publicly or in a federated manner. A collaboration may add value in contrast to individually trained models but does not account for data privacy. Providing data privacy may come at the price of limited performance compared to models trained on all data and must be evaluated in the context of practical needs. Yet, federated learning focuses on predictive models (Li et al. 2020) but is not limited to that.

7.2.3 Prescriptive Analytics

With a particular focus on combining analytics and operations management models, the research field of prescriptive analytics arises, in general, as well as aligned with federated learning. Instead of building and solving complex operations management models through (manual designed) approximations, recently, the abundance of data, in combination with advances in computer science and statistics, has led to a shift in operations management research (Mišic and Perakis 2020). Using historical data and related features to support operational decisions is becoming increasingly essential to improve decisions and address natural uncertainties (Gambella, Ghaddar, and Naoum-Sawaya 2020; Kraus, Feuerriegel, and Oztekin 2020).

Facing volatile situations with increased uncertainties, prescriptive analytics approaches may thus become even more critical to provide reliable decision support. Concerning the fundamental principle of prescriptive approaches —typically, an unknown joint distribution of uncertain parameters and available data is assumed to model future uncertainty and approximate an optimization problem (Bertsimas, McCord, and Sturt 2019)—the challenge of reflecting on changed circumstances arises, too. Taking additional (heterogeneous) data into account and applying the methods depicted in this thesis may improve prescriptive decisions concerning changed circumstances and provide a wide field of future research. Due to the combination of analytics and operations management models, future research could evaluate differences between incorporating additional data in the individual models and shed light on how decisions are affected. Extending this to the field of federated learning and comparing the performance of (separated) prescriptive models with federated models for assessing whether decisions are affected by (not) sharing data would also add to future research.

7.2.4 Adoption Analytics Services

Besides increasing analytics research from a technical perspective, the behavioral aspect of analytics-enabled information systems and services provides interesting future directions. Technology acceptance models (TAM; Davis, Bagozzi, and Warshaw 1989) and extensions, such as the unified theory of acceptance and user technology (UTAUT) model (Venkatesh et al. 2003), can be used to explain the interaction between users and, e.g., information systems or services through behavioral constructs. Combining such constructs in a valuable way then adds to the explanation of why users adopt a system and can guide the design of future systems.

The change in adoption as simple information systems are extended with analytics capabilities, e.g., for data-driven decision-making, could be particularly interesting. Identifying driving constructs for analytics adoption would provide insights and potentially guide the development of future analytics-enabled information systems. An additional aspect of such analyses may focus on the effect of explainability of analytics models, which is in line with the growing research field of explanatory artificial intelligence (Doshi-Velez and Kim 2017; Meske et al. 2021; Arrieta et al. 2020). A second aspect relates to the adoption of prescriptive services. As pointed out by Grover, Kar, and Dwivedi (2020) “not much of work has been undertaken in the area of using AI on a real-time basis in operations management.” With systems that go beyond decision support towards automated actions, analyzing factors for adoption and its change, from support to action, would provide valuable research.

7.3 Practical Implications

Increasing digitalization of organizations facilitates using analytics-enabled information systems to improve business processes. This thesis sheds light on the design of such systems through a step-wise extension from descriptive to prescriptive analytics and, in parallel, points out the evolution from an informational to an operational system. Any system extension can provide value, but it must ultimately be realized in practice, which can be challenging for various reasons (Chapter 2). In particular, quantifying the value of informational-focused systems can be challenging (Chapter 3). Operational-focused systems can overcome this (Chapters 4 and 6) but necessitate interactions with complex processes, and thus the systems' design becomes particularly important (Chapter 5). Having addressed these challenges, there are two key pillars for the success I depicted during this thesis—patience, and communication. Either development or deployment of analytics-enabled systems takes time, and outcomes often do not realize immediately. Thus, it is fundamental to communicate realistic objectives from an analytics scope and timely perspective. Then the foundations for implementing analytics projects in organizations are in place, and the actual purpose of process improvement can be pursued.

List of Figures

List of Tables

Bibliography

Aalst, Wil van der, Arya Adriansyah, and Boudewijn van Dongen. 2012. "Replaying history on process models for conformance checking and performance analysis". *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2 (2): 182–192.

Abbasi, Ahmed, Suprateek Sarker, and Roger HL Chiang. 2016. "Big data research in information systems: Toward an inclusive research agenda". *Journal of the association for information systems* 17 (2): 3.

Abowd, Gregory D, Anind K Dey, Peter J Brown, Nigel Davies, Mark Smith, and Pete Steggles. 1999. "Towards a better understanding of context and context-awareness". In *International symposium on handheld and ubiquitous computing*, 304–307. Springer.

Acimovic, Jason, and Vivek F Farias. 2019. "The Fulfillment-Optimization Problem". In *Operations Research & Management Science in the Age of Analytics*, 218–237. INFORMS.

Acimovic, Jason, and Stephen C Graves. 2015. "Making better fulfillment decisions on the fly in an online retail environment". *Manufacturing & Service Operations Management* 17 (1): 34–51.

Adams, Jan Niklas, Sebastiaan J van Zelst, Lara Quack, Kathrin Hausmann, Wil M. P. van van der Aalst, and Thomas Rose. 2021. "A Framework for Explainable Concept Drift Detection in Process Mining". *arXiv arXiv:2105.13155*.

Aguirre, Santiago, Carlos Parra, and Marcos Sepulveda. 2017. "Methodological proposal for process mining projects". *International Journal of Business Process Integration and Management* 8 (2): 102–113.

Ahmad, Azizah. 2015. "Business intelligence for sustainable competitive advantage". In *Sustaining competitive advantage via business intelligence, knowledge management, and system dynamics*. Emerald Group Publishing Limited.

Alaei, Saeed, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2013. "The online stochastic generalized assignment problem". In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 11–25. Springer.

Albareda-Sambola, Maria, Maarten H van Der Vlerk, and Elena Fernandez. 2006. "Exact solutions to a class of stochastic generalized assignment problems". *European journal of operational research* 173 (2): 465–487.

Albareda-Sambola, Maria, and Elena Fernandez. 2000. "The stochastic generalised assignment problem with Bernoulli demands". *Springer* 8 (2): 165–190.

Alcacer, Vitor, and Virgilio Cruz-Machado. 2019. "Scanning the industry 4.0: A literature review on technologies for manufacturing systems". *Engineering Science and Technology, an International Journal*.

Alkawsi, Gamal Abdulnaser, Norashikin Ali, Abdulsalam Salihu Mustafa, Yahia Baashar, Hitham Alhussian, Ammar Alkahtani, Sieh Kiong Tiong, and Janaka Ekanayake. 2021. "A hybrid SEM-neural network method for identifying acceptance factors of the smart meters in Malaysia: Challenges perspective". *Alexandria Engineering Journal* 60 (1): 227–240.

Almeida Marodin, Giuliano, and Tarcisio Abreu Saurin. 2015. "Managing barriers to lean production implementation: context matters". *International Journal of Production Research* 53 (13): 3947–3962.

Andrews, Robert, Suriadi Suriadi, Moe Wynn, Arthur HM ter Hofstede, and Sean Rothwell. 2018. "Improving patient flows at St. Andrew's War Memorial Hospital's emergency department through process mining". In *Business Process Management Cases*, 311–333. Springer.

Armengaud, Eric, Michael Fruhwirth, Martin Rothbart, Martin Weinzerl, and Georg Zembacher. 2020. "Digitalization as an Opportunity to Remove Silo-Thinking and Enable Holistic Value Creation". *Systems Engineering for Automotive Powertrain Development*: 1–28.

Arrieta, Alejandro Barredo, Natalia Diaz-Rodriguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, et al. 2020. “Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI”. *Information Fusion* 58:82–115.

Augusto, Adriano, Raffaele Conforti, Marlon Dumas, Marcello La Rosa, Fabrizio Maria Maggi, Andrea Marrella, Massimo Mecella, and Allar Soo. 2018. “Automated discovery of process models from event logs: Review and benchmark”. *IEEE Transactions on Knowledge and Data Engineering* 31 (4): 686–705.

Ausiello, Giorgio, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela, and Marco Protasi. 2012. *Complexity and approximation: Combinatorial optimization problems and their approximability properties*. Springer Science & Business Media.

Baardman, Lennart, Setareh Borjian Boroujeni, Tamar Cohen-Hillel, Kiran Panchamgam, and Georgia Perakis. 2020. “Detecting customer trends for optimal promotion targeting”. *Manufacturing & Service Operations Management*.

Babiceanu, Radu F, and Remzi Seker. 2016. “Big Data and virtualization for manufacturing cyber-physical systems: A survey of the current status and future outlook”. *Computers in Industry* 81:128–137.

Baines, Tim S, Howard W Lightfoot, and John M Kay. 2009. “Servitized manufacture: practical challenges of delivering integrated products and services”. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 223 (9): 1207–1215.

Ban, Gah-Yi, Jeremie Gallien, and Adam J Mersereau. 2019. “Dynamic procurement of new products with covariate information: The residual tree method”. *Manufacturing & Service Operations Management* 21 (4): 798–815.

Ban, Gah-Yi, and Cynthia Rudin. 2019. “The big data newsvendor: Practical insights from machine learning”. *Operations Research* 67 (1): 90–108.

Barua, Anitesh, Prabhudev Konana, Andrew B. Whinston, and Fang Yin. 2004. “An Empirical Investigation of Net-Enabled Business Value”. *MIS Quarterly* (USA) 28, no. 4 (): 585–620.

Barua, Anitesh, Charles H Kriebel, and Tridas Mukhopadhyay. 1995. "Information technologies and business value: An analytic and empirical investigation". *Information systems research* 6 (1): 3–23.

Becker, Jörg, Christoph Mathas, and Axel Winkelmann. 2009. *Geschäftsprozessmanagement*. Springer-Verlag.

Becker, Till, and Wacharawan Intoyoad. 2017. "Context aware process mining in logistics". *Procedia Cirp* 63:557–562.

Becker, Till, Michael Lütjen, and Robert Porzel. 2017. "Process maintenance of heterogeneous logistic systems—a process mining approach". In *Dynamics in Logistics*, 77–86. Springer.

Beheshti, Amin, Boualem Benatallah, and Hamid Reza Motahari-Nezhad. 2018. "ProcessAtlas: A scalable and extensible platform for business process analytics". *Software: Practice and Experience* 48 (4): 842–866.

Benatallah, Boualem, Sherif Sakr, Daniela Grigori, Hamid Reza Motahari-Nezhad, Moshe Chai Barukh, Ahmed Gater, Seung Hwan Ryu, et al. 2016. *Process analytics: concepts and techniques for querying and analyzing process data*. Springer.

Benscoter, Bud. 2012. "How to Identify and Analyze Problems in Your Organization". In *The Encyclopedia of Human Resource Management*, 290–294. John Wiley & Sons, Ltd.

Bergstra, James, and Yoshua Bengio. 2012. "Random Search for Hyper-Parameter Optimization". *Journal of Machine Learning Research* 13:281–305.

Berti, Alessandro, Sebastiaan J van Zelst, and Wil M. P. van der Aalst. 2019. "Process mining for python (PM4PY): bridging the gap between process-and data science". *arXiv preprint arXiv:1905.06169*.

Bertram, Patrick, Christian Kränzler, Pascal Rübel, and Martin Ruskowski. 2020. "Development of a Context-Aware Assistive System for Manual Repair Processes-A Combination of Probabilistic and Deterministic Approaches". *Procedia Manufacturing* 51:598–604.

Bertsimas, Dimitris, Vishal Gupta, and Nathan Kallus. 2018a. "Data-driven robust optimization". *Mathematical Programming* 167 (2): 235–292.

— . 2018b. "Robust sample average approximation". *Mathematical Programming* 171 (1): 217–282.

Bertsimas, Dimitris, and Nathan Kallus. 2018. *From Predictive to Prescriptive Analytics*. arXiv: `1402.5481 [stat.ML]`.

— . 2020a. "From predictive to prescriptive analytics". *Management Science* 66 (3): 1025–1044.

— . 2020b. "The Power and Limits of Decision Making with Confounded Data: The Case of Pricing". *arXiv preprint*.

Bertsimas, Dimitris, Nathan Kallus, Alexander M Weinstein, and Ying Daisy Zhuo. 2017. "Personalized diabetes management using electronic medical records". *Diabetes care* 40 (2): 210–217.

Bertsimas, Dimitris, and Nihal Koduri. 2021. "Data-driven optimization: A reproducing kernel hilbert space approach". *Operations Research* 70 (1): 454–471.

Bertsimas, Dimitris, Christopher McCord, and Bradley Sturt. 2019. "Dynamic optimization with side information". *arXiv preprint arXiv:1907.07307*.

Bertsimas, Dimitris, Allison O'Hair, Stephen Relyea, and John Silberholz. 2016. "An analytics approach to designing combination chemotherapy regimens for cancer". *Management Science* 62 (5): 1511–1531.

Bertsimas, Dimitris, and Aurelie Thiele. 2006. "Robust and data-driven optimization: modern decision making under uncertainty". In *Models, methods, and applications for innovative decision making*, 95–122. INFORMS.

Bhuiyan, Nadia, and Amit Baghel. 2005. "An overview of continuous improvement: from the past to the present". *Management decision*: 761–771.

Bilgeri, Dominik, Heiko Gebauer, Elgar Fleisch, and Felix Wortmann. 2019. "Driving process innovation with IoT field data". *MIS Quarterly Executive* 18:191–207.

Blower, Sally M, and Hadi Dowlatabadi. 1994. "Sensitivity and uncertainty analysis of complex models of disease transmission: an HIV model, as an example". *International Statistical Review/Revue Internationale de Statistique*: 229–243.

Bolt, Alfredo, and Wil M. P. van der Aalst. 2015. "Multidimensional process mining using process cubes". In *Enterprise, Business-Process and Information Systems Modeling*, 102–116. Springer.

Bolton, Ruth N, Janet R McColl-Kennedy, Lilliemay Cheung, Andrew Gallan, Chiara Orsingher, Lars Witell, and Mohamed Zaki. 2018. "Customer experience challenges: bringing together digital, physical and social realms". *Journal of Service Management*.

Bonawitz, Keith, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konecny, Stefano Mazzocchi, H Brendan McMahan, et al. 2019. "Towards federated learning at scale: System design". *arXiv preprint arXiv:1902.01046*.

Borkowski, Michael, Walid Fdhila, Matteo Nardelli, Stefanie Rinderle-Ma, and Stefan Schulte. 2019. "Event-based failure prediction in distributed business processes". *Information Systems* 81:220–235.

Boyes, Hugh, Bil Hallaq, Joe Cunningham, and Tim Watson. 2018. "The industrial internet of things (IIoT): An analysis framework". *Computers in Industry* 101:1–12.

Bozkaya, Melike, Joost Gabriels, and Jan Martijn van der Werf. 2009. "Process diagnostics: a method based on process mining". In *2009 International Conference on Information, Process, and Knowledge Management*, 22–27. IEEE.

Braun, Virginia, and Victoria Clarke. 2006. "Using thematic analysis in psychology". *Qualitative research in psychology* 3 (2): 77–101.

Breiman, Leo. 2001. "Random forests". *Machine learning* 45 (1): 5–32.

Breiman, Leo, Jerome Friedman, Charles J Stone, and Richard A Olshen. 1984. *Classification and regression trees*. CRC press.

Breuker, Dominic, Martin Matzner, Patrick Delfmann, and Jörg Becker. 2016. "Comprehensible Predictive Models for Business Processes." *MIS Quarterly* 40 (4): 1009–1034.

Brinkkemper, Sjaak. 1996. "Method engineering: engineering of information systems development methods and tools". *Information and Software Technology* 38 (4): 275–280.

Brodersen, Kay Henning, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M Buhmann. 2010. "The balanced accuracy and its posterior distribution". In *2010 20th International Conference on Pattern Recognition*, 3121–3124. IEEE.

Brougham, David, and Jarrod Haar. 2018. "Smart technology, artificial intelligence, robotics, and algorithms (STARA): Employees' perceptions of our future workplace". *Journal of Management & Organization* 24 (2): 239–257.

Brown, Tom B, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. "Language models are few-shot learners". *arXiv preprint arXiv:2005.14165*.

Brownlee, Jason. 2018. *Better Deep Learning: Train Faster, Reduce Overfitting, and Make Better Predictions*. Machine Learning Mastery.

— . 2020. *Imbalanced Classification with Python: Better Metrics, Balance Skewed Classes, Cost-Sensitive Learning*. Machine Learning Mastery.

— . 2017. "Long short-term memory networks with python". *Machine Learning Mastery*.

Brunk, Jens, Johannes Stottmeister, Sven Weinzierl, Martin Matzner, and Jörg Becker. 2020. "Exploring the effect of context information on deep learning business process predictions". *Journal of Decision Systems*: 1–16.

Brynjolfsson, Erik. 1993. "The productivity paradox of information technology". *Communications of the ACM* 36 (12): 66–77.

Brynjolfsson, Erik, and Lorin M Hitt. 1998. "Beyond the productivity paradox". *Communications of the ACM* 41 (8): 49–55.

Brynjolfsson, Erik, Lorin M Hitt, and Heekyung Hellen Kim. 2011. "Strength in numbers: How does data-driven decisionmaking affect firm performance?" *Available at SSRN 1819486*.

Camm, Jeffrey D, James J Cochran, Michael J Fry, and Jeffrey W Ohlmann. 2020. *Business Analytics*. Cengage AU.

Canizo, Mikel, Isaac Triguero, Angel Conde, and Enrique Onieva. 2019. "Multi-head CNN–RNN for multi-time series anomaly detection: An industrial case study". *Neurocomputing* 363:246–260.

Cardin, Olivier. 2019. “Classification of cyber-physical production systems applications: Proposition of an analysis framework”. *Computers in Industry* 104:11–21.

Chae, Bongsug, David Olson, and Chwen Sheu. 2014. “The impact of supply chain analytics on operational performance: a resource-based view”. *International Journal of Production Research* 52 (16): 4695–4710.

Chai, Tianfeng, and Roland R Draxler. 2014. “Root mean square error (RMSE) or mean absolute error (MAE)?–Arguments against avoiding RMSE in the literature”. *Geoscientific model development* 7 (3): 1247–1250.

Chalabi, Zaid, David Epstein, Claire McKenna, and Karl Claxton. 2008. “Uncertainty and value of information when allocating resources within and between healthcare programmes”. *European journal of operational research* 191 (2): 530–539.

Chen, Hsinchun, Roger HL Chiang, and Veda C Storey. 2012. “Business Intelligence and Analytics: From Big Data to Big Impact.” *MIS quarterly* 36 (4): 1165–1188.

Cheng, Yang, and John Johansen. 2016. “The servitisation of manufacturing function: empirical case studies”. *International Journal of Manufacturing Technology and Management* 30 (6): 369–391.

Chollet, François. 2018. *Deep Learning mit Python und Keras: Das Praxis-Handbuch vom Entwickler der Keras-Bibliothek*. Frechen: MITP-Verlags GmbH & Co. KG.

Chu, Leon Yang, J George Shanthikumar, and Zuo-Jun Max Shen. 2008. “Solving operational statistics via a Bayesian analysis”. *Operations Research Letters* 36 (1): 110–116.

Clark, Shawn M, Dennis A Gioia, David J Ketchen Jr, and James B Thomas. 2010. “Transitional identity as a facilitator of organizational identity change during a merger”. *Administrative Science Quarterly* 55 (3): 397–438.

Côrte-Real, Nadine, Pedro Ruivo, and Tiago Oliveira. 2020. “Leveraging internet of things and big data analytics initiatives in European and American firms: Is data quality a way to extract business value?” *Information & Management* 57 (1): 103141.

Cunha Mattos, Talita da, Flavia Maria Santoro, Kate Revoredo, and Vanessa Tavares Nunes. 2014. “A formal representation for context-aware business processes”. *Computers in Industry* 65 (8): 1193–1214.

Cuzzocrea, Alfredo, Francesco Folino, Massimo Guarascio, and Luigi Pontieri. 2019. “Predictive monitoring of temporally-aggregated performance indicators of business processes against low-level streaming events”. *Information Systems* 81:236–266.

Da Silva, A. M., and M. C.C. Baranauskas. 2000. “The Andon system: Designing a CSCW environment in a lean organization”. In *Proceedings - 6th International Workshop on Groupware, CRIWG 2000*, 130–133. Institute of Electrical / Electronics Engineers Inc.

Dai, Zihang, Hanxiao Liu, Quoc V Le, and Mingxing Tan. 2021. “CoAtNet: Marrying Convolution and Attention for All Data Sizes”. *arXiv preprint arXiv:2106.04803*.

Daneshvar Kakhki, Mohammad, and Vidyaranya B Gargeya. 2019. “Information systems for supply chain management: a systematic literature analysis”. *International Journal of Production Research* 57 (15-16): 5318–5339.

Davenport, Thomas, and Jeanne Harris. 2017. *Competing on analytics: Updated, with a new introduction: The new science of winning*. Harvard Business Press.

Davis, Fred D, Richard P Bagozzi, and Paul R Warshaw. 1989. “User Acceptance of Computer Technology: A Comparison of Two Theoretical Models”. *Management science* 35 (8): 982–1003.

Denisov, Vadim, Elena Belkina, and Dirk Fahland. 2018. “BPIC’2018: Mining concept drift in performance spectra of processes”. In *8th International Business Process Intelligence Challenge*.

Dhuieb, Mohamed Anis, Florent Laroche, and Alain Bernard. 2016. “Context-awareness: a key enabler for ubiquitous access to manufacturing knowledge”. *Procedia CIRP* 41:484–489.

Di Francescomarino, Chiara, Marlon Dumas, Marco Federici, Chiara Ghidini, Fabrizio Maria Maggi, and Williams Rizzi. 2016. "Predictive Business Process Monitoring Framework with Hyperparameter Optimization". In *Advanced Information Systems Engineering*, ed. by Selmin Nurcan, Pnina Soffer, Marko Bajec, and Johann Eder, 361–376. Cham: Springer International Publishing.

Di Francescomarino, Chiara, Chiara Ghidini, Fabrizio Maria Maggi, Giulio Petrucci, and Anton Yeshchenko. 2017. "An eye into the future: leveraging a-priori knowledge in predictive business process monitoring". In *International Conference on Business Process Management*, 252–268. Springer.

Diba, Kiarash. 2019. "Towards a comprehensive methodology for process mining". In *Proceedings of the 11th Central European Workshop on Services and their Composition, Bayreuth*, 9–12.

Doan, AnHai, Alon Halevy, and Zachary Ives. 2012. *Principles of data integration*. Elsevier.

Dombrowski, Uwe, Thomas Richter, and Philipp Krenkel. 2017. "Interdependencies of Industrie 4.0 & Lean Production Systems: A Use Cases Analysis". *Procedia Manufacturing* 11:1061–1068.

Dong, Xin Luna, and Divesh Srivastava. 2013. "Big data integration". In *2013 IEEE 29th international conference on data engineering (ICDE)*, 1245–1248. IEEE.

Doshi-Velez, Finale, and Been Kim. 2017. "Towards a Rigorous Science of Interpretable Machine Learning". *arXiv preprint arXiv:1702.08608*.

Drohomeretski, Everton, Sergio E Gouvea da Costa, Edson Pinheiro de Lima, and Paula Andrea da Rosa Garbuio. 2014. "Lean, Six Sigma and Lean Six Sigma: an analysis based on operations strategy". *International Journal of Production Research* 52 (3): 804–824.

Dubey, Rameshwar, Angappa Gunasekaran, Stephen J Childe, David J Bryde, Mihalis Giannakis, Cyril Foropon, David Roubaud, and Benjamin T Hazen. 2020. "Big data analytics and artificial intelligence pathway to operational performance under the effects of entrepreneurial orientation and environmental dynamism: A study of manufacturing organisations". *International Journal of Production Economics* 226:107599.

Dudek, Grzegorz. 2015. "Pattern similarity-based methods for short-term load forecasting–Part 1: Principles". *Applied Soft Computing* 37:277–287.

Dudek, Grzegorz, and Paweł Pełka. 2021. "Pattern similarity-based machine learning methods for mid-term load forecasting: A comparative study". *Applied Soft Computing* 104:107223.

Al-Dulaimi, Ali, Soheil Zabihi, Amir Asif, and Arash Mohammadi. 2019. "A multimodal and hybrid deep neural network model for remaining useful life estimation". *Computers in Industry* 108:186–196.

Dumas, Marlon, Marcello La Rosa, Jan Mendling, and Hajo A Reijers. 2013. *Business process management*. Springer.

— . 2018a. "Introduction to business process management". In *Fundamentals of Business Process Management*, 1–33. Springer.

— . 2018b. *Fundamentals of Business Process Management*. Berlin, Heidelberg: Springer.

Dyer, Martin, and Alan Frieze. 1992. "Probabilistic analysis of the generalised assignment problem". *Mathematical Programming* 55 (1): 169–181.

Eck, Maikel L van, Xixi Lu, Sander JJ Leemans, and Wil M. P. van der Aalst. 2015. "PM^2: a process mining project methodology". In *International Conference on Advanced Information Systems Engineering*, 297–313. Springer.

Edgington, Theresa M, TS Raghu, and Ajay S Vinze. 2010. "Using process mining to identify coordination patterns in IT service management". *Decision Support Systems* 49 (2): 175–186.

Eggers, Julia, and Andreas Hein. 2020. "Turning Big Data into Value: A Literature Review on Business Value Realization from Process Mining." In *ECIS*.

Eggers, Julia, Andreas Hein, Markus Böhm, and Helmut Krcmar. 2021. "No Longer Out of Sight, No Longer Out of Mind? How Organizations Engage with Process Mining-Induced Transparency to Achieve Increased Process Awareness". *Business & Information Systems Engineering*: 1–20.

Ehrendorfer, Matthias, Juergen Mangler, and Stefanie Rinderle-Ma. 2021. "Assessing the Impact of Context Data on Process Outcomes During Runtime". In *International Conference on Service-Oriented Computing*, 3–18. Springer.

Eisenhardt, Kathleen M. 1989. "Building theories from case study research". *Academy of management review* 14 (4): 532–550.

Eisenhardt, Kathleen M, and Melissa E Graebner. 2007. “Theory building from cases: Opportunities and challenges”. *Academy of management journal* 50 (1): 25–32.

Elmachtoub, Adam N, and Paul Grigas. 2021. “Smart “predict, then optimize””. *Management Science*.

ElMaraghy, Hoda, Günther Schuh, Waguih ElMaraghy, Frank Piller, Paul Schönsleben, Mitchell Tseng, and Alain Bernard. 2013. “Product variety management”. *Cirp Annals* 62 (2): 629–652.

Emamjome, Fahame, Robert Andrews, and Arthur HM ter Hofstede. 2019. “A case study lens on process mining in practice”. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, 127–145. Springer.

Erasmus, Jonnro, Irene Vanderfeesten, Konstantinos Traganos, and Paul Grefen. 2020. “Using business process models for the specification of manufacturing operations”. *Computers in Industry* 123:103297.

Errarhout, A, S Kharraja, and C Corbier. 2016. “Two-stage Stochastic Assignment Problem in the Home Health Care”. 8th IFAC Conference on Manufacturing Modelling, Management and Control MIM 2016, *IFAC-PapersOnLine* 49 (12): 1152–1157.

Evans, James R, and Carl H Lindner. 2012. “Business Analytics: The Next Frontier for Decision Sciences”. *Decision Line* 43 (2): 4–6.

Evermann, Joerg, Jana-Rebecca Rehse, and Peter Fettke. 2016. “A deep learning approach for predicting process behaviour at runtime”. In *International Conference on Business Process Management*, 327–338. Springer.

Eversheim, Walter. 2013. *Prozessorientierte Unternehmensorganisation: Konzepte und Methoden zur Gestaltung „schlanker “Organisationen*. Springer-Verlag.

Fahey, Will, Paul Jeffers, and Paula Carroll. 2020. “A business analytics approach to augment six sigma problem solving: A biopharmaceutical manufacturing case study”. *Computers in Industry* 116:103153.

Farias Jr, Ismael R de, Ellis L Johnson, and George L Nemhauser. 2000. “A generalized assignment problem with special ordered sets: a polyhedral approach”. *Mathematical Programming* 89 (1): 187–203.

Feldman, Jacob, Alice Paul, and Huseyin Topaloglu. 2019. "Assortment optimization with small consideration sets". *Operations Research* 67 (5): 1283–1299.

Feldman, Jon, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pal. 2009. "Online ad assignment with free disposal". In *International workshop on internet and network economics*, 374–385. Springer.

Feng, Qi, and J George Shanthikumar. 2018. "How research in production and operations management may evolve in the era of big data". *Production and Operations Management* 27 (9): 1670–1684.

Ferreira, Kris Johnson, Bin Hong Alex Lee, and David Simchi-Levi. 2016. "Analytics for an online retailer: Demand forecasting and price optimization". *Manufacturing & Service Operations Management* 18 (1): 69–88.

Feurer, Matthias, Aaron Klein, Katharina Eggensperger, Jost Tobias Springenberg, Manuel Blum, and Frank Hutter. 2019. "Auto-sklearn: efficient and robust automated machine learning". In *Automated Machine Learning*, 113–134. Springer, Cham.

Fischer, Marcus, Florian Imgrund, Christian Janiesch, and Axel Winkelmann. 2019. "Directions for future research on the integration of SOA, BPM, and BRM". *Business Process Management Journal*.

Fisher, Marshall L, and Ramchandran Jaikumar. 1981. "A generalized assignment heuristic for vehicle routing". *Networks* 11 (2): 109–124.

Flath, Christoph M., and Nikolai Stein. 2018. "Towards a Data Science Toolbox for Industrial Analytics Applications". *Computers in Industry* 94:16–25.

Folino, Francesco, Massimo Guarascio, and Luigi Pontieri. 2012. "Context-aware predictions on business processes: an ensemble-based solution". In *International Workshop on New Frontiers in Mining Complex Patterns*, 215–229. Springer.

Gambella, Claudio, Bissan Ghaddar, and Joe Naoum-Sawaya. 2020. "Optimization problems for machine learning: A survey". *European Journal of Operational Research*.

Garcia, Cleiton dos Santos, Alex Meincheim, Elio Ribeiro Faria Junior, Marcelo Rosano Dallagassa, Denise Maria Vecino Sato, Deborah Ribeiro Carvalho, Eduardo Alves Portela Santos, and Edson Emilio Scalabrin. 2019. "Process mining techniques and applications – A systematic mapping study". *Expert Systems with Applications* 133 (): 260–295.

Ghobakhloo, Morteza, and Masood Fathi. 2020. "Corporate survival in Industry 4.0 era: the enabling role of lean-digitized manufacturing". *Journal of Manufacturing Technology Management* 31, no. 1 (): 1–30.

Giesecke, Kay, Gui Liberali, Hamid Nazerzadeh, J George Shanthikumar, and Chung Piaw Teo. 2018. "Call for Papers—Management Science—Special Issue on Data-Driven Prescriptive Analytics". *Management Science* 64 (6): 2972–2972.

Gilchrist, Alasdair. 2016. *Industry 4.0: the industrial internet of things*. Springer.

Glaeser, Chloe Kim, Marshall Fisher, and Xuanming Su. 2019. "Optimal retail location: Empirical methodology and application to practice: Finalist–2017 m&som practice-based research competition". *Manufacturing & Service Operations Management* 21 (1): 86–102.

Glowalla, Paul, Christoph Rosenkranz, and Ali Sunyaev. 2014. "Evolution of IT use: A case of business intelligence system transition". In *International Conference on Information Systems (ICIS) 2014 Proceedings*.

Goerigk, Marc, and Jannis Kurtz. 2020. "Data-Driven Robust Optimization using Unsupervised Deep Learning". *arXiv preprint arXiv:2011.09769*.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.

Grisold, Thomas, Jan vom Brocke, Steven Gross, Jan Mendling, Maximilian Röglinger, and Katharina Stelzl. 2021. "Digital Innovation and Business Process Management : Opportunities and Challenges as Perceived by Practitioners". *Communications of the Association for Information Systems*.

Grisold, Thomas, Jan Mendling, Markus Otto, and Jan vom Brocke. 2020a. "Adoption, use and management of process mining in practice". *Business Process Management Journal*.

Grisold, Thomas, Bastian Wurm, Jan Mendling, and Jan vom Brocke. 2020b. “Using process mining to support theorizing about change in organizations”. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*.

Grover, Purva, Arpan Kumar Kar, and Yogesh K Dwivedi. 2020. “Understanding artificial intelligence adoption in operations management: insights from the review of academic literature and social media discussions”. *Annals of Operations Research*: 1–37.

Grover, Varun, Roger HL Chiang, Ting-Peng Liang, and Dongsong Zhang. 2018. “Creating strategic business value from big data analytics: A research framework”. *Journal of Management Information Systems* 35 (2): 388–423.

Gruszczyński, K. 2019. “Enhancing business process event logs with software failure data”. *ECONTECHMOD: An International Quarterly Journal on Economics of Technology and Modelling Processes* 8.

Gupta, Daya, and Naveen Prakash. 2001. “Engineering methods from method requirements specifications”. *Requirements Engineering* 6 (3): 135–160.

Gürdür, Didem, Jad El-khoury, and Martin Törngren. 2019. “Digitalizing Swedish industry: What is next?: Data analytics readiness assessment of Swedish industry, according to survey results”. *Computers in Industry* 105:153–163.

Gursoy, Dogan, Oscar Hengxuan Chi, Lu Lu, and Robin Nunkoo. 2019. “Consumers acceptance of artificially intelligent (AI) device use in service delivery”. *International Journal of Information Management* 49:157–169.

Gust, Gunther, Christoph M. Flath, Tobias Brandt, Philipp Ströhle, and Dirk Neumann. 2016. “Bringing Analytics Into Practice: Evidence From the Power Sector”. In *Proceedings of the 37th International Conference on Information Systems (ICIS)*.

Gust, Gunther, Dirk Neumann, Christoph M. Flath, Tobias Brandt, and Philipp Ströhle. 2017. “How a traditional company seeded new analytics capabilities”. *MIS Quarterly Executive* 16 (3): 215–230.

Hallerbach, Alena, Thomas Bauer, and Manfred Reichert. 2008. “Context-based configuration of process variants”. In *3rd International Workshop on Technologies for Context-Aware Business Process Management (TCoB), Barcelona*, 31–40.

Han, Jiawei, Jian Pei, and Micheline Kamber. 2011. *Data mining: concepts and techniques*. Elsevier.

Hastie, Trevor, Robert Tibshirani, and J. H. Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*. 2nd ed. New York, NY: Springer.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep residual learning for image recognition". In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 770–778.

Heinrich, Kai, Patrick Zschech, Christian Janiesch, and Markus Bonin. 2021. "Process data properties matter: Introducing gated convolutional neural networks (GCNN) and key-value-predict attention networks (KVP) for next event prediction with deep learning". *Decision Support Systems* 143:113494.

Hernandez-Resendiz, Jaciel David, Edgar Tello-Leal, Heidy Marisol Marin-Castro, Ulises Manuel Ramirez-Alcocer, and Jonathan Alfonso Mata-Torres. 2021. "Merging Event Logs for Inter-organizational Process Mining". In *New Perspectives on Enterprise Decision-Making Applying Artificial Intelligence Techniques*, 3–26. Springer.

Holsapple, Clyde, Anita Lee-Post, and Ram Pakath. 2014. "A Unified Foundation for Business Analytics". *Decision Support Systems* 64:130–141.

Holtzblatt, Karen, and Hugh Beyer. 1997. *Contextual design: defining customer-centered systems*. Elsevier.

— . 1993. "Making customer-centered design work for teams". *Communications of the ACM* 36 (10): 92–103.

Horn, Max, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. 2020. "Set functions for time series". In *International Conference on Machine Learning*, 4353–4363. PMLR.

Hull, Giles Hindle, Martin Kunc, Michael Mortensen, Asil Oztekin, and Richard Vidgen. 2018. "Call for Papers—Business Analytics: Defining the Field and Identifying a Research Agenda". *European Journal of Operations Research*.

Jacobs, Jesse V, Lawrence J Hettinger, Yueng-Hsiang Huang, Susan Jeffries, Mary F Lesch, Lucinda A Simmons, Santosh K Verma, and Joanna L Willetts. 2019. "Employee acceptance of wearable technology in the workplace". *Applied ergonomics* 78:148–156.

Jaech, Aaron, Baosen Zhang, Mari Ostendorf, and Daniel S Kirschen. 2018. “Real-time prediction of the duration of distribution system outages”. *IEEE Transactions on Power Systems* 34 (1): 773–781.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2017. *An Introduction to Statistical Learning: with Applications in R*. Ed. by G. Casella, S. Fienberg, and I. Olkin. New York, NY: Springer.

Janiesch, Christian, and Jörn Kuhlenkamp. 2018. “Enhancing business process execution with a context engine”. *Business Process Management Journal* 25 (6): 1273–1290.

Janiesch, Christian, Martin Matzner, and Oliver Müller. 2011. “A Blueprint for Event-Driven Business Activity Management”. In *International Conference on Business Process Management (BPM 2011)*, ed. by Stefanie Rinderle-Ma, Farouk Toumani, and Karsten Wolf, 17–28. Berlin, Heidelberg: Springer Berlin Heidelberg.

Jia, Yuting, Qinqin Zhang, Weinan Zhang, and Xinbing Wang. 2019. “Community-gan: Community detection with generative adversarial nets”. In *The World Wide Web Conference*, 784–794.

Jiang, Haoming, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019. “Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization”. *arXiv preprint arXiv:1911.03437*.

Jokonowo, Bambang, Jan Claes, Riyanarto Sarno, and Siti Rochimah. 2018. “Process mining in supply chains: A systematic literature review”. *International Journal of Electrical and Computer Engineering* 8 (6): 4626–4636.

Joshi, Kailash. 1991. “A model of users’ perspective on change: the case of information systems technology implementation”. *MIS quarterly*: 229–242.

Kagermann, Henning, Johannes Helbig, Ariane Hellinger, and Wolfgang Wahlster. 2013. *Recommendations for implementing the strategic initiative INDUSTRIE 4.0: Securing the future of German manufacturing industry; final report of the Industrie 4.0 Working Group*. Forschungsunion.

Kahani, Mohsen, Behashid Behkamal, et al. 2021. “Concept drift detection in business process logs using deep learning”. *Signal and Data Processing* 17 (4): 33–48.

Kannan, Rohit, Güzin Bayraksan, and James R Luedtke. 2020. "Data-driven sample average approximation with covariate information". *Optimization Online*.

Karlsen, Jan Terje. 2002. "Project stakeholder management". *Engineering Management Journal* 14 (4): 19–24.

Kassner, Laura, Pascal Hirmer, Matthias Wieland, Frank Steimle, Jan Königsberger, and Bernhard Mitschang. 2017. "The social factory: connecting people, machines and data in manufacturing for context-aware exception escalation". In *Proceedings of the 50th Hawaii International Conference on System Sciences*.

Kerpedzihev, Gorgi Dimov, Ulrich Matthias König, M Roglinger, and Michael Rosemann. 2021. "An exploration into future business process management capabilities in view of digitalization". *Business & Information Systems Engineering* 63 (2): 83–96.

Kim, Hee-Woong, and Atreyi Kankanhalli. 2009. "Investigating user resistance to information systems implementation: A status quo bias perspective". *MIS quarterly*: 567–582.

Kim, Sun Hye, and Fani Boukouvala. 2019. "Machine learning-based surrogate modeling for data-driven optimization: a comparison of subset selection for regression techniques". *Optimization Letters*: 1–22.

Kim, Taenyun, and Hayeon Song. 2020. "The Effect of Message Framing and Timing on the Acceptance of Artificial Intelligence's Suggestion". In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–8.

Kleywegt, Anton J, Alexander Shapiro, and Tito Homem-de-Mello. 2002. "The sample average approximation method for stochastic discrete optimization". *SIAM Journal on Optimization* 12 (2): 479–502.

Klumpp, Matthias, Marc Hesenius, Ole Meyer, Caroline Ruiner, and Volker Gruhn. 2019. "Production logistics and human-computer interaction—state-of-the-art, challenges and requirements for the future". *The International Journal of Advanced Manufacturing Technology* 105 (9): 3691–3709.

Knoll, Dino, Gunther Reinhart, and Marco Prüglmeier. 2019. "Enabling value stream mapping for internal logistics using multidimensional process mining". *Expert Systems with Applications* 124:130–142.

Kogan, Konstantin, Eugene Khmelnitsky, and Toshihide Ibaraki. 2005. "Dynamic generalized assignment problems with stochastic demands and multiple agent–task relationships". *Journal of Global Optimization* 31 (1): 17–43.

Kogan, Konstantin, and Avraham Shtub. 1997. "DGAP-the dynamic generalized assignment problem". *Annals of Operations Research* 69:227–239.

Kohli, Rajiv, and Sarv Devaraj. 2004. "Realizing the Business Value of Information Technology Investments: An Organizational Process." *MIS Quarterly Executive* 3 ().

Kohzadi, Nowrouz, Milton S. Boyd, Bahman Kermanshahi, and Iebeling Kaastra. 1996. "A comparison of artificial neural network and time series models for forecasting commodity prices". *Neurocomputing* 10 (2): 169–181.

Komer, Brent, James Bergstra, and Chris Eliasmith. 2019. "Hyperopt-sklearn". In *Automated Machine Learning*, 97–111. Springer, Cham.

Konecny, Jakub, H Brendan McMahan, Felix X Yu, Peter Richtarik, Ananda Theertha Suresh, and Dave Bacon. 2016. "Federated learning: Strategies for improving communication efficiency". *arXiv preprint arXiv:1610.05492*.

Kovacs, György. 2020. "Combination of Lean value-oriented conception and facility layout design for even more significant efficiency improvement and cost reduction". *International Journal of Production Research*: 1–21.

Kratsch, Wolfgang, Jonas Manderscheid, Maximilian Röglinger, and Johannes Seyfried. 2020. "Machine Learning in Business Process Monitoring: A Comparison of Deep Learning and Classical Approaches Used for Outcome Prediction". *Business & Information Systems Engineering*.

Kraus, Mathias, Stefan Feuerriegel, and Asil Oztekin. 2020. "Deep learning in business analytics and operations research: Models, applications and managerial implications". *European Journal of Operational Research* 281 (3): 628–641.

Kregel, Ingo, Dietmar Stemann, Julian Koch, and André Coners. 2021. "Process Mining for Six Sigma: Utilising Digital Traces". *Computers & Industrial Engineering* 153:107083.

Krumeich, Julian, Dirk Werth, and Peter Loos. 2016. "Prescriptive Control of Business Processes - New Potentials Through Predictive Analytics on Big Data in the Process Manufacturing Industry". *Business & Information Systems Engineering (BISE)* 58 (4): Online First.

Kuckartz, Udo, and Stefan Rädiker. 2019. *Analyzing qualitative data with MAXQDA*. Springer.

Kuhn, Heinrich. 1995. "A heuristic algorithm for the loading problem in flexible manufacturing systems". *International Journal of Flexible Manufacturing Systems* 7 (3): 229–254.

Kundakcioglu, O. Erhun, and Saed Alizamir. 2009. "Generalized assignment problemGeneralized Assignment Problem". In *Encyclopedia of Optimization*, ed. by Christodoulos A. Floudas and Panos M. Pardalos, 1153–1162. Boston, MA: Springer US.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning". *Nature* 521 (7553): 436.

Lee, CKH, GTS Ho, KL Choy, and GKH Pang. 2014. "A RFID-based recursive process mining system for quality assurance in the garment industry". *International journal of production research* 52 (14): 4216–4238.

Lenz, Juergen, Valerio Pelosi, Marco Taisch, Eric MacDonald, and Thorsten Wuest. 2020. "Data-driven Context Awareness of Smart Products in Discrete Smart Manufacturing Systems". *Procedia Manufacturing* 52:38–43.

Li, Guangming, and Wil MP van der Aalst. 2017. "A framework for detecting deviations in complex event logs". *Intelligent Data Analysis* 21 (4): 759–779.

Li, Shancang, Li Da Xu, and Shanshan Zhao. 2015. "The internet of things: a survey". *Information Systems Frontiers* 17 (2): 243–259.

Li, Tian, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. "Federated learning: Challenges, methods, and future directions". *IEEE Signal Processing Magazine* 37 (3): 50–60.

Lin, Yuxiao, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. 2021. "BertGCN: Transductive Text Classification by Combining GCN and BERT". *arXiv preprint arXiv:2105.05727*.

Liyanage, Liwan H, and J George Shanthikumar. 2005. “A practical inventory control policy using operational statistics”. *Operations Research Letters* 33 (4): 341–348.

Lopez-Leyva, J. A., A. Molina-Inzunza, P. Navarro-Paz, S. Verduzco-Unzon, and M. Yañez. 2020. “Customized Smart Andon System to Improve the Efficiency of Industrial Departments”. *NISCAIR-CSIR* 79 (1).

Lopez, Robert, Peter ED Love, David J Edwards, and Peter R Davis. 2010. “Design error classification, causation, and prevention in construction engineering”. *Journal of performance of constructed facilities* 24 (4): 399–408.

Lorenz, Rafael, Julian Senoner, Wilfried Sihn, and Torbjørn Netland. 2021. “Using process mining to improve productivity in make-to-stock manufacturing”. *International Journal of Production Research*: 1–12.

Lustig, Irv, Brenda Dietrich, Christer Johnson, and Christopher Dziekan. 2010. “The Analytics Journey”. *Analytics Magazine*, no. 6: 11–13.

Lyu, Jrjung, Chia Wen Liang, and Ping-Shun Chen. 2020. “A Data-Driven Approach for Identifying Possible Manufacturing Processes and Production Parameters That Cause Product Defects: A Thin-Film Filter Company Case Study”. *IEEE Access* 8:49395–49411.

Ma, Sai, and Fulei Chu. 2019. “Ensemble deep learning-based fault diagnosis of rotor bearing systems”. *Computers in Industry* 105:143–152.

Maanen, John van. 1979. “Reclaiming qualitative methods for organizational research: A preface”. *Administrative science quarterly* 24 (4): 520–526.

Macdonald, John R, and Thomas M Corsi. 2013. “Supply chain disruption management: Severe events, recovery, and performance”. *Journal of Business Logistics* 34 (4): 270–288.

Maggi, Fabrizio Maria, Chiara Di Francescomarino, Marlon Dumas, and Chiara Ghidini. 2014. “Predictive Monitoring of Business Processes”. In *International Conference on Advanced Information Systems Engineering*, ed. by Matthias Jarke, John Mylopoulos, Christoph Quix, Colette Rolland, Yannis Manolopoulos, Haralambos Mouratidis, and Jennifer Horkoff, 457–472. Cham: Springer International Publishing.

Maisenbacher, Marco, and Matthias Weidlich. 2017. “Handling Concept Drift in Predictive Process Monitoring.” *SCC* 17:1–8.

Maita, Ana Rocıo Cardenas, Lucas Corrêa Martins, Carlos Ramon Lopez Paz, Laura Rafferty, Patrick CK Hung, Sarajane Marques Peres, and Marcelo Fantinato. 2018. "A systematic mapping study of process mining". *Enterprise Information Systems* 12 (5): 505–549.

Malega, Peter. 2014. "Escalation management as the necessary form of incident management process". *J Emerg Trends Comput Inf Sci* 5 (6): 641–646.

Mans, Ronny, Hajo Reijers, Hans Berends, Wasana Bandara, and Rogier Prince. 2013. "Business process mining success". In *Proceedings of the 21st European Conference on Information Systems (ECIS).*

Markus, M Lynne. 1983. "Power, politics, and MIS implementation". *Communications of the ACM* 26 (6): 430–444.

Markus, M Lynne, and Daniel Robey. 1983. "The organizational validity of management information systems". *Human relations* 36 (3): 203–225.

Marquez-Chamorro, Alfonso Eduardo, Manuel Resinas, and Antonio Ruiz-Cortes. 2017. "Predictive monitoring of business processes: a survey". *IEEE Transactions on Services Computing* 11 (6): 962–977.

Martin, Niels, Dominik A Fischer, Georgi D Kerpedzhiev, Kanika Goel, Sander JJ Leemans, Maximilian Röglinger, Wil MP van der Aalst, Marlon Dumas, Marcello La Rosa, and Moe T Wynn. 2021. "Opportunities and challenges for process mining in organizations: results of a Delphi study". *Business & Information Systems Engineering* 63 (5): 511–527.

Matschewsky, Johannes, Marianna Lena Kambanou, and Tomohiko Sakao. 2018. "Designing and providing integrated product-service systems–challenges, opportunities and solutions resulting from prescriptive approaches in two industrial companies". *International Journal of Production Research* 56 (6): 2150–2168.

Mehdiyev, Nijat, Joerg Evermann, and Peter Fettke. 2017. "A multi-stage deep learning approach for business process event prediction". In *Proceedings - 2017 IEEE 19th Conference on Business Informatics, CBI 2017*, 1:119–128. Institute of Electrical / Electronics Engineers (IEEE).

— . 2020. "A novel business process prediction model using a deep learning method". *Business & information systems engineering* 62 (2): 143–157.

Melville, Nigel, Kenneth Kraemer, and Vijay Gurbaxani. 2004. “Review: Information Technology and Organizational Performance: An Integrative Model of IT Business Value”. *MIS quarterly* 28 (2): 283–322.

Meske, Christian, Enrico Bunde, Johannes Schneider, and Martin Gersch. 2021. “Explainable artificial intelligence: objectives, stakeholders, and future research opportunities”. *Information Systems Management*: 1–11.

Mine, H, M Fukushima, K Ishikawa, and I Sawa. 1983. “An algorithm for the assignment problem with stochastic side constraints”. *Memoirs of the Faculty of Engineering, XLV (part 4)*.

Mišic, Velibor V, and Georgia Perakis. 2020. “Data analytics in operations management: A review”. *Manufacturing & Service Operations Management* 22 (1): 158–169.

Mo, Hyunho, Federico Lucca, Jonni Malacarne, and Giovanni Iacca. 2020. “Multi-Head CNN-LSTM with Prediction Error Analysis for Remaining Useful Life Prediction”. In *2020 27th Conference of Open Innovations Association (FRUCT)*, 164–171. IEEE.

Mohamad, Effendi, Mohd Soufhwee Abd Rahman, Teruaki Ito, and Azrul Azwan Abd Rahman. 2019. “Framework of Andon Support System in Lean Cyber-Physical System Production Environment”. *The Proceedings of Manufacturing Systems Division Conference* 2019 (0): 404.

Mohd Salleh, Noor Akma, Fiona Rohde, and Peter Green. 2017. “Information systems enacted capabilities and their effects on SMEs’ information systems adoption behavior”. *Journal of Small Business Management* 55 (3): 332–364.

Monostori, L., B. Kadar, T. Bauernhansl, S. Kondoh, S. Kumara, G. Reinhart, O. Sauer, G. Schuh, W. Sihn, and K. Ueda. 2016. “Cyber-physical systems in manufacturing”. *CIRP Annals* 65, no. 2 (): 621–641.

Moon, Junhyung, Gyuyoung Park, and Jongpil Jeong. 2021. “POP-ON: Prediction of Process Using One-Way Language Model Based on NLP Approach”. *Applied Sciences* 11 (2): 864.

Mooney, John G, Vijay Gurbaxani, and Kenneth L Kraemer. 1996. “A process oriented framework for assessing the business value of information technology”. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems* 27 (2): 68–81.

Mounira, Zerari, and Boufaida Mahmoud. 2010. "Context-aware process mining framework for Business Process flexibility". In *Proceedings of the 12th International Conference on Information Integration and Web-based Applications & Services*, 421–426.

Muehlen, Michael zur, and Robert Shapiro. 2015. "Business Process Analytics". In *Handbook on Business Process Management 2*, 243–263.

Müller, Oliver, Iris Junglas, Stefan Debortoli, and Jan vom Brocke. 2016. "Using text analytics to derive customer service management benefits from unstructured data". *MIS Quarterly Executive* 15 (4): 243–258.

Müller, Rainer, Matthias Vette, Leenhard Hörauf, Christoph Speicher, and Dirk Burkhard. 2017. "Lean Information and Communication Tool to Connect Shop and Top Floor in Small and Medium-sized Enterprises". *Procedia Manufacturing* 11 (): 1043–1052.

Munoz-Gama, Jorge, et al. 2016. *Conformance checking and diagnosis in process mining*. Springer.

Naedele, Martin, Hong Mei Chen, Rick Kazman, Yuanfang Cai, Lu Xiao, and Carlos V.A. Silva. 2015. "Manufacturing execution systems: A vision for managing software development". *Journal of Systems and Software* 101 (): 59–68.

Näslund, Dag. 2008. "Lean, six sigma and lean sigma: fads or real process improvement methods?" *Business process management journal*.

Neff, Gina, Anissa Tanweer, Brittany Fiore-Gartland, and Laura Osburn. 2017. "Critique and contribute: A practice-based framework for improving critical data studies and data science". *Big data* 5 (2): 85–97.

Nemati, Hamid R, and Christopher D Barko. 2003. "Key factors for achieving organizational data-mining success". *Industrial Management & Data Systems*.

Ning, Chao, and Fengqi You. 2019. "Optimization under uncertainty in the era of big data and deep learning: When machine learning meets mathematical programming". *Computers & Chemical Engineering* 125:434–448.

Notz, Pascal M, and Richard Pibernik. 2021. "Prescriptive analytics for flexible capacity management". *Management Science* 68 (3): 1756–1775.

O'Malley, Tom, Elie Bursztein, James Long, François Chollet, Haifeng Jin, Luca Invernizzi, et al. 2019. *Keras Tuner*. `https://github.com/keras-team/keras-tuner`.

Oberdorf, Felix, Kevin McFall, and Joachim Kempkes. 2018. "A Gradient Descent Based Efficiency Calculation Method With Learning Rate Adaption". In *Proceedings of the 36th International Conference Science in Practice (SIP 2018)*, 5–13.

Oberdorf, Felix, Kevin McFall, Sebastian Moros, and Joachim Kempkes. 2018. "A gradient descent based method for maximum efficiency calculation". In *Proceedings of the 2018 International IEEE Conference and Workshop in Óbuda on Electrical and Power Engineering (CANDO-EPE)*, 201–206.

Oberdorf, Felix, Sebastian Moros, and Joachim Kempkes. 2017. "MATLAB based automated and parameterized 3D-FEA modeling in ANSYS Maxwell". In *Proceedings of the 35th International Conference Science in Practice (SIP 2017)*.

Oberdorf, Felix, Myriam Schaschek, Nikolai Stein, and Christoph M. Flath. 2021a. "Neural Process Mining: Multi-Headed Predictive Process Analytics in Practice". In *Proceedings of the 29th European Conference on Information Systems (ECIS)*.

— . 2022a. "Success factors for process mining—A multiple case study". Working Paper.

Oberdorf, Felix, Myriam Schaschek, Nikolai Stein, Richard Pibernik, and Christoph M. Flath. 2022b. "Data-Driven Approximate Dynamic Stochastic Programming for Maintenance Job Assignment in Manufacturing". Working Paper.

Oberdorf, Felix, Myriam Schaschek, Sven Weinzierl, Nikolai Stein, Martin Matzner, and Christoph M. Flath. 2023. "Predictive End-to-End Enterprise Process Network Monitoring". Business & Information Systems Engineering (BISE).

Oberdorf, Felix, Nikolai Stein, and Christoph M. Flath. 2021. "Analytics-enabled escalation management: System development and business value assessment". *Computers in Industry* 131.

— . 2020. "Data-Driven Cycling Policy Guidance using GIS". In *Proceedings of the 41st International Conference on Information Systems (ICIS)*.

Oberdorf, Felix, Nikolai Stein, Nicolas Walk, Matthias Griebel, and Christoph M. Flath. 2020. "ADR for Big-Data IT Artifact Development: An Escalation Management Example". In *Proceedings of the 41st International Conference on Information Systems (ICIS)*.

Oberdorf, Felix, Peter Wolf, Myriam Schaschek, and Nikolai Stein. 2021b. "Strategic Decision Support System for Fleet Investments in the Vaccine Supply Chain". In *Proceedings of the 42nd International Conference on Information Systems, (ICIS)*.

Öncan, Temel. 2007. "A survey of the generalized assignment problem and its applications". *INFOR: Information Systems and Operational Research* 45 (3): 123–141.

Ongena, Guido, and Pascal Ravesteyn. 2019. "Business process management maturity and performance: A multi group analysis of sectors and organization sizes". *Business Process Management Journal*.

Oprea, Eugen Marius, Mihnea Alexandru Moisescu, and Simona Iuliana Caramihai. 2021. "Context Awareness in Enterprise Systems Design". In *Proceedings of the 23rd International Conference on Control Systems and Computer Science (CSCS)*, 280–286. IEEE.

Ostrom, Amy L, Darima Fotheringham, and Mary Jo Bitner. 2019. "Customer acceptance of AI in service encounters: understanding antecedents and consequences". In *Handbook of Service Science, Volume II*, 77–103. Springer.

Papers with Code. 2021. *Browse State-of-the-Art*. https://paperswithcode.com/sota. [Online; accessed 23-September-2021].

Park, Gyunam, and Minseok Song. 2020. "Predicting performances in business processes using deep neural networks". *Decision Support Systems* 129:113191.

Pasquadibisceglie, Vincenzo, Annalisa Appice, Giovanna Castellano, and Donato Malerba. 2019. "Using convolutional neural networks for predictive process analytics". In *2019 International Conference on Process Mining (ICPM)*, 129–136. IEEE.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. "Scikit-learn: Machine Learning in Python". *Journal of Machine Learning Research* 12:2825–2830.

Peffers, Ken, Tuure Tuunanen, and Björn Niehaves. 2018. *Design science research genres: introduction to the special issue on exemplars and criteria for applicable design science research*.

Penas, Olivia, Regis Plateaux, Stanislao Patalano, and Moncef Hammadi. 2017. "Multi-scale approach from mechatronic to Cyber-Physical Systems for the design of manufacturing systems". *Computers in Industry* 86:52–69.

Polyvyanyy, Artem, Chun Ouyang, Alistair Barros, and W. M. P. van der Aalst. 2017. "Process querying: Enabling business intelligence through query-based process analytics". *Decision Support Systems* 100:41–56.

Powell, Warren B. 1996. "A stochastic formulation of the dynamic assignment problem, with an application to truckload motor carriers". *Transportation Science* 30 (3): 195–219.

— . 2007. *Approximate Dynamic Programming: Solving the curses of dimensionality*. Vol. 703. John Wiley & Sons.

Qin, Jian, Ying Liu, and Roger Grosvenor. 2016. "A categorical framework of manufacturing for industry 4.0 and beyond". *Procedia Cirp* 52:173–178.

Răileanu, Silviu, Florin Anton, Theodor Borangiu, Silvia Anton, and Maximilian Nicolae. 2018. "A cloud-based manufacturing control system with data integration from multiple autonomous agents". *Computers in Industry* 102:50–61.

Ram, Sundaresan, and Jagdish N Sheth. 1989. "Consumer resistance to innovations: the marketing problem and its solutions". *Journal of consumer marketing*.

Rama-Maneiro, Efrén, Juan C. Vidal, and Manuel Lama. 2020a. *Deep Learning for Predictive Business Process Monitoring: Review and Benchmark*.

— . 2020b. "Deep learning for predictive business process monitoring: Review and benchmark". *arXiv preprint arXiv:2009.13251*.

Rebuge, Alvaro, and Diogo R Ferreira. 2012. “Business process analysis in healthcare environments: A methodology based on process mining”. *Information systems* 37 (2): 99–116.

Reddy, Aala Santhosh. 2016. “Why IoT Analytics Are a Manufacturer’s Most Important Tool”. *Harvard Business Review*.

Rogge-Solti, Andreas, Wil M. P. van der Aalst, and Mathias Weske. 2013. “Discovering stochastic petri nets with arbitrary delay distributions from event logs”. In *International Conference on Business Process Management*, 15–27. Springer.

Romero, David, Thorsten Wuest, Johan Stahre, and Dominic Gorecky. 2017. “Social factory architecture: social networking services and production scenarios through the social internet of things, services and people for the social operator 4.0”. In *IFIP International Conference on Advances in Production Management Systems*, 265–273. Springer.

Rosemann, Michael, Jan Recker, and Christian Flender. 2008. “Contextualisation of business processes”. *International Journal of Business Process Integration and Management* 3 (1): 47–60.

Rosemann, Michael, Jan Recker, Christian Flender, and Peter-Daniel Ansell. 2006. “Understanding context-awareness in business process design”. In *Proceedings of the 17th Australasian Conference on Information Systems*, 1–10. Australasian Association for Information Systems.

Ross, G Terry, and Richard M Soland. 1975. “A branch and bound algorithm for the generalized assignment problem”. *Mathematical programming* 8 (1): 91–103.

Rubin, Greg, Ajay George, DJ Chinn, and Clive Richardson. 2003. “Errors in general practice: development of an error classification and pilot study of a method for detecting errors”. *BMJ Quality & Safety* 12 (6): 443–447.

Rumelhart, David E, Geoffrey E Hinton, and Ronald J Williams. 1985. *Learning internal representations by error propagation*. Tech. rep. California Univ San Diego La Jolla Inst for Cognitive Science.

Rüßmann, Michael; et al. 2015. “Future of Productivity and Growth in Manufacturing”. *Boston Consulting*, no. April.

Russom, Philip. 2011. *TDWI Best Practices Report: Big Data Analytics*. Tech. rep. TDWI.

Sahni, Sartaj, and Teofilo Gonzalez. 1976. "P-complete approximation problems". *Journal of the ACM (JACM)* 23 (3): 555–565.

Saldaña, Johnny. 2015. *The coding manual for qualitative researchers*. Sage.

Sanchez-Marquez, Rafael, Jose Miguel Albarracın Guillem, Eduardo Vicens-Salort, and Jose Jabaloyes Vivas. 2020. "Diagnosis of quality management systems using data analytics–A case study in the manufacturing sector". *Computers in Industry* 115:103183.

Sanders, Nada R., and Ram Ganeshan. 2015. "Special Issue of Production and Operations Management on "Big Data in Supply Chain Management"". *Production and Operations Management* 24 (3): 519–520.

Schaschek, Myriam, Felix Oberdorf, Nikolai Stein, Axel Winkelmann, and Christoph M. Flath. 2022. "A framework for context-aware Process Mining: An action-driven manufacturing use case". In *Review at the 30th European Conference on Information Systems (ECIS)*.

Scheepers, Helana, and Rens Scheepers. 2008. "A process-focused decision framework for analyzing the business value potential of IT investments". *Information Systems Frontiers* 10 (3): 321–330.

Scheer, August-Wilhelm. 2013. *ARIS—vom Geschäftsprozess zum Anwendungssystem*. Springer-Verlag.

Schmenner, Roger W. 2012. *Getting and staying productive: applying swift, even flow to practice*. Cambridge University Press.

Schmenner, Roger W, and Morgan L Swink. 1998. "On theory in operations management". *Journal of operations management* 17 (1): 97–113.

Schryen, Guido. 2010. "Preserving knowledge on IS business value". *Business & Information Systems Engineering* 2 (4): 233–244.

Schuh, Günther, Andreas Gützlaff, Sven Cremer, and Marco Schopen. 2020. "Understanding process mining for data-driven optimization of order processing". *Procedia Manufacturing* 45:417–422.

Schuh, Günther, Gunther Reinhart, Jan-Philipp Prote, Frederick Sauermann, Julia Horsthofer, Florian Oppolzer, and Dino Knoll. 2019. "Data mining definitions and applications for the management of production complexity". *Procedia CIRP* 81:874–879.

Schwegmann, Bernd, Martin Matzner, and Christian Janiesch. 2013. "A Method and Tool for Predictive Event-Driven Process Analytics". In *Proceedings of the 11th International Conference on Wirtschaftsinformatik (WI 2019)*.

Seidl, Thomas. 2021. "Concept Drift Detection on Streaming Data with Dynamic Outlier Aggregation". In *Process Mining Workshops: ICPM 2020 International Workshops*, 406:206. Springer Nature.

Sein, Maung K, Ola Henfridsson, Sandeep Purao, Matti Rossi, and Rikard Lindgren. 2011. "Action design research". *MIS Quarterly* 35 (1): 37–56.

Sen, Suvrajeet, and Yunxiao Deng. 2018. "Learning enabled optimization: Towards a fusion of statistical learning and stochastic programming". *INFORMS Journal on Optimization (submitted)*.

Senderovich, Arik, Chiara Di Francescomarino, Chiara Ghidini, Kerwin Jorbina, and Fabrizio Maria Maggi. 2017. "Intra and inter-case features in predictive process monitoring: A tale of two dimensions". In *International Conference on Business Process Management*, 306–323. Springer.

Senderovich, Arik, Chiara Di Francescomarino, and Fabrizio Maria Maggi. 2019. "From knowledge-driven to data-driven inter-case feature encoding in predictive process monitoring". *Information Systems* 84:255–264.

Senoner, Julian, Torbjorn Netland, and Stefan Feuerriegel. 2021. "Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing". *Management Science*.

Shapiro, Alexander. 2003. "Monte Carlo sampling methods". *Handbooks in operations research and management science* 10:353–425.

Shmoys, David B, and Eva Tardos. 1993. "An approximation algorithm for the generalized assignment problem". *Mathematical programming* 62 (1): 461–474.

Shraga, Roee, Avigdor Gal, Dafna Schumacher, Arik Senderovich, and Matthias Weidlich. 2019. “Inductive Context-aware Process Discovery”. In *Proceedings of the 2019 International Conference on Process Mining (ICPM)*, 33–40. IEEE.

— . 2020. “Process discovery with context-aware process trees”. *Information Systems*: 101533.

Snee, Ronald D. 2010. “Lean Six Sigma–getting better all the time”. *International Journal of Lean Six Sigma*.

Snoek, Jasper, Hugo Larochelle, and Ryan P Adams. 2012. “Practical bayesian optimization of machine learning algorithms”. *Advances in neural information processing systems* 25.

Song, Minseok, Christian W Günther, and Wil MP van der Aalst. 2008. “Trace clustering in process mining”. In *Proceedings of the International conference on business process management*, 109–120. Springer.

Sonnenberg, Christian, and Jan vom Brocke. 2012. “Evaluations in the science of the artificial–reconsidering the build-evaluate pattern in design science research”. In *Proceedings of the International Conference on Design Science Research in Information Systems*, 381–397. Springer.

Sousa, Rafael Gaspar de, Sarajane Marques Peres, Marcelo Fantinato, and Hajo Alexander Reijers. 2021. “Concept drift detection and localization in process mining: an integrated and efficient approach enabled by trace clustering”. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, 364–373.

Spenrath, Yorick, and Marwan Hassani. 2020. “Predicting Business Process Bottlenecks In Online Events Streams Under Concept Drifts.” In *ECMS*, 190–196.

Spoerl, David R, and R Kevin Wood. 2004. “A stochastic generalized assignment problem”. *academia.edu*.

Srinivasan, V, and Gerald L Thompson. 1973. “An algorithm for assigning uses to sources in a special class of transportation problems”. *Operations Research* 21 (1): 284–295.

Stein, Nikolai, Jan Meller, and Christoph M. Flath. 2018. “Big data on the shopfloor: sensor-based decision-support for manual processes”. *Journal of Business Economics* 88 (5): 593–616.

Stein, Nikolai, Felix Oberdorf, and Jonas Pirner. 2020. "Convolutional Neural Networks for Survey Response Classification". In *Proceedings of the 26th Americas Conference on Information Systems (AMCIS)*.

Stierle, Matthias, Sven Weinzierl, Maximilian Harl, and Martin Matzner. 2021. "A technique for determining relevance scores of process activities using graph-based neural networks". *Decision Support Systems* 144:113511.

Sunk, Alexander, Peter Kuhlang, Thomas Edtmayr, and Wilfried Sihn. 2017. "Developments of traditional value stream mapping to enhance personal and organisational system and methods competencies". *International Journal of Production Research* 55 (13): 3732–3746.

Suriadi, Suriadi, Moe T Wynn, Chun Ouyang, Arthur HM ter Hofstede, and Nienke J van Dijk. 2013. "Understanding process behaviours in a large insurance company in Australia: A case study". In *International Conference on Advanced Information Systems Engineering*, 449–464. Springer.

Sutton, Robert I, and Barry M Staw. 1995. "What theory is not". *Administrative science quarterly*: 371–384.

Syed, Rehan, Sander JJ Leemans, Rebekah Eden, and Joos ACM Buijs. 2020. "Process Mining Adoption". In *International Conference on Business Process Management*, 229–245. Springer.

Tax, Niek, Natalia Sidorova, and Wil M. P. van der Aalst. 2019. "Discovering more precise process models from event logs by filtering out chaotic activities". *Journal of Intelligent Information Systems* 52 (1): 107–139.

Tax, Niek, Ilya Verenich, Marcello La Rosa, and Marlon Dumas. 2017. "Predictive business process monitoring with LSTM neural networks". In *International Conference on Advanced Information Systems Engineering*, 477–492. Springer.

Teinemaa, Irene, Marlon Dumas, Anna Leontjeva, and Fabrizio Maria Maggi. 2018. "Temporal stability in predictive process monitoring". *Data Mining and Knowledge Discovery* 32 (5): 1306–1338.

Theis, Julian, and Houshang Darabi. 2019. "Decay replay mining to predict next process events". *IEEE Access* 7:119787–119803.

Thiede, Malte, Daniel Fuerstenau, and Ana Paula Bezerra Barquet. 2018. “How is process mining technology used by organizations? A systematic literature review of empirical studies”. *Business Process Management Journal*.

Thoben, Klaus-Dieter, Stefan Wiesner, and Thorsten Wuest. 2017. ““Industrie 4.0” and smart manufacturing-a review of research issues and application examples”. *International journal of automation technology* 11 (1): 4–16.

Thong, James YL, Chee-Sing Yap, and KS Raman. 1996. “Top management support, external expertise and information systems implementation in small businesses”. *Information systems research* 7 (2): 248–267.

Tiwari, Ashutosh, Chris J Turner, and Basim Majeed. 2008. “A review of business process mining: state-of-the-art and future trends”. *Business Process Management Journal*.

Toktas, Berkin, Joyce W Yen, and Zelda B Zabinsky. 2006. “Addressing capacity uncertainty in resource-constrained assignment problems”. *Computers & operations research* 33 (3): 724–745.

Trevor, Hastie, and Tibshirani Robert. 2001. *Friedman J Jerome H. The elements of statistical learning*.

Trieu, Van-Hau. 2017. “Getting value from Business Intelligence systems: A review and research agenda”. *Decision Support Systems* 93:111–124.

Tukey, John W. 1977. *Exploratory Data Analysis*. Pearson.

Turner, Chris J, Ashutosh Tiwari, Richard Olaiya, and Yuchun Xu. 2012. “Process mining: from theory to practice”. *Business Process Management Journal*.

Valles, Adan, Jaime Sanchez, Salvador Noriega, and Berenice Gómez Nuñez. 2009. “Implementation of Six Sigma in a manufacturing process: A case study”. *International Journal of Industrial Engineering* 16 (3): 171–181.

van der Aalst, W. M. P. 2016a. “Data science in action”. In *Process mining*, 3–23. Springer.

— . 2011. “Process mining: discovering and improving Spaghetti and Lasagna processes”. In *2011 IEEE Symposium on Computational Intelligence and Data Mining*, 1–7. IEEE.

van der Aalst, W. M. P., et al. 2012a. "Process Mining Manifesto". In *Business Process Management Workshops*, ed. by Florian Daniel, Kamel Barkaoui, and Schahram Dustdar, 169–194. Berlin, Heidelberg: Springer Berlin Heidelberg.

van der Aalst, W. M. P., Arya Adriansyah, Ana Karla Alves De Medeiros, Franco Arcieri, Thomas Baier, Tobias Blickle, Jagadeesh Chandra Bose, Peter van Den Brand, Ronald Brandtjen, Joos Buijs, et al. 2011a. "Process Mining Manifesto". In *International Conference on Business Process Management*, 169–194. Springer.

van der Aalst, Wil M. P. 2019. "A practitioner's guide to process mining: Limitations of the directly-follows graph". *Procedia Computer Science* 164:321–328.

— . 2013. "Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining". In *Asia-Pacific conference on business process management*, 1–22. Springer.

— . 2016b. *Process Mining - Data Science in Action*. Berlin, Heidelberg: Springer.

van der Aalst, Wil M. P., A Adriansyah, AKA De Medeiros, F Arcieri, T Baier, and T Blickle. 2012b. "Process mining manifesto". In *International Conference on Business Process Management. Berlin, Germany: Springer Heidelberg. BPM*, vol. 11.

van der Aalst, Wil M. P., M Helen Schonenberg, and Minseok Song. 2011. "Time prediction based on process mining". *Information systems* 36 (2): 450–475.

van der Aalst, Wil, Arya Adriansyah, Ana Medeiros, Franco Arcieri, Thomas Baier, Tobias Blickle, Jagadeesh Chandra Bose R.P., Peter Brand, Ronald Brandtjen, Joos Buijs, Andrea Burattin, Josep Carmona, Malú Castellanos, Jan Claes, Jonathan Cook, Nicola Costantini, Francisco Curbera, Ernesto Damiani, Massimiliano de Leoni, and Moe Wynn. 2011b. "Process Mining Manifesto". In *Lecture Notes in Business Information Processing*, 99:169–194.

van Dongen, B. F., R. A. Crooy, and W. M.P. van der Aalst. 2008. "Cycle time prediction: When will this case finally be finished?" In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, 319–336. Berlin, Heidelberg: Springer.

Vater, Johannes, Lars Harscheidt, and Alois Knoll. 2019. “Smart manufacturing with prescriptive analytics”. In *Proceedings of the 2019 8th International Conference on Industrial Technology and Management (ICITM)*, 224–228. IEEE.

Venable, John, Jan Pries-Heje, and Richard Baskerville. 2012. “A comprehensive framework for evaluation in design science research”. In *Proceedings of the International conference on design science research in information systems*, 423–438. Springer.

Venkatesh, Viswanath, Michael G Morris, Gordon B Davis, and Fred D Davis. 2003. “User acceptance of information technology: Toward a unified view”. *MIS quarterly*: 425–478.

Vera-Baquero, Alejandro, Ricardo Colomo-Palacios, and Owen Molloy. 2013. “Business Process Analytics Using a Big Data Approach”. *IT Professional* 15 (6): 29–35.

Verbitskiy, Sergey, and Viacheslav Vyshegorodtsev. 2021. “ERANNs: Efficient Residual Audio Neural Networks for Audio Pattern Recognition”. *arXiv preprint arXiv:2106.01621*.

Verenich, Ilya, Marlon Dumas, Marcello La Rosa, Fabrizio Maria Maggi, and Irene Teinemaa. 2019. “Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring”. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10 (4): 1–34.

Verweij, Bram, Shabbir Ahmed, Anton J Kleywegt, George Nemhauser, and Alexander Shapiro. 2003. “The sample average approximation method applied to stochastic routing problems: a computational study”. *Computational optimization and applications* 24 (2): 289–333.

Vidgen, Richard, Sarah Shaw, and David B Grant. 2017. “Management challenges in creating value from business analytics”. *European Journal of Operational Research* 261 (2): 626–639.

vom Brocke, Jan, Marie-Sophie Baier, Theresa Schmiedel, Katharina Stelzl, Maximilian Röglinger, and Charlotte Wehking. 2021a. “Context-Aware Business Process Management”. *Business & Information Systems Engineering*: 1–18.

vom Brocke, Jan, Mieke Jans, Jan Mendling, and Hajo A Reijers. 2021b. *A Five-Level Framework for Research on Process Mining*.

— . 2020. "Call for Papers—Business & Information System Engineering—Special Issue on "Process Mining at the Enterprise Level"". *Business & Information Systems Management* 62 (2): 185–187.

vom Brocke, Jan, and Michael Rosemann. 2014. *Handbook on business process management 1: Introduction, methods, and information systems*. Springer.

Wahid, Nur Ahmad, Taufik Nur Adi, Hyerim Bae, and Yulim Choi. 2019. "Predictive Business Process Monitoring–Remaining Time Prediction using Deep Neural Network with Entity Embedding". *Procedia Computer Science* 161:1080–1088.

Walk, Harro. 2010. "Strong laws of large numbers and nonparametric estimation". In *Recent developments in applied probability and statistics*, 183–214. Springer.

Wan, Renzhuo, Shuping Mei, Jun Wang, Min Liu, and Fan Yang. 2019. "Multivariate temporal convolutional network: A deep neural networks approach for multivariate time series forecasting". *Electronics* 8 (8): 876.

Wang, Jiaxing, Sibin Gao, Zhejun Tang, Dapeng Tan, Bin Cao, and Jing Fan. 2021. "A context-aware recommendation system for improving manufacturing process modeling". *Journal of Intelligent Manufacturing*: 1–22.

Wang, Jinjiang, Yulin Ma, Laibin Zhang, Robert X Gao, and Dazhong Wu. 2018. "Deep Learning for Smart Manufacturing: Methods and Applications". *Journal of Manufacturing Systems* 48:144–156.

Wang, Lihui, Martin Törngren, and Mauro Onori. 2015. "Current status and advancement of cyber-physical systems in manufacturing". *Journal of Manufacturing Systems* 37:517–527.

Wang, Rui, Bicheng Li, Shengwei Hu, Wenqian Du, and Min Zhang. 2019a. "Knowledge graph embedding via graph attenuated attention networks". *IEEE Access* 8:5212–5224.

Wang, Shan, William Yeoh, Gregory Richards, Siew Fan Wong, and Younghoon Chang. 2019b. "Harnessing business analytics value through organizational absorptive capacity". *Information & Management* 56 (7): 103152.

Weinzierl, Sven, Sebastian Dunzer, Sandra Zilker, and Martin Matzner. 2020. "Prescriptive business process monitoring for recommending next best actions". In *International conference on business process management*, 193–209. Springer.

Widmer, Gerhard, and Miroslav Kubat. 1996. "Learning in the presence of concept drift and hidden contexts". *Machine learning* 23 (1): 69–101.

Wuest, Thorsten, Daniel Weimer, Christopher Irgens, and Klaus-Dieter Thoben. 2016. "Machine Learning in Manufacturing: Advantages, challenges, and applications". *Production & Manufacturing Research* 4 (1): 23–45.

Yang, Heetae, Hwansoo Lee, and Hangjung Zo. 2017. "User acceptance of smart home services: an extension of the theory of planned behavior". *Industrial Management & Data Systems*.

Yang, Jian, Xiangtong Qi, and Gang Yu. 2005. "Disruption management in production planning". *Naval Research Logistics* 52 (5): 420–442.

Yang, Qiang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. "Federated machine learning: Concept and applications". *ACM Transactions on Intelligent Systems and Technology (TIST)* 10 (2): 1–19.

Yeshchenko, Anton, Fernando Durier, Kate Revoredo, Jan Mendling, and Flavia Santoro. 2018. "Context-aware predictive process monitoring: The impact of news sentiment". In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*, 586–603. Springer.

Yiu, LM Daphne, Andy CL Yeung, and TC Edwin Cheng. 2020. "The impact of business intelligence systems on profitability and risks of firms". *International Journal of Production Research*: 1–24.

Zelst, Sebastiaan J. van, Felix Mannhardt, Massimiliano de Leoni, and Agnes Koschmider. 2020. "Event abstraction in process mining: literature review and taxonomy". *Granular Computing*: 1–18.

Zerari, Mounira, and Mahmoud Boufaida. 2011. "Dynamic context-aware Business Process flexibility: an artefact-based approach using process mining". *International Journal of Business Intelligence and Data Mining* 6 (4): 345–361.

Zhang, Bin, Shaohui Zhang, and Weihua Li. 2019. "Bearing performance degradation assessment using long short-term memory recurrent network". *Computers in Industry* 106:14–29.

Zhang, Cai Wen, and Hoon Liong Ong. 2007. "An efficient solution to biobjective generalized assignment problem". *Advances in Engineering Software* 38 (1): 50–58.

Zhang, Yingfeng, Shan Ren, Yang Liu, and Shubin Si. 2017. "A big data analytics architecture for cleaner manufacturing and maintenance processes of complex products". *Journal of Cleaner Production* 142:626–641.

Zhang, Zhen, Jiajun Bu, Martin Ester, Jianfeng Zhang, Chengwei Yao, Zhi Yu, and Can Wang. 2019. "Hierarchical graph pooling with structure learning". *arXiv preprint arXiv:1911.05954*.

Zhao, Bendong, Huanzhang Lu, Shangfeng Chen, Junliang Liu, and Dongya Wu. 2017. "Convolutional neural networks for time series classification". *Journal of Systems Engineering and Electronics* 28 (1): 162–169.

Zhu, Yongchun, Fuzhen Zhuang, Jindong Wang, Jingwu Chen, Zhiping Shi, Wenjuan Wu, and Qing He. 2019. "Multi-representation adaptation network for cross-domain image classification". *Neural Networks* 119:214–221.

Zschech, Patrick. 2018. "A taxonomy of recurring data analysis problems in maintenance analytics". In *Proceedings of the 26th European Conference on Information Systems (ECIS)*.

Zur Muehlen, Michael, and Robert Shapiro. 2015. "Business process analytics". In *Handbook on Business Process Management 2*, 243–263. Springer.

Appendix

A List of Publications

Oberdorf, Felix, Kevin McFall, and Joachim Kempkes. 2018. “A Gradient Descent Based Efficiency Calculation Method With Learning Rate Adaption”. In *Proceedings of the 36th International Conference Science in Practice (SIP 2018)*, 5–13.

Oberdorf, Felix, Kevin McFall, Sebastian Moros, and Joachim Kempkes. 2018. “A gradient descent based method for maximum efficiency calculation”. In *Proceedings of the 2018 International IEEE Conference and Workshop in Óbuda on Electrical and Power Engineering (CANDO-EPE)*, 201–206.

Oberdorf, Felix, Sebastian Moros, and Joachim Kempkes. 2017. “MATLAB based automated and parameterized 3D-FEA modeling in ANSYS Maxwell”. In *Proceedings of the 35th International Conference Science in Practice (SIP 2017)*.

Oberdorf, Felix, Myriam Schaschek, Nikolai Stein, and Christoph M. Flath. 2021a. “Neural Process Mining: Multi-Headed Predictive Process Analytics in Practice”. In *Proceedings of the 29th European Conference on Information Systems (ECIS)*.

— . 2022a. “Success factors for process mining—A multiple case study”. Working Paper.

Oberdorf, Felix, Myriam Schaschek, Nikolai Stein, Richard Pibernik, and Christoph M. Flath. 2022b. “Data-Driven Approximate Dynamic Stochastic Programming for Maintenance Job Assignment in Manufacturing”. Working Paper.

Oberdorf, Felix, Myriam Schaschek, Sven Weinzierl, Nikolai Stein, Martin Matzner, and Christoph M. Flath. 2023. “Predictive End-to-End Enterprise Process Network Monitoring”. Business & Information Systems Engineering (BISE).

Oberdorf, Felix, Nikolai Stein, and Christoph M. Flath. 2021. “Analytics-enabled escalation management: System development and business value assessment”. *Computers in Industry* 131.

— . 2020. “Data-Driven Cycling Policy Guidance using GIS”. In *Proceedings of the 41st International Conference on Information Systems (ICIS).*

Oberdorf, Felix, Nikolai Stein, Nicolas Walk, Matthias Griebel, and Christoph M. Flath. 2020. “ADR for Big-Data IT Artifact Development: An Escalation Management Example”. In *Proceedings of the 41st International Conference on Information Systems (ICIS).*

Oberdorf, Felix, Peter Wolf, Myriam Schaschek, and Nikolai Stein. 2021b. “Strategic Decision Support System for Fleet Investments in the Vaccine Supply Chain”. In *Proceedings of the 42nd International Conference on Information Systems, (ICIS).*

Schaschek, Myriam, Felix Oberdorf, Nikolai Stein, Axel Winkelmann, and Christoph M. Flath. 2022. “A framework for context-aware Process Mining: An action-driven manufacturing use case”. In *Review at the 30th European Conference on Information Systems (ECIS).*

Stein, Nikolai, Felix Oberdorf, and Jonas Pirner. 2020. “Convolutional Neural Networks for Survey Response Classification”. In *Proceedings of the 26th Americas Conference on Information Systems (AMCIS).*

B Checklist for Interviews

The Interview Guide is not a questionnaire, but a checklist. The sequence of questions may vary, and follow-up questions may be asked.

Companies' employee guide

- Please describe your current responsibilities and work situation.
- In which tasks and processes are you involved?
- How does access to process data facilitate your work?
- What are your motives for using process mining?
- How do you estimate the potential of process mining in the context of internal processes?
- What business value do you see through process mining?
- How do you see the acceptance of information systems (with process mining) in the company?
- How do you assess challenges for the continuous use of process mining?
- How do you evaluate the results of process mining?
- What success factors do you see for establishing process mining?
- Why do you see ... as a challenge or success for process mining? [*With mentioned and identified challenges/factors from literature*]

Process mining expert and consultant guide

- Please describe your current responsibilities, clients, and work situation.
- What are your motives for using process mining?

- What business value do you see through process mining?
- How do you evaluate the results of process mining?
- What success factors do you see for establishing process mining? [*Without providing process mining business value framework*]
- How can the process mining business value framework be adopted by your clients? [*With provided process mining business framework*]
- Can you provide additional real-world scenarios where the framework could have been helpful?

C Multi-Headed Neural Network Architecture

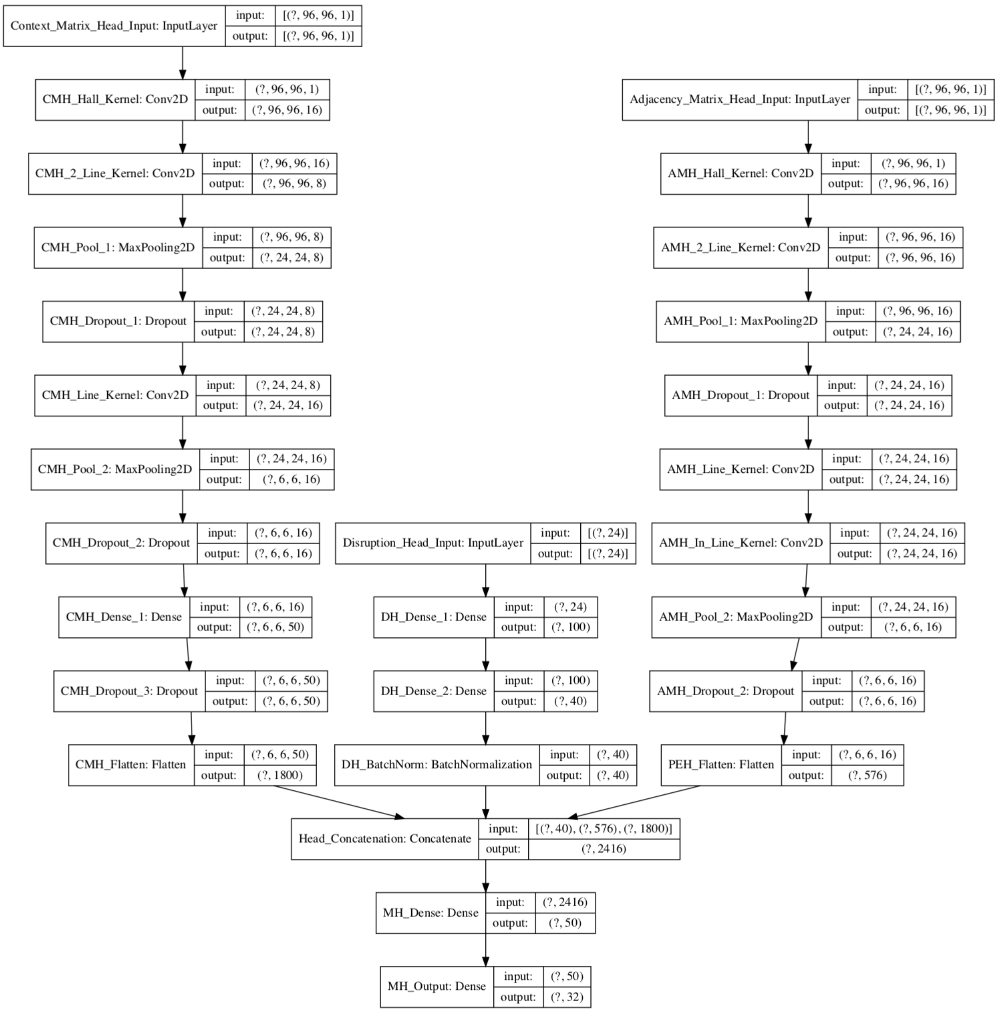

Figure C.1: Visualization of the three-headed NN structure for disruption type classification. ? refers to the TensorFlow representation for the chosen batch-size of 32.