

Decision Mining in the Rail Industry

A Case Study on an Application of Decision Mining for Process Model Enhancement and Conformance Checking in the Context of an Industrial Wheelset Revision Process

Master Thesis
Business Informatics
Bart Hoornstra



Decision Mining in the Rail Industry

A Case Study on an Application of Decision Mining
for Process Model Enhancement and Conformance
Checking in the Context of an Industrial Wheelset
Revision Process

by

Bart Hoornstra

Master Thesis
to obtain the degree of Master of Science (MSc)
in Information Science
at Utrecht University,
to be defended publicly on Monday, February 12, 2024 at 15:30.

Student number:	1633171
Faculty:	Faculty of Science, Graduate School of Natural Sciences
Department:	Information and Computing Sciences
Research group:	Business Process Management & Analytics
Project duration:	February 1, 2023 – February 1, 2024
First supervisor:	Dr. G.C. (Inge) van de Weerd Utrecht University
Second examiner:	Dr. I.M. (Iris) Beerepoot Utrecht University
External supervisor:	Drs. J.F. (Joris) Mens N.V. Nederlandse Spoorwegen

*The archived thesis is confidential and cannot be made public without explicit permission from NS.
This is the redacted version suitable for publication.*

Cover: Machine Learning Process, Quardia, Shutterstock, ID: 1938746143

An electronic version of this thesis is available online at
<https://studenttheses.uu.nl/>



**Utrecht
University**



Preface and acknowledgments

The completion of this thesis marks the end of my academic journey that started in September 2016 at the Vrije Universiteit, Amsterdam. It continued in September 2020 with the Master program Business Informatics at Utrecht University that is now coming to an end, and probably this also ends my academic career. The process that spanned roughly the past 7.5 years cannot be described in a few words, but it has been everything from extremely rewarding to exceedingly challenging at a variety of different levels. However, I sincerely feel that I have made the right decisions and that I would not have missed the journey. I have discovered new things about myself as a person, have deeply educated myself, experienced personal growth, broadened my horizon, and met so many amazing people over the course of the years. For some of those, now is of course the time to express my feelings of sincere gratitude.

First of all, I would like to thank my first supervisor Dr. Inge van de Weerd for her guidance throughout this project. Your pragmatic approach, positively critical and appropriately reflective attitude greatly helped shape this research. In addition, it helped me learn about some of my strengths and weaknesses and how to deal with them. Furthermore, I would like to thank Dr. Iris Beerepoot for being my second examiner. Your enthusiasm for my research, helpful tips, and constructive criticism were all very much appreciated and motivating. Lastly, I also want to thank Dr. ir. Xixi Lu for the interesting exchanges of thoughts that we had about this research.

Next, I would like to thank Drs. Joris Mens for being my external supervisor during my internship at the Dutch Railways (NS). Your enthusiastic though still relaxed attitude greatly helped me get acquainted with the practicalities of the company and how to properly fit the project within the limits of what was possible. Furthermore, I would like to thank all members of the NS Innovation Platform — Stijn, Rembert, Rainish, Roxanne, Germo and Melvin — for their interest in my research and the fruitful discussions we had, especially before and after my intermediate presentations.

In addition, I would like to explicitly thank Mick Pas, Patrick Vinke, and Mees Waal from the wheelset revision department (RLW) of NS. Your no-nonsense, practical attitude greatly helped accelerate this research. Without your commitment, enthusiasm, expertise, and valuable contributions, this research would not have been what it is now.

Finally, returning to personal notes, I would like to thank my girlfriend Anne-Lot — my wonderful Wuppie — for her everlasting trust and support in both good and bad times. Without you by my side, all this would have been a much more cumbersome experience. Lastly, I want to thank my family, family in law, friends, other relatives, and fellow students for their support and moments of reflection. Also, a special word of thanks to my sister in law Maartje for proofreading this thesis.

*Life is like riding a bicycle.
To keep your balance you must keep moving.*

—Albert Einstein (1879 – 1955)

*Bart Hoornstra
Utrecht, February 2024*

Abstract

The field of process mining has evolved into a multidisciplinary research paradigm in which relevant data mining practices led to new avenues of analysis and better process understanding. However, the role of decisions in relation to the modeling and analysis of processes is still significantly unexplored. The nascent discipline of decision mining aims to fill this void by separating concerns, while a more holistic approach is also being considered. This research took the latter approach, as there is nowadays an abundance of process data available that can potentially uncover captivating details without the need of a separate project. Following the design science research methodology, a methodological framework for integrated process and decision mining was developed. It is based on the synthesis of an established process mining project methodology and a review of the state-of-the-art in decision mining approaches. The extended methodology was applied and evaluated with a case study of the industrial wheelset revision process at the Dutch national railway company, which included a focus group with several experts. The results demonstrated that the addition of a decision perspective to process models allows for better process understanding and therefore is potentially useful for an organization that is executing a process mining project. In addition, the evaluation identified a new form of conformance checking that can be used to validate whether the process was executed correctly in accordance with the decisions taken. Future work could aim to increase generalizability by applying the framework in other contexts. Furthermore, the framework could be validated further beyond the sole evaluation of the resulting artifacts and insights.

Keywords: Decision Mining · Decision Point Analysis · Process Mining · Process Model Enhancement · Conformance Checking · Design Science Research Methodology · Industrial Wheelset Revision Process

Contents

Preface and acknowledgments	i
Abstract	ii
Contents	v
List of Figures	vii
List of Tables	vii
Nomenclature	viii
1 Introduction	1
1.1 Problem statement	2
1.2 Research gap	3
1.3 Research aims and objectives	3
2 Theoretical background	5
2.1 Business Process fundamentals	5
2.2 Business Process Modeling	6
2.2.1 Petri nets	6
2.2.2 Business Process Model and Notation (BPMN)	7
2.3 Business Process Management and Information Systems	8
2.4 Process Mining	9
2.5 Decision Management & Modeling	10
2.5.1 Business rules	11
2.5.2 Business logic	11
2.5.3 Decision Model and Notation (DMN)	11
2.6 Decision Mining	13
2.6.1 Decision logs	14
2.6.2 Towards holistic process and decision mining	15
3 Research design	18
3.1 Methodology	18
3.1.1 Problem identification and motivation	19
3.1.2 Definition of the objectives for a solution	19
3.1.3 Design and development	20
3.1.4 Demonstration and evaluation	20
3.1.5 Instantiation within a project methodology	20
3.1.6 Communication	21
3.2 Systematic literature review	22
3.2.1 Research questions	23
3.2.2 Inclusion and exclusion criteria	23
3.2.3 Quality criteria	24
3.3 Single-case study	24
3.3.1 Context	24
3.3.2 Cases	25
3.3.3 Data collection	25
3.3.4 Data analysis	26
3.3.5 Ethical considerations	26
3.4 Validity threats	27
3.4.1 Systematic literature review	27
3.4.2 Single-case study	29

4	Systematic literature review	31
4.1	Overview and characteristics of the SLR procedure	31
4.2	Identification of decision points with contextual process data	31
4.3	Enhancement of process analyses with decision information	32
4.4	Evaluation criteria for decision-enhanced process models	33
5	Methodological framework	35
5.1	Rationale and details of the design procedure	35
5.2	An overview of the methodological framework: PM ² xDM	35
5.3	Stage 1: Planning	36
5.4	Stage 2: Extraction	37
5.5	Stage 3: Data processing	37
5.6	Stage 4: Mining and analysis	38
5.7	Stage 5: Evaluation	39
5.8	Stage 6: Process improvement and support	40
6	Case study results	41
6.1	Rationale	41
6.2	Description	42
6.2.1	Process description and characteristics	42
6.2.2	Wheelset types, axles, and components	43
6.3	Implementation	43
6.3.1	Stage 1: Planning	43
6.3.2	Stage 2: Extraction	45
6.3.3	Stage 3: Data processing	47
6.3.4	Stage 4: Mining and analysis	48
6.3.5	Stage 5: Evaluation	54
6.3.6	Stage 6: Process improvement and support	60
6.4	Case study summary	60
7	Discussion	62
7.1	Implications	62
7.2	Contributions	63
7.2.1	Scientific	63
7.2.2	Practical	63
7.3	Retrospective analysis of validity threats	63
7.4	Remaining challenges and limitations	64
7.5	Future research opportunities	64
8	Conclusion	65
	References	67
A	SLR PRISMA 2020 report	78
B	Overview of decision mining approaches	79
C	PDDs, Activity and Concept tables	82
D	Case study protocol	87
D.1	Preamble	87
D.2	General	87
D.3	Procedures	87
D.4	Research instruments	87
D.5	Data analysis guidelines	88
D.6	Appendix	88
E	Focus group and interview consent form	89
F	Focus group and interview protocol	91
F.1	Outline	91
F.2	Welcome	91

F.3	Introduction	91
F.4	Explanation of process mining and decision mining	92
F.5	Process characteristics	92
F.6	Process modeling techniques	92
F.7	Evaluation of baseline process models	92
F.8	Evaluation of enhanced process models	93
F.9	Evaluation of derived insights and process improvement potential	93
F.10	Closing	93
G	Source code	94
H	Petri net with data for wheelset type 328	98

List of Figures

2.1	A basic and abstract example of a Petri net [41].	7
2.2	An example of a non-free choice workflow net [42].	7
2.3	A subset of the core BPMN constructs [47].	8
2.4	The relation between BPMN and the two levels of DMN illustrated in a single condensed model [82].	12
2.5	The decision mining quadrant, adapted from De Smedt <i>et al.</i> [18].	15
2.6	The approaches to converge decision and process mining, quadrant proposed by De Smedt <i>et al.</i> [18].	16
2.7	An overview of traditional process mining and its extensions, courtesy of Beerepoort <i>et al.</i> [90].	17
3.1	The DSRM and its implementation specific to this research project, adapted from Peffers <i>et al.</i> [27].	19
3.2	The Interview Protocol Refinement (IPR) cycle [117].	27
5.1	An initial overview of PM ² xDM its decision-related activities in relation to the stages of PM ² [99].	36
6.1	An overview of the physical factory layout of the wheelset revision process (in Dutch).	43
6.2	A schematic overview of the components of the IT architecture surrounding the wheelset revision process.	45
6.3	The initial model as shown in Disco after applying the data processing steps, 100% of activities and paths visible.	49
6.4	The model created in Disco after filtering for wheelset type 328, 100% of activities and paths visible.	50
6.5	Example of a decision point in the Petri net for wheelset type 328 where activity 337 is an optional examination step.	51
6.6	Example of decision points in the Petri net for wheelset type 328 with multiple choices and a repeatable step.	52
6.7	Example of an annotated decision point for the optional examination step in the Petri net for wheelset type 328.	53
6.8	Example of an annotated decision point with multiple exclusive conditions in the Petri net for wheelset type 328.	53
6.9	Example of annotated decision points with logical OR/AND and contradictions in the Petri net for wheelset type 328.	53
6.10	Part of the Petri net of wheelset type 328 where a livelock seems to be introduced by the guard expressions.	58
6.11	An overview of the PM ² xDM framework after implementation within the case study and the evaluation.	60
6.12	Excerpt of the enhanced Petri net for wheelset type 328 that shows both the conditions and path frequencies.	61
A.1	Overview of the search-screen phases and the resulting amount of papers, template by Haddaway <i>et al.</i> [196].	78
C.1	The PDD model of the PM ² methodology by Van Eck <i>et al.</i> [99].	83
C.2	The PDD model of the extended PM ² xDM methodology.	86
E.1	The consent form for the focus group or interview that was signed beforehand by the participants.	90

List of Tables

2.1	A rudimentary example of an event log, with supplemental <i>resource</i> attribute.	9
2.2	A basic example of a decision log.	14
2.3	The earlier example of an event log, however, the resource is now embedded in a complex set of attributes.	14
3.1	An overview of the mappings between the research questions and the different research methods used.	19
3.2	The dimensions and criteria for the evaluation within the case study context, adapted from Prat <i>et al.</i> [94].	26
4.1	Overview of data types and the frequency of applied learning algorithms for decision mining.	32
4.2	Overview of evaluation criteria and metrics identified within the literature.	34
6.1	An overview of the project team members of this research project, details accurate as of December 1st, 2023.	45
6.2	The columns as present in the initial data export CSV file for the year 2022.	46
6.3	An overview of the mapping between the data set attributes and the process mining concepts.	47
6.4	An overview of the characteristics of the baseline event log created with Disco.	48
B.1	The typology of the different approaches for decision mining.	79
C.1	Activity table for the PDD of PM ² , descriptions based on [59, 99].	84
C.2	Concept table for the PDD of PM ² , descriptions based on [59, 99].	84
F.1	A mapping of the evaluation criteria with the questions in the focus group / interview protocol, based on [195].	93

Nomenclature

Abbreviations

Abbreviation	Definition
BAM	Business Activity Monitoring
BI	Business Intelligence
BP	Business Process
BPA	Business Process Analysis
BPI	Business Process Intelligence
BPM	Business Process Management
BPMI	Business Process Modeling Initiative
BPMN	Business Process Model and Notation
BPMS	Business Process Management System
BPR	Business Process Re-engineering
BPfM	Business Performance Management
CoE	Center of Excellence
CRISP-DM	CRoss Industry Standard Process for Data Mining
CSF	Critical Success Factor
DFG	Directly-Follows Graph
DM	Decision Mining
DMN	Decision Model and Notation
DPA	Decision Point Analysis
DPM	Data Process Model
DPN	Petri Net with Data
DRD	Decision Requirements Diagram
DRG	Decision Requirements Graph
DSRM	Design Science Research Methodology
DSS	Decision Support System
EBSE	Evidence-Based Software Engineering
ERP	Enterprise Resource Planning
FEEL	Friendly Enough Expression Language
GPT	Generative Pre-trained Transformer
GSM	Guard-Stage-Milestone
IEC	International Electrotechnical Commission
IS	Information System
ISO	International Organization for Standardization
LLM	Large Language Model
NLP	Natural Language Processing
OMG	Object Management Group
PM	Process Mining
PM ²	Process Mining Project Methodology
SLR	Systematic Literature Review
UML	Unified Modeling Language
WFMS	WorkFlow Management System



Introduction

It was already in the early twentieth century that Frederick W. Taylor recognized that scientific management techniques required the accumulation and analysis of detailed work-related data [1], but these opportunities for increasing organizational effectiveness were then still limited by the technology of the time [2]. Nevertheless, his proposed principles of managing business processes analytically from a task-based perspective, to optimize the efficiency of human and physical resource utilization, can still be seen as fundamental building blocks of modern-age Business Process Management (BPM). And although we cannot deny that Taylor his work has been highly influential — and still is relevant — to organizational management practices nowadays, the contextual environment where it could be applied to is evolving rapidly, and his complementary ideas on reward and punishment have also been put in new perspectives in the meantime [3, 4].

With the advent of large-scale computing systems and low-cost storage, organizations became capable of producing and collecting vast amounts of data that is not necessarily usable *information* yet [2]. Unsurprisingly, the subsequent ability to analyze meaningful and relevant data, and convert this data into information and knowledge, and ultimately take timely action to favorably influence an organization is a key competitive advantage nowadays [2]. Although there is an abundant amount of data and information available, it becomes exponentially harder to distinguish the relevant from the irrelevant bits, and carry out meaningful forms of analysis. This phenomenon is referred to as “information overload”, where the sheer volume of data and information at some point even makes it impossible to effectively use it [5]. Therefore, emerging research into business intelligence and data analytics is increasingly focusing on technologies that are capable of dealing with high-velocity, high-volume types of data [6]. A concrete example is the domain of text analytics, where useful information and knowledge is extracted from large bodies of unstructured or semi-structured text [6].

Emerging technologies make it nowadays possible to process, query, and extract relevant information from data sets that are impossible to process manually. Analogous to what happens within the realm of natural language, the perspective of (big) data analytics is also shifting its paradigm towards mining more or less structured types of data streams [6]. Although the goal still is to “analyze critical business data to help an enterprise better understand its business and market and make timely business decisions” [6], the context is no longer only limited to structured content that is subsequently subjected to a classical statistical analysis. Areas of analysis have been expanding into the process perspective using more structured forms of sequential process execution data (process mining) [7], and mining of patterns contained within temporal aspects (time series analysis). Even more recently, this is being expanded with extracting relevant information from contextual high-speed data streams and sensor data [6]. In sum, the current era of big data analytics demonstrates that there are ample innovative opportunities to be leveraged that convert data into informational value to businesses. This can either be accomplished by analyzing existing types of data, such as large text bodies, in new ways using emerging technologies, as well as extracting new types of information and knowledge from forms of data that were not available earlier, such as decision process execution logs and streaming sensor data.

Chen *et al.* [6] further acknowledge that business intelligence and data analytics technologies can help organizations to “leverage opportunities presented by abundant data and domain-specific analytics.” Additionally, organizations that are top-performing typically employ agile, high-speed decision-making

based on rigorous analysis, and further use these insights for improving their day-to-day operations as well as to guide future strategies [8]. However, business analytics implementations do not inherently create value [9], especially since the technologies should be merely seen as tools — not drivers — that aid in dealing with information overload [5]. The rapid advances within information and communication technology have led to the rather paradoxical condition that, even though there is an abundance of available information, it is nowadays presumably more difficult to extract the relevant and useful information when needed [5]. However, the *potential* value already justifies organizational investments in such technologies to a large extent [9]. Although business analytics practices do not create business value from its direct effects, such effects, like improved information flows, ultimately combat information overload. The most significant contributions to improved organizational performance come from the less tangible indirect effects, such as improved customer knowledge and satisfaction [9]. Therefore, organizations should transform their decision-making processes to accommodate the use of business analytics [9]. The application of these techniques subsequently enables superior decision-making within organizations, which in turn leads to improved organizational performance [9]. Given that business analytics and decision-making are so strongly intertwined in relation to organizational value creation through the use of information and data, the remainder of this work will focus on the aspect of data-driven decision-making within organizations. More specifically, this work approaches data-driven decision-making from a business process analysis and management perspective, as processes are fundamental building blocks of any organization. Therefore, this study aims to investigate how to leverage data from non-process-aware information systems to create relevant and useful insights into the decision points that are present within a (business) process.

This chapter further elaborates on the scientific relevance and positioning of the research project associated with this thesis. First, it summarizes the related research gap that was discovered in the literature. Subsequently, it clarifies the research aim and objectives of this study.

1.1. Problem statement

Information systems are not only capable of managing and automating processes, but at the same time they also generate information about the underlying processes through which an organization performs its work [10]. This kind of information is usually referred to as an *event log*, which in its most essential form provides an ordered sequence of activities that actually happened during the process execution. The act of distilling a structured process description (e.g. a process model) from these logs such that it is consistent with the observed dynamic behavior, is referred to as *process mining* [7]. Process mining in its essence allows not only for the investigation of causal relations between activities, but additional data attributes for example also enable the investigation of performance (timestamps) and workload (resources) [7]. Therefore, both information availability [11] as well as subsequent data and event log quality [12] are still some of the highly relevant critical success factors (CSFs) when performing such process mining projects in general. With the abundance of data available nowadays, as mentioned earlier in this section, it becomes increasingly relevant to critically assess and evaluate event log quality [13]. While significant research has been carried out to address these latter aspects for event logs in the realm of process mining in general [14–16], only until very recently limited attempts have been done on enhancing event logs with data from the context of the process execution [17]. Nevertheless, the results demonstrate that it is worthy of continued investigation as the decisions — and in turn the process flow — can be influenced by event and case attributes [17].

Furthermore, the field of decision mining recently gained more widespread attention within the scientific community [18]. This development is grounded in the idea that at least some separation of concerns between business logic (rules, decisions) and processes should be achieved for the appropriate balance between flexibility, compliance, efficiency, and effectiveness of supporting information systems [19]. While processes and decisions are intertwined by nature, there are several addressable issues observed at their intersection. Firstly, when a process model incorporates too detailed decision paths, it becomes more or less a decision tree represented as a cluttered process model. These unnecessarily convoluted process models are hard to reuse and maintain [18]. Secondly, process models that contain hard-coded business rules, impose changes to the control-flow representation in case that the rules change at some point, even though the process itself is not adjusted. Thirdly, decisions might be the driver behind the activities and workflows of all process stakeholders, and as such they should be modeled separately to accurately document the related knowledge, and to allow for reuse beyond

a single process. Fourthly, a process might be the execution of a complex decision in itself, where the relationships between decisions should be explicitly modeled such that decision making can be facilitated by an optimal process. Fifth and final, processes that are highly dynamic, human-centric, and non-standardized could benefit from declarative process modeling where the principles are the same, but each case is genuinely distinct [19].

In sum, the aforementioned issues indicate that there does not exist a one-size-fits-all solution to integrate business logic with process knowledge. Nevertheless, an improved separation of concerns can be both beneficial to process model comprehensibility, as well as to provide new insights into underlying business logic and the relation with process optimization. De Smedt *et al.* [18] state that we should strive for a more holistic integration of process and decision models, which in turn allows for a better representation of the interplay between decisions and its dynamic influence on the control-flow [18]. Therefore, ideally one does not only want to explain routing and decision points in terms of control-flow [20], but be able to dynamically discover a decision model that covers the full process execution span [21]. This allows organizations to better analyze and further optimize their decision-making processes, which seems fundamental to yield significant improvements in organizational performance [9]. Although this is a promising research direction, we are still far out from such a holistic approach. Therefore, this research aims to pave the way forward by investigating the fundamental relation between decisions and processes from a data-centric perspective. While multiple definitions of process-related data exist (e.g. auxiliary, contextual, exogenous), we have chosen to consistently follow the definition of *endogenous* data [17]. This means the data follows the typical structure of an event log, for example from a system that also registers successful activity execution. The opposite is data that is extracted from the context of the process execution, which we refer to as *exogenous* data.

1.2. Research gap

Early seminal works on decision mining [22] and the related discovery of more data-aware process models [23] have proven that these enriched models describe reality better. At that point in time, there was no established formal notation for the modeling of decisions, hence the choice to extend an already common notation. In the mean time, more expressive — though less formal — modeling languages matured in terms of specification and implementation, such as BPMN, which in turn later gave rise to complementary extensions, such as DMN. While these integrated models allow for a more comprehensive view of real-world execution of processes, the discovery and consistent integration of such models still proves challenging [18, 20, 21, 24].

Several approaches that employ contextual data from the process environment to enrich both process mining and decision mining models have proven to be of significant value [17, 25, 26]. However, further research is needed to investigate the requirements for properly aligning the fit between the available input data and the required angle of analysis [18]. Furthermore, the aforementioned approaches all share the *modus operandi* of deriving decision points as a secondary step after discovering the control-flow structure. The current state of the art demonstrates that an opposite approach has been significantly under-explored. In this case, the overall logic of the decisions is derived first. This logic could either be derived from case-specific attributes that evolve and are available to query throughout the process, or specific decision data could be readily available. In a subsequent step, the available sequence information is employed to determine where and how sub-steps in the decision logic are related to the control-flow of the process [18]. The research gap presented by this alternative data-centric decision-oriented approach, is what this thesis project aims to fill.

1.3. Research aims and objectives

The primary aim of this thesis is to increase the understanding of the role of decisions within processes, at the intersection of process mining and decision mining research. As a secondary aim, this study explores the role and usage of endogenous process data to uncover and further analyze these decisions. These are subsequently not only presumably used as a starting point for process model and control-flow discovery, but also as way to complement and extend process models in a new suggested form of process enhancement.

The study at hand aims to construct an artifact based on theory, that can ultimately be valorized in practice. Therefore, iterative refinements are deemed necessary, thus the project comprises an initial design, and several validation, evaluation, and redesign cycles. Such an incremental study can be

classified as a design science project. Hence, it will be constructed using the widely-adopted design science research project methodology by Peffers *et al.* [27], as elaborated further in Chapter 3. There exists a universal template that can help with concisely formulating the associated research objective [28]. The execution of that template makes the Main Research Objective (MRO) to be as follows:

MRO: *Improve* the representation of the influence of decisions within process models *by* the design and empirical validation of a methodological framework for integrated decision and process mining *such that* endogenous process data can be used in process and decision mining analyses *in order to* present a more realistic perspective on real-world processes within their respective modeling artifacts.

This research project comprises four different components that altogether help reach the objective and address the research gap identified earlier. Firstly, an overview of related work and the theoretical background is built upon an investigation of the literature, as well as a complementary systematic literature review on the state-of-the-art regarding the specific niche of the usage of auxiliary data within process and decision mining practices. Secondly, based on this theoretical foundation, a methodological framework will be designed and developed, and relevant evaluation criteria and metrics will be selected from those that are well established in the scientific body of literature. Thirdly, the methodology will be demonstrated by application within a process mining project that follows an established methodology. Finally, a real-world evaluation of the framework in terms of the resulting artifacts and insights will be carried out, by means of a single case study. This will involve the application of the methodology in an actual business problem context at the Dutch national railway company *NS*, using the initial evaluation criteria. Subsequently, validation is carried out in conjunction with the relevant stakeholders. All of these are experts from the associated (business) domain, or experts with relevant knowledge of the process itself.

2

Theoretical background

This chapter presents the theoretical background in terms of context, concepts, and definitions, that forms the basis of this research project. First, the essentials of business processes and their management are treated, as well as how this was shaped by information systems. Second, the essence of related process modeling and the subsequent emergence of process mining is explicated. Third, the practice of decision management and its relation to business rules and logic is presented, to further define the landscape of decision modeling. Fourth, the emerging field of decision mining and its characteristics are disseminated. Finally, a synthesis of the literature that most closely relates to this work is presented. That work finds itself at the intersection of process mining and decision mining, and how these analyses can be enriched with endogenous data.

2.1. Business Process fundamentals

There is no consensus on how to formally define a business process within the literature [29]. However, at least “it contains purposeful activities, it is carried out collaboratively by a group, it often crosses functional boundaries, and it is invariably driven by outside agents or customers” [30]. More briefly, a business process is “the set of internal activities performed to serve a customer” [31], or more abstractly that it is “a sequence of activities which transform inputs into outputs” [29], or even more generic that it is a “set of partially ordered activities intended to reach a goal” [32]. In an encompassing assembly, a business process is a collection of activities that takes one or more kinds of input and creates an output that is of value to the customer. It has a clearly defined goal and is affected by events occurring in the external world or in other processes [32].

Some put a strong emphasis on that business processes entail *how* work is done within an organization, instead of focusing on *what* is being done [33, 34]. This is especially true now that business processes no longer solely focus on production work, but on the organization as a whole, for example also including office work. Therefore, both a more abstract definition of a business process, as well as emphasis on the workflow in the process itself instead of a focus the desired output for a certain market or customer, aligns well with business process management and business process modeling. After all, a model still is an abstracted view of the real world to make it better interpretable or understandable in a certain context [29].

As the description of organizations in terms of their business processes shifted away from the typical production perspective, the relation between the nature of business processes and their respective modeling approaches needed a more pluralistic definition [35]. Therefore, we first need to examine how business process re-engineering evolved into business process management and how the latter is supported by business process modeling. Essentially, re-engineering was re-engineered itself to take a broader perspective on processes, and follow an iterative approach to improvement over radical change, as is also common nowadays in software development [35]. Business process management can therefore be seen as a contingent and holistic approach over a radical, mechanistic approach of change, where information technology is more an enabler of change, over a driver of change in itself. As such, business processes did not fit a single archetypal definition anymore. Therefore, four perspectives on processes were identified [35]:

- **Deterministic machines:** a fixed sequence of well-defined activities or tasks performed by 'human machines' that convert inputs into outputs in order to accomplish clear objectives.
- **Complex dynamic systems:** a set of complex, dynamic and interactive features. An open loop system that adapts to a changing environment and is based on interaction and dynamic behavior.
- **Interacting feedback loops:** a closed loop system with intrinsic control. A network of interactions between internal structure and policies. The flows are regulated by policies (decisions), which represent explicit statements of actions to be taken in order to achieve a desired result.
- **Social constructs:** different perceptions constructed by various individuals and groups as a result of different frames of interpretation, shaped by beliefs, values, expectations and previous experience.

The most fundamental graphical notation of processes that contain behavior like choice, iteration and concurrent execution is a Petri net [36]. From a visual perspective, it has similar expressive and communicative value as flow charts, while the mathematical foundation behind its definition allows for rigorous analysis of the modeled systems using formal expressions and equations [37]. Therefore, Petri nets strike a balance between theory and practice. In contrast, BPMN provides a richer set of objects and notations. It is more expressive in terms of syntax, but it lacks the formal semantics of Petri nets and is therefore prone to modeling errors [38]. Albeit a trade-off, this aligns with the direction in which research in business process modeling is heading: towards a more integrated approach where the resulting models provide value in different areas of business interest [39]. As such, business processes should be examined from a combination of the four perspectives mentioned in the previous paragraph. Some of the analysis types which are elaborated later in this chapter, clearly put emphasis on one or more of these perspectives. However, as both Petri nets and BPMN are non-deterministic, it seems rather trivial to observe that modeling deterministic processes is not possible without complementary features. This is partially mitigated by extensions to these modeling approaches, and thereby one tries to find a balance between construct overload, redundancy, and excess, while at the same time increasing the coverage of the different analysis perspectives [38].

2.2. Business Process Modeling

This section provides the fundamentals of, and further expands on, the process modeling techniques that are most relevant to this research project, namely Petri nets and BPMN.

2.2.1. Petri nets

Petri nets are a graphical notation and precise mathematical modeling tool used to describe and analyze the flow of information and control within concurrent systems [37]. Petri nets are especially useful when parallel and asynchronous behavior is present, but there are constraints on the concurrence, precedence, or frequency of these occurrences, as is very common in workflows, data exchanges, information systems or even biochemical reactions [40].

Petri nets are directed bipartite graphs, that consist of a set of places, transitions, and arcs that connect them. Places represent conditions or states, transitions represent events or actions, and arcs represent the flow of tokens, which can be interpreted as resources, signals, or data [37]. Petri nets can capture both the structure and behavior of a system, allowing for the analysis of properties such as reachability, liveness, and deadlock-freeness. Furthermore, Petri nets can be combined with other modeling techniques, such as process algebra, to enable more expressive and precise modeling [37].

Figure 2.1 shows an elementary and abstract example of a Petri net with four places (A, B, C, D), five transitions (a, b, c, d, e) and twelve arcs. The Petri net contains two tokens, one in place A and one in place D, and therefore this Petri net also constitutes a *marking*. A marking represents the state of the modeled system at a particular point in time. When a transition fires, it always consumes one token from all inputs. As such, a transition is *enabled* (able to fire) if there is at least one token available to be consumed at *all* of its input arcs. Upon firing, the transition also produces new tokens at all of its output arcs. The behavior of the system is therefore limited by the availability and distribution of tokens, where a given marking can also be an initial or final marking [41].

The most relevant property of Petri nets to this research, is that it enables the modeling of *choices* apart from sequential, parallel or iterative behavior. In traditional workflow management and its respective modeling practice, Petri nets have shown to be of great value [42]. The formal semantics in addition




Figure 2.1: A basic and abstract example of a Petri net [41].

to the graphical nature, the ability to represent states over events, and the availability of a variety of (mathematical) analysis techniques are all good reasons to use Petri nets as the foundation of workflow modeling. For this reason, we should note the concept of workflow nets [43]. Workflow nets essentially are classical Petri nets with three additional constraints:

- It has a clear starting point i , a place that does not accept any tokens.
- It has a clear ending point o , a place where tokens are never consumed from.
- All other places and transitions are on a path from i to o .




Figure 2.2: An example of a non-free choice workflow net [42].

However, in terms of representing choice, some workflows exhibit behavior that requires unwanted combinations of constructs. For example, shown in Figure 2.2, when parallel behavior is mixed with choice, this could lead to the choice being influenced by the sequential order of the earlier activities. Under optimal circumstances, the routing of a case should be independent of the order in which tasks are executed [42]. A non-free choice workflow net poses restrictions to the available avenues of analysis, and therefore should be avoided. Nevertheless, workflow nets are used frequently in practice, and the semantics of some of the most widely used modeling languages, such as UML Activity Diagrams and BPMN, are converging towards Petri nets [43]. Furthermore, BPMN models can be converted into Petri nets for analysis and simulation purposes, and there are methods to transform one into the other with direct mapping relations to investigate certain mathematical properties [44]. BPMN provides a more encompassing set of elements with respect to syntax, but is also inherently less strict in terms of formal semantics, which provides the necessary rationale for these types of translations.

2.2.2. Business Process Model and Notation (BPMN)

The Business Process Model and Notation (BPMN) is the de-facto standard modeling language to graphically capture business processes [45]. It provides a formal but still expressive way of capturing business processes, as such that it is easily understandable by both end users as well as domain experts. The current version of the BPMN standard is significantly different from the initial release version [46]. It not only provides more syntactical constructs, it also formalizes the execution semantics for all BPMN elements and defines an extensibility mechanism for both process model extensions and graphical extensions [45]. Furthermore, it refines event composition and correlation, extends the definition of human interactions, and defines choreography and conversation models. On the other hand, it still manages to reduce inconsistencies and ambiguities that were present in earlier versions [45].

The latest version of the BPMN standard acknowledged the relevance of data to process modeling, analysis and execution. In fact, data are no more part of the artifacts but are a separate element category, including data input/output, collections of data objects, data stores and messages [45]. In addition to process and collaboration models, different diagram types were added to provide more meaningful models with respect to exchange and flows of data and information, such as choreographies and conversation diagrams.

Figure 2.3 shows the essential BPMN constructs. Each of these has a Petri net counterpart, where for example a start event is modeled as a (silent) transition between two places. Parallel gateways are modeled as a transition from a single to two distinct places, where XOR gateways branch from a single place to two distinct transitions and corresponding places for each. An example of a BPMN model, and how this can be extended with additional modeling constructs, is found in Section 2.5.3.

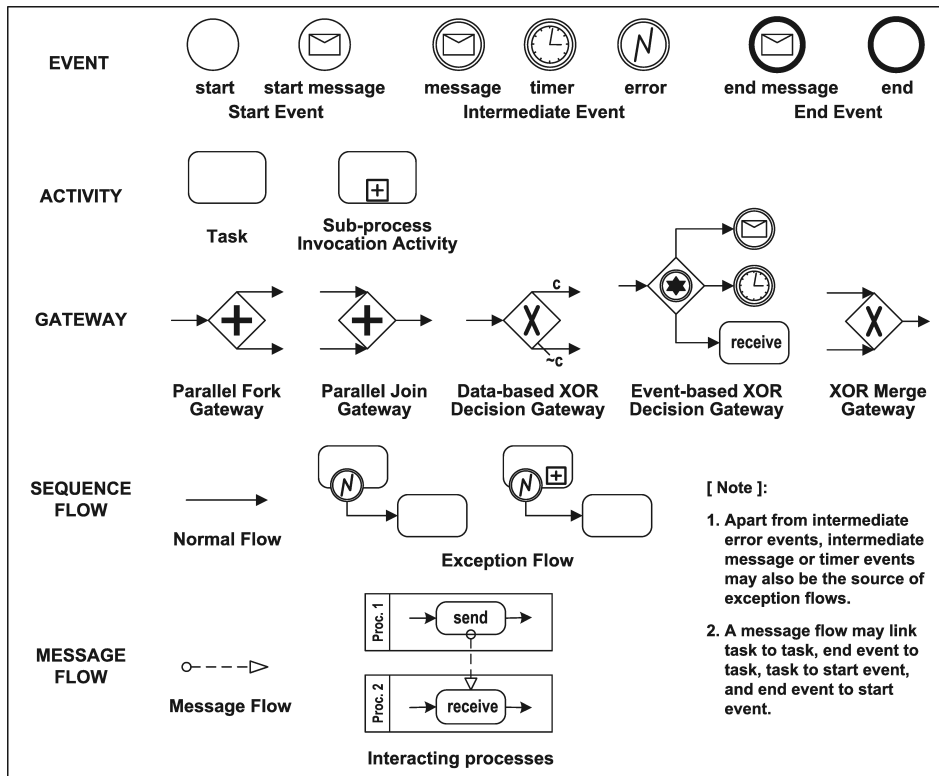


Figure 2.3: A subset of the core BPMN constructs [47].

2.3. Business Process Management and Information Systems

In roughly the past three decades, the ever-increasing adoption of information systems has not only led to the automation of already structured business processes, but it also introduced radical changes into other fundamental business procedures [48]. Instead of employing information systems merely to aid business process execution, companies live and thrive by information systems, and are even shaped and structured around them. Information systems therefore are an essential part of the strategy of every successful company [49], and there is a significant relation between information technology adoption and organizational change [50].

Now that information systems are omnipresent in both the executive and supportive role for business processes within organizations, new opportunities for process analysis and optimization have emerged. Following the workflow wave of the nineties, business process management (BPM) emerged as an overarching discipline, that we define as *supporting business processes using methods, techniques, and software to design, enact, control, and analyze operational processes involving humans, organizations, applications, documents and other sources of information* [51]. It is worth noting that this definition is purposely limited to operational processes, as the systems within the BPM sphere need to be process-aware [52]. Therefore, the processes at hand should be tangible or explicit in their execution in at least some form, otherwise

there is not enough information available to provide any kind of support.

These opportunities for process-support through software triggered the evolution of traditional workflow management, enterprise resource planning, and customer relationship management systems into integrated software suites that are commonly referred to as business process management systems (BPMSs) [52]. A BPMS is defined as *a generic software system that is driven by explicit process designs to enact and manage operational business processes* [51]. As traditional workflow management systems focused solely on one-time process design, configuration, and process enactment, little to no attention was paid to possibilities regarding process diagnosis through simulation, verification and validation [52]. BPMSs do inherently provide these functionalities, and the majority of modern information systems are capable of providing BPMS-like functions, or at least provide records of what happened during day-to-day usage. In turn, these BPM capabilities, and especially logging functionalities, provide clear opportunities for in-depth process analysis and subsequent improvement, which falls under the umbrella term of business process analysis (BPA) [51].

These logging functionalities sparked the idea that process diagnosis could also take place by investigating the execution data (logs) [51]. This would not only provide performance metrics about the process, but it could also demonstrate the actual process models [53]. These models provide a reliable ground truth, as they are not drawn-up from an ideal perspective by business experts, rather discovered in practice [54]. This not only underpins the relevance of evidence-based BPM [55], but also hints at opportunities for feedback loops, in an integral and iterative approach to continuous analysis and improvement [56]. Nowadays, these practices are commonly referred to as *process mining* [57].

2.4. Process Mining

In the past twenty years, process mining emerged and matured as scientific discipline within the realm of BPM, with a focus on process analysis using event data [58]. Whereas traditional data mining and analysis techniques tend to zoom in on a single step or part of the process, process mining follows a holistic and integral approach. The widespread adoption of information systems allows for analysis of the process from end-to-end, given that event data is available for different aspects of the process that are at least somehow software-supported [58]. Therefore, processes do not need to be IT-centric themselves.

The aim of process mining is to *discover, monitor and improve real processes by extracting knowledge from event logs* [55]. The smallest unit of examination is an *event*, where each event refers to an *activity* within the process (e.g. a single step that has been completed). Each event belongs to a particular *case*, which is one execution of the process, sometimes referred to as *process instance*. It is particularly important that all events are *ordered* sequentially, either by a numerical property or for example by a *timestamp*. In addition, each event could contain more information such as the *resource* involved with the activity or additional *data attributes* about conditions, the state or execution of the process. All events from a set of process instances combined form an *event log* [58].

Table 2.1: A rudimentary example of an event log, with supplemental *resource* attribute.

index	case_id	event_id	timestamp	activity	resource
1	1	1	2023-02-11T10:36:53+00:00	ticket_created	mary
2	1	2	2023-02-11T10:56:58+00:00	ticket_assigned	peter
3	1	3	2023-02-11T11:34:23+00:00	repair_planned	john
4	1	4	2023-02-12T15:34:23+00:00	repair_completed	john
5	1	5	2023-02-12T16:42:23+00:00	customer_informed	mary
6	1	6	2023-02-12T16:59:59+00:00	ticket_closed	peter
7	2	1	2023-02-13T09:45:56+00:00	ticket_created	mary
8	2	2	2023-02-13T10:15:51+00:00	ticket_assigned	john
9	2	3	2023-02-13T12:34:56+00:00	problem_resolved	peter
10	2	4	2023-02-13T13:37:01+00:00	customer_informed	mary
11	2	5	2023-02-13T13:45:56+00:00	ticket_closed	peter

In Table 2.1 above, a fictitious example of a rudimentary event log from a help-desk process is shown. In this example, there are two distinct cases with six and five activities that are performed respectively. If timestamps would not be available, the *event_id* could still be used to determine the sequential order of events within a single process instance. This example contains the minimal properties needed for process mining, except for the *resource* property which already supports other forms of analysis.

Three types of process mining activities are commonly identified, that are possible with said event logs: discovery, conformance checking and enhancement (or extension). Firstly, process discovery is the creation of a model solely based on the observed events. Secondly, conformance checking deals with the verification if an event log complies with an (existing) process model, and the other way around. Finally, as opposed to conformance checking, process enhancement does not compare a model with reality [58]. Instead, it tries to change, correct, extend or enrich the already existing model. This can either be already accomplished by examining timestamps and calculating time differences to demonstrate service times, and to indicate possible bottlenecks. Additionally, one could include the resource attribute to for example identify resources that are underutilized, frequently execute related activities, or lead to specific or unwanted behavior. These different activities in turn correlate with four dominant analysis perspectives within the process mining paradigm [59]:

- **Control-flow:** concerns the (sequential) order of activities within the process execution.
- **Time:** investigates e.g. the throughput, service and waiting times based on different calculations with timestamp data.
- **Organizational:** considers organizational factors, such as the utilization of (human) resources within the process. Also denominated as the resource perspective.
- **Data:** examines attributes and values that are specific to a certain case, such as data attributes that vary throughout the process. Sometimes also referred to as the case or (process) instance perspective.

These perspectives are not mutually exclusive but rather of complementary nature. The control-flow perspective is usually regarded as the most fundamental, as it produces the as-is process model from the execution data. However, the other perspectives have proven to also provide highly relevant avenues of analysis [59].

2.5. Decision Management & Modeling

Decision management is a set of methodologies and technologies used to automate and improve decision-making processes within an organization. It involves the use of data analysis, business rules, business logic, and decision models to facilitate informed and consistent decision making [60]. In the context of business processes, decision management is a critical component that helps organizations streamline their decision-making processes and increase efficiency [32]. In most cases, business processes involve multiple decision points that require input from various stakeholders. Thus, decision management is critical in enabling organizations to make effective decisions at each point of the business process. Logically, an holistic approach to integrating decision management into business processes can help organizations achieve better results by reducing decision-making cycle times, minimizing errors, and increasing overall efficiency [61].

From a historical perspective on decision management, we should note the existence of a separate long-lasting discipline of research in and development of Decision Support Systems (DSSs), that finds its origins in the availability of modern computing technology [62]. While a DSS is just another form of business intelligence (BI) technology to aid an organization in decision-making [62], it is different in the sense that it is mainly focused on supporting and improving the managerial perspective on decisions [63]. Although BPM and BI initiatives can be of complementary nature in the sense that they both contribute to organizational performance management [64], the difference in perspective cannot be ignored. As this research explicitly positions itself within the context of business processes, the investigation of system-supported decision-making as an isolated form of business intelligence is considered out of scope. Nevertheless, the avenue of research into how process and decision mining can assist in decision making is still an open area of exploration in this research project.

2.5.1. Business rules

In the context of business, a *business rule* can be defined as a specific statement or guideline that dictates how an organization conducts its operations, processes transactions, and interacts with its customers, partners, and stakeholders [65]. Business rules are typically created to ensure consistency, accuracy, and compliance with legal and regulatory requirements [66].

Business rules can cover a wide range of topics, such as pricing, discounts, payment terms, delivery schedules, quality standards, data privacy, security, and risk management [67]. They can be expressed in various forms, such as natural language, decision tables, decision trees, or computer code, depending on their complexity and level of automation [67].

The development and management of business rules can be facilitated by specialized software tools, such as rule engines, rule editors, and rule repositories, which enable business analysts and domain experts to define, test, and modify rules in a collaborative and transparent manner [68]. The application of business rules can also be automated through integration with business processes, workflows, and applications, enabling real-time decision-making and improved operational efficiency [69].

2.5.2. Business logic

In contrast to business rules, business logic refers to the overall framework of principles and processes that guide the decision-making and operations of an organization [70]. Business logic encompasses its goals, strategies, and operating principles, and it may involve a combination of business rules, best practices, and industry standards [70]. The key difference between business rules and business logic is that business rules are formal, specific, granular guidelines that dictate how individual transactions or interactions should be handled (business policy) [65], while business logic is a broader framework that guides the overall strategy and decision-making of an organization [70].

From an information systems perspective, business logic is the part of a system which determines how data is transformed or calculated, and how it is routed to people or other software systems [65]. Anything that is a process or procedure can be deemed business logic, and anything that is neither a process nor a procedure is a business rule [71]. For example, the screening of a new employee is a workflow or process, consisting of steps to be taken, whereas stating that every new employee must undergo screening is a business rule. Furthermore, business logic is procedural whereas business rules are declarative [72].

The study of business rules and business logic typically involves examining how organizations develop and implement these guidelines and frameworks [70], how they use technology and data to automate and enforce them [73], and how they balance the need for consistency and compliance with the need for flexibility and adaptability in a dynamic marketplace [74].

2.5.3. Decision Model and Notation (DMN)

The relevance of explicit business rule and logic management to improved business outcomes and increased organizational competitiveness [75], is leading to an increased integration of decisions and business processes within the business process management paradigm [76]. It is not up until quite recently that there has been a proper integration between business process and decision management [77]. Meanwhile, process modeling languages have often been abused for either implicit or explicit decision modeling [78], or decisions were represented using separate entities, such as knowledge models and ontologies [78]. Not only does this leave the overall view of a decision and its interplay with other decision and data requirements dispersed and hard to maintain [78], but it also results in overly complex models that suffer from reduced comprehensibility [79] and that are hard to maintain [80]. Therefore, the process logic should be isolated from the underlying rules and decisions, to an extent that one preferably reaches a true separation of concerns [80]. Decisions are then no longer implicitly defined in process structures, and subsequently any modification of the decision logic does not necessarily need to be reflected in the process model [77].

In response to the need for declarative business process modeling [81], and to actually break the ground for a more holistic business process and decision modeling approach, the OMG introduced the concept of the Decision Model and Notation (DMN) in 2015, as a complement to the BPMN modeling language [82]. DMN constitutes a decision model that comprises two complementary layers, namely *decision requirements* and *decision logic* [77]. Decision requirements can be captured in one or more Decision Requirement Graphs (DRGs) that together form a Decision Requirement Diagram (DRD). Each decision requirement can have an underlying specification of decision logic, where DMN provides

a language for further specifying that in the form of FEEL (Friendly Enough Expression Language), and corresponding notations (boxed expressions and decision tables) [78, 82]. It is possible to link a BPMN model to its DMN complement by means of *decision activities*, where a decision is associated with the activity in which it takes place [82]. The associated decision model further specifies the decision requirements and the inner decision logic of the activity [77]. In essence, DMN is the *functional* — and thus also *declarative* — counterpart of *procedural, imperative* BPMN, to efficiently capture and structure decision logic [81].

Figure 2.4 shows a rather simple BPMN model of an application process, accompanied by a DMN model that further specifies a specific decision activity that provides input for the further process routing. Instead of implicitly hiding this decision within the control-flow, the specifics are made explicit by this composite activity. The DRD at the decision requirements level shows the input parameters for the actual routing decision. The rectangular boxes represent the actual decisions, which is the act of determining an output from a number of inputs, using decision logic which may reference one or more business knowledge models [82]. A decision requires some form of input to determine its output, which can be other decisions, input data (boxes with rounded corners), business knowledge models, knowledge sources, or a combination. Business knowledge models can recursively have other business knowledge models as their input, an expression thereof, or an authoritative knowledge source [82].

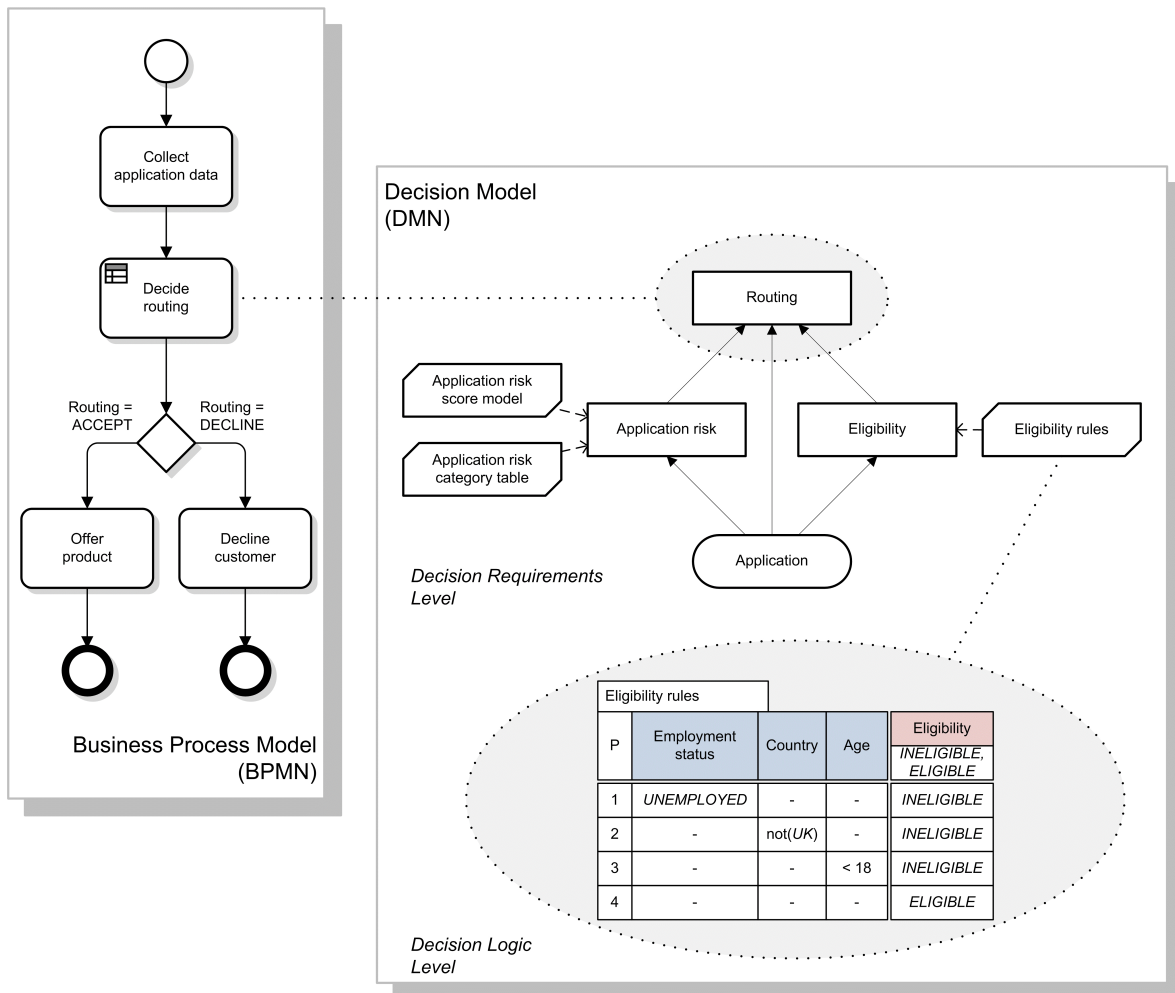


Figure 2.4: The relation between BPMN and the two levels of DMN illustrated in a single condensed model [82].

2.6. Decision Mining

A significant proportion of the most widely available and used process mining tools and techniques, tend to mainly focus on the classical control-flow perspective [83], while the influence of data attributes on routing of a case in a process was not widely investigated earlier [22]. Within the processes considered, and their respective process models, usually a considerable amount of branching or alternate paths is present, which represent different types of decisions in operational (workflow-like) processes [84]. The different types of decisions that can be examined using process mining are as follows [84]:

- **Design-time decisions:** explicit decisions made during design and initial modeling of a process. These decisions are inherently embedded into specifications and models that were created to specify the process. An example would be that some activity always needs to be executed after another due to regulatory requirements.
- **Configuration-time decisions:** decisions related to the configuration and customization of a process or (software) system specifically tailored to the organizational context in which it is used. A specific ERP system might for example comprise several modules which are either enabled or disabled and therefore allow certain functionality to be used.
- **Control-time decisions:** more of an ad-hoc type of decisions to manage running processes within their context. These decisions are therefore contextually dependent, and although they can evolve over time, they relate to the process as a whole instead of individual process instances (cases). This could for example be the allocation of capacity or particular prioritization of cases to overcome bottlenecks, due to extraordinary high demands in certain periods of the year.
- **Run-time decisions:** decisions that are taken at the level of individual process instances (cases). This is the type of decision that is most commonly found in process models. For example a purchase order that exceeds a certain value threshold, will take a path through the process that require additional authorizations. The properties (e.g. certain data attributes) of a particular case usually determine this kind of decisions.

The focus area of this work is on run-time decisions, as these exhibit the closest relation to case data attributes. However, process mining is relevant to all of the above types of decisions, as it can e.g. uncover deviations from a-priori design decisions (process discovery), non-compliance and configuration issues (conformance checking), and bottlenecks (process extension and enhancement) [84]. With respect to the branches related to run-time decisions, these are usually explicitly preceded by decision points [85], where the underlying decision rules are based on, or represented as, non-arbitrary values of data attributes. These external details, supplementary to an event log, can reveal useful information for process optimization and analysis [17, 25, 86]. The extraction and exploitation of these additional data attributes for determining the branching characteristics of discovered process models is referred to as decision point analysis (DPA) [23]. If the decisive attributes are used to further enhance the process models and to make predictions on future decisions, we refer to the more general practice of decision mining [22]. However, please note that the terms decision point analysis and decision mining might be used interchangeably.

DPA in conjunction with exogenous data can be employed by annotating events in an event log, by first linking, slicing, and transforming the supplemental data [17]. Linking is the process of assigning the relevant exogenous data to a particular trace. Returning a subset of the data for each distinct event is referred to as slicing, and a transformation function is applied to this subset which returns transformed attributes. Next, each event is annotated with the resulting attributes. Subsequently, a discovery function yields a Petri Net with Data (DPN) that is extended with guard conditions (preconditions) to include external factors not represented in the endogenous event log, resulting in an exogenous DPN (xDPN). The event log is then aligned with the xDPN, and a visualization presents the subset of exogenous data set relevant for each transition in a so-called traceback xDPN [17]. The resulting event log can be leveraged by a data-aware process discovery technique to annotate the decision points with preconditions based on the exogenous data in a discovered process model.

An observed limitation of DPA is that it only considers single attribute values for each decision point [25]. In practice, however, data attributes appear in a variety of shapes, where most process environments nowadays comprise at least multiple *endogenous* data streams in the form of transaction logs or audit trails, which contain only internal data directly related to the process execution and the ultimate goal of the process [17]. Additionally, this is complemented with *exogenous* data, data that

is not related directly to the process but rather to the context the process is executed in [17]. This information describes the environment as accurately as possible over time, for example by consecutive interval-based measurements, opposed to the single point-in-time emission of events that represents the endogenous process mining data. These attributes could for example be continuous measurements of sensor values throughout the process [17]. This data is usually known as time series data, and it can exhibit patterns and intervals that can be relevant to the enhancement of decision mining [86]. However, it can be challenging to assess the suitability of certain data assets within this context [87]. Furthermore, there are different approaches to incorporate interval-based time series data into event logs [25, 26, 88]. Ideally, one wants to retain the possibility to apply existing process mining tools and techniques, and strive for a holistic and integrated decision and process model representation [18].

While decision mining and process mining should not be considered as a single field, there is a significant amount of overlap [89]. Recent work has illustrated that investigation of how data influences the workflow can be valuable and of complementing nature to existing process mining analyses [21]. However, variables or data attributes that are used by activities can be influenced by decisions throughout the workflow without an actual effect on the activity execution order. Therefore, there need not be a correlation between the workflow data and the control-flow of the activities. The majority of current decision mining techniques still take the control-flow as a starting point, and alternate approaches have not yet been properly addressed in the literature [21].

2.6.1. Decision logs

In Section 2.4, the concepts of process mining and event logs were introduced. Within the realm of decision mining, it is necessary to also note that there exists the concept of a *decision log* or *decision event log*. Table 2.2 shows an example of such a decision log, with a case identifier, two criteria and a conclusion (status). While this certainly allows us to deduce some of the rules that might underpin the conclusion, there are no sequential patterns that are needed for application of existing process mining techniques and algorithms. Although the decisions in the log could be linked to specific cases in a separate event log by their case identifier, the dimension of analysis is limited to a single item per case. This makes it impossible to uncover decision patterns that change over time, and also does not allow to link decisions to different activities within the process. An integrated example of an event log, with

Table 2.2: A basic example of a decision log.

index	case_id	grade	verified	status
1	1	8	true	pass
2	2	4	true	fail
3	3	6	false	fail
4	4	7	true	pass

Table 2.3: The earlier example of an event log, however, the resource is now embedded in a complex set of attributes.

index	case_id	event_id	timestamp	activity	attributes
1	1	1	2023-02-11T10:36:53+00:00	ticket_created	{{res=mary;attr1=x1};attr2=y6;attr3=z3}
2	1	2	2023-02-11T10:56:58+00:00	ticket_assigned	{{res=peter;attr1=x1};attr2=y7;attr3=z4}
3	1	3	2023-02-11T11:34:23+00:00	repair_planned	{{res=john;attr1=x1};attr2=y8;attr3=z5}
4	1	4	2023-02-12T15:34:23+00:00	repair_completed	{{res=john;attr1=x1};attr2=y9;attr3=z6}
5	1	5	2023-02-12T16:42:23+00:00	customer_informed	{{res=mary;attr1=x2};attr2=y1;attr3=z7}
6	1	6	2023-02-12T16:59:59+00:00	ticket_closed	{{res=peter;attr1=x2};attr2=y5;attr3=z4}
7	2	1	2023-02-13T09:45:56+00:00	ticket_created	{{res=mary;attr1=x1};attr2=y6;attr3=z3}
8	2	2	2023-02-13T10:15:51+00:00	ticket_assigned	{{res=john;attr1=x1};attr2=y7;attr3=z4}
9	2	3	2023-02-13T12:34:56+00:00	problem_resolved	{{res=peter;attr1=x1};attr2=y1;attr3=z9}
10	2	4	2023-02-13T13:37:01+00:00	customer_informed	{{res=mary;attr1=x2};attr2=y1;attr3=z7}
11	2	5	2023-02-13T13:45:56+00:00	ticket_closed	{{res=peter;attr1=x2};attr2=y6;attr3=z3}

attributes that could represent some form of a decision structure, is shown in Table 2.3. In this case, the attributes are encoded in a single item with multi-dimensional key-value pairs. The evolution of these parameters is retained over the whole process execution span of each case, and as such it is possible

to identify relations between certain activities and combinations of parameters that exist. In this case, by structure alone, we observe that *attr1* specifies some attribute about the resource. However, the attributes can be shaped differently, such as in regular columns in a tabular format, but this is only a syntactical difference that does not yield much semantic value. Such an integrated example of an event log with (decision) attributes nevertheless provides a first step towards a more holistic approach of process and decision mining [21].

2.6.2. Towards holistic process and decision mining

To be able to accurately position decision mining research within the different data analysis and business process management paradigms, a decision mining quadrant was proposed as shown in Figure 2.5 [18]. The quadrant supports two dimensions, namely the decision control-flow relation (vertical) and decision model maturity dimension (horizontal), where each of those poses constraints to the necessary input data and the associated techniques [18]. On the vertical axis, this distinguishes whether or not the decision making is driven by a decision flow. In terms of process mining, the left hand side of the model captures the control-flow first approach, where decisions are an integral part of the process itself and therefore the resource, performance, or general data perspective, are out of scope (Q2). In Q1, this is exhibited from a data mining perspective, where consequentially the dynamic aspects of the decision process are ignored as these approaches usually represent either isolated single-stage or multi-stage decisions [18]. In the example in Q1, this is represented as a decision tree result being inserted into a neural network.

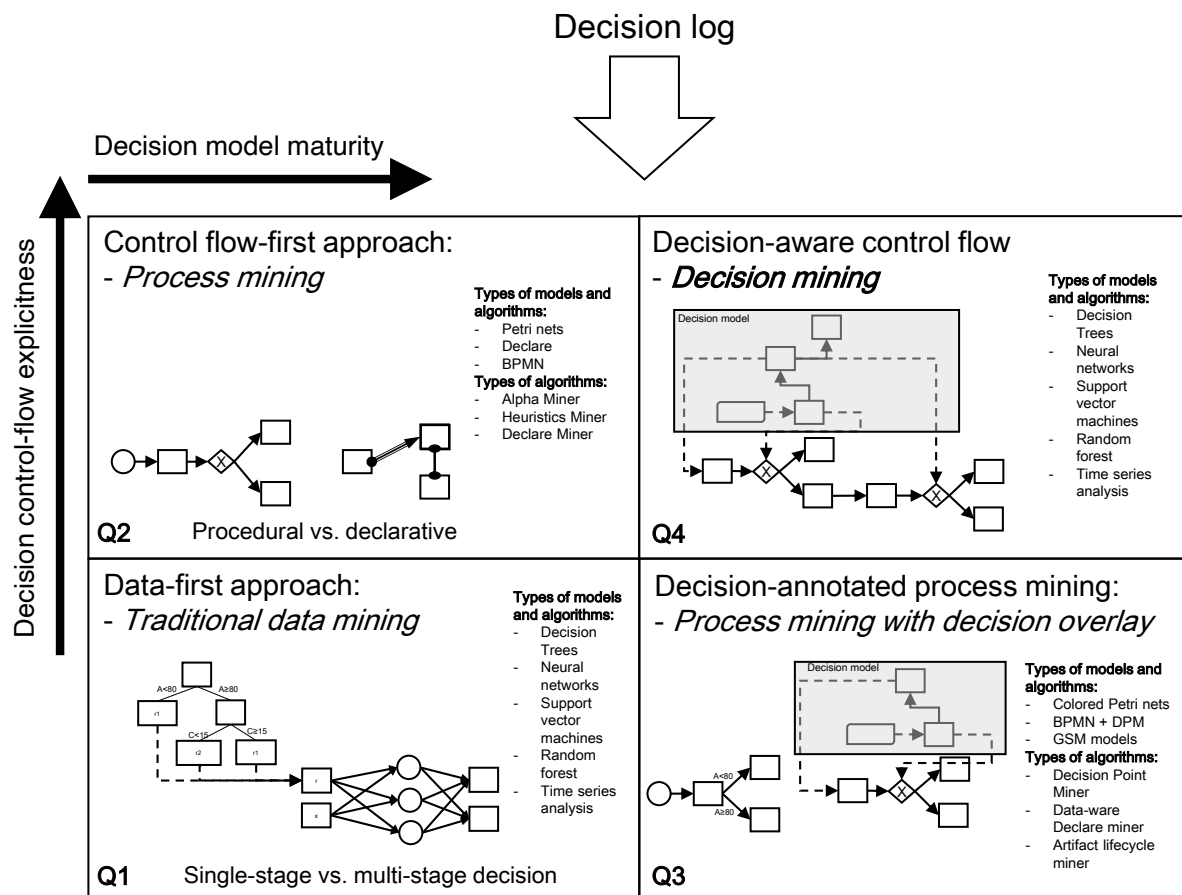


Figure 2.5: The decision mining quadrant, adapted from De Smedt *et al.* [18].

The horizontal dimension of the quadrant relates to the employment of a decision model (maturity) [18]. A decision model is inherently present at the right hand quadrants of the model, while the left hand side either captures the decisions within the control-flow or a single decision is dissected further with a traditional data mining approach. The right hand side of the quadrant therefore differs vertically

in terms of holism. Q3 offers a decision mining overlay as an annotation to a process model, where not the whole workflow is considered. This allows for inconsistencies to occur between the process model and the decision model [18]. The right upper corner stipulates the most holistic view, where the decision model is tightly integrated with the process model and it covers all decision points [21]. This additionally could enforce the need for alignment and consistency, which in turn could benefit the accuracy and possibly also the richness of the resulting models [24].

If we consider the the most holistic perspectives, we should focus on the right hand side of the quadrant. Q3 and Q4 actually enable the use of a mix of event-based and instance based data, where Q3 prioritizes the initial discovery of a control-flow and subsequently enrich that with the decision points that influenced the workflow based on the instance attributes. On the other hand, Q4 takes a fully integrated approach where there are no longer fixed decision points but rather decision inputs that are reused throughout the model [18]. In this type of approach, the control-flow that exists over the correlation of the event data is also mined, and represented as an additional layer with full closure, giving rise to a truly holistic model representation.

Figure 2.6 demonstrates the dimensional differences between Q3 and Q4 approaches. Both approaches are capable of providing a dynamic type of analysis, however the Q4 approach assumes the explicit availability of a decision log. This research aims to bridge the gap between Q3 and Q4 approaches. Therefore, the data perspective of Q3 in terms of sole availability of a case-based event log is combined with the analysis capabilities of the Q4 approach. Instead of relying on the explicit availability of a decision log, the exogenous data of the process is used as a substitution to subsequently derive the decision model constructs and elements that result in a full decision model extension to the process model.




Figure 2.6: The approaches to converge decision and process mining, quadrant proposed by De Smedt *et al.* [18].

As mentioned by De Smedt *et al.* [18], an inverse approach has been under-explored up until now. While the starting point traditionally is the extraction of a control-flow model from event data with a decision perspective placed on top, we could also consider this the other way around. In this approach, a decision model is enhanced with sequence based constructs from the data semantics, or it is holistically linked to a process model as in Q4 [18]. However, the need for a more comprehensive or even inverse approach to decision mining in relation to the usual control-flow-discovery-first method, also poses additional constraints to the available data. The usage of endogenous process data available within the process for either of these approaches, poses a research gap in itself. With our development of a methodological framework to leverage the potential of this data within this context, we also aim to at least partially address this gap.

In conclusion, decision mining is one of several emerging extensions of traditional process mining that has gained interest in the scientific community. Traditional process mining focused predominantly

on the discovery of the order of activities of a single process from event data. However, the increasing availability of other types of data expanded the discipline to new purposes and perspectives [90]. In a nutshell, decision mining investigates the influence of data attributes on the choices made in a process, thus representing the data perspective of process mining. Figure 2.7 shows a handy overview of the subfields of process mining. Although it is not meant to be complete, it provides a comprehensive picture of how the different disciplines relate to each other. For a more elaborate explanation and details of the concepts presented, refer to the conference paper by Beerepoot *et al.* [90].

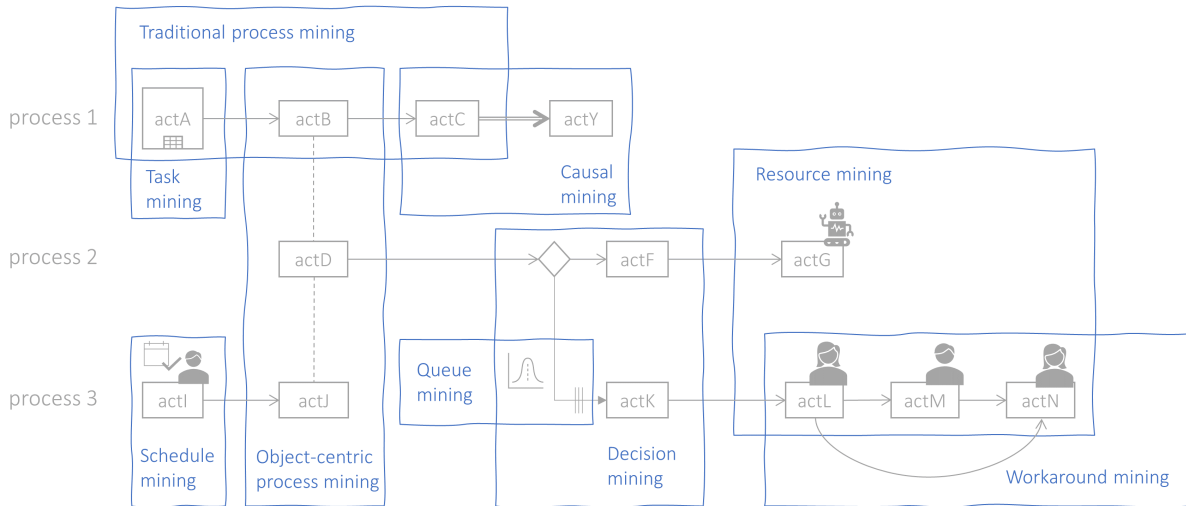


Figure 2.7: An overview of traditional process mining and its extensions, courtesy of Beerepoot *et al.* [90].

3

Research design

In this chapter, we present our proposed research approach. This approach entails a combination of research methods following an established methodology to fulfill its aim. The following sections describe the research design, research questions, and their relation with the different methods. We conclude this chapter with an overview of the threats to validity, and the associated mitigations.

3.1. Methodology

This project follows the design science research methodology (DSRM) proposed by Peffers *et al.* [27]. Design science is a particularly well-suited approach when the artifact under consideration will be constructed and evaluated in an iterative fashion, where the associated cycles of validation and rework are part of the design and development phase [91]. Next to that, the design science approach allows for the application of several well-renowned frameworks that cater towards objective evaluation strategy designs for IS artifacts [92–94]. The aforementioned evaluation phase is of paramount importance to provide proof of the actual value of the proposed approach and associated methodological framework. Additionally, it provides a feedback loop to another design iteration. The overall DSRM process therefore contains six distinct phases: problem identification and motivation, objectives of a solution, design and development, demonstration, evaluation, and communication [27]. Each of these different phases require different, sometimes overlapping, research methods for their execution, that ultimately contribute to the solution for a problem by designing and evaluating an artifact within the relevant context. The IS artifact that is being created is a methodological framework for the application of endogenous process data into a process mining workflow or project methodology.

Adhering to the DSRM provides an opportunity to contribute significantly to knowledge and ensures consistency, integrity, and scientific rigor among different IS research projects [95]. In terms of knowledge contribution, the project at hand can be considered an *improvement*, where a new solution is developed for a known problem [95]: consistent integration and visualization of decision flows in process models. On the other hand, it can also be seen as an *exaptation*, where a known solution is extended to new problems [95]: extending a process mining methodology to leverage additional analysis potential with endogenous data.

The overall execution of the DSRM implementation aims to achieve the Main Research Objective (MRO), where each of the phases can be employed to answer one or more of the Supporting Research Questions (SRQs). How the different phases of the DSRM relate to the phases of the research project is shown in Figure 3.1 below, where each of them is explicated further in the upcoming sections. The remainder of this chapter provides an overview of the respective DSRM phases, with the related research questions and how these will be answered. Although our research objective was already formulated during the introduction in Chapter 1, for convenience it is presented here again.

MRO: *Improve* the representation of the influence of decisions within process models *by* the design and empirical validation of a methodological framework for integrated decision and process mining *such that* endogenous process data can be used in process and decision mining analyses *in order to* present a more realistic perspective on real-world processes within their respective modeling artifacts.




Figure 3.1: The DSRM and its implementation specific to this research project, adapted from Peffers *et al.* [27].

3.1.1. Problem identification and motivation

This phase entails the identification of the research problem, and the associated value that is stipulated by solving the problem, which inherently provides the motivation for pursuing this study. The problem statement, research gap, and contextual background knowledge were all identified by means of a narrative literature review [96], and the resulting knowledge is disseminated in earlier Chapters 1 and 2 respectively.

SRQ1: *How do the disciplines of process mining and decision mining relate to each other in terms of their context, fundamental concept definitions and data requirements?*

3.1.2. Definition of the objectives for a solution

The definition of solution objectives not only relies on inferences that can be drawn from the problem statement, but they also need to be solidified in terms of quantitative and qualitative meaning [27]. Therefore, there is also a need to further structurally investigate the existing body of knowledge by means of a systematic literature review [97]. This not only aids in being able to accurately position this research, in addition to defining what sets this research apart from earlier work, but also to investigate well-grounded ways of quantitative and qualitative evaluation of the artifact under study. In turn, this is beneficial to the subsequent iterative phases of design and development, validation, demonstration, and evaluation as the output can serve as relevant input for design and evaluation strategies. The specifics of the systematic literature review, e.g. the protocol, is defined in Section 3.2.

SRQ2: *Which methods or techniques already exist to detect decision points within a process using event data attributes?*

SRQ3: *What is the current state of research in enhancement of process models with decision information?*

SRQ4: *What are relevant criteria and metrics to evaluate process models enhanced with decision information?*

Table 3.1: An overview of the mappings between the research questions and the different research methods used.

SRQ	NLR	SLR	DSRM	Case Study	Focus Group
1	✓				
2		✓			
3		✓			
4		✓			
5			✓	✓	✓
6			✓	✓	
7			✓	✓	
8			✓	✓	
9			✓	✓	✓

3.1.3. Design and development

This phase comprises the actual creation of the artifact, in this case a *methodological framework* [27]. To determine the desired functionality, its architecture, and ultimately to design a first draft iteration of the inherent methodology, the synthesis of the output of the performed SLR serves as the primary input for this phase. The output of this phase is an experimental version of the methodology, which will be demonstrated and evaluated, where iterative refinement could occur after both subsequent phases, up to an extent where even the objectives could need additional refinement.

SRQ5: *How to design an effective methodology to enhance process models with decision information?*

SRQ6: *How to convert event data attributes into decision conditions?*

SRQ7: *How to integrate decision information into process models?*

SRQ8: *Where does the approach and methodological framework under investigation fit into the application of a process mining project methodology?*

3.1.4. Demonstration and evaluation

The demonstration phase of the DSRM should be considered as the first instantiation of the methodology in practice and is directly followed by the evaluation. As this is a functional demonstration that is carried out after the creation of the artifact in a real-world context, it is considered an *naturalistic, ex post* form of evaluation [93].

The evaluation of the newly created artifact is conducted in the form of a single-case study. This provides insights into the effectiveness and applicability of the methodology, and is therefore of paramount importance to the succeeding of this research project. Although the research project as a whole can already be seen as a *mixed methods approach* [98], the relevance of this phase as form of naturalistic evaluation [93] makes that the case study comprises multiple distinct projects in itself:

- **Process model discovery:** following the (adapted) PM² methodology [99], the processes under investigation will be examined and the reference process model will be determined.
- **Endogenous data discovery:** using a subset of relevant procedures from the CRISP-DM approach [100], the entry points of endogenous data that could serve as input for our proposed method will be investigated.
- **Methodology implementation:** the proposed methodology will be applied to the discovered event logs and the endogenous data that has been gathered in the previous two sub-projects. The output will be a model extension for the reference process model.
- **Methodology evaluation:** using the evaluation metrics and criteria defined earlier, the methodology will be evaluated within the real-world context of the case study.
- **Meta-evaluation:** by means of a focus group with process, domain, and business experts, the relevance of the methodology and value of the model extension will be validated within the real-world process context.

The output of the latter two sub-projects can serve as input for an additional iteration of methodological design (redesign). A further detailed dissemination of the case study elements and its context is presented in Section 3.3.

3.1.5. Instantiation within a project methodology

As mentioned earlier, within the demonstration and evaluation phase of the DSRM, elements of the CRISP-DM [100] and PM² [99] project methodologies will be employed to execute the project. As the final component of the demonstration phase, the following steps are sequentially executed to examine where our methodology fits best within the workflow of such other methodologies.

1. **Understand the existing process mining methodology:** before you can instantiate a method, you need to have a good understanding of the methodology that you plan to adapt. This involves studying the principles, stages, and techniques of the methodology.

2. **Identify the project requirements:** you need to identify the specific requirements of the project, including the goals, scope, and data sources. This information will help you determine the changes you need to make to the methodology, and where to make them.
3. **Determine the method instantiation approach:** there are different approaches to method instantiation, and you need to select the one that is most suitable for the project. For example, you can modify existing techniques, add new techniques, or combine existing techniques in a new way.
4. **Adapt the methodology:** using the approach you have selected, you can start adapting the methodology to fit the project requirements. This may involve changing the order of stages, modifying techniques to work with specific data sources, or creating new techniques to address unique challenges.
5. **Validate the instantiation:** once you have adapted the methodology, you need to validate the instantiation to ensure that it is effective and efficient in achieving the goals of the project. This may involve testing the instantiation on a sample of the data or comparing it to alternative methodologies.
6. **Document the instantiation:** finally, you need to document the instantiation so that it can be shared with other project stakeholders and used as a reference for future projects. The documentation should include the rationale behind the changes made to the methodology, the specific adaptations made, and any testing or validation results.

SRQ9: *How does the proposed methodology perform in a **real-world** organizational context?*

3.1.6. Communication

The communication phase of this project comprises ten different deliverables. These are in presumed chronological order:

1. **Long proposal:** a written document that explicates the background and context of the problem and the associated research approach. It concludes the long proposal phase of the overall research project.
2. **Colloquium presentation:** a mandatory presentation of the research project long proposal in a colloquium session of the Master Business Informatics, to an audience of fellow students, researchers and professors. To gather formal feedback on the project from the academic audience from mostly a high-level perspective.
3. **Intermediate stakeholder presentation:** a company internal presentation that presents the research design, including an overview of the relevant projects that are analyzed within the case study, and the intended valorization of the artifact for the organization.
4. **Intermediate scientific presentation:** a presentation of the research project proposal to the primary supervisor, second examiner, external supervisor, and other interested members of the research group. The goal is to have a brainstorm session, and to gather an additional round of ideas and (informal) feedback to further scope the research project.
5. **Executive summary:** a written statement or small report (e.g. in the form of a (internal) white paper) that summarizes the outcomes that are relevant in light of the internship project associated with this thesis. It describes the aspects that are especially relevant to the organizational context of the research project, and how the main takeaways of the case study can be utilized to its full potential within the organization.
6. **Final stakeholder presentation:** a company internal presentation that accompanies the executive summary, with an overview of the relevant project outcomes with the involved stakeholders within the case study context.
7. **Thesis document:** the written research report that covers all aspects of the research project, conforming to the academic guidelines of what it should entail.
8. **Thesis defense:** the formal public defense of the thesis, including a presentation and plenary discussion of the matter at hand, with an audience comprising the involved researchers (supervisors and examiners), external stakeholders, and other interested parties.
9. **Scientific publication:** a scientific article to be submitted to one or more of the leading conferences on business process management and process mining.

10. **Open data set:** if possible and feasible, the real-world data set that is used for the demonstration, validation, and evaluation of the framework under investigation should be made available as reusable open data for further academic research. Of course, this is at the sole discretion of the supplier of the data (the organization associated with the case study), and if appropriate measures for anonymization and pseudonymization of the data can be applied without nullifying the purpose of the data in the first place.

3.2. Systematic literature review

According to Kitchenham *et al.* [97], research within software engineering should adopt an approach of Evidence-Based Software Engineering (EBSE). In this context, *evidence* is defined as “a synthesis of best quality scientific studies on a specific topic or research question” [97], and the preferred method to gather this is by means of a Systematic Literature Review (SLR). While an SLR is not only a methodologically rigorous review of research results, opposed to ad-hoc selection of literature common to narrative reviews to present an initial contextual overview, but it also supports the creation of evidence-based guidelines for practitioners to use and to provide appropriate software engineering solutions in a specific context [97]. According to the secondary study by Xiao and Watson [101], any type of SLR comprises at least the following steps:

1. **Formulate the research question:** define the scope of the review and determine the concise research question that the review will address.
2. **Develop inclusion and exclusion criteria:** specify the criteria that will be used to determine which studies are relevant to the review.
3. **Search for relevant studies:** this involves identifying and obtaining the relevant studies through a systematic search process, which includes using appropriate search terms, databases, and search engines.
4. **Evaluate the quality of the studies:** an assessment of the quality and relevance of the studies that have been identified using predetermined criteria, such as study design, sample size, and data analysis methods.
5. **Extract data from the studies:** this involves extracting relevant data from the selected studies using a standardized data extraction form.
6. **Synthesize the findings:** summarizing and synthesizing the findings of the selected studies, highlighting patterns, themes, and gaps in the literature.
7. **Interpret the results:** the findings in relation to the research question and context need to be interpreted, and subsequently this leads to drawing conclusions based on the evidence.
8. **Report the review:** the review concludes with documenting the review process, including the methods used, the findings, and any limitations of the review.

In addition to defining a protocol and following a systematic approach, we opt to use a tool-supported method to assist with searching, evaluating, and data extraction and synthesis [102]. Although we opt to apply a similar methodology and phases as described by Bandara *et al.* [102] in detail, our approach with respect to the actual tools used in the process will be different. Instead of using separate, non-integrated, tools for the different aspects of the review, like NVivo for the analysis and EndNote for the article management, our study will primarily be conducted using Elicit¹.

Elicit is an online tool that is specifically tailored to support scientific literature reviews, where it has been configured to allow for a non-biased, objective, search for the most relevant papers to a specific research question [103]. While it also employs Generative Pre-Trained Transformer (GPT) language models, it has been designed to consistently reproduce the same results, as opposed to ChatGPT that is now widely under investigation [104]. As ChatGPT its responses are conversational and context-dependent in nature, they are more difficult to validate and reproduce. Even though Elicit should be as unbiased as technically possible, it remains a tool that assists — but is not meant to replace — researchers and domain experts. One might state that it becomes even more of paramount importance for them to critically assess and reflect on (the relevance of) the results. However, the systematic literature review is not the main objective of this study in itself, its main goal is to fill gaps

¹<https://elicit.org/>

in the literary context. Furthermore, it aims to provide design building blocks from and evidence for the performance of similar proposed methodologies, and to investigate optimal ways and metrics for evaluation.

As this research project as a whole is subjected to considerable time constraints, completeness is not the primary purpose of the literature review. Nevertheless, Elicit uses Semantic Scholar² as its searching back-end, and it has proven to cover nearly 99% of the papers present in high-quality secondary studies within software engineering [105]. It therefore seems reasonable to follow the aforementioned novel tool-supported approach to alleviate some of the cumbersome and time-consuming aspects of such an endeavor. This is of course only a valid argument as long as the followed protocol has still been properly documented, which we will summarize in the upcoming subsections, to ensure validity, and especially consistent reproducibility.

3.2.1. Research questions

The SLR at hand aims to answer SRQ2, SRQ3, and SRQ4, as defined in Subsection 3.1.2. A novelty of using such a tool, that is semantically aware of the similarities between keywords, is that it allows for natural input of the search query, opposed to complicated conjunctions of different boolean expressions to present a composite search string. As search queries, we will separately input the relevant keywords that relate to the aforementioned SRQs into Elicit. For each query, the top 42 papers will be considered as input for our filtering and selection process. This number is not arbitrarily selected, but Elicit presents additional results in batches of seven. The discriminative power of the ranking algorithm is assumed to be exhausted after incorporating twice the amount of articles, as only the top twenty results have been ranked in two stages by distinct algorithms within Elicit³.

Elicit does not support Boolean queries. Therefore, the search queries for the research questions are defined as string concatenations as follows:

- **SRQ2:** detect+decision+point+process+method+technique+data+auxiliary+contextual+exogenous
 - Abstract keyword requirement filter: decision+process+mining
- **SRQ3:** process+model+enhancement+decision+data+auxiliary+contextual+exogenous
 - Abstract keyword requirement filter: decision+process+mining
- **SRQ4:** metric+criteria+evaluate+model+quality+decision+process+data+auxiliary+contextual+exogenous+dmn+bpmn
 - Abstract keyword requirement filter: decision+process+mining

3.2.2. Inclusion and exclusion criteria

The following operational inclusion and exclusion criteria to answer all of the related SRQs have been defined. Literature will be included:

- **IC1:** That is in the field of information systems or software engineering research.
- **IC2:** That at least mentions the domain of business process management, business process modeling, process mining, decision point analysis, decision modeling or decision mining in their abstract.
- **IC3:** That relates to the employment of exogenous process data.
- **IC4:** That proposes and evaluates a method, tool or other artifact using predetermined evaluation metrics and criteria.

The SLR will explicitly disregard literature that:

- **EC1:** Is written in any other language than English.
- **EC2:** Has been published before 2001, as the scientific fields of process mining and decision mining were non-existent before that time.
- **EC3:** Is not any kind of scientific publication. This includes items such as gray literature, white papers, business reports, news articles and slide decks.

²<https://www.semanticscholar.org/>

³<https://elicit.org/faq#appendix-how-does-elicit-work>

3.2.3. Quality criteria

As mentioned in the exclusion criteria above, this SLR will regard all possible types of scientific publications, including but not limited to books, articles, journal items, conference proceedings, and other works such as Master theses or PhD dissertations. Given that the aim of the SLR is to pragmatically find papers that have proposed specific goal-oriented artifacts at the still nascent intersection of process and decision mining, it makes no sense to explicitly put an arbitrary lower bound on the amount of citations an article should have to determine relevance nor quality. Furthermore, the number of citations is not a key quality indicator in dimensions apart from scientific impact and relevance [106]. Nevertheless, the work should preferably be peer-reviewed, which usually is the case for conference proceedings and journal articles. Furthermore, the following general quality criteria are evaluated in a binary way (true or false) for each resulting work:

- **QC1:** The study should present a clearly formulated problem statement or research gap.
- **QC2:** The study clearly describes the metrics and criteria, and the context of evaluation (real-world, artificial).
- **QC3:** The findings are presented and explained in an understandable way, supported by quantitative and/or qualitative evidence.
- **QC4:** The study presents a discussion on its findings, or (systematically) addresses its limitations.

3.3. Single-case study

A single-case study is performed as part of this research project to evaluate our approach and methodological framework in a real-world context. The case study protocol is found in Appendix D and follows the method proposed by Yin [107], not only for the purpose of scientific rigor, but also because it stipulates the combination of quantitative and qualitative evidentiary sources [108]. The case study protocol is designed according to the guidelines put forward by Pervan and Maimbo [109], as these are specifically tailored to the case study method within the context of information systems research. The case has multiple units of analysis (the process) within its context. Additionally, the overarching organization is a common factor. Therefore, the study is considered an embedded, single-case design [107]. Although the study at hand primarily adopts the Yinian perspective on case study research [107], as this research takes place within the context of software engineering, we also have to make note of the seminal work by Höst *et al.* [110]. As the work by Höst *et al.* [110] is largely based on the earlier revisions of the work by Yin [107], there is a significant amount of overlap between either of the approaches. Apart from that case studies should exhibit a clear research question, have clearly defined cases and use multiple sources of data collection, a synthesis of the most important aspects yields the following key characteristics [107, 110]:

- The case should be representative of a relevant phenomenon in software engineering, such as a software development process or a software product, or it should be selected based on its ability to provide insights into a theoretical or practical issue.
- Data analysis should be based on a framework, method, or other artifact that is relevant to the research question. Pattern-matching logic can be applied, which involves comparing the data with the existing theory.
- Findings should be validated through triangulation, i.e. the use of multiple sources of data *and* multiple methods of data analysis (e.g. quantitative and qualitative). One could also employ rival explanations for validation, i.e. the consideration of alternative interpretations of the data.

This section further disseminates the context in which the case study is conducted, the procedures employed for data collection, data analysis, and how the findings are presumably synthesized. Please note that the DSRM employs an iterative approach that fits case study research particularly well. As a consequence, the case study details may slightly evolve as the research project progresses further over time.

3.3.1. Context

The embedded single-case study will in its entirety be conducted at the Dutch national railway company, fully denominated “N.V. Nederlandse Spoorwegen” which indicates that it is a public limited company.

Nevertheless, the company is usually referred to in its shorthand naming form as NS.⁴ A brief overview of relevant facts and figures of the largest rail transport organization in the Netherlands⁵:

- A complex and diverse organization that not only concerns the actual transport of travelers by train (NS Operations), but also the exploitation of train stations and its facilities like public toilets and snack shops. Additionally, there are complementary travel services that support the last miles of door-to-door journeys, such as short-term bike rental (OV-fiets), secure bike parking, and luggage lockers. As such, a wide variety of processes and related data is available throughout the organization.
- NS has embraced innovation and wants to become a more data-driven organization. A Tech6 has been defined with the six most promising innovative technologies that are being considered and implemented, where process mining is one of those, among others such as real-time asset monitoring with sensor data, 5G networking, 3D printing and extended reality.
- A Center of Excellence (CoE) for the sole purpose of disseminating process-centric ways of working and process modeling was only recently established. With a vast amount of processes being present throughout the organization, its goal is to not only document all business processes and procedures, but also investigate opportunities where insights from process analysis allow for further optimizations and increased operational efficiency. Consequentially, process mining has not yet matured within the organization. Therefore, the organization is looking for innovative ways to assist them with exploiting process mining opportunities.

The context of the largest rail transport organization in the Netherlands is not only selected for reasons of convenience, but it is also societally relevant as it supports almost a million travelers with their train journeys each day. Furthermore, such a complex organization accommodates a significant amount of different processes and their respective contexts. It therefore also has shown to offer a unique and interesting perspective on challenges and opportunities related to process mining project approaches in practice [111]. Additionally, several projects in the past have demonstrated that it can presumably serve rather well as a testbed for process mining applications [112–114].

3.3.2. Cases

Due to organizational constraints, a single case was ultimately selected for inclusion. A total of four cases have originally been elected for inclusion, to provide an appropriate variety with respect to context. However, this was considered unfeasible within the constraints of the project. The selected cases shares the minimum requirements that an internal process mining endeavor has at least been attempted. As such, documentation, (reference) models, and supplemental materials are all readily available to some extent. Eventually, case B was left as the only available case to investigate. However, the list of cases originally considered was as follows:

- **Case A:** the incident handling process of a service desk, regarding service disruptions and maintenance of a variety of assets at railway stations, such as escalators, elevators, and toilets.
- **Case B:** the wheel-set overhaul process in a robotized factory setting. The system was designed and implemented to generate rich and detailed event log data throughout the end-to-end process.
- **Case C:** the unplanned withdrawal of train sets from the active fleet. Allows for investigation of the relation with the maintenance process and the diagnostic data that is systematically collected from the trains.
- **Case D:** the maintenance process of the OBIS (On-Board Information System) that is present in the majority of train types. For example the displays that provide travelers with information on the current and upcoming train stations.

3.3.3. Data collection

The data collection process for the single-case study at hand comprises a threefold strategy. Firstly, the gathering of contextual documents of any process mining projects that have already been carried out within the project context. These can appear in the form of (archival) documents, reports and presentations, relevant meta-data about the project execution (e.g. the method that has been followed),

⁴<https://www.ns.nl/en/about-ns/who-are-we/history>

⁵https://www.nsannualreport.nl/FbContent.ashx/pub_1001/downloads/v230414160311/NS_annualreport_2022.pdf

and resulting reference process models. Second, the retrieval of the raw data-set(s) that have been used throughout the project, where the proposed methodology will be applied to. The quantitative aspect of the evaluation will be executed using the same criteria as the initial demonstration. The execution of the case study entailed the quantitative aspect of this research. Finally, validation and evaluation data was collected through a focus group [115] with relevant stakeholders, such as business and domain experts. The focus group with the experts was used to evaluate the output in terms of artifacts and insights of the methodology after it has been applied. The emphasis will be on criteria that evaluate environmental factors, such as the consistency with the organization, output understandability, and utility in a business context [94]. This presumably allows one to also further evaluate the inherent business value proposition of the proposed methodological framework. An overview of the additional adapted criteria for the case study context is shown in Table 3.2. As the focus group followed a structured approach, a focus group protocol [115] was developed with initial and follow-up questions relevant to the different contexts and the respective predetermined evaluation criteria [94]. The proposed collection of data from various sources using different methodologies allows for increased validity and reliability of case study-based research through so-called *triangulation* of evidence [116].

Table 3.2: The dimensions and criteria for the evaluation within the case study context, adapted from Prat *et al.* [94].

Dimension	Criterion	Sub-criterion	Explanation
Environment	Consistency-People	Utility	Does the artifact provide relevant practical value to an individual?
Environment	Consistency-People	Understandability	Are the structure and procedures of the artifact easy to understand?
Environment	Consistency-People	Ease of use	Is the artifact easy to employ within an overall methodology?
Environment	Consistency-Organization	Utility	Does the artifact provide relevant practical value to the organization?
Environment	Consistency-Organization	Understandability	Is the output of the artifact comprehensible within the organizational context?
Environment	Consistency-Organization	Fit	How well does the artifact align with the organizational environment?
Evolution	Robustness	-	Is the artifact resistant to changes in the environment?

Focus group session procedure

The composition of the focus group in terms of participants can be found in Table 6.1 in Chapter 6. For the focus group session, a separate protocol was developed based on the relevant evaluation criteria [115]. The questions within the protocol were created according to the guidelines of Roberts [117]. The questions were crafted as meticulously as possible by the researcher to satisfy the intended goals. The questions were discussed in advance of the focus group with one of the experts. This could be considered a pilot activity, as shown in the Interview Protocol Refinement cycle in Figure 3.2 [117]. However, no additional iteration of the IPR cycle was performed due to time constraints. The audio of the focus group was recorded using the video conferencing system available in the meeting room (Microsoft Teams). The researcher took notes and the audio recordings were saved and uploaded to an online tool called Amberscript⁶, so that they could be consulted afterward for clarification. In addition, the information was stored in the case study database and kept for reference purposes and further analysis.

3.3.4. Data analysis

In addition to the quantitative evaluation in the form of the case study implementation, the qualitative components of the case study need to be analyzed as well. Given the diverse amount of evidence types that is being collected, the descriptive and explanatory power of the case study can be harnessed by means of extensive synthesis [118]. Therefore, the approach at hand comprises within-case analysis to construct interpretations across the different types of evidence, should there be considerable differences between them [118]. Additionally, narrative synthesis is conducted to form the basis of the case study reports and to elaborate on the chain of evidence present in the qualitative data. Furthermore, it is used to provide logical rationalizations based on the results of the evaluation, primarily being the results of the focus group with the different relevant stakeholders from the business and domain context [118].

3.3.5. Ethical considerations

As is common in many areas of research that involves working with sensitive data or human participants, case study research also requires ethical considerations at design time [116]. At the commencement of this project, agreements have been settled between the host organization NS, Utrecht University, and

⁶<https://www.amberscript.com/en>

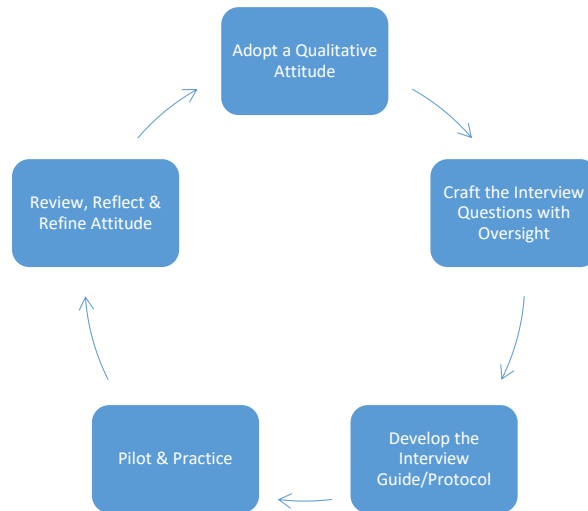


Figure 3.2: The Interview Protocol Refinement (IPR) cycle [117].

the involved researcher, in the form of a Work Place Agreement (WPA). Additionally, NS has provided a separate confidentiality agreement that is applicable to internships, and has been signed by both parties (organization and the researcher). If specific requests for the usage of organizational data are made (e.g. for process mining analyses), an internal Data Usage Board (DUB) of NS will assess the legitimacy of such a request, and approve or deny it based on the provided rationale. Furthermore, they will provide constraints and requirements for the processing of the data, such as needed measures with respect to data anonymization or minimization. With respect to the participants involved in the different interviews, the principles of informed consent apply. This will be adhered to by presenting an informed consent form to the participants to sign beforehand, of which the design will be presented in one of the appendices. Finally, some parts of the research, e.g. part of the case study protocol and its associated data, might remain under embargo or is kept strictly confidential at the sole discretion of the involved organization.

3.4. Validity threats

This section describes a synthesis of the relevant validity threats and possible remedies and mitigation strategies, that are applicable to this research project. With respect to validity, there are two different perspectives: one that focuses on the secondary study in the form of the SLR, and the other that refers to the primary study component in the form of the multiple-case study. The validity threats that are relevant to the SLR, are treated from the perspectives presented in two seminal tertiary studies on validity threats to SLRs within software engineering research [119, 120]. With respect to the multiple case study component, the validity threats are addressed conform the intersection of guidelines presented by Runeson and Höst [116], Yin [107], and Wohlin *et al.* [121].

3.4.1. Systematic literature review

Any SLR suffers from a distinct variety of validity and reliability threats, and it can be a challenge to rigorously assess them [119]. Most of which are mitigated by encompassing search strategies [119], the upfront determination of well-defined criteria [120], and accurate documentation and execution of the review protocol [120]. The threats and their respective proposed mitigations are identified and structurally presented in the upcoming subsections.

Conclusion validity

Within SLRs, the conclusions drawn from the actual review can be deemed valid when they are also reproducible [120]. Therefore, one should accurately document and report the followed procedures for searching, selection, and analysis. Furthermore, to avoid search and selection biases, a meta-search engine was employed for discovery, instead of a manual selection of several digital libraries, and the research question is used as the actual search query. Any bias with respect to publication has

been overcome by considering unpublished scientific works and not regarding a minimum number of citations. However, any form of gray literature has been excluded, which might exclude relevant articles from practice. Finally, we should note that we only consider articles in English, as this is considered the standard for written works in information and computing science. Although this introduces a language bias, the insignificant amount of works in other languages should not pose a real threat to conclusion validity.

Internal validity

As internal validity is closely related to conclusion validity, similar biases and their mitigation are deemed the most important. As an SLR does not measure any statistically significant relationship between certain variables, it is the process itself that should be rigorously designed and free of selection biases [122]. By means of accurate specification of the research procedure and using an objective tool-set, this threat is attempted to be properly mitigated. Additionally, personal biases could be present in the researcher executing the procedure. These could be mitigated by involving at least one additional researcher within the process. However, given the objectives of the SLR and the constraints on the allotted time-frame, this seem to be too much of an effort to apply.

Construct validity

The outcomes of the SLR are intended as input for the design cycle iterations of the DSRM. In this sense, construct validity entails the identification of the correct operational measures for the concepts under investigation, which in turn relates to the objective selection of study objects that they originate from. Logically, there is a significant amount of overlap with internal validity assessment. Nevertheless, most significant threats that arise from the formulation, search, and selection procedures have already been covered earlier. In addition, construct validity threats mostly arise from the usage of incomprehensible venues or databases, inappropriate inclusion and exclusion criteria, time-span limitations, or the lack of expert evaluation [119]. Given the limited scope and our meta-search method, lack of venues does not seem to pose a considerable threat or at least it has been partially mitigated. Furthermore, the criteria for in- and exclusion have been accurately defined, and we also consider recent works. Nevertheless, a more longitudinal study approach might derive richer insights. The lack of expert evaluation is not considered a threat here, as a cross-synthesis of the evidence for evaluation methods is applied, and experts are consulted within the case study context.

External validity

The scope of the SLR is quite narrow on purpose and therefore it inherently exhibits a low generalizability, being constrained to the specific domain of information systems and software engineering research. In addition, it is impossible to generalize beyond the primary studies that are considered [119]. However, this is not an overly large concern as it serves as specific input to our proposed method design that hopefully is applicable to a wider array of contexts. Nevertheless, we should acknowledge that the investigation of such a nascent body of scientific literature additionally restricts the generalizability further, as some of the method implementations could be specifically tailored to their respective context. At least the cross-comparison and identification of similarities and shared characteristics among the investigated works might partially mitigate this concern.

Reliability

Reliability within the scope of an SLR is mostly procedural of nature, as consistency and stability of the results mostly depends on the execution of the same steps with respect to searching, selection, and filtering of the body of literature under investigation [120]. Both internal consistency as well as test-retest reliability are guaranteed by sufficiently detailing the procedures that are carried out. As exhaustion and saturation are not primary objectives of the SLR, this does not seem to be an overly large concern. With respect to test-retest reliability, we should note that the review considers all work available at the time of writing. Therefore, to ensure repeatability over time, one should explicitly consider the time-frame from the defined criteria, as in the meantime new publications might have appeared. Inter-rater reliability is not applicable to this review, as it concerns only a single author. What this entails with respect to internal validity has already been explicated earlier.

3.4.2. Single-case study

Within the realm of case study research, there is a common trade-off to be made between the level of realism of the context that the phenomenon under investigation is situated in, and the associated amount of generalization that is still possible with the gathered insights [116]. Nevertheless, combining complementary — quantitative and qualitative — sources of evidence to obtain multiple perspectives on the phenomenon is the best overall strategy to ensure validity and reliability of case studies [123]. How this fits in, and relates to, the categorization of these different threats is further explained in the subsequent part of this chapter.

Conclusion validity

Conclusion validity within case study research can be reached by maintaining a chain of evidence [107], which usually is achieved by means of triangulation [116]. While ensuring validity with respect to the conclusion can actually be seen as a goal of the case study research procedure itself, as the case study is the complementary naturalistic evaluation to the synthetic evaluation in the form of a simulation. There are two aspects of the case study component of this research that are affected by threats to conclusion validity, namely the applicability to, and representativeness of, the evaluation criteria with respect to the selected cases, and the meta-evaluation conducted by means of semi-structured interviews. For the first threat, the mitigation strategy is to employ a wide variety of well-established evaluation metrics and criteria that have been applied to similar cases, as identified within the SLR. With respect to the semi-structured interviews, the mitigation tactic is to have multiple disjoint interview sessions for the respective cases, and as such a form of triangulation can be applied. Nevertheless, given the time constraints, we should acknowledge that the sample sizes are still relatively small and it is hard to account for personal and organizational biases present with the involved participants.

Internal validity

Within case study research, internal validity threats relate to the instrumentation, history and maturation, and selection and researcher biases [122]. Regarding the instrumentation, validity across the cases is ensured by employing the same evaluation criteria as identified by the SLR, and developing a concise case study protocol with predetermined analysis tools [124]. The risks with respect to history and maturation are deemed low, given that the project is executed within a relatively short time-span of three to four months. With respect to the selection bias, we should acknowledge that the case selection process has been carried out from a convenience perspective. However, the cases have been selected such that they exhibit both similar and dissimilar characteristics in a way that can be evaluated with the same criteria and procedures, but still are sufficiently different. Finally, researcher bias might be present throughout the study. This will be partially mitigated by the iterative nature of the method design and the meta-evaluations (interviews), where inputs from different sources are triangulated. Unfortunately, having a second researcher individually perform the analysis is beyond the scope and time constraints of this project. However, evaluation with supervisors and other peers should provide some relief to this threat.

Construct validity

Construct validity threats within case study research arise from insufficient upfront conceptualization and operationalization of the concepts under study [121]. This can be avoided by defining clear research questions and objectives, and by not studying the same phenomenon from different angles [122]. To mitigate these threats within this particular research, a clear case study protocol with transparent metrics and criteria has been defined which allows the reader to reconstruct the path from questions to conclusions in the form of a chain of evidence [107]. Additionally, given that a multiple-case study is performed, cross-case synthesis [118] will be applied between the different cases to perform triangulation and in turn increase objectivity of the constructs under study.

External validity

Case studies have been known to emphasize external validity at the expense of internal and construct validity [122]. Nevertheless, a nested multiple-case study within one organization does not provide a solid foundation for generalization. However, the characteristics of the selected cases — processes in this context — can be of more general nature and thus also present insights that are relevant to other organizations of similar size, complexity, and structure, and in turn allow for at least some form

of analytical generalization [122]. To achieve this, a cross-case analysis of at least four similar case studies should be performed [122], which is of course beyond the scope of this research but offers opportunities for future research directions. To furthermore mitigate remaining threats to external validity, the rationale and context are clearly defined within the to-be defined case study protocol [109].

Reliability

Threats to the reliability of a case study can roughly be divided into two categories, on one hand inter-case and intra-case reliability, and on the other hand biases within data collection and analysis [122]. Most of these are merely a composite of threats from the aforementioned categories, where biases within data collection and analysis can be avoided by defining and following a predefined case study protocol and documenting the results in a case study database. In other words, observing complete transparency with respect to analysis and reporting of the facts. With respect to ensuring inter-case reliability, the same methods should be applied to the different cases. However, intra-case reliability is harder to maintain, as the context of the case (e.g. the organization) and the interpretations of the researcher are subject to evolution and change over time. This can not only happen due to influence of the case study execution and results itself, but also by dynamic or environmental factors that transcend the scope and boundaries of any case study.



Systematic literature review

This chapter summarizes the results of the SLR component of this research, which aims to answer SRQ2, SRQ3, SRQ4 through enumeration and synthesis of related work. Answering these questions should provide knowledge of the detection of decision points within processes using different forms of endogenous and exogenous data, how process analyses and models can be enriched with this additional perspective, and how this can formally be evaluated using relevant criteria and metrics. This in turn provides input for the subsequent design and development, validation, and evaluation cycles of the DSRM that result in the initial version of the methodological framework.

4.1. Overview and characteristics of the SLR procedure

The SLR acts as a review that serves as background for an empirical study, as opposed to stand-alone pieces of work [101]. It has been executed in line with the EBSE tradition [97], and the results are processed in accordance to the procedure of framework synthesis [101]. Subsequently, the synthesis and the extracted information from the relevant set of papers is used for inception of the initial methodological framework, and to provide the applicable evaluation criteria and associated metrics. From a procedural perspective, the review is reported using the PRISMA 2020 guidelines for systematic reviews and meta-analyses [125]. A graphical overview of the results of the searching, screening and selection procedure is provided in Appendix A.

4.2. Identification of decision points with contextual process data

The (visual) identification of the actual decision points within any given process by means of a process model is a rather trivial operation. With respect to Petri nets, a decision point is defined as a place with more than one outgoing transition [126]. In a BPMN model, this is represented as a branching path that follows a XOR-gateway. In terms of the most simple notation, a Directly-Follows Graph (DFG), this would entail an activity that is followed by more than one distinct activities. However, DFGs are known to be limited in their capability of clearly expressing divergent and convergent behavior, and the related forms of interleaving semantics (e.g. loops versus intended routing logic, or any other type of mixed behavior) [127].

Although the awareness of decision points within a process could demonstrate a logical starting point for an in-depth process analysis, it is not the sole presence of these points that provides significant analytical value [128]. For that, more behind-the-scenes information about each decision point needs to be investigated, preferably based on data that relates with the underlying decision parameters [25]. As explained in Chapter 2, it could therefore be beneficial to separately derive the decision structure from the actual (exogenous) data that reflects the values of these parameters, and thus inherently provides the decision criteria. Several approaches that apply a similar workflow have been systematically identified within the literature. These are subsequently categorized by application domain, analysis goal, data type, data structuredness level, and the applied algorithm. A full overview of the investigated approaches and an associated typology formed by an enumeration of their relevant properties is provided in a convenient tabular format in Appendix B.

Several observations can be made from the 26 research papers under scrutiny that propose a decision mining approach that employs a form of endogenous or exogenous data. First, there does not seem to exist a single business domain or context that is overly represented. The application domains range from medical and health care settings, IT incident management and ticketing, financial loan and credit assessments, up to logistics and industrial maintenance and manufacturing processes. Surprisingly, six approaches demonstrate an application in a medical or health care related domain. Although this domain is known to pose additional challenges within the realm of process mining in general, it also provides opportunities in terms of data availability [129]. Therefore, the domain itself should not be the primary constraint to the applicability of decision mining, rather the availability (and quality) of the data that are used to extract the decision structures.

Second, the examined endogenous and exogenous data types exhibit typical structures such as numerical values, booleans (true/false), categorical attributes, and date-timestamps. In most cases, the approaches provide a method in which data points are converted into attributes at case, event or activity level, which subsequently allows the application of common process analysis and data mining techniques. However, the approaches where a singular relation between the data attribute and a case concept does not exist, warrant the implementation of self-developed or more complex types of decision classification algorithms. The majority of the approaches implement decision tree classification. While some studies are opaque in terms of which decision tree classifier has been employed, the most common variant is the J48 decision tree classifier, which is an implementation of the C4.5 algorithm. An overview of the frequency of occurrence of the relation between available data types and the applied learning algorithms as discovered within the literature is provided in summarizing Table 4.1.

Table 4.1: Overview of data types and the frequency of applied learning algorithms for decision mining.

	Boolean	Categorical	DateTime	Numerical	Text
Association rules	1	-	-	-	-
Decision trees	6	5	3	13	2
Deep learning	-	1	-	1	-
Evolutionary	-	1	1	1	-
Logistic regression	1	1	-	1	1
Probabilistic	-	1	-	-	-
Proprietary	3	2	2	3	1
Text mining	-	-	-	-	1

Third and last, only four of the 26 approaches employ a form of text data as a possible source for decision patterns. Three of them analyze attributes that are stored in an attribute-like textual format, and only one of them demonstrates that complex business rules in knowledge-intensive processes can be discovered from unstructured logs in the form of natural text [130]. However, the applied text mining procedure was not very sophisticated. It involved loading the raw text data into the R statistical software suite and applying a manual search for keywords using regular expressions, an elaborate form of pattern matching. Given that a significant amount of tacit business process knowledge is still documented using written text, this therefore provides an interesting avenue for further research. Additionally, with recent significant advances in Large Language Models (LLMs), the opportunities for Natural Language Processing (NLP) on large bodies of unstructured text have already led to a wealth of new applications [131]. As shown in this part of the review, the state-of-the-art within research on decision mining has not widely considered the application of advanced forms of text mining, let alone NLP with LLMs. However, due to constraints within the organizational context, the development of our methodology will not focus on harnessing the power of novel language models to better use and understand free text objects as exogenous data within this field of research. Instead, existing endogenous data within the process will be used to further enhance the process mining activities.

4.3. Enhancement of process analyses with decision information

The enrichment of process models with decision information is a practice that focuses on the modeling and analysis of decision-making processes embedded in organizational workflows. The integration of these perspectives aims to optimize and enhance business processes through better informed decision-

making. The remainder of this section provides an overview of the work that was included in this part of the SLR.

The aggregation of individual decision-making models is explored by Petrusel [132]. This research is fundamental in understanding how different decision-making approaches, when aggregated, can enrich and optimize process models. The diversity in decision-making and its impact on process efficiency is emphasized, suggesting a model that accommodates these variations for more effective overall decision-making processes. However, this approach used explicit decision logs to explicitly mine a separate decision model and focused on complex decisions. This stipulates the need for a large amount of data to be extracted to create a representative decision model.

The approach of Hasić *et al.* [133] introduces a novel concept of embedding decision-making into process models through a service-oriented architecture. It enables organizations to treat decision-making as an adaptable service, offering a modular and flexible framework that can be tailored to specific process requirements. The proposed architecture is particularly effective in environments where decisions need to be quickly adapted to changing conditions. Similarly, the work by Hasić *et al.* [134] builds on the idea of closely integrating decision-making with process models. The framework leverages process mining techniques to extract decision-related data, providing a richer, data-informed view of how decisions impact processes. This approach helps to create more nuanced and effective process models, particularly in complex decision-making scenarios. Thabet *et al.* [135] utilize process mining to extend business process models with a focus on cost analysis. Their context-based approach provides a comprehensive view of how decision-making related to costs can be integrated into process models, offering insights for more cost-effective process management.

Not only analysis but also process improvement is an area of research where decision information is used. Brzychczy *et al.* [136] explore the use of data analytics to improve machinery utilization. This study shows how data-driven decisions can lead to significant improvements in process models, particularly in terms of efficiency and resource utilization. Shahzad and Zdravkovic [137] also focus on the role of decision making in improving the efficiency of business processes. They propose a methodology where decision-making is central to identifying and implementing process improvements. This approach aligns decision-making with strategic goals, leading to more effective and goal-oriented process models.

The decision perspective can also be used to further enhance visualizations. Farooqui *et al.* [138] discuss a methodology for visualizing manufacturing processes. This approach emphasizes the transformation of raw data from manufacturing floors into comprehensible process models, integrating decision-making elements for better visualization and analysis. It helps identify key decision points and their impact on structured manufacturing processes. Pereira Detro *et al.* [139] illustrate the application of process mining and semantic reasoning in customizing process models in healthcare. This research highlights the importance of personalized decision making in complex environments such as healthcare, where customization of the process model using the decision perspective can lead to improved patient outcomes and operational efficiency.

The enrichment of process analyses and models with decision information, as explored in aforementioned papers, offers a multifaceted view of the importance of informed decision-making in various sectors. From service-oriented architectures to healthcare, the integration of decision-making processes enhances efficiency, adaptability, and strategic alignment. The collective insights from these studies provide a robust foundation for understanding and implementing decision-enriched process models in organizational contexts.

4.4. Evaluation criteria for decision-enhanced process models

In the realm of business process management and decision support systems, the quality and value of process models enhanced with decision information are crucial. This review systematically examines various scholarly contributions, analyzing the criteria and metrics they propose for the evaluations of their proposed implementations. The metrics and resulting criteria that occur most frequently are used to guide the determination of the metrics and criteria to be used for the evaluation of the artifacts and insights of this research project.

Table 4.2: Overview of evaluation criteria and metrics identified within the literature.

Paper	Contextually implemented evaluation strategy	Criteria and metrics
Osei-Bryson [140]	Context relevance, Applicability of results	Relevance, Applicability
Ivanchikj <i>et al.</i> [141]	Clarity, Accuracy in visual representations	Clarity, Accuracy
Liu and Salvendy [142]	Visual clarity, Interpretability of decision trees	Clarity, Understandability
Liu <i>et al.</i> [143]	Data accuracy, Completeness, Relevance	Accuracy, Completeness, Relevance
Yang <i>et al.</i> [144]	Recommendation relevance, Timeliness	Relevance, Timeliness
Combi <i>et al.</i> [145]	Temporal accuracy, Representation of time-sensitive decisions	Accuracy, Representativeness
Raheja <i>et al.</i> [146]	Predictive accuracy, Data source integration	Accuracy
Bruha [147]	Rule simplicity, Generalizability, Predictive performance	Simplicity, Generalizability, Performance
Al-Salim and Abdoli [148]	Process efficiency, Error reduction, Continuous improvement	Efficiency, Error rate, Performance
Gallardo <i>et al.</i> [149]	Business objectives alignment, Clarity in requirement modeling	Applicability, Clarity
Ghosh and Maiti [150]	Defect reduction, Process stability, Efficiency gains	Performance, Efficiency
Wątróbski <i>et al.</i> [151]	Method adaptability, Comprehensiveness of criteria	Adaptability, Comprehensiveness
Žnidaršic <i>et al.</i> [152]	Model adaptability, Accuracy over time, Responsiveness to new information	Adaptability, Accuracy, Robustness
Tsang <i>et al.</i> [153]	Adaptability to process changes, Long-term performance enhancement	Robustness
Corrales <i>et al.</i> [154]	Data veracity, Relevance, Timeliness	Accuracy, Relevance
Gu and Baxter [155]	Accuracy and efficiency of linkage processes	Accuracy, Efficiency
Mortada and Yacout [156]	Predictive accuracy, Logical coherence of decisions	Accuracy, Coherence
Peeters [157]	Accuracy of log interpretation, Robustness of decision logic	Accuracy, Robustness
Alizamini <i>et al.</i> [158]	Rule precision, Handling of data uncertainty	Precision, Robustness

To sum up, this review presents diverse methods for assessing process models that are enhanced with decision information. The main factors to consider include the quality of the data, the accuracy of the model, adaptability, the clarity of the visualization, the representation of time and timeliness, and the alignment with the business goals (applicability). These criteria and measurements offer a comprehensive basis for evaluating the effectiveness and usefulness of decision-enhanced process models that accommodate various business and technological contexts.

5

Methodological framework

This chapter presents the reference fundamentals and first version of the methodological framework that resulted from the initial design and development iteration of the DSRM. The result is based on the initial PM² methodology in conjunction with the SLR findings on decision mining approaches. Explicit attention was also paid to the types of endogenous data that were used in relation to potential avenues of analysis within the case study. Furthermore, this chapter further defines the scope and clarifies the actual implementation of the evaluation metrics and criteria that were broadly defined earlier.

5.1. Rationale and details of the design procedure

The methodological framework is constructed as an adaptation and extension of the PM² methodology by Van Eck *et al.* [99]. The rationale for this choice is four-fold. From a project perspective, decision mining can be of complementary nature to process mining. Depending on the project characteristics, it might therefore be mandatory to (partially) execute a process mining project to identify project suitability and research questions fit to answer with decision mining. The practical nature of the framework is also helpful in that regard. On the data-technical side, several resources in terms of data and project outputs might be reused for decision mining, such as process models and identified decision points. Third, the close involvement and collaboration of the project team and stakeholders throughout the project is acknowledged, which is quintessential for contextual analyses such as decision mining, especially in complex projects [159]. Finally, the methodology is highly iterative, which benefits both the refinement of initial research questions as well as the conversion of the findings into actual process improvements [99].

First, a common understanding of the reference project methodology must be established formally. Therefore, the PM² methodology is modeled as a Process-Deliverable Diagram (PDD), a metamodeling language based on UML [160]. The model is used to clearly demonstrate the dependencies, constraints, and execution order of the six distinct stages, the underlying activities, and their relationship to the deliverables. As this is a data-centric project, emphasis is placed on the required adaptations and extensions of the stages related to data extraction, processing, mining, and analysis. This corresponds to stages two, three and four of PM² [99], however, the other phases will also be covered less extensively. Second, for each of the respective phases, a method fragment is created in the same PDD notation. These fragments are used to formally document and practically illustrate overlap and differences with the reference methodology. For each stage, the requirements are specified in terms of data, procedures, and tool implementations. Third and last, the framework evaluation is carried out by implementing the adapted stages in the context of the case study. This also concerns the development and implementation of technical tools to execute the methodology given the context of the project. The evaluation of this research project is not considered a separate stage of the methodology. It is rather the end-to-end implementation of the framework itself and its complementary technical artifact in a real-world context.

5.2. An overview of the methodological framework: PM²xDM

The proposed methodological framework aims to provide rigor, clues, and guidelines on how to tailor and apply specific features of decision mining in a process mining project. The underlying goal is to

further enhance the project objectives with respect to improving process performance or confirming conformance. This is achieved, for example, by performing established process mining activities in conjunction with decision-point analysis. Such an investigation could, among others, potentially uncover the root causes for certain process behavior, such as the inception of certain process variants. Given that the method is based on PM² [99], we have decided to politely honor its origins and therefore designated our artifact as Process Mining Project Methodology with Decision Mining or PM²xDM in short. The *x* not only signifies the extended approach, but it coincidentally also resembles a part of the possible notation of a decision point in BPMN (in this case, a decision represented as an exclusive OR-gate). Figure 5.1 demonstrates a summary of the main stages of PM² and the relationship with the decision mining activities within each stage. The following subsections summarize the contents of the original methodology and describe what the adaptations and extensions for decision mining entail for each stage in further detail. The complete model and the associated activity and concept tables are found in Appendix C.

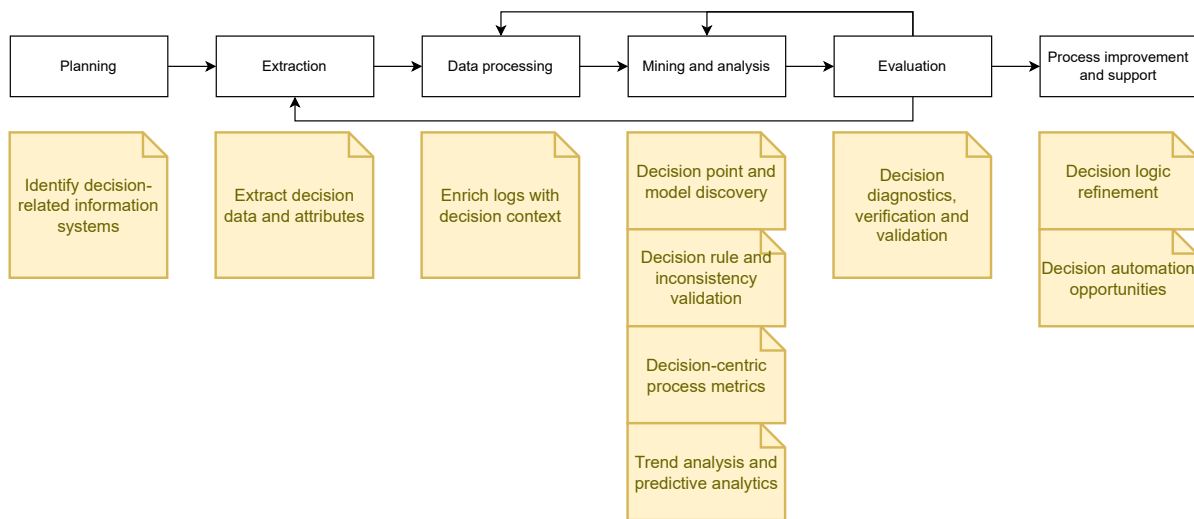


Figure 5.1: An initial overview of PM²xDM its decision-related activities in relation to the stages of PM² [99].

5.3. Stage 1: Planning

The planning stage mainly revolves around the initialization of the process mining project and the determination of research questions in relation to the intended objectives. In addition, related business processes are selected, information systems involved are determined, and the project team is composed [99, 161]. To integrate decision mining, the following adaptations and extensions are proposed.

1. **Select business processes:** the considerations with respect to the selection of business processes do not need to be significantly altered. Both types of challenges related to the process characteristics and data quality also apply when an approach is extended in the context of decision mining [86]. And although this may add one or more additional layers of complexity, an approach enhanced with decision mining enhanced could also yield useful insights in terms of auditing, control, and verification [162]. In turn, this allows one to broaden the considered range of eligible types of processes, and therefore also consider more human-oriented processes that are less structured and exhibit higher levels of complexity or flexibility. Those can be characterized as *knowledge-intensive processes* [163]. Depending on the context, it may be necessary to distinguish between knowledge-intensive and decision-making activities [164]. However, the separation of concerns stipulated by decision mining is already a step forward to uncover how process execution is influenced by explicit decisions rather than tacit or explicit knowledge [134].
2. **Identify information systems:** instead of being embedded in the business process selection phase, this part should be recognized as a separate phase within this stage of the methodology. Information that can be used for decision mining, for example on related decision parameters and outcomes, may not necessarily be available as an integral part of the information system

that stores the primary process execution data. In typical cases where an advanced BPMS is not implemented, decision management or other supporting systems could exist as separate entities within the IT landscape of an organization. Therefore, the foremost aim of this step is to carefully examine which ISs are related to the process. This should include both systems that support the primary process execution and systems of secondary nature.

3. **Identify research questions:** the scope of research questions to be identified is broader, under the assumption that decision mining is applied. While PM^2 defines that research questions should be answered using *event data* only [99], the extended methodology allows the development of questions that can be answered using *decision data*. These data can take the form of an explicit decision log with a relation to the cases, it could be available in structured or unstructured form in a separate information system, or it should be derivable from features or attributes that already exist in the event data. Therefore, the previous step is of paramount importance, as some decision-related research questions can only be answered using data from a specific information system. The aspects of a process to which these additional questions pertain are the same, as well as that they can vary in their level of abstractness. However, the decision mining perspective offers opportunities for integrated analysis, such as decision conformance checking. The attribute values that are used to analyze and cluster process deviations could then be used for root cause analysis.
4. **Compose project team:** the composition effort for the project team does not differ significantly. However, a closer participation of process participants could be needed to gather knowledge about decision points and their relation to process execution. Additionally, if data from different sources are needed, more emphasis should be placed on information system experts that are part of the project team.

5.4. Stage 2: Extraction

The extraction stage aims to extract and collect relevant event data and related documentation, such as process descriptions and existing process models. It involves three distinct activities related to the scope of the project, data extraction, and process knowledge exchange [99, 165]. Within $PM^2 \times DM$, the following adaptations and extensions are proposed.

1. **Determine scope:** the scope of data extraction should be extended such that either the event data includes *dynamic* attributes that are related to activity outputs or case outcomes, or decision data from other sources should be included in the scope.
2. **Extract event data:** the extraction of event data should include *dynamic* attributes that change over time or are related to the outcome of the case. Attributes that are *static* within a single case could be related to the outcome, but do not have an explanatory value for possible decision points within the process. However, a different case notion could yield another result, which will be treated in the next stage.
3. **Extract decision data:** in our methodology, extraction of decision data is a separate step. As approaches to automatic discovery and holistic integration are still nascent [18], the necessity assessment and extraction have to be executed manually [134]. If decision-related information systems have been identified earlier, these should be consulted for information extraction given the scope determined earlier.
4. **Transfer process knowledge:** the knowledge exchange that can be performed simultaneously with the previous three steps, extends into an additional dimension. This not only considers knowledge about the process but also about the decisions related to the process. Classification of decisions that are worthy of investigation could be made according to the typology introduced in Section 2.6.

5.5. Stage 3: Data processing

The actual construction of event logs from the event data obtained earlier is the main purpose of this stage [99]. The event data are reshaped in different ways so that different event logs can be derived for further analysis [166], which in turn can be enriched with additional information [167]. Additionally, existing process models can be used for advanced filtering operations [168, 169]. The proposed adaptations of and extensions to activities of this stage for decision mining are as follows.

1. **Create views:** the incorporation of decision mining into the data processing workflow does not stipulate a specific perspective on or a restriction to the selection and definition of case notions and event classes. However, a different assignment of case notions and event classes could have a significant impact on the representation of the process, and thus also on the relevant decision context of the process that is inherent to that particular view. Therefore, each distinct view that is created should be consistently associated with exactly one decision context. In addition, the creation of views should be executed such that the maximum amount of original data attributes related to cases and events is retained.
2. **Aggregate events:** the aggregation of events in an *is-a* relationship could provide pointers to variations within an event class that are explainable in terms of a decision attribute or outcome. For example, whether a thorough or superficial assessment is necessary could be dictated by an attribute that, in turn, could be the result of an earlier activity. On the contrary, however, the *part-of* aggregation could lead to the obfuscation of such structures, whether it be an explicit choice or a sub-process that is the actual execution of the decision. Therefore, the latter kind of simplification should be applied with caution, as this might lead to information loss that could otherwise have been beneficial to decision mining.
3. **Enrich logs:** with respect to log enrichment two primary methods are discussed within the original methodology, namely deriving or computing attributes or events from the log itself, and adding external data. In the previous two phases, attributes might have been added at the case and event level, given that they could be directly mapped to either of those. Some of these attributes could represent the decision context, such as initial assessment parameters or process outcomes. Regarding the external data that is used to enrich the logs, two types of data can be identified. There might exist parameters that explain the decision rules that lead to an eventual outcome of a decision, as well as data that explicitly represents the decision outcome itself. One should make sure that the event logs capture these outcomes either implicitly or explicitly. Additionally, decision metadata could also be used to enrich the logs, this could include information on who made the decision, what system or tool was used, or any rules or guidelines that were applied.
4. **Filter logs:** the event logs can be filtered according to certain criteria to acquire a different perspective on the data, or to reduce overall complexity by focusing on a specific part of the process. The original methodology defines three general types of filtering:
 - *Slice and dice* or *attribute filtering* can be used to generically filter all cases that match an attribute value, such as the involvement of a specific resource or a specific event that happened. Within a decision mining context, this could already be employed to do a rudimentary form of attribute-based decision conformance checking. For example, all cases that have a certain combination of characteristic parameters should all have followed the same path from a control-flow perspective. Deviations could serve as input for further analysis or additional research questions.
 - *Variance based* filtering is used to group similar traces to reduce the complexity of the process under scrutiny. Similarly to event aggregation, this should be applied with care, as the complexity of the discovered process might be related to an underlying decision context. Overly simplifying the process might prevent the correct mapping of influential parameters with decision logic if that is explicitly embedded within the events.
 - *Compliance based* is a type of filtering that explicitly discards traces or events based on rules or the fit with a given (normative) process model. Similar to the attribute filtering, this could be based on decision attributes to remove incompatible traces. However, application in this stage could also remove (parts of the) decision context that might lead to interesting insights in the subsequent stage.

5.6. Stage 4: Mining and analysis

In this stage, the event logs that have been created in the former stage are the primary input and thus subjected to process mining techniques. The purpose is to gain insights with respect to process performance and conformance in light of the research questions [170]. Explorative techniques can be applied to iteratively refine research questions, optionally in conjunction with process discovery to uncover the control flow if no overall view is readily available. If a-priori (normative) process models are available, these can be used for conformance checking and enhancement [171]. However, the discovered

fact-based process models can also be enhanced with additional perspectives [55, 167]. The output of this stage is the set of performance and compliance findings that answer the research questions.

The execution of decision mining within a process mining analysis yields an additional perspective for the four primary activities from the original methodology. The implications are discussed for each of these activities. It is important to note that decision point analysis usually requires that a fact-based process model based on the event logs has already been discovered. The relevant decision points can be identified based on the branching characteristics present within the discovered model.

1. **Process discovery:** apart from discovering the control-flow of the process at hand, the decision context adds the discovery of other aspects:
 - **Decision point discovery:** the initial enumeration of all branching points within the process. For each of those, it should be determined if they are worthy of further analysis, e.g. by investigating the associated case variants and if attributes or other data points are available to potentially describe their decision context. Decision points can be either implicit, where a path is stipulated by one or more attributes, or explicit, where an activity is the actual execution of the decision and subsequent branching occurs.
 - **Decision model discovery:** depending on the complexity of the decision structure behind a decision point, a separate decision model might need to be created. This helps to offload the process model from the additional complexity that might be needed to describe the different paths in terms of their attributes.
2. **Conformance checking:** the decision perspective adds the following opportunities for decision-based conformance checking:
 - **Decision rule validation:** in addition to validating the alignment between the actual process execution and the modeled process, it is now also possible to check if the decisions made in real-life align with the expected or prescribed decision rules.
 - **Inconsistent decision detection:** detect process instances where similar conditions led to different outcomes, revealing inconsistencies in (automated) decision-making.
 - **Root cause analysis:** when process deviations are observed, decision mining can be used to determine whether the process diverged from the anticipated model due to the decision logic or its incorrect application.
3. **Enhancement:** the process model can be enhanced with a decision perspective, either by annotating the decision points with the attributes that determine the specific path or by a separate decision model that covers one or more decision points and explains their conditional routing.
4. **Process analytics:** several types of additional process analytics emerge from the application of decision mining:
 - **Decision-centric metrics:** new metrics from the decision context, such as accuracy and consistency, can be measured and analyzed. In addition, these metrics can be related to more traditional process metrics such as frequency and duration.
 - **Trend analysis:** depending on data availability, it could be possible to investigate how decision logic, outcomes, and relationship to process execution evolve over time, similar to concept drift detection and analysis [172]. These shifts in decision-making patterns can serve as input for improvements to process execution, decision-making or even both.
 - **Predictive analytics:** future decision outcomes based on specific conditions can be predicted using predictive models, given that historical decision data is substantially available. Combined with the available data on control-flow execution, this could not only predict what will happen next [173], but also the eventual outcome of a case [174].

5.7. Stage 5: Evaluation

The evaluation stage concerns the relationship of the analysis findings with improvement ideas so that project goals can be achieved. Artifacts and findings from the mining and analysis stage are used to generate these ideas or lead to new research questions. This is done in close cooperation between process analysts, and domain and business experts, as unexpected analysis results could otherwise be difficult to interpret [159, 175]. With respect to decision mining, the adaptations and extensions are proposed as follows.

1. **Diagnose:** the diagnostic activities of the evaluation stage are similar to those applied in a regular process mining project without decision mining. The results in terms of decision points and models should be correctly interpreted, the segregation of expected from unexpected or interesting outcomes with respect to the decision context should be investigated, and decision-related research questions might be refined or newly created. The latter could be based solely on interesting results from the analysis of the decision context, but also in conjunction with traditional process mining findings.
2. **Verify and validate:** the addition of decision mining extends the verification and validation activities into the decision context, in the following fashion:
 - **Verification:** the findings, both expected and especially the unexpected, should be checked for correctness against the data and implementations as they have been designed and documented upfront. This could entail, for example, the verification of findings in relation to their decision attributes and the respective decision outcomes, as well as the verification of findings that are based on decision data sourced from external systems. Knowledge of the decision structures might require more involvement of domain experts that are well informed about the subject matter.
 - **Validation:** the feedback loop with stakeholders and other experts to validate their claims with the findings is the final part of the central iterative cycle of data processing, analysis and evaluation. As decision mining can uncover implicit business rules or guidelines that were not explicitly known, it is crucial to engage with them to understand the meaning of the findings and the optimization potential that could be leveraged from possible exceptions.

5.8. Stage 6: Process improvement and support

In the final stage, the learned insights are used to implement changes in the process [176]. As a starting point for these modifications, the improvement ideas from the previous evaluation stage can be used. The feasibility of operational support can depend on the structuredness of the process [55]. Decision mining suggests the following adaptations and extensions.

1. **Implement improvements:** the improvements are not limited to changes to the process itself, but also to the decision context that influences the process or is part of the process. Examples of other areas for improvement are:
 - **Decision logic refinement:** inefficiencies or inconsistencies that have been detected in the decision logic can be repaired, or the process might be changed in such a way that these can be avoided. Traditional process redesign heuristics could also be applied in this context, such as the introduction of an additional path based on a triage activity that incorporates an earlier decision outcome [177].
 - **Decision automation opportunities:** the gained insights into business rules and logic could signify opportunities for decision automation.
2. **Support operations:** the insights into the decision context add another perspective to the area of operational support. Useful operational analysis of problematic running cases or recommended actions could arise from root causes in the decision context of the process. Furthermore, a decision support system could be implemented or enhanced to assist (human) operators in making consistent and informed decisions throughout the process execution.

6

Case study results

This chapter describes the details and results of the case study that is used to implement and evaluate the proposed methodological framework. Chapter 5 describes the integrated process and decision mining framework that is based on the PM² methodology [99] in combination with methods from the extant body of literature on decision mining. First, a more detailed description of the rationale behind and the context of the case study is given. Second, the researcher implemented the framework according to the stages of the PM² methodology, and the related implementation details are disseminated. Third, the framework is evaluated in situ by means of a focus group with several different experts and stakeholders. Finally, the implications to the framework of its implementation and evaluation are described in more detail.

6.1. Rationale

To date, the application of process mining techniques in industrial manufacturing processes is still limited, in contrast to traditional statistics and data mining practices [178]. However, more recent work has shown significant potential value that can be realized in these types of processes, which generally exhibit high complexity and variety [179, 180]. Similarly to process mining for healthcare, complexity and variety can make process analysis extra challenging [129]. However, this stresses the need for deliberate selection of the analysis perspective and adequate scoping. Therefore, it is also essential to note that the complexity of a process cannot be characterized from a single perspective [181]. There exist four main perspectives through which the complexity of a process can be analyzed. A process can be complex in terms of the number of activities, elaborate control-flow constructs, related data flows and activity parameters, or the resources involved [181]. As elaborated in Chapter 2, decision mining aims to explain or reduce this complexity by bridging the gap between the control-flow and data-flow. It does so by providing relevant insights into the influence of data flows on the case routing within the process, hence further explaining what is either implicitly or explicitly driving the control-flow.

The selection of the process at hand for the application and evaluation of the methodological framework is driven by three main factors. Primarily, the process is heavily data-based and decision-driven. Furthermore, it has sufficient complexity, so analysis of a part of the process could already yield interesting insights. This is useful given its size, and therefore an analysis of the whole end-to-end process would be unfeasible within the time frame allocated to this study. Secondly, the process was also selected for convenience reasons, as the availability of both stakeholders and real-world process data was deemed optimal for this specific process, compared to other candidate processes that have been evaluated for inclusion in the case study. Finally, the process seems to be one of a kind. Although there are certainly other locations around the world where wheelsets are being refurbished, presumably none of them are designed like this particular instance. This is an advantage in the sense that it offers unique opportunities for process optimization through process and decision mining. Unfortunately, it also limits the generalizability of the research results, as they might be very closely related to this specific process. However, the objective is to generate insights that are still useful in the broader context of the application of decision mining in industrial manufacturing processes. In the next section, the specific process is described in more detail.

6.2. Description

The Dutch Railways (“NS”) have an organizational division (“NS Engineering”, in Dutch: “NS Techniek”) that is primarily responsible for repairs, maintenance, and modernization of rolling stock inventory. At the industrial site of NS Train Modernization (NSTM) in Haarlem, complete train sets are being refurbished at multiple times in their decades-long lifespan. Some of the individual parts that are necessary to keep trains rolling require more frequent maintenance. The respective maintenance is carried out in individual plants due to their complex revision process. The revision of bogies, and in particular the wheelsets that make up each bogie assembly, are done on-site but in separate production locations. These complementary processes are out of the scope of this research.

The wheelset overhaul process (“Revision of Moving Parts”, in Dutch: “Revisie LoopWerken” or *RLW* in short) is a unique asset of its own and takes care of inspection and refurbishment of wheelsets and respective components. In the past decade, the decision to transition from tired wheels to forged or casted *monoblock* wheels also required a redesign of the process. In 2017, this culminated in the completion of a new sustainable factory building where the redesigned operational process was commissioned. From then on, the process was fully supported by an IS and became highly automated and roboticized, similar to modern industrial car manufacturing plants. In this case, the IS is called the Manufacturing Execution System (MES)¹, based on an implementation of the Production Execution Manager² from the AspenTech aspenONE industrial automation suite. The system orchestrates the entire process, ranging from registering parts measurements, deciding routing through the plant, controlling equipment and transport cranes, and tracking (manual) tasks of on-site workers and production engineers³.

6.2.1. Process description and characteristics

The current operational process was originally designed as a fully linear production process with integrated buffering capabilities. Buffering in this respect means that parts can wait at a certain station for other parts of the process to finish or components that need to become available. However, the fact that the core process comprises several independent and non-interacting subprocesses stipulated the decision to relax this strict linear requirement [182]. This allowed for optimization of the execution sequence of some of the distinct subprocesses, which has already been performed earlier using data from the MES. Figure 6.1 visually depicts the physical layout of the process.

The initialization steps of a wheelset comprises a cleaning routine and work preparation steps, such as removal of bearings and the inspection of the gearbox. More notably, this last step also includes determining the material plan based on the outcomes of the pre-screening routine. The material plan prescribes which treatments to apply and thus the route the wheelset, the axle, and other parts should follow throughout the factory. The actual wheelset revision process that follows consists of five notable steps. First, the wheelset is disassembled into parts. This means that wheels and optionally disc brakes are removed. Wheelsets with gearboxes will also have those removed, and the gearbox itself enters a separate revision process. Second, the axles are decoated, cleaned, measured, and inspected using non-destructive observation techniques (ultrasonic or magnetic). Depending on the result of the inspection, such as whether cracks or other damage have been found, the axle could be rejected or needs to undergo one or more rounds on a lathe machine. The axle can also be exchanged in place for a new part. Third, the respective axle enters the conservation station, where two layers of coating are applied and a subsequent drying process of up to eight hours is executed. Motor axles can skip this step, as they have a gearbox mounted on the axle. Fourth, the on-press station is where the various components are gathered and reassembled. At least the wheels and optionally the disc brakes are pushed onto the axle. If this operation somehow fails, the wheelset exits the process and moves to the manual revision process in a separate factory. The fifth and final step again deals with measurements of the combination of wheels, disc brakes, and axles. If needed, disc brakes can be flattened, or wheels can be re-profiled. For non-gearbox axles, balance is also tested as possible rejection criterion. Subsequently, if everything went smoothly, some remaining parts, such as bearings, are mounted, and the gearbox is filled with fresh oil. Then, a final quality check is executed, and the process ends.

¹<https://www.ict.eu/en/projects/successful-installation-and-commissioning-mes-ns-train-modernisation>

²<https://www.aspentech.com/en/products/msc/aspens-production-execution-manager>

³https://www.ict.eu/sites/corporate/files/files/ICT-GROUP_CasestudyNS-MES_UK_Dig-1.pdf

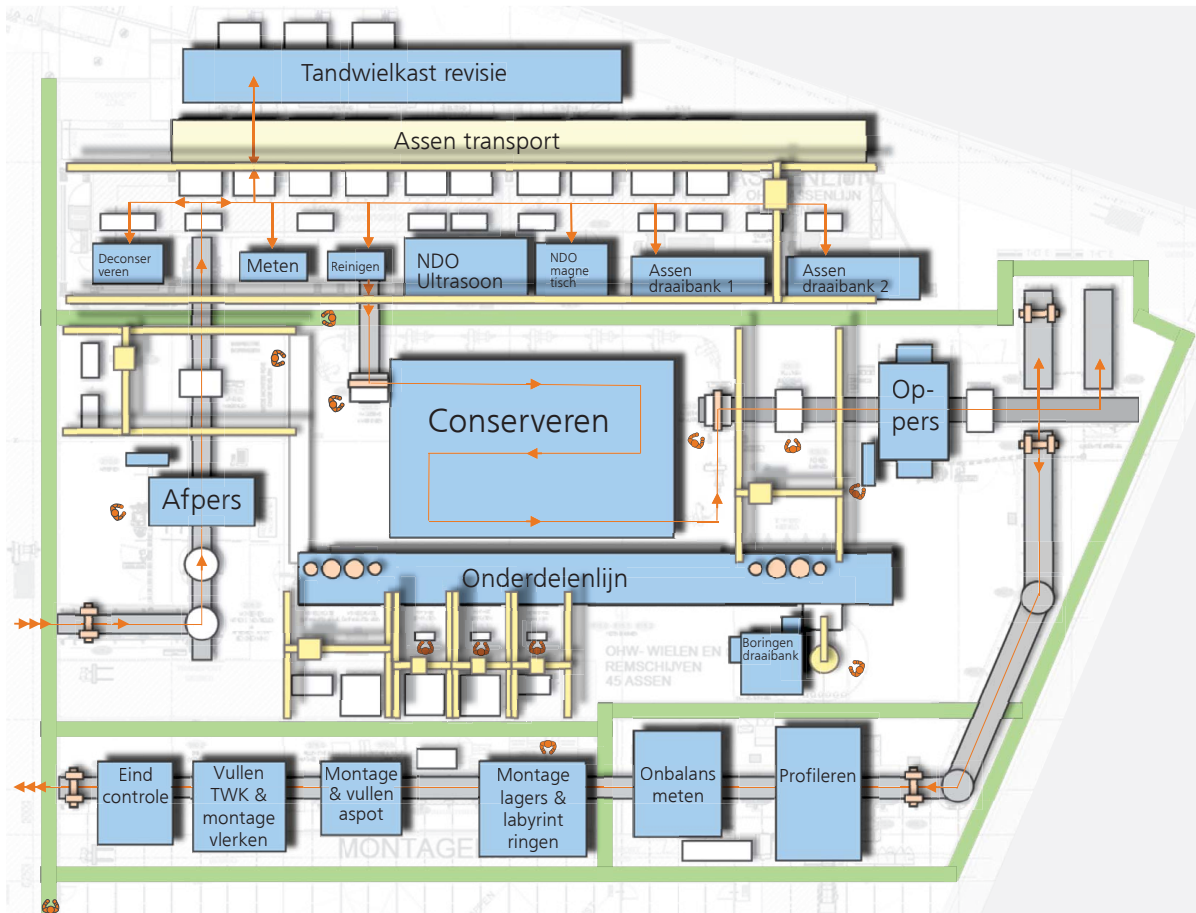


Figure 6.1: An overview of the physical factory layout of the wheelset revision process (in Dutch).

6.2.2. Wheelset types, axles, and components

The process currently handles the treatment of 24 different types of wheelsets, all of which have a numerical type indication. Each different type of train is fitted with one or more types of wheelsets in each bogie, which are typical for a certain type of trainset. Primarily, two types of wheelset are distinguished. Motor wheelsets are actively driven; therefore, they necessarily have a gearbox and brake plates that are mounted at the outer ends of the wheels. Optionally, these wheelsets can be fitted with a cardan drum. All others are referred to as running wheelsets, which are not connected to the train its traction motor, and thus are only passively driven. This type can be fitted with brake plates on the wheels or with two or three brake discs on the axle. Each distinct wheelset type dictates the steps that need to be executed in the process, depending on its components. This is recorded in the aforementioned material plan that is determined upon the insertion of the wheelset into the factory.

6.3. Implementation

This section details the implementation of the steps of the methodological framework. Based on the different stages, the way of working as applied in the context of this particular case study is further explained. Furthermore, the intermediate results of the execution of each stage are disseminated.

6.3.1. Stage 1: Planning

Select business processes

The selection of the process revolved around three of the criteria outlined by Van Eck *et al.* [99], focusing on process adaptability, data availability and quality, and stakeholder commitment. Firstly, adaptability and changeability was confirmed by its IS-supported flexibility in object routing through the plant. This level of adaptability suggests the practical application potential of the findings. Secondly, the

traceability and automation of the process ensured ample availability of high-quality data without privacy concerns, distinguishing it from previously considered projects for inclusion within the case study. Lastly, the complex process required significant stakeholder participation to ensure sufficient knowledge exchange for project success within the time constraints.

Furthermore, explicit attention was paid to the applicability of the decision mining perspective. Although the process itself does not strictly classify as knowledge work, it could still be characterized as knowledge-intensive, as it has an explicit decision-driven process structure [163]. Decision-making within the process is not the accumulation of knowledge evaluation by a human process participant; rather, it is the evaluation of rules and guidelines. In practice, these are the measurements and properties that stipulate certain actions and operations in the process which translate into a routing. Additionally, exceptions could occur throughout the process which change the course of action (e.g. addition or skipping of certain steps) and thus dictate an alternate control-flow. The verification of adherence to expected routing given preassessment outcomes, component types, and intermediate results — in what would be decision conformance checking — seems to be an avenue of decision mining that could yield interesting insights into process performance. Therefore, the process seems to be well eligible for evaluation of the framework and to test the applicability of decision mining.

Identify information systems

The identification of information systems was explicitly carried out to uncover the landscape and architecture of the information systems that support the process. Through a combination of document analysis and a series of informal meetings with the product owner of the MES system, the IT architecture was identified as shown in Figure 6.2. Central orchestration of the process occurs through MES, which dispatches the work items to the different work places and gathers the result of the actions that have been executed along with (machine) measurement results (if applicable). The logistic movements of the wheelsets are exchanged with an integrated logistic and financial system, whereas the summaries of the revision activities are exchanged with the reporting system. This is the central asset tracking and reporting system of NS, which is also used to prove if maintenance requirements have been met. The resources and inventory of parts and materials are managed by an ERP system, where MES also exchanges data with respect to the stock levels, parts, and materials that have been used. This also includes the main components of a wheelset, such as wheels and braking discs. Finally, there is a configuration management system that does not interface with MES. However, it stores the relevant documents (work instructions) that describe the routines and procedures that are performed at the different work stations. These documents also contain information on the approval/rejection criteria for axles and other parts of a wheelset. This information is only stored as unstructured written text documents and is therefore not electronically accessible to MES. MES though has its own internal repository with routing logic and the associated measurement results, criteria, and tolerances that determine the route of a wheelset throughout the revision process. This repository was not accessible throughout this project. However, this is not necessarily problematic, as one of the goals of this research is to actually uncover these constructs from the resulting process data.

Identify research questions

The research goals were developed interactively in an iterative collaborative effort to address not only the main objective of this research, but also a real business need with associated practical relevance. Initially, an idea was evaluated to verify the compliance of the execution of the work processes with the work instruction documents from the Infor PLM configuration management system. However, the process data from MES is abstracted at the work place level and does not contain data about subroutines and procedures that are, for example, employed by a mechanic or engineer. Therefore, the goals were decided as follows.

- Create an overview (e.g. process models) of the overall process that accurately depicts the control-flow and the related possible behavioral constructs, using an appropriate modeling technique. The high-level evaluation criteria are understandability, quality, relevance, and generalizability.
- Compare complex constructs in discovered behavior (e.g. sequential, parallel, optional, and exclusive) with the real process and knowledge thereof (conformance checking). Subsequently, analyze and discuss any deviations that presumably should not occur and, if possible, further investigate these cases and their root causes.

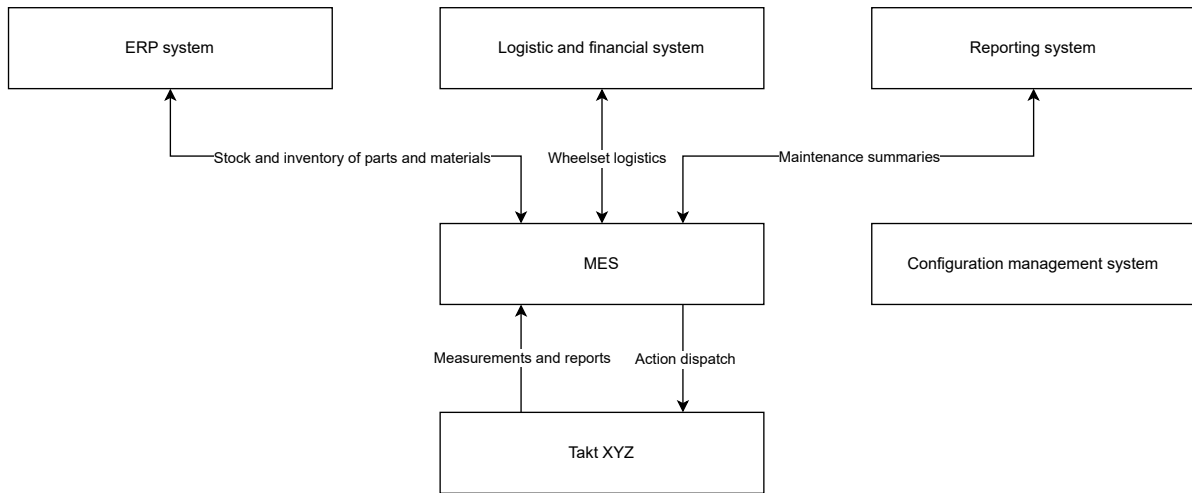


Figure 6.2: A schematic overview of the components of the IT architecture surrounding the wheelset revision process.

- Enrich the artifacts with the decision information as far as is present in the data using the proposed methodology. Evaluate the findings (confirmation), quality, relevance, and generalizability in the context of the results presented by the execution of this activity.

Compose project team

The project team was composed in such a way that all the different stages could be supported by the same team members, from initialization to final evaluation of the methodology and its results. Given that the process is unique, complex, and highly specialized, the participation of domain experts was of paramount importance. As this project was subject to considerable time constraints, it would be impossible for the process analyst — in this case the researcher — to gather sufficient domain and process knowledge to conduct a fruitful analysis, let alone rigorously evaluate the findings. The evaluation of the project was carried out with the respective team, and the numbering of experts is consistent throughout this document, including the supplemental transcriptions. The roles were categorized according to the definition of Van Eck *et al.* [99], so the project team was composed as shown in Table 6.1. Professional experience was defined as years of active participation in the labor market in general, not necessarily in this specific field. Tenure was limited to a position at NS that is related to the wheelset revision process.

Table 6.1: An overview of the project team members of this research project, details accurate as of December 1st, 2023.

Expert ID	Project Team Role(s)	Organizational Role	Tenure	Prof. Exp.	Process Mining Exp.	Focus Group
Exp1	Business Owner, Business Expert, System Expert	Manager Engineering, Product Owner MES	5 years	10 years	Some projects	Yes
Exp2	Business Expert	Production Engineer	4 years	8 years	None until now	Yes
Exp3	Business Expert	Trainee Engineering	4 months	2 years	None until now	Yes
Exp4	Process Analyst (researcher, moderator)	Intern Innovation Process Mining	6 months	12 years	1,5 years academic	Yes
Exp5	Business Owner, Business Expert	Manager Engineering	1,5 years	7 years	-	No

6.3.2. Stage 2: Extraction

Determine scope

Determining the scope of data extraction was a two-stage process. The first step was to evaluate the available and required granularity of the data within the ISs and the relevant period to extract. With respect to granularity, the process data should capture the actual physical process well. After an initial brainstorm session with an expert, the intent was to compare the workflows with the data at the level of the individual steps that engineers execute at a single workstation. These procedures and related decision criteria are formally documented in associated work instruction documents, and this could provide an interesting form of conformance checking. However, it turned out that there was no data available to work with at this level of granularity. Therefore, it was jointly decided that the data at the level of the transitions between the workstations would be used to further analyze the process.

Regarding the time period to consider, the data was available from 2020 onward. After careful consideration, the years 2020 and 2021 were ignored due to possible deviations caused by the COVID-19 crisis. The most recent data for a full calendar year would be most appropriate. Given that for the running year 2023 not all data were available at the time of analysis, it was therefore decided that the full year 2022 would be used for further analysis. Separate data on the intake process of the wheel sets were also available. However, these were not considered for this project, as it is a completely isolated process that takes place before the actual revision. In addition, it has a different case identifier, making the analysis more difficult. If an integrated analysis of both processes is demanded in the future, a possibility would be to link them together based on the unique serial number of wheel sets involved.

The second step dealt with evaluating the richness and suitability of the data for decision mining. This meant investigating the characteristics of the data in terms of the available attributes at the case and event level. In an excerpt of the example data set of 2023, both types were present. Given the IT architecture identified earlier, the MES environment does not seem to contain external systems that could contain relevant decision information. Therefore, the analysis is limited to only data exports from the MES itself.

Extract event data

The event data from MES for the full year 2022 was supplied as a single Comma Separated Value (CSV) file of 597,6 MB in size. It contains 10.871.409 rows divided into six columns. Table 6.2 shows the columns of the data structure and the respective descriptions. Data were presented in long format, rather than in wide format, where the attributes of each workstation are listed as key-value pairs. Each order also has a "Common" characteristic, where the attributes valid for each instance of the subtype Wheel Set Revision (WSR) are stored. Given the impossibility of creating a proper event log from these data in a long format, this stipulated the first data preprocessing activities. To maintain a well-documented reproducible workflow, a Jupyter Notebook [183] was initialized to develop and execute the different steps in the Python code. For data reshaping, the Pandas library and its associated concept of DataFrames were used [184].

Table 6.2: The columns as present in the initial data export CSV file for the year 2022.

Column	Description
MESPONUMMER	MES Production Order number
Takt	Workstation number
Takt instantie	Workstation instance sequence number
Karakteristiek	Characteristic attribute key name
Karakteristiek instantie	Characteristic instance sequence number
Karakteristiek waarde	Characteristic attribute value

Extract decision data

The decision data that could explain routing decisions were assumed to be present in the hundreds of descriptive attributes in the log. Therefore, no separate decision data extraction (e.g. from an external information system) was performed. The criteria stored in the MES that are used to decide routing based on measurement outcomes would be interesting to verify against actual routing decisions. However, it was not viable to extract those from the MES at this point in time.

Transfer process knowledge

Transferring process knowledge usually is performed simultaneously with data extraction, as well as iteratively throughout subsequent stages. The initial round of preprocessing steps that only reshaped the data from long to wide format was quite trivial. However, the final part of the extraction stage was concerned with the actual mapping of the data attributes to the process mining concepts. This required the participation of experts to determine the appropriate constructs. Therefore, this initial round was condensed for the better part into two interactive sessions and a handful of additional email conversations to not overly burden the business experts.

Furthermore, additional documentation was collected in the form of earlier process mining project outcomes (reports and process models), hand-crafted flowcharts of the process, a document with

extensive knowledge on wheelsets, and a description of the different workstations. The a-priori process models have been discussed but have proved to be of only minor value according to the experts, as they were either difficult to interpret or did not capture the process sufficiently well.

6.3.3. Stage 3: Data processing

In this stage and the subsequent mining and analysis stage, a combination of three tools was used:

- Fluxicon Disco⁴ 3.6.7 for exploration of the data sets and creation/manipulation of event logs
- ProM 6.13 [185] for further exploration and process model generation other than DFGs
- PM4Py 2.7.4 [186] with Scikit-learn [187] for model generation and decision mining activities

The source code for data preparation and subsequent processing, mining, and analysis is available in Appendix G.

Create views

The relevant process mining concepts were extracted from the list of characteristic attribute key names. All other attributes were maintained separately as attributes at the case or activity level, so that they could be used in subsequent decision mining activities. This resulted in the simplified mapping shown in Table 6.3. The data contain start and completion timestamps, and thus analysis opportunities that make use of life cycle information can be executed, such as active/idle time, waiting time, and case utilization [188]. Furthermore, given that the activity is defined as the unique workstation number, there are multiple start and completion timestamps available.

Table 6.3: An overview of the mapping between the data set attributes and the process mining concepts.

Column	Description	Concept
MESPONUMMER	MES Production Order number	Case ID
Takt	Workstation number	Activity
Takt instantie	Workstation instance sequence number	Attribute
Karakteristiek instantie	Characteristic instance sequence number	Attribute
Starttijd	MES Activity Start	Start Timestamp
Eindtijd	MES Activity End	Complete Timestamp
StarttijdEQU	EQU Equipment Start (Ready)	Start Timestamp
EindtijdEQU	EQU Equipment End	Complete Timestamp
StarttijdMOB	MOB Hand Terminal Start	Start Timestamp
EindtijdMOB	MOB Hand Terminal End	Complete Timestamp

Aggregate events

Aggregation of events was not applied to the event log. To maintain granularity to potentially uncover decision structures, collapsing similar events into their event class could potentially remove useful information. On the other hand, grouping similar events in an *is-a* relationship was not considered necessary.

Enrich logs

Initial exploration of the resulting event log yielded a complex spaghetti-like process model that was considered unusable for further analysis [189]. Note that the process map in Disco was explored with both detail levels set to 100%, which means that all activities and paths are visible. However, on closer inspection, impossible behavior was observed. Activities that certainly could not take place at the same time were depicted as if they were happening in parallel. After further analysis of the data and a follow-up discussion with an expert, it turned out that the equipment in the process can emit a start event before the MES indicates that the actual activity should start. Therefore, the equipment start event can precede the start event of the main activity. An additional factor was that Disco considers the outer windows of start and complete timestamps for the related activity. This caused sequential workstations to also be shown as overlapping, while it is logically impossible for a physical axle to be in two stations

⁴Courtesy of an academic license, more information available at <https://fluxicon.com/disco/>

at the same time. Although it would be possible to disregard the supplemental timestamps altogether, it was decided to maintain this information as it might be useful in subsequent analyses. Therefore, additional activities with these timestamps were added to the log with the same workstation number but with EQU and MOB suffixes, respectively.

Furthermore, the expert raised the point that not everyone involved in the project could be fully familiar with all the different plain workstation numbers as activity names. Therefore, a list was collaboratively compiled that contains comprehensible descriptions of all workstation numbers. This was included as an attribute of the respective activity, so that it could be included in a composite name. Disco conveniently offers the feature of compiling an activity by merging several attributes upon import of raw data.

Filter logs

The data set contains cases that started in 2021 but ended in 2022. As the analysis should focus on that year only, the cases that extend before January 1, 2022 have been excluded using the "contained in timeframe" trimming feature. This led to the exclusion of 70 of the 1969 cases. Subsequently, the log was filtered to contain only the events that are emitted by MES, and not by machines (EQU) or hand terminals (MOB) to prevent incorrect parallel behavior, as described earlier. Furthermore, the expert stated that the separate processes for gearbox revision and component lines should be ignored. This was achieved by filtering only activities in the range of 0-599 resulting a reduction from 126 to 69 activities.

The resulting data set exhibits the global characteristics shown in Table 6.4 and the initial Disco DFG is presented in Figure 6.3. The most important observation with respect to the characteristics of the process is that the process has a large variance. This is signified by the ratio between cases and variants, where there exist a significant number of rare variants. An explanation for this might be that the options that differ throughout the execution of a single case can vary due to the configuration of certain types of wheelsets and exceptions that might occur, such as a measurement result that falls outside the tolerance range. All of this gives rise to the introduction of new process variants. Although the process exhibits large variance, even at 100% fitness a readable process model is presented. However, it still visually resembles a spaghetti-like model, although portrayed sideways it would look more like lasagna [189].

Table 6.4: An overview of the characteristics of the baseline event log created with Disco.

Characteristic	Value
Events	98.583
Cases	1.899
Activities	69
Attributes	510
Variants	262
First event	03.01.2022 07:24:28
Last event	07.12.2022 15:03:38
Median case duration	10 days
Mean case curation	12,7 days

6.3.4. Stage 4: Mining and analysis

Process discovery

The discovery of an initial DFG of the wheelset revision process, as shown earlier, was carried out using Disco, which internally uses a proprietary implementation of the Fuzzy miner [190]. Most of the "spaghetti" paths shown are very rare; however, Disco does not offer a way to filter *only* these rare paths. Instead, the overall paths should be reduced. Additionally, the variation filter is not suitable, as it only filters exceptional cases in a single act to keep the mainstream behavior. This would remove 94% of the variants, 83% of the cases, and 85% of the events, which makes a further analysis rather pointless.

After a discussion with an expert, a viable solution would be to filter the data set according to the different types of wheelset. Each wheelset stipulates certain actions based on the components it contains, in addition to the results of the measurements and diagnostics that are conducted throughout the process. Therefore, the possible different types of necessary treatment combined with the intermediate results could lead to cluttering of the model. To validate whether the model would be improved, a




Figure 6.3: The initial model as shown in Disco after applying the data processing steps, 100% of activities and paths visible.

filter was applied in Disco on the most common type of wheelset (328). On the one hand, given that a significant amount of cases were still retained, noise could still be an issue if that was the underlying cause. On the other hand, noise could also be present only in cases of a certain type of wheelset. This was an accepted side effect of working at a fitness level of 100%.

The model shown in Figure 6.4 indeed shows a lasagna model of the process for wheelset type 328. This confirms that either the different wheelset type compositions introduce infrequent behavior or that noise is present only for specific wheeltypes. Another interesting option that was not explored at this point is to cluster the wheelset types by similar composition. This could then yield information on which wheelset types have deviating treatment procedures.

Although the DFGs generated in Disco provided an initial overview of the characteristics and "minability" of the process, a requirement for the decision mining activities was the use of Petri nets as reference models, so that a decision tree could be mapped onto the net using Token-Based Replay (TBR) [191]. The resulting model would then become a Petri net with data (DPN), as explained in Section 2.6 of this thesis. The log contains start and complete events, and thus life cycle information [188]. Furthermore, the high proportion of variants indicates that the log contains infrequent behavior [192]. Therefore, the most obvious choice for the mining algorithm to use was the Inductive miner [193]. This algorithm detects infrequent behavior, is robust against noise, and has the capability of discovering concurrent and interleaving behavior using life cycle information. Finally, it is implemented as a plug-in in ProM, and is also available as a module in PM4Py.

In ProM, mining a Petri net on the event log that contained all axle types with the IMflc (Inductive Miner - infrequent and life cycle) resulted only in crashes due to a Java error. Adjusting the noise threshold to a higher value did not produce a working result. Presumably, the search space for the mining algorithm became too large with the included temporal information. Therefore, the exploration in ProM was continued with the filtered log for the most common wheelset type 328. This similarly




Figure 6.4: The model created in Disco after filtering for wheelset type 328, 100% of activities and paths visible.

produced a lasagna-like Petri net of the revision process for this wheelset type at 100% fitness. At the default fitness level of 80%, the model improved significantly in terms of readability, since two main blocks exhibiting concurrent behavior disappeared. These were caused by a few deviating cases that exhibit a different execution order, which in turn makes the model accommodate any execution order for the related activities. However, for the upcoming decision mining analysis, a fitness of 100% is necessary to investigate if deviating paths originating from the decision points can be explained by attributes of associated activities.

Decision point discovery

The subsequent analysis took place in a Jupyter Notebook with the Pandas library [184] and PM4Py [186]. The filtered event log for wheelset type 328 was exported from Disco and it was subsequently imported as a Pandas DataFrame. Initially, the discovery of a Petri net with 100% fitness using the Inductive miner resulted in a model full of loops. An investigation in the documentation unveiled that the PM4Py implementation of the Inductive mining algorithm does support infrequent behavior, but not the inclusion life cycle information. Therefore, the algorithm assumed that all activities were repeated at least once because of separate start and completion events. By filtering only events with life cycle transition status *complete*, this problem was overcome. However, this also leads to an unfortunate loss of information. For the detection of concurrent or interleaving behavior, timestamps could no longer be used to distinguish temporal overlap from distinct execution order.

In the DFG presented earlier in Figure 6.4, all events with more than one outgoing arc could represent a decision point. On the one hand, most alternate paths have a very low frequency of occurrence. However, these could be interesting from a compliance perspective, as a skipped step might violate an explicit rule. If an explanation is to be found in the respective attributes, this could give a clue for root cause analysis. On the other hand, the exceptions that occur more frequently could be clustered based on their attributes to investigate whether a pattern is present for all these cases or if several subclasses exist within these deviations.

In a similar way, in a Petri net a decision point is represented as a place with two or more outgoing arcs. An example is shown in Figure 6.5, where an optional examination step can be executed. The place before the optional activity has an arc to either a silent transition that does not relate to any event in the log, or to the actual activity. In the former case, the activity is skipped. Otherwise, in the latter case, the token is consumed by the actual activity. Silent transitions are therefore needed to generate sound workflow nets. As the silent transitions do not represent an activity present in the event log, any respective conditions are always related to the actual activity in this optional choice construct.




Figure 6.5: Example of a decision point in the Petri net for wheelset type 328 where activity 337 is an optional examination step.

Decision points of a more elaborate nature have also been identified. Figure 6.6 shows a place after a silent transition with three output arcs that represent exclusive choices. On the top-most path, activity 345 (imbalance measurement) is required to be executed at least once, but can be repeated multiple times. For all distinct paths, conditions might be discoverable from the data attributes that function as guards for the respective transitions to occur. Automated discovery of a complementary *decision model*, such as a DMN as presented earlier in Section 2.5.3 of this work, was unfortunately considered beyond the scope of this research at this point. Currently, no tooling was available to easily implement such a discovery activity. However, if time permitted, it would be possible to create such a model manually using the conditions discovered.

Conformance checking

The seminal notion of conformance checking is essentially to strike a balance between fitness and appropriateness of (real-world) behavior and a documented or discovered process model [171]. In other words, the model should not accommodate unwanted behavior, but it should still be able to replay all




Figure 6.6: Example of decision points in the Petri net for wheelset type 328 with multiple choices and a repeatable step.

appropriate traces present in the event log, given that these relate to the "happy flow" of the process. However, the application of decision mining to a process leads to a conflict of interest, as exceptional cases could be explained by a circumstance that is encapsulated in a particular combination of attributes. Apart from validating the discovered decision rules and investigating inconsistent decisions, a possible cause could thus be explained by the decision attributes. Therefore, the model needs to at least accommodate that particular trace for the TBR replay to work. Given that in this project a fitness level of 100% for the most part did not render incomprehensible models, this could be a criterion to consider for the application of decision mining. For less structured processes, this inhibits the potential for analysis of unexpected or unwanted behavior. Nevertheless, if a well-fitting model is discoverable, there still exists potential to enhance models of the ideal process using decision mining.

In this case study, a one-shot final evaluation strategy was applied due to time constraints. Apart from data preparation and processing iterations, no formal conformance checking activities were executed at this point. Instead, the resulting models were evaluated in a single evaluation episode in the form of a focus group with the business experts that participated throughout the project.

Enhancement

The enhancement step involves the extension and improvement of the process model with additional information; in this case, any explanatory conditions signified by the attribute values present in the event log. A decision tree for the attribute classification problem is generated by clustering the attribute values, similar to the approach of De Leoni *et al.* [167]. The underlying decision tree is evaluated for each case that is replayed on the model using the TBR approach [191]. Any characteristic that has a correlation with a certain path can become attached as a guard condition to the respective transition if it corresponds with the leaf of the decision tree. The decision tree classifier performs automatic feature selection from the available attributes. The initial execution yielded the discovery of conditions related to individual wheelsets, such as serial numbers or anonymized attributes. As these inherently do not provide any explanatory value, a short discussion with an expert was conducted. Subsequently, it was decided to remove these specific features from the list of available attributes.

Figure 6.7 shows the same optional examination step as discussed earlier. However, now annotated with a guard expression. The expression states that the examination step is only carried out if the "UnLoadBaan" parameter has a value lower than or equal to 2.5. At first glance, it seems like an arbitrary value. However, upon inspection of the possible values in the data, it turned out to be a discrete variable. Given that the decision tree classifier has interpreted this variable as continuous, the threshold value is not rounded as an integer. The possible values found in the data are 1, 2, 3 and 5. After a discussion with an expert, this variable indicates the track that is chosen for the axle to follow. This track depends on the measurement result in activity 335 and logically results in a parameter change from activity 336 to optionally activity 337. Values 1 and 2 are the examination tracks, which essentially means that the wheelset is set aside for further inspection and treatment. This can also show a delay, as these treatments are usually postponed to the next workday. Value 3 is the approval track and 5 is the track for wheelsets that have already been rejected earlier. If we cast the condition as an integer, then this correlation with the followed path in the control flow makes perfect sense. All wheelsets with values one or two pass through the examination step, and any other skips this activity.

In Figure 6.8, a more complex decision construct that was discovered is shown. Three outgoing arcs and one incoming arc are connected to this place, where the latter indicates a loop back from a later stage. Double ampersands in the guard expression indicate a logical *AND*, which means that both conditions should hold for the transition to be enabled. Additionally, any token that is produced at this place from a later stage will also be again evaluated against the guard expressions. The variable "Takt instantie" indicates the sequence number if a wheelset has been passed through the treatment stage for




Figure 6.7: Example of an annotated decision point for the optional examination step in the Petri net for wheelset type 328.

an additional time. When the variable "Gerepareerd_JA" evaluates to true, a repair has been conducted. As such, if it is a second instance, activity 333 will follow anyway; otherwise, it depends on the criterion of whether the wheelset has been repaired or not. In this specific case, an interesting observation is that the addition of guard expressions in combination with a loop actually introduces a deadlock. If the attribute values do not enable any transition, a token might get stuck at the respective place. However, for analysis purposes, this is intentional, as it marks an execution that should not be possible. In the paths that follow from this example, some of the conditions appear again in several subcompositions to determine further branching. Please consult the complete model in Appendix H for further details.

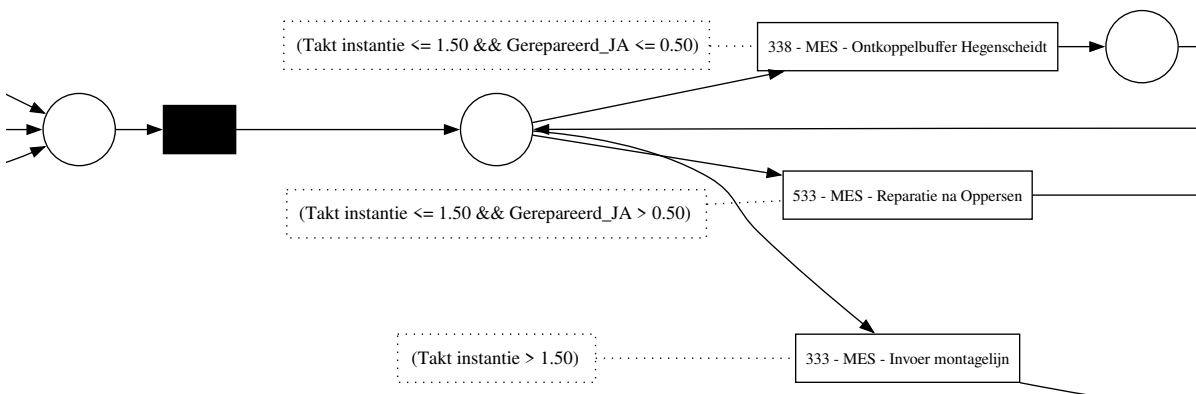


Figure 6.8: Example of an annotated decision point with multiple exclusive conditions in the Petri net for wheelset type 328.

Furthermore, guard expressions are also discovered that contain a logical OR and also a combination with an AND construct as shown in Figure 6.9. The condition for the silent transition that loops back to the place shown in Figure 6.8 is guarded for firing by an expression that evaluates to true if the instance has an integer value less than two and the wheelset has been repaired, or the instance is valued two or higher. In other words, the loop back is enabled if this is the first treatment iteration of a wheelset and it has been repaired, or it is entering the process again and therefore the second treatment of that particular wheelset. Finally, this part also contains an expression that resembles a semantic contradiction. However, it programmatically states that the instance sequence number is either one or lower, or greater than two. Upon inspection of the underlying data and filtering for the presence of activity 590, it turned out that the instance attribute value does not explain whether activity 590 is executed or not. The classifier found that this attribute was explanatory, but there was no consistent condition during replay that could be assigned to each cluster and that would also satisfy both available routes.




Figure 6.9: Example of annotated decision points with logical OR/AND and contradictions in the Petri net for wheelset type 328.

Process analytics

As discussed in Chapter 5, data mining techniques — including decision mining — and other types of visual analytics can be used to answer specific research questions related to certain characteristics of the process [167]. The addition of decision mining in a process mining project adds a dimension to that, in the sense that existing metrics or analyses can be enriched with the decision context. In this project, the upcoming evaluation of the enhanced models with decision information was essentially concerned as a form of process analytics. Therefore, the proposed process analytics activities, such as decision trend analysis, have not been formally performed. However, these possibilities have been discussed as part of the focus group in the evaluation stage. Subsequently, a follow-up request from one of the experts was to further enrich the models enhanced with the decision perspective with information on the execution frequency of the different paths.

6.3.5. Stage 5: Evaluation

In the proposed methodological framework, no significant changes have been made in the structure of the evaluation episodes as presented in the original PM² methodology. The focus area of interest were the enhanced artifacts and the additional insights that could be generated from them. Although evaluation can take several different forms, this research project bound to time constraints stipulated a final one-shot evaluation strategy in the form of a focus group with the involved experts. This was in addition to several intermediate discussions that arose during the collaborative analysis iterations. Altogether in an effort to investigate the application of the framework under development and subsequently use the learnings to improve it.

The focus group with four participants including the researcher lasted slightly longer than two hours and was conducted according to a predefined protocol found in Appendix F. The best practices and instructions for both the procedure as well as the design of the protocol were applied as presented by Krueger and Casey [115]. For the process of conducting the focus group itself, the guidelines of Saunders *et al.* [194] were also taken into account. The results are grouped thematically for convenience and presented in the following subsections.

Process characteristics

In the planning phase of PM², assessing the *changeability* of a process is crucial to determine the viability of a process mining project [99]. Taking on a project is futile if the process lacks adaptability. Therefore, it was valuable to explore its adaptability characteristics, its physical and logical layout correlation with event data abstraction, and how its configurability fits within the decision-making context. This analysis aimed to elucidate the process architecture, providing insight into the applicability, feasibility, and generalizability of the process and decision mining techniques.

The changeability of the process turned out to be directly related to the physical and logical dimensions of the process architecture. From a physical perspective, the structure of the process is largely dictated by the laws of physics. It makes perfect sense that it is impossible to drill a hole in a wheel that matches the diameter of the axle until both have been measured and are physically available at the same location. Exp1 stated that the logical abstraction of the process in the data is fully customizable in the software that controls the process. This underpins the idea that any event log is an interpretation of the real-world process, where choices in the software configuration might skew the view on how the process is actually executed in practice. The sole logical constraint is that the different steps can be executed only in a linear forward-oriented fashion. However, primarily, the constraints that stem from the physical composition of a wheelset are still in the lead. Exp1 illustrated this in a slightly sarcastic way by saying: *"We can configure the data in every way possible. But if you ask me how to build a car, well, then you do not start by inflating the tires."* Exp2 slightly disagreed and added that there are also constraints stipulated by the as-is physical layout of the factory, where Exp1 regarded these constraints solely as "money-based". In other words, significant changes to the physical layout of the factory are usually avoided, as they presumably are costly. Any physical change to the process will nevertheless also stipulate changes to the logical configuration in the software.

Regarding the relationship between the decisions taken and the changeability and configurability of the process, three significant points were discussed. First, the routing logic is fully defined in the software system. For example, if a work station is duplicated, they can be defined as mutually exclusive. To achieve something similar, an additional assessment step could be introduced that determines the next activity. Second, the assessment of the necessity of any step is made upon invocation of the activity.

To illustrate this, let us say that a wheelset has been rejected in step two, and therefore all steps up to step six can be skipped because step six contains a treatment that is *only* executed for rejected wheelsets. Each of the intermediate steps will be considered anyway, but as long as the execution criteria are not met, they will not actually be carried out. Therefore, a specific characteristic of this process is that all criteria are evaluated *locally* at each activity, even if the actual criterion was determined in an earlier activity. Thus, it could be that the information that explains a decision point later in the process was obtained in a different activity. This could prove problematic for decision point analysis if the respective attributes are not coupled with the activity that directly precedes the branching point. In addition, these decisions are only taken ad hoc, and there is subsequently no forecasting employed of what potentially needs to happen in the future. Third and last, it was noted that a part of the complexity comes from the fact that in the MES there is no notion of routes or paths of a wheelset through the process. Exp1 explained that *"We only know the considerations of what to do for each operation based on the input that is collected until that point in time. If you have then produced a wheel set and look back, you accidentally turn out to have a route."* In essence, the process is a partial game of imperfect information that unfolds as it progresses.

This last point ultimately stressed an important distinction in the typology of these processes, namely that of industrial *production* opposed to *revision* processes. In a production process of a new item, there is usually a blueprint available, and at each point in time, it is known what is required in terms of materials and what needs to be done to get to the end result. For a revision process, this does not have to be the case. In this process, the information gradually becomes available during execution. Also, measurements taken at some point might prove invalid at a later stage. Exp1 illustrated this specific trait of a revision process by reiterating the analogy of manufacturing a car. *"When you build a car, you know that you are going to equip it with four seats. What we do is grab an existing car, remove the seats, determine that three of them are good, and thus we only need a fourth one. We do not have a path that states that we by default order four seats; we only have a path that we will assess the seats, and then some four steps later we will order only the rejected seats. [...] That is a completely different vision, which is specific to a revision process. I think that all revision processes are dealing with this."*

Key takeaways

- The architecture of the process is largely dictated by the physical composition of a wheelset
- Physical changeability of the process (e.g. factory layout) is limited due to costs
- Logical configuration options are virtually unlimited
- Routing decisions within the process are fully defined in software
- Criteria for execution of each activity are evaluated locally, without any forecasting
- Decisions are based on ongoing assessments and possibly changing information
- Production processes have a clear blueprint and known requirements at each stage
- Revision processes tend to have the information and needs revealed progressively

Process modeling techniques

The part on general process modeling techniques of the focus group had two primary goals. First, identify the preferences for modeling techniques with business experts and the appropriateness of the respective modeling paradigms in this context. Second, familiarize the participants with the use of Petri nets to represent their process. The latter was compulsory because it was unfortunately not possible to convert the Petri nets with data (DPNs) into BPMN while also retaining the guard expressions.

First, an excerpt from the wheelset revision process was presented to the experts modeled as a DFG, a Petri net, and a BPMN model. Two out of the three participants immediately mentioned that the BPMN representation was the most clear in terms of syntax. Upon asking why, Exp1 expounded that *"because it explicitly displays if something is executed in parallel or not, and it explicitly states where a decision is taken."* Exp3 stated that *"with a DFG, you need to deduce where it happens in parallel and where a decision is taken."* Exp2 was slightly in doubt between the DFG and BPMN, but nevertheless added that *"the Petri net is the hardest to read."* After guiding the experts through interpreting the behavior, consensus was reached that the BPMN models are the most readable and clear in terms of how to correctly interpret the displayed process behavior.

Second, the notion of tokens in the Petri net was briefly explained and what that mechanism entails to read such a model. Subsequently, the example of a non-free choice workflow net from Figure 2.2

was presented. Presumably, since it was now an abstract example, Exp1 initially misinterpreted that $t5$ would directly follow $t1$, but therefore $t3$ must be executed first. After some additional clarification regarding the enablement of transitions in relation to the location of the tokens, the operation of this non-trivial Petri net was clear to the experts. Exp3 correctly summarized that in this case *"the choice in the process makes it that if you execute one activity, the other is not possible to execute anymore."* At this point, a rudimentary understanding of Petri nets was achieved so that we could proceed with the analysis of the actual models of the wheelset revision process.

Key takeaways

- BPMN models are most clear in terms of readability and understandability
- DFGs make it hard to deduce what actually happens and obscure behavior
- Petri nets have a steeper learning curve, but are usable after proper introduction

Baseline process models

The evaluation of the baseline process models involved the presentation of the Petri nets for the different wheelset types without any additional information. Due to time constraints in the focus group, the discussion was focused on the model of wheelset type 328 that was also used for the evaluation of the subsequent model enhancement. Instead of treating all models at the surface level, it was decided in the interest of time to treat the most frequently occurring wheelset type in detail. Furthermore, the different data preparation operations and the general characteristics of the event log, as well as any assumptions made, were also discussed, as recommended by Koorn *et al.* [195].

Regarding the general characteristics of the event log, such as the number of cases, activities, median, and mean case duration, the experts agreed that all the properties presented seemed plausible in relation to how they experience the real-world process. A remark was made about the calculation of the median and mean case durations and whether they were accounted for by business hours. It was explained that this was not the case since the analysis did not focus on aspects such as throughput times. If such an analysis were to be made in the future, this is something to pay explicit attention to.

During the walk-through discussion of the model, it was mentioned whether the model could be annotated with the frequencies of the different paths. This could help distinguish paths that occur frequently from those that are followed more rarely. A short investigation in PM4Py showed that this is possible to implement in PM4Py for both the baseline and the enhanced Petri nets. In addition, this would also help identify the particular conditions that occur the most frequently.

The general feeling was that for the most part the baseline model quite truthfully represents the process as it is executed in practice. However, a part of the model allows for behavior that is impossible in practice, such as the option of parallel behavior that is physically impossible or behavior where the execution order is not enforced even though activities need to complete before the process can proceed. It is nevertheless important to note that the notion of parallelism in this respect is different from the notion of true concurrency. The activities need to complete, but the order is not specified. That means that they could happen exactly at the same time, but they might also execute independently, not stating which one of them goes first. For a more realistic portrayal of the process, these should be as accurate and strict as possible. However, this depends on the application of the analysis. The discussion was expanded to possible types of analysis and the relationship with the fitness of the models. On the one hand, this could be to have a correct overview of the process for documentation and training purposes, where you want to focus on the intended happy flow and disregard deviations. Subsequently, it was discussed that there are also paths that occur less frequently and that are still valid, and thus should be visible if the goal were to accurately document the process. On the other hand, Exp1 suggested that *"another product could be that, by this process mining exercise, we validate if all our products have been produced according to the prescribed process rules. [...] Then you should be able to see under which condition branching has occurred."* This hints at an advanced form of conformance checking, where the addition of the decision perspective could shed light on whether a deviation is actually a violation of procedures and prescriptions.

Key takeaways

- The general characteristics of the event log seemed plausible in relation to the real process

- A further analysis of the time perspective should pay explicit attention to the business hours
- The evaluated baseline model quite truthfully portrays the process, except for the last third
- One possible application mentioned was to document the process for training of new engineers
- Another mentioned application was to validate if all wheelsets have been treated correctly
- For both these applications, the trade-offs in terms of model fitness are different
 - For documentation, fitness should be set such that only the valid paths are shown
 - For conformance checking, all variants including rare deviations should be shown

Enhanced process models

The evaluation of the enhanced process model of wheelset type 328 was conducted similarly to the baseline model. The model was presented, and the discussion focused on the decision points that were present in the model and the discovered guard expressions with which those points were automatically annotated. The second part focused again on the possible applications and to gather feedback on how the artifacts could be improved to make them more accurate, appropriate, and useful.

Unfortunately, most of the guards discovered for the decision points in the model did not correspond to any sensible explanation of the path that was followed. An example of a condition that was considered valid by the experts is that of the optional examination step shown in Figure 6.7. However, the discussion that followed indicated that this was merely a trivial finding. The "UnLoadBaan" attribute is the result of a choice made after assessment by an operator, depending on whether the earlier on-press activity 335 was executed correctly. Therefore, it is logical that a strong correlation was identified between that attribute and the route followed, since the chosen path of the axle determines the activity that will follow. In the data, the attributes representing the outcome that stipulates which choice needs to be made are determined in an earlier activity than where the actual decision point is located. Consequently, the fact that the decision miner only considers the attributes locally at the decision point is a crucial shortcoming in actually identifying the underlying decision criteria.

In the discussion that followed from this finding, this possible area of improvement was further dissected. The information that explains where the information that the decision is based on can be found is registered in the written work instruction documents, but also in the MES software script that determines the routing. For this semantic enrichment to be implemented, the data could be enriched with an attribute that states for a decision point at which other activity the actual decision attributes are to be found. Another approach would be to always consider the attributes of n amount earlier activities than at the actual decision point to determine what influenced the actual choice. However, this might not be feasible when working with the classical notion of a decision tree projected onto the model. Furthermore, Exp1 suggested that it could also be an improvement to eliminate the attribute that has the strongest correlation but does not have any explanatory value for the choice. However, therefore, the data of an activity earlier than the decision point should nevertheless be considered.

The evaluation continued with the modeling construct shown in Figure 6.10. Exp1 remarked about the discovered conditions that *"it is really difficult to read, but it is what we do."* However, when the inspection continued, an interesting observation was made; the guard expressions in this part of the Petri net actually introduced a livelock. When the instance attribute is valued two or higher, this enables activity 333 but as long as that remains, it loops back to this initial decision point, and the continuation of the process with activity 396 firing is not possible. The experts stated that the value of the activity instance attribute never changes to a lower value. However, inspection of cases with an instance number greater than two revealed that it actually changes to a lower value for the latter part of the process.

The discussion then continued that activities with a different instance sequence number should actually be regarded as a separate version of that activity, which would imply that the model would probably change significantly. This could also explain why this attribute caused the introduction of a livelock, as it should be regarded as part of the activity itself and not the activity context. Further inspection of the data after the focus group revealed two different patterns with respect to the activity instance attribute. It is either consistent during the part of the process that is repeated, but changes back to value one for the final four activities of the process, or the value two is only emitted once for activity 333 which is the initiation of the rework. The problematic guard expression for activity 396 also appeared to have a cause, as the duration of that activity overlaps with a significant part of subsequent activities. The instance value for this activity was still one, and hence the guard expression. Therefore, it is incorrectly placed sequentially in the model, since life cycle information is ignored due to the PM4Py

implementation of the inductive miner. However, the suggestion of pertaining the instance attribute as part of the activity might be a solution, in combination with the retention of life cycle information. This different abstraction of the activity instance might be more insightful. Regarding what this adds to the analysis goals, Exp1 stated in a concise way that *"we want to understand the process, not the physical stations"* and Exp2 in turn confirmed that this adaptation would probably paint a better picture of the process. An additional analysis iteration with the information mentioned taken into account should be adequate to investigate whether this yields the desired result.



Figure 6.10: Part of the Petri net of wheelset type 328 where a livelock seems to be introduced by the guard expressions.

Key takeaways

- Most guard expressions for decision points were found irrelevant or lacking a sensible explanation
- A valid condition was identified, which actually revealed its dependence on earlier activities
- The decision miner its local consideration of attributes failed to capture the actual decision criteria
- A consideration is to add information which earlier activities and attributes should be considered
- A livelock problem in the Petri net was identified, caused by certain guard expressions
- Activity instances should become separate activities, potentially altering the model significantly
- The process flow should be understood logically rather than in terms of physical stations

Insights and process improvement potential

The final part of the evaluation dealt with what has been presented so far in terms of artifacts and insights, what the relevance has been so far, and what should be done in a subsequent analysis iteration. Exp1 explained earlier that his experiences with process mining had not been that fruitful up to now, but said that: *"Look, I am cautiously a bit positive that you are already showing more than what I have seen so far in process mining by adding those decisions. These are not yet the right decisions, and I get confused because now at the end and then again in the loops you are not handling the instances properly."* In the follow-up discussion, it turned out that the activity instance attribute might be crucial in distinguishing the logical from the physical execution of the process, as this signifies that some parts of the process can be repeated as part of a valid execution. However, the representation of this behavior in the data without paying attention to the instance attribute led to a model that exhibited less precision for the respective rework loops. Furthermore, wheelset repair activity 396 has a significantly longer duration and overlaps in the data with subsequent activities. Therefore, the model displays this activity as occurring later in time. This was the case because the inductive miner in PM4Py does not consider life cycle information but only completion events of the activities.

Another improvement that was discussed actually refers to semantically enriching the model by applying more elaborate feature engineering to the features that are eventually fed into the decision tree classifier. Exp1 phrased it as: *"If you were to stay here any longer, I would ask you to discuss this model again with an engineer and ask them, per activity where we see a choice, to indicate in the data which elements determine that choice according to the engineer. [...] You will then have a much better idea of what will happen. What he has now done is to look for the strongest correlation. [...] This feels actually a little bit dumb in the sense that it has been given too much freedom."* However, part of the research objective was to investigate whether this knowledge of the process could be extracted without upfront supplying elaborate semantic information about the process. Subsequently, it was discussed that attributes that are the result of, e.g., an interpretation of a measurement should be eliminated in favor of the attributes that relate to the actual decision criteria.

The final part of the discussion dealt with the handling of loops and rework in process mining analyses in general, where the sentiment was that this is still lacking adequate support. The detection of loops is a known problem for process mining algorithms. It was also touched upon that conditionally repeating a part of the process has a specific representation in terms of process mining concepts. In this specific process, it has been decided to have an instance sequence number to represent this, which could be translated into part of the activity concept. In other processes, especially industrial revision

processes, these could be represented differently, such as starting a new case or replacing values in the data. Nevertheless, if repeating a part of the process is a quintessential operation, for example, if requirements are not yet met, then explicit attention should be paid to within modeling and analysis.

The discussion ended with an outlook on what a follow-up product of the proposed concept would look like. The primary related business need that was presented was related to a control and validation mechanism such that it can be formally proven if an arbitrary yield of wheelsets has been produced according to the regulations. This is relevant, as Exp1 illustrates: *"Apparently we went through 262 different processes to deliver a wheelset. So, how do we know that all 262 variations have been valid and have produced a sound product? How can you guarantee that? [...] How can you adequately assess 262 different variations? [...] I think this should be possible if your model is a bit more accurate."* Exp2 added to that: *"Yes, this [concept] could then definitely help with that."*

Key takeaways

- Cautious optimism was expressed about the inclusion of decision points in process models
- Concerns were raised about improper handling of instance attributes in specific segments
- Activity instances were found to differentiate between logical and physical process executions
- There is a need for more elaborate feature engineering to enhance the decision tree classifier
- It was suggested to involve engineers in identifying data elements that influence decision-making
- Attributes directly related to decision criteria should be used, rather than derived interpretations
- Inadequacy of current process mining tools in handling loops and rework effectively was discussed
- The business need to validate whether wheelset production meets regulations was identified
- The potential of the concept to aid in assessment and validation of variations was acknowledged

In the following subsections, it is briefly described how the evaluation stage with the focus group was related to the activities that were identified in the original PM² methodology, as well as what it additionally involved for the decision mining aspect of the extended PM²xDM methodology.

Diagnose

The diagnosis of the findings began with the correct interpretation of the results in terms of the discovered model. Therefore, a joint understanding of the syntax and semantics of the Petri nets was first established, as this was the preferred modeling paradigm within this research project. Subsequently, a collaborative walk-through of the model was facilitated to correctly interpret the discovered model for the most common wheelset type 328. For any interesting or unusual results that were displayed in the model, a further in-depth review and discussion was conducted to identify possible causes. The identified results were collected to be used in new iterations of the analysis to further refine the model. Regarding the diagnosis of the identified decision points and related criteria, the same steps were repeated, but emphasis was placed on the syntax and semantics of the decisions.

Verify and validate

The verification of the results against the data and the implementation of the system was carried out for the discovered model, the decision points, and the associated attributes. For verification purposes, the characteristics and architecture of the process were discussed before diving into the actual modeled representations of the process. Some choices, such as how conditional repetition within the process is represented as an attribute within the data, could have severe consequences on how the process is eventually modeled. The same notion holds for the decision points in relation to the case and activity attributes. Some of them are the resulting interpretation of measurement data gathered at a different activity, which indicates that a local approach might not produce sensible results.

Regarding the validation aspect, which involves comparing the findings with the claims of the process stakeholders, some interesting discussions ensued with the domain experts. The modeled representation of the process exhibited some constructs that were physically impossible in practice. However, follow-up investigation showed that this was presumably caused by a small number of cases where data manipulation occurred due to a manual override. The remaining discrepancies were mostly caused by the data representation and abstraction of the actual process, which could be solved in an additional analysis iteration.

The validation of the decision attributes represented in the enhanced model eventually resulted in a discussion on how the decisions are taken in practice, versus how they are represented in the data. This offered insights on the important considerations of how an improved approach should regard also non-local information, either by considering earlier activities or artificially creating an attribute for that purpose. Furthermore, the discussion yielded an application for decision-enhanced models and analysis, where it could be employed for a novel form of conformance checking. This could be used to prove that, for example, a rejected wheelset was conditionally treated again until the applicable tolerances would have been met.

6.3.6. Stage 6: Process improvement and support

This stage was neither intended nor feasible to implement for this research project due to time constraints. However, at the time of writing some practical ideas for implementation are still in the works. The first subsequent iteration will focus on the refinement of the models so that they can be used to document the process. Next, the idea of production validation with the enhanced models will be further explored.

6.4. Case study summary

- A process mining project was initiated according to the PM² methodology
- Additional steps and extensions were executed to demonstrate decision mining
- Several intermediate iterations were needed to create a fitting baseline model
- An additional round of refinement is needed for the models to serve as documentation
- The baseline model was enhanced with guard expressions based on attributes
- The models were evaluated with the experts in a single evaluation episode (focus group)
- Additional knowledge was gathered about the process for further iterations
- Applications of the models and decision context enhancements have been discussed

Figure 6.11 shows the resulting overview of the activities of the PM²xDM methodology as they have been executed, omitted, deemed unfeasible, are still ongoing or out of scope for this research.

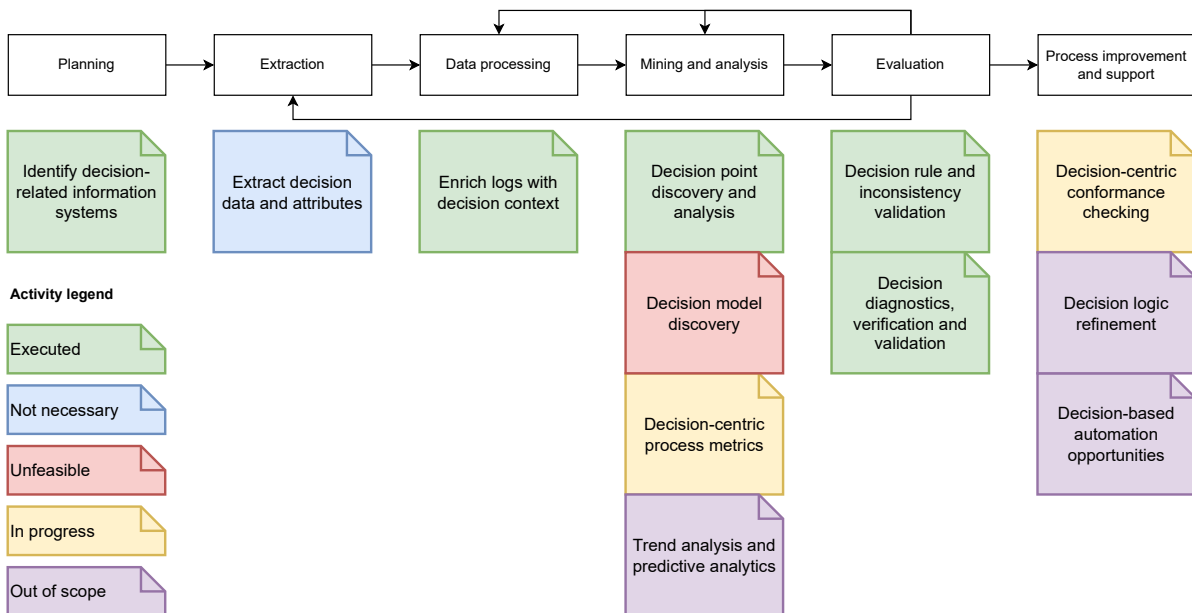


Figure 6.11: An overview of the PM²xDM framework after implementation within the case study and the evaluation.

During the focus group, the need arose to also include frequency information in the Petri net to distinguish frequent from infrequent paths in relation to the decision criteria. Figure 6.12 shows such an example of the section of the Petri net for wheelset type 328 that includes the guard expression annotations with additional frequency information for the respective paths.




Figure 6.12: Excerpt of the enhanced Petri net for wheelset type 328 that shows both the conditions and path frequencies.



Discussion

The objectives of this thesis were to create a better understanding of the role of decisions within processes and to explore whether and how process models and associated activities could be improved by leveraging the data attributes present in the event data. In previous chapters, the development of the methodology, its implementation, and evaluation have been presented to achieve this objective. This chapter describes the implications of the findings in this research project and how these relate to existing research, as well as the contributions from both a scientific and a practical perspective. In addition, a final assessment of the validity threats identified earlier will be carried out and any remaining challenges and limitations will be discussed. Finally, some of the opportunities for future research are identified and disseminated.

7.1. Implications

First and foremost, this research has demonstrated the potential relevance and applicability of decision mining within a process mining project. Enhanced process models were produced using case and activity attribute data. These models were produced with only limited initial semantic knowledge about the process. An analysis of the decision points within the process aided by such visualizations demonstrated an interesting starting point for further applications, such as richer process documentation that shows under which conditions certain paths are taken. In addition, an enhanced form of conformance checking could be developed using these models. Validation of whether the production of assets has been performed in accordance with the required guidelines and regulations could be supported using these artifacts. This implies that, depending on the project goals, it is worthwhile to assess the suitability for decision mining analysis. However, the quality of the condition annotations was not yet at a usable level. Therefore, improvements should be made to the input data and the decision mining algorithm. On the one hand, more elaborate feature engineering and reduction of the feature space are areas of optimization. On the other hand, the attributes from nonlocal activities should be considered, e.g. by enriching activities with attributes from earlier activities or a symbolic link that states the attributes of which other activities should be considered at a certain decision point.

Furthermore, it was investigated what and how activities should be carried out and what they entail in terms of suitable process characteristics and data requirements to pursue a relevant and meaningful decision-mining analysis. A significant observation was that it should be possible to obtain a sufficiently readable process model at fitness levels greater than 80% to be able to perform a meaningful analysis. An argument for this is that if specific deviations are not present in the model, these will also not be annotated with the conditions under which they occur. Therefore, this type of analysis is less applicable to processes that are only loosely structured or exhibit a very high degree of variation. This is in line with the analysis challenges posed by knowledge-intensive processes [163] or processes that accommodate a wide variety of different needs, such as healthcare processes [129].

7.2. Contributions

7.2.1. Scientific

The scientific contributions of this research are twofold. First, this research explored a potential avenue for a more holistic integration between process and decision mining, as suggested by De Smedt *et al.* [20]. Although it seemed unfeasible with the present tools and techniques to discover a fully integrated model of control flow and decisions [18], it nevertheless supports the notion that the underutilized data perspective of process mining can provide relevant insights [17, 59]. The methodology was implemented within a case study in a real-world context, and the resulting artifacts and insights were validated and evaluated within a focus group with the relevant experts. Although at least an additional iteration would be necessary to actually valorize the concept in practice, it is consistent with the literature that the concept can be potentially relevant and insightful [24].

Second, the foundational PM² methodology [99] has been extended with a decision mining component. The resulting methodological framework in the form of PM²xDM identifies which decision mining activities are related to the stages of the original methodology. The synthesis of the common activities based on the literature and the practical implementation not only helps to increase our common understanding of the intersection between process and decision mining, but also helps in shaping future research opportunities for the respective activities that have been defined.

7.2.2. Practical

One of the key success factors for a process mining project is following a structured approach [11]. Therefore, from a practical perspective, the proposed methodological framework can also help practitioners execute the decision mining component of a process mining project in a systematic way. Furthermore, the fact that it is based on and integrated with a generic process mining project methodology means that it can be included in an existing project if it aligns with the project goals. This in turn helps optimize efficient resource usage, as it does not require the creation of a distinct project as is the case with classical data mining projects that serve similar purposes [140].

7.3. Retrospective analysis of validity threats

Although validity threats have been extensively treated *ex ante* in Chapter 3, this section briefly reflects in hindsight on the five validity dimensions of the research project as a whole. The threats to the internal validity of this work are partially mitigated by following a systematic approach documented in Chapter 3 and the case study protocol in Appendix D. The procedures within the case study were partially executed by the researcher himself, possibly resulting in researcher bias. However, the focus group evaluation was carried out in a real-life setting with the researcher executing a documented protocol and, therefore, as naturalistic as possible.

To ensure construct validity, the theoretical background, concepts, metrics, and evaluation criteria have been based on a combination of extensive narrative and systematic reviews of the literature. Thus, the constructs defined for measurement using the research questions have been crafted as meticulously as possible. The external validity of this work concerns the generalizability of the findings in other contexts. On the one hand, the case study was conducted within a particular process context within a single company. Hence, some aspects of the methodology might be inadvertently tailored to this specific process. These aspects were investigated and reported as process characteristics. On the other hand, it has been acknowledged that the insights are specific to this process, and consequently, these should not be generalized. However, the steps and activities of the framework under scrutiny cover a wider area of applications. Although the framework itself was not specifically evaluated, a focus group was conducted to evaluate the implementation. Therefore, the conclusion validity of this work is deemed sufficient, but is still limited by the fact that a single case study was executed.

Finally, threats to the reliability of this study are mitigated in several ways. First, the case study was performed in a real-life context with actual data, where the procedures were based on established literature and have been extensively described. Second, all related materials are presented either in the main body of this document or in the appendices, except for classified materials which are stored in the case study repository of the researcher, such as raw data and focus group transcripts. Third and last, the nature of the case study stipulated several iterations of interim validation that have also been documented, which in turn provides an inherent form of repeatability.

7.4. Remaining challenges and limitations

A research project is certainly never perfect, and neither is this one. Therefore, there are several remaining challenges and limitations that must be acknowledged. First, the project was carried out in the convenient context of an internship project, which limits the generalizability of the research. To improve the methodological framework, the implementation should preferably be repeated with several other case studies in different contexts. Second, due to time constraints, the methodology itself was only executed and evaluated based on its artifacts and insights. However, it was not validated and evaluated in terms of procedures, for example, with process analysts. This could yield useful information on other dimensions, such as ease of use. Third, the visualizations presented were limited to Petri nets due to constraints within the tooling, whereas the focus group results showed that others, such as BPMN, might be more appropriate in practical applications where experts with limited process modeling knowledge are involved. Fourth and last, the decision mining results could be improved by more semantical enrichment of the attributes, so that it could better cope with activity attributes that are not local to a decision point and categorically typed attributes.

7.5. Future research opportunities

Future work could build on this research in several ways. First and foremost, the $PM^2 \times DM$ framework should be repeatedly applied in different environments and contexts to develop a more robust context-agnostic version. Such follow-up experiments could, in addition, contain a part that also pays special attention to the execution of the methodology itself by process analysts. Second, research could focus on developing a toolkit that integrates several of the decision mining assessment steps and activities of the framework into a single software package, for a more straightforward application within a process mining project. Furthermore, research could also focus on enabling additional interoperability between visualizations, such as the conversion of Petri nets with data into BPMN diagrams that retain these conditions.



Conclusion

This chapter briefly summarizes how the supporting research questions have been answered throughout this thesis, ultimately leading to the fulfillment of the main research objective.

SRQ1: *How do the disciplines of process mining and decision mining relate to each other in terms of their context, fundamental concept definitions and data requirements?*

Decision mining is a form of data mining that analyzes decision points in a given process. It tries to uncover the influence of data attributes on the choices within a process, regardless of whether these implicitly or explicitly stipulate separate branches in the control-flow of a process. If an explicit decision log is present, a decision model can be discovered that visualizes relations between the decision criteria. Combined with process mining, it can enhance process models with a decision perspective. This can be employed to find correlations between paths and criteria and, for example, to execute conformance checking that incorporates these criteria. A full elaboration on the relation is found in Chapter 2 of this thesis.

SRQ2: *Which methods or techniques already exist to detect decision points within a process using event data attributes?*

Several techniques have been implemented as form of decision mining shown in the literature review, ranging from association rule mining, deep learning to text mining. The technique most often applied is the use of decision tree algorithms that cluster attributes from the respective process paths in terms of attribute values. Additional information is found in Chapter 4 and the full list of assessed approaches is found in Appendix B.

SRQ3: *What is the current state of research in enhancement of process models with decision information?*

In Chapter 4 a brief overview is presented of approaches that enhance process models with decision information. Most importantly, there exists a large variance in the different applications that have been investigated. However, all approaches that were part of the review seem to have sought to improve decision-making within or around the process, instead of representing the criteria for the decisions within the process model itself. This distinguishes our approach from those that have already been investigated.

SRQ4: *What are relevant criteria and metrics to evaluate process models enhanced with decision information?*

The contextual implementations of the evaluation strategies in the papers considered in the literature review have been extracted. The most frequently used criteria are accuracy, relevance/applicability, and clarity/understandability. Subsequently, these have been used to synthesize the evaluation criteria for the artifacts within our research. An overview of the criteria for each article is found in Table 4.2.

SRQ5: *How to design an effective methodology to enhance process models with decision information?*

The execution of the case study was the implementation of a process mining project guided by the formally adopted methodology in the form of PM².

SRQ6: *How to convert event data attributes into decision conditions?*

For this research, a decision tree classifier was used to cluster the event log attributes and subsequently map these onto the decision points of a Petri net using Token Based Replay (TBR).

SRQ7: *How to integrate decision information into process models?*

A Petri net visualization of the process was enriched with the output of the decision tree. For each decision point where an applicable condition was found, it was annotated with a guard expression that states the Boolean condition under which the respective transition can fire. The resulting artifact is a Petri net with data (DPN).

SRQ8: *Where does the approach and methodological framework under investigation fit into the application of a process mining project methodology?*

How the different stages relate to the activities of the PM² methodology after implementation within the real-world case study is shown in Figure 6.11. This figure essentially shows the activities of the extension to PM² that was developed. Further details on the integrated PM²xDM methodology are found in the respective PDDs of both the original and extended methodology in Appendix C.

SRQ9: *How does the proposed methodology perform in a **real-world** organizational context?*

The methodology was evaluated by means of implementation within the case study. A final evaluation was conducted in a focus group that treated the resulting artifacts and insights. Most of the annotated decision points did not provide plausible conditions or additional value. However, with an additional analysis iteration, the results could already be significantly improved. This would at least entail more elaborate feature engineering. Other suggested improvements, such as the inclusion of attributes from activities that are not local to the decision point, will probably require more development and research. Please refer to Chapter 6 for more details on the implementation of the case study and its results.

MRO: *Improve* the representation of the influence of decisions within process models *by* the design and empirical validation of a methodological framework for integrated decision and process mining *such that* endogenous process data can be used in process and decision mining analyses *in order to* present a more realistic perspective on real-world processes within their respective modeling artifacts.

In terms of the main research objective, the methodological framework PM²xDM was developed based on the established process mining project methodology PM². It allowed us to enrich a Petri net process model with conditions based on the event data attributes, converting it into a Petri net with data (DPN). This research has shown that visualization of decisions in process models is *potentially* useful to organizations implementing a process mining project. Additionally, it helps to present a more realistic perspective on the process during discovery, and it allows for enhanced activities, such as decision-based conformance checking.

References

- [1] F. W. Taylor, *The principles of scientific management*. Harper & Brothers, New York & London, 1911.
- [2] D. Bumblauskas, H. Nold, P. Bumblauskas, and A. Igou, "Big data analytics: transforming data to action," *Business Process Management Journal*, vol. 23, no. 3, pp. 703–720, 2017.
- [3] C. M. Giannantonio and A. E. Hurley-Hanson, "Frederick Winslow Taylor: Reflections on the relevance of the principles of scientific management 100 years later," *Journal of Business & Management*, vol. 17, no. 1, 2011.
- [4] D. A. Wren, "The Centennial of Frederick W. Taylor's The Principles of Scientific Management: A Retrospective Commentary," *Journal of Business & Management*, vol. 17, no. 1, 2011.
- [5] A. Edmunds and A. Morris, "The problem of information overload in business organisations: a review of the literature," *International Journal of Information Management*, vol. 20, no. 1, pp. 17–28, 2000.
- [6] H. Chen, R. H. Chiang, and V. C. Storey, "Business intelligence and analytics: From big data to big impact," *MIS quarterly*, pp. 1165–1188, 2012.
- [7] W. M. P. Van der Aalst and A. J. Weijters, "Process mining: a research agenda," *Computers In Industry*, vol. 53, no. 3, pp. 231–244, 2004.
- [8] S. LaValle, E. Lesser, R. Shockley, M. S. Hopkins, and N. Kruschwitz, "Big data, analytics and the path from insights to value," *MIT Sloan Management Review*, 2010.
- [9] R. Sharma, S. Mithas, and A. Kankanhalli, "Transforming decision-making processes: a research agenda for understanding the impact of business analytics on organisations," *European Journal of Information Systems*, vol. 23, no. 4, pp. 433–441, 2014.
- [10] S. Zuboff, "Automate/informate: The two faces of intelligent technology," *Organizational Dynamics*, 1985.
- [11] A. Mamudu, W. Bandara, M. T. Wynn, and S. J. Leemans, "A Process Mining Success Factors Model," in *Business Process Management: 20th International Conference, BPM 2022, Münster, Germany, September 11–16, 2022, Proceedings*, Springer, 2022, pp. 143–160.
- [12] R. R. Mans, H. Reijers, H. Berends, W. Bandara, and R. Prince, "Business process mining success," in *Proceedings of the 21st European Conference on Information Systems*, Association for Information Systems (AIS), 2013, pp. 1–13.
- [13] M. O. Kherbouche, N. Laga, and P.-A. Masse, "Towards a better assessment of event logs quality," in *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, IEEE, 2016, pp. 1–8.
- [14] S. Suriadi, R. Andrews, A. H. ter Hofstede, and M. T. Wynn, "Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs," *Information systems*, vol. 64, pp. 132–150, 2017.
- [15] D. A. Fischer, K. Goel, R. Andrews, C. G. J. van Dun, M. T. Wynn, and M. Röglinger, "Enhancing event log quality: Detecting and quantifying timestamp imperfections," in *Business Process Management: 18th International Conference, BPM 2020, Seville, Spain, September 13–18, 2020, Proceedings 18*, Springer, 2020, pp. 309–326.
- [16] B. R. van Wensveen, "Estimation and analysis of the quality of event log samples for process discovery," M.S. thesis, 2020.
- [17] A. Banham, S. J. Leemans, M. T. Wynn, and R. Andrews, "xPM: a Framework for Process Mining with Exogenous Data," in *Process Mining Workshops: ICPM 2021 International Workshops, Eindhoven, The Netherlands, October 31–November 4, 2021, Revised Selected Papers*, Springer International Publishing Cham, 2022, pp. 85–97.

- [18] J. De Smedt, S. K. vanden Broucke, J. Obregon, A. Kim, J.-Y. Jung, and J. Vanthienen, "Decision mining in a broader context: An overview of the current landscape and future directions," in *Business Process Management Workshops: BPM 2016 International Workshops, Rio de Janeiro, Brazil, September 19, 2016, Revised Papers 14*, Springer, 2017, pp. 197–207.
- [19] J. Vanthienen, F. Caron, and J. De Smedt, "Business rules, decisions and processes: five reflections upon living apart together," in *Proceedings SIGBPS Workshop on Business Processes and Services (BPS'13)*, 2013, pp. 76–81.
- [20] J. De Smedt, F. Hasić, S. K. vanden Broucke, and J. Vanthienen, "Towards a holistic discovery of decisions in process-aware information systems," in *Business Process Management: 15th International Conference, BPM 2017, Barcelona, Spain, September 10–15, 2017, Proceedings 15*, Springer, 2017, pp. 183–199.
- [21] J. De Smedt, F. Hasić, S. K. vanden Broucke, and J. Vanthienen, "Holistic discovery of decision models from process execution data," *Knowledge-Based Systems*, vol. 183, p. 104866, 2019.
- [22] A. Rozinat and W. M. P. van der Aalst, "Decision mining in ProM," *Business Process Management*, vol. 4102, pp. 420–425, 2006.
- [23] A. Rozinat, R. Mans, M. Song, and W. M. P. Van der Aalst, "Discovering colored Petri nets from event logs," *International Journal on Software Tools for Technology Transfer*, vol. 10, pp. 57–74, 2008.
- [24] L. Janssens, E. Bazhenova, J. De Smedt, J. Vanthienen, and M. Denecker, "Consistent Integration of Decision (DMN) and Process (BPMN) Models.," in *CAiSE forum*, Citeseer, vol. 1612, 2016, pp. 121–128.
- [25] R. Dunkl, S. Rinderle-Ma, W. Grossmann, and K. Anton Fröschl, "A Method for Analyzing Time Series Data in Process Mining: Application and Extension of Decision Point Analysis," in *Information Systems Engineering in Complex Environments: CAiSE Forum 2014, Thessaloniki, Greece, June 16-20, 2014, Selected Extended Papers 26*, Springer, 2015, pp. 68–84.
- [26] B. Scheibel and S. Rinderle-Ma, "Decision mining with time series data based on automatic feature generation," in *Advanced Information Systems Engineering: 34th International Conference, CAiSE 2022, Leuven, Belgium, June 6–10, 2022, Proceedings*, Springer, 2022, pp. 3–18.
- [27] K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A design science research methodology for information systems research," *Journal of Management Information Systems*, vol. 24, no. 3, pp. 45–77, 2007.
- [28] R. J. Wieringa, *Design Science Methodology for Information Systems and Software Engineering*. Springer, 2014.
- [29] A. Lindsay, D. Downs, and K. Lunn, "Business processes—attempts to find a definition," *Information and software technology*, vol. 45, no. 15, pp. 1015–1019, 2003.
- [30] M. A. Ould, *Business processes: modelling and analysis for re-engineering and improvement*. Wiley, 1995.
- [31] I. Jacobson, M. Ericsson, and A. Jacobson, *The object advantage: business process reengineering with object technology*. ACM Press/Addison-Wesley Publishing Co., 1994.
- [32] M. Hammer and J. Champy, *Reengineering the corporation: A manifesto for business revolution*. Harper Collins, New York, 1993, p. 233.
- [33] T. H. Davenport, *Process innovation: reengineering work through information technology*. Harvard Business Press, 1993.
- [34] H.-E. Eriksson and M. Penker, *Business Modeling With UML: Business Patterns at Work*. John Wiley & Sons, Inc., New York, 1998.
- [35] N. Melão and M. Pidd, "A conceptual framework for understanding business processes and business process modelling," *Information Systems Journal*, vol. 10, no. 2, pp. 105–129, 2000.
- [36] C. A. Petri, "Fundamentals of a Theory of Asynchronous Information Flow.," in *IFIP congress*, vol. 62, 1962, pp. 386–390.
- [37] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE*, vol. 77, no. 4, pp. 541–580, 1989.

- [38] J. Recker, M. Rosemann, M. Indulska, and P. Green, "Business process modeling—a comparative analysis," *Journal of the Association for Information Systems*, vol. 10, no. 4, p. 1, 2009.
- [39] M. Indulska, J. Recker, M. Rosemann, and P. Green, "Business process modeling: Current issues and future challenges," in *Advanced Information Systems Engineering: 21st International Conference, CAiSE 2009, Amsterdam, The Netherlands, June 8-12, 2009. Proceedings 21*, Springer, 2009, pp. 501–514.
- [40] J. L. Peterson, "Petri nets," *ACM Computing Surveys (CSUR)*, vol. 9, no. 3, pp. 223–252, 1977.
- [41] J. Desel and G. Juhás, "“What Is a Petri Net?” Informal Answers for the Informed Reader," in *Unifying Petri Nets: Advances in Petri Nets*, Springer, 2001, pp. 1–25.
- [42] W. M. P. Van der Aalst, "The application of Petri nets to workflow management," *Journal of Circuits, Systems, and Computers*, vol. 8, no. 01, pp. 21–66, 1998.
- [43] K. M. van Hee, N. Sidorova, and J. M. van der Werf, "Business process modeling using petri nets," in *Transactions on Petri Nets and Other Models of Concurrency VII*, Springer, 2013, pp. 116–161.
- [44] I. Raedts, M. Petkovic, Y. S. Usenko, J. M. E. van der Werf, J. F. Groote, and L. J. Somers, "Transformation of BPMN Models for Behaviour Analysis.," *MSVVEIS*, vol. 2007, pp. 126–137, 2007.
- [45] M. Chinosi and A. Trombetta, "BPMN: An introduction to the standard," *Computer Standards & Interfaces*, vol. 34, no. 1, pp. 124–134, 2012.
- [46] OMG, *Business Process Model and Notation (BPMN), Version 2.0.2*, Object Management Group, Dec. 2013. [Online]. Available: <http://www.omg.org/spec/BPMN/2.0.2>.
- [47] R. M. Dijkman, M. Dumas, and C. Ouyang, "Formal semantics and analysis of BPMN process models using Petri nets," *Queensland University of Technology, Tech. Rep*, pp. 1–30, 2007.
- [48] V. Gurbaxani and S. Whang, "The impact of information systems on organizations and markets," *Communications of the ACM*, vol. 34, no. 1, pp. 59–73, 1991.
- [49] "The winning organization," *Fortune*, pp. 50–60, Sep. 26, 1988.
- [50] R. L. Ackoff, "Management misinformation systems," *Management Science*, vol. 14, no. 4, B147–B156, Dec. 1967.
- [51] W. M. P. van der Aalst, A. H. M. ter Hofstede, and M. Weske, "Business Process Management: A Survey," in *Business Process Management: International Conference, BPM 2003, Eindhoven, The Netherlands, June 26–27, 2003 Proceedings 1*, Springer, 2003, pp. 1–12.
- [52] W. M. P. van der Aalst, "Business Process Management: A personal view," *Business Process Management Journal*, vol. 10, no. 2, 2004.
- [53] W. M. P. Van der Aalst, T. Weijters, and L. Maruster, "Workflow mining: Discovering process models from event logs," *IEEE transactions on knowledge and data engineering*, vol. 16, no. 9, pp. 1128–1142, 2004.
- [54] A. A. Kalenkova, W. M. P. van der Aalst, I. A. Lomazova, and V. A. Rubin, "Process mining using BPMN: relating event logs and process models," in *Proceedings of the ACM/IEEE 19th International Conference on Model Driven Engineering Languages and Systems*, 2016, pp. 123–123.
- [55] W. M. P. van der Aalst, *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011, vol. 2.
- [56] C. Larman and V. R. Basili, "Iterative and incremental developments: A brief history," *Computer*, vol. 36, no. 6, pp. 47–56, 2003.
- [57] W. M. P. van der Aalst, B. F. Van Dongen, J. Herbst, L. Maruster, G. Schimm, and A. J. M. M. Weijters, "Workflow mining: A survey of issues and approaches," *Data & Knowledge Engineering*, vol. 47, no. 2, pp. 237–267, 2003.
- [58] W. M. P. van der Aalst, "Process mining: Overview and opportunities," *ACM Transactions on Management Information Systems (TMIS)*, vol. 3, no. 2, pp. 1–17, 2012.
- [59] W. M. P. van der Aalst, *Process Mining: Data Science in Action*. Springer, 2016, vol. 2.

- [60] J. F. Yates, *Decision management: How to assure better decisions in your company*. John Wiley & Sons, 2003.
- [61] J. Ghattas, P. Soffer, and M. Peleg, "Improving business process decision making based on past experience," *Decision Support Systems*, vol. 59, pp. 93–107, 2014.
- [62] F. Burstein, C. W. Holsapple, and D. J. Power, "Decision support systems: a historical overview," *Handbook on Decision Support Systems 1: Basic themes*, pp. 121–140, 2008.
- [63] D. Arnott and G. Pervan, "Eight key issues for the decision support systems discipline," *Decision Support Systems*, vol. 44, no. 3, pp. 657–672, 2008.
- [64] V. B. Vukšić, M. P. Bach, and A. Popovič, "Supporting performance management with business process management and business intelligence: A case analysis of integration and orchestration," *International journal of information management*, vol. 33, no. 4, pp. 613–619, 2013.
- [65] T. Morgan, *Business rules and information systems: aligning IT with business goals*. Addison-Wesley Professional, 2002.
- [66] O. Levina, O. Holschke, and J. Rake-Revelant, "Extracting business logic from business process models," in *2010 2nd IEEE International Conference on Information Management and Engineering*, IEEE, 2010, pp. 289–293.
- [67] P. Kardasis and P. Loucopoulos, "Expressing and organising business rules," *Information and Software technology*, vol. 46, no. 11, pp. 701–718, 2004.
- [68] J. Boyer, H. Mili, J. Boyer, and H. Mili, *Agile business rule development*. Springer, 2011.
- [69] K. Taveter and G. Wagner, "Agent-oriented enterprise modeling based on business rules," in *Conceptual Modeling—ER 2001: 20th International Conference on Conceptual Modeling Yokohama, Japan, November 27–30, 2001 Proceedings 20*, Springer, 2001, pp. 527–540.
- [70] B. Von Halle and L. Goldberg, *The decision model: a business logic framework linking business and technology*. CRC Press, 2009.
- [71] M. Zur Muehlen and M. Indulska, "Modeling languages for business processes and business rules: A representational analysis," *Information systems*, vol. 35, no. 4, pp. 379–390, 2010.
- [72] S. Goedertier and J. Vanthienen, "Declarative process modeling with business vocabulary and business rules," in *On the Move to Meaningful Internet Systems 2007: OTM 2007 Workshops: OTM Confederated International Workshops and Posters, AWeSOMe, CAMS, OTM Academy Doctoral Consortium, MONET, OnToContent, ORM, PerSys, PPN, RDDS, SWS, and SWS 2007, Vilamoura, Portugal, November 25-30, 2007, Proceedings, Part I*, Springer, 2007, pp. 603–612.
- [73] J. Shao and C. Pound, "Extracting business rules from information systems," *BT Technology Journal*, vol. 17, no. 4, pp. 179–186, 1999.
- [74] M. Wang and H. Wang, "From process logic to business logic—A cognitive approach to business process management," *Information & Management*, vol. 43, no. 2, pp. 179–193, 2006.
- [75] J. Taylor and J. Purchase, *Real-world decision modeling with DMN*. Meghan-Kiffer Press Tampa, 2016.
- [76] T. Biard, A. Le Mauff, M. Bigand, and J.-P. Bourey, "Separation of decision modeling from business process modeling using new "Decision Model and Notation" (DMN) for automating operational decision-making," in *Risks and Resilience of Collaborative Networks: 16th IFIP WG 5.5 Working Conference on Virtual Enterprises, PRO-VE 2015, Albi, France,, October 5-7, 2015, Proceedings 16*, Springer, 2015, pp. 489–496.
- [77] E. Bazhenova, F. Zerbato, B. Oliboni, and M. Weske, "From BPMN process models to DMN decision models," *Information Systems*, vol. 83, pp. 69–88, 2019.
- [78] L. Janssens, J. De Smedt, and J. Vanthienen, "Modeling and enacting enterprise decisions," in *Advanced Information Systems Engineering Workshops: CAiSE 2016 International Workshops, Ljubljana, Slovenia, June 13-17, 2016, Proceedings 28*, Springer, 2016, pp. 169–180.
- [79] H. van der Aa, H. Leopold, K. Batoulis, M. Weske, and H. A. Reijers, "Integrated process and decision modeling for data-driven processes," in *Business Process Management Workshops: BPM 2015, 13th International Workshops, Innsbruck, Austria, August 31–September 3, 2015, Revised Papers 13*, Springer, 2016, pp. 405–417.

- [80] K. Batoulis, A. Meyer, E. Bazhenova, G. Decker, and M. Weske, "Extracting decision logic from process models," in *Advanced Information Systems Engineering: 27th International Conference, CAiSE 2015, Stockholm, Sweden, June 8-12, 2015, Proceedings 27*, Springer, 2015, pp. 349–366.
- [81] S. Goedertier, J. Vanthienen, and F. Caron, "Declarative business process modelling: principles and modelling languages," *Enterprise Information Systems*, vol. 9, no. 2, pp. 161–185, 2015.
- [82] OMG, *Decision Model and Notation (DMN), Version 1.3*, Object Management Group, Feb. 2021. [Online]. Available: <http://www.omg.org/spec/DMN/1.3>.
- [83] M. de Leoni and F. Mannhardt, "Decision Discovery in Business Processes," in *Encyclopedia of Big Data Technologies*, 2019.
- [84] F. Burstein, C. W Holsapple, and W. M. P. van der Aalst, "Decision support based on process mining," *Handbook on Decision Support Systems 1: Basic Themes*, pp. 637–657, 2008.
- [85] P. Wohed, W. M. P. van der Aalst, M. Dumas, A. H. ter Hofstede, and N. Russell, "On the suitability of BPMN for business process modelling," in *Business Process Management: 4th International Conference, BPM 2006, Vienna, Austria, September 5-7, 2006. Proceedings 4*, Springer, 2006, pp. 161–176.
- [86] A. Rozinat, R. Mans, and W. Van Der Aalst, "Mining CPN models: discovering process models with data from event logs," in *7th Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools (CPN'06), October 24-26, 2006, Aarhus, Denmark*, University of Aarhus, 2006, pp. 57–76.
- [87] C. A. Amaral, M. Fantinato, H. A. Reijers, and S. M. Peres, "Enhancing completion time prediction through attribute selection," in *Information Technology for Management: Emerging Research and Applications: 15th Conference, AITM 2018, and 13th Conference, ISM 2018, Held as Part of FedCSIS, Poznan, Poland, September 9–12, 2018, Revised and Extended Selected Papers 16*, Springer, 2019, pp. 3–23.
- [88] N. Tax, N. Sidorova, R. Haakma, and W. M. P. van der Aalst, "Event abstraction for process mining using supervised learning techniques," in *Proceedings of SAI Intelligent Systems Conference (IntelliSys) 2016: Volume 1*, Springer, 2018, pp. 251–269.
- [89] R. de Jong, S. Leewis, and M. Berkhout, "Decision Mining versus Process Mining: a Comparison of Mining Methods," in *2021 5th International Conference on Software and e-Business (ICSEB)*, 2021, pp. 28–32.
- [90] I. Beerepoot *et al.*, "A Window of Opportunity: Active Window Tracking for Mining Work Practices," in *2023 5th International Conference on Process Mining (ICPM)*, IEEE, 2023, pp. 57–64.
- [91] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," *Management Information Systems Quarterly*, vol. 28, no. 1, p. 6, 2004.
- [92] K. Peffers, M. Rothenberger, T. Tuunanen, and R. Vaezi, "Design science research evaluation," in *Design Science Research in Information Systems. Advances in Theory and Practice: 7th International Conference, DESRIST 2012, Las Vegas, NV, USA, May 14-15, 2012. Proceedings 7*, Springer, 2012, pp. 398–410.
- [93] J. Venable, J. Pries-Heje, and R. Baskerville, "A comprehensive framework for evaluation in design science research," in *Design Science Research in Information Systems. Advances in Theory and Practice: 7th International Conference, DESRIST 2012, Las Vegas, NV, USA, May 14-15, 2012. Proceedings 7*, Springer, 2012, pp. 423–438.
- [94] N. Prat, I. Comyn-Wattiau, and J. Akoka, "Artifact evaluation in information systems design-science research — A holistic view," in *PACIS 2014: 18th Pacific Asia Conference on Information Systems*, 2014, p. 23.
- [95] S. Gregor and A. R. Hevner, "Positioning and presenting design science research for maximum impact," *MIS Quarterly*, pp. 337–355, 2013.
- [96] M. Juntunen and M. Lehenkari, "A narrative literature review process for an academic business research thesis," *Studies in Higher Education*, vol. 46, no. 2, pp. 330–342, 2021.

- [97] B. Kitchenham, O. P. Brereton, D. Budgen, M. Turner, J. Bailey, and S. Linkman, "Systematic literature reviews in software engineering—a systematic literature review," *Information and Software Technology*, vol. 51, no. 1, pp. 7–15, 2009.
- [98] R. B. Johnson, A. J. Onwuegbuzie, and L. A. Turner, "Toward a definition of mixed methods research," *Journal of Mixed Methods Research*, vol. 1, no. 2, pp. 112–133, 2007.
- [99] M. L. Van Eck, X. Lu, S. J. Leemans, and W. M. P. Van Der Aalst, "PM2: a process mining project methodology," in *Advanced Information Systems Engineering: 27th International Conference, CAiSE 2015, Stockholm, Sweden, June 8-12, 2015, Proceedings*, Springer, 2015, pp. 297–313.
- [100] C. Shearer, "The CRISP-DM model: the new blueprint for data mining," *Journal of Data Warehousing*, vol. 5, no. 4, pp. 13–22, 2000.
- [101] Y. Xiao and M. Watson, "Guidance on conducting a systematic literature review," *Journal of Planning Education and Research*, vol. 39, no. 1, pp. 93–112, 2019.
- [102] W. Bandara, S. Miskon, and E. Fieft, "A systematic, tool-supported method for conducting literature reviews in information systems," in *ECIS 2011 proceedings [19th European conference on information systems]*, AIS Electronic Library (AISeL)/ Association for Information Systems, 2011, pp. 1–13.
- [103] Ought, *Elicit: The AI Research Assistant*, Feb. 22, 2023. [Online]. Available: <https://elicit.org>.
- [104] B. D. Lund and T. Wang, "Chatting about ChatGPT: how may AI and GPT impact academia and libraries?" *Library Hi Tech News*, 2023.
- [105] A. Hannousse, "Searching relevant papers for software engineering secondary studies: Semantic Scholar coverage and identification role," *IET Software*, vol. 15, no. 1, pp. 126–146, 2021.
- [106] D. W. Aksnes, L. Langfeldt, and P. Wouters, "Citations, citation indicators, and research quality: An overview of basic concepts and theories," *Sage Open*, vol. 9, no. 1, p. 2 158 244 019 829 575, 2019.
- [107] R. K. Yin, *Case Study Research and Applications: Design and Methods*, 6th ed. Sage, 2018.
- [108] B. Yazan, "Three approaches to case study methods in education: Yin, Merriam, and Stake," *The Qualitative Report*, vol. 20, no. 2, pp. 134–152, 2015.
- [109] G. Pervan and M. Maimbo, "Designing a case study protocol for application in IS research," in *Proceedings of the Ninth Pacific Asia Conference on Information Systems (PACIS)*, PACIS, 2005, pp. 1281–1292.
- [110] M. Höst, A. Rainer, P. Runeson, and B. Regnell, *Case study research in software engineering: Guidelines and examples*. John Wiley & Sons, 2012.
- [111] K. Smit and J. Mens, "Process mining in the rail industry: a qualitative analysis of success factors and remaining challenges," in *BLEED 2019 Proceedings*, 2019.
- [112] P. Kecman and R. M. Goverde, "Process mining of train describer event data and automatic conflict identification," *Computers in Railways XIII: Computer System Design and Operation in the Railway and Other Transit Systems*, vol. 127, p. 227, 2013.
- [113] E. Bandis, M. Petridis, and S. Kapetanakis, "Business process workflow mining using machine learning techniques for the rail transport industry," in *Artificial Intelligence XXXV: 38th SGAI International Conference on Artificial Intelligence, AI 2018, Cambridge, UK, December 11–13, 2018, Proceedings 38*, Springer, 2018, pp. 446–451.
- [114] E. Bandis, M. Petridis, and S. Kapetanakis, "Predictive process mining using a hybrid CBR approach for the rail transport industry," *ICCBR 2018*, p. 175, 2018.
- [115] R. A. Krueger and M. A. Casey, "Focus group interviewing," *Handbook of practical program evaluation*, pp. 506–534, 2015.
- [116] P. Runeson and M. Höst, "Guidelines for conducting and reporting case study research in software engineering," *Empirical Software Engineering*, vol. 14, pp. 131–164, 2009.
- [117] R. E. Roberts, "Qualitative interview questions: Guidance for novice researchers.," *Qualitative Report*, vol. 25, no. 9, 2020.


- [118] D. S. Cruzes, T. Dybå, P. Runeson, and M. Höst, "Case studies synthesis: a thematic, cross-case, and narrative synthesis worked example," *Empirical Software Engineering*, vol. 20, pp. 1634–1665, 2015.
- [119] X. Zhou, Y. Jin, H. Zhang, S. Li, and X. Huang, "A map of threats to validity of systematic literature reviews in software engineering," in *2016 23rd Asia-Pacific Software Engineering Conference (APSEC)*, IEEE, 2016, pp. 153–160.
- [120] A. Ampatzoglou, S. Bibi, P. Avgeriou, M. Verbeek, and A. Chatzigeorgiou, "Identifying, categorizing and mitigating threats to validity in software engineering secondary studies," *Information and Software Technology*, vol. 106, pp. 201–230, 2019.
- [121] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in Software Engineering*. Springer Science & Business Media, 2012.
- [122] M. Gibbert, W. Ruigrok, and B. Wicki, "What passes as a rigorous case study?" *Strategic Management Journal*, vol. 29, no. 13, pp. 1465–1474, 2008.
- [123] C. Quintão, P. Andrade, and F. Almeida, "How to Improve the Validity and Reliability of a Case Study Approach?" *Journal of Interdisciplinary Studies in Education*, vol. 9, no. 2, pp. 264–275, 2020.
- [124] B. Kitchenham, L. Pickard, and S. L. Pfleeger, "Case studies for method and tool evaluation," *IEEE software*, vol. 12, no. 4, pp. 52–62, 1995.
- [125] M. J. Page *et al.*, "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews," *International Journal of Surgery*, vol. 88, p. 105 906, 2021.
- [126] G. Park, A. Küsters, M. Tews, C. Pitsch, J. Schneider, and W. M. P. van der Aalst, "Explainable predictive decision mining for operational support," in *International Conference on Service-Oriented Computing*, Springer, 2022, pp. 66–79.
- [127] W. M. P. Van Der Aalst, *A practitioner's guide to process mining: Limitations of the directly-follows graph*, 2019.
- [128] M. De Leoni and W. M. P. Van Der Aalst, "Data-aware process mining: Discovering decisions in processes using alignments," in *Proceedings of the 28th annual ACM symposium on applied computing*, 2013, pp. 1454–1461.
- [129] J. Munoz-Gama *et al.*, "Process mining for healthcare: Characteristics and challenges," *Journal of Biomedical Informatics*, vol. 127, p. 103 994, 2022.
- [130] J. Campos, P. Richetti, F. A. Baião, and F. M. Santoro, "Discovering business rules in knowledge-intensive processes through decision mining: An experimental study," in *Business Process Management Workshops: BPM 2017 International Workshops, Barcelona, Spain, September 10-11, 2017, Revised Papers 15*, Springer, 2018, pp. 556–567.
- [131] J. Wei *et al.*, "Emergent abilities of large language models," *arXiv preprint arXiv:2206.07682*, 2022.
- [132] R. Petrusel, "Aggregating individual models of decision-making processes," in *Advanced Information Systems Engineering: 24th International Conference, CAiSE 2012, Gdansk, Poland, June 25-29, 2012. Proceedings 24*, Springer, 2012, pp. 47–63.
- [133] F. Hasić, J. De Smedt, S. vanden Broucke, and E. Serral, "Decision as a service (DaaS): a service-oriented architecture approach for decisions in processes," *IEEE Transactions on Services Computing*, vol. 15, no. 2, pp. 904–917, 2020.
- [134] F. Hasić, J. De Smedt, and J. Vanthienen, "Developing a modelling and mining framework for integrated processes and decisions," in *On the Move to Meaningful Internet Systems. OTM 2017 Workshops: Confederated International Workshops, EI2N, FBM, ICSP, Meta4eS, OTMA 2017 and ODBASE Posters 2017, Rhodes, Greece, October 23–28, 2017, Revised Selected Papers*, Springer, 2018, pp. 259–269.
- [135] D. Thabet, S. Ayachi Ghannouchi, and H. Hajjami Ben Ghezala, "A process mining-based solution for business process model extension with cost perspective context-based cost data analysis and case study," in *Computer Information Systems and Industrial Management: 17th International Conference, CISIM 2018, Olomouc, Czech Republic, September 27-29, 2018, Proceedings 17*, Springer, 2018, pp. 434–446.

- [136] E. Brzychczy, P. Gackowiec, and M. Liebetrau, "Data analytic approaches for mining process improvement—machinery utilization use case," *Resources*, vol. 9, no. 2, p. 17, 2020.
- [137] K. Shahzad and J. Zdravkovic, "A decision-based approach to business process improvement," in *2010 IEEE International Conference on Systems, Man and Cybernetics*, IEEE, 2010, pp. 810–818.
- [138] A. Farooqui, K. Bengtsson, P. Falkman, and M. Fabian, "From factory floor to process models: A data gathering approach to generate, transform, and visualize manufacturing processes," *CIRP Journal of Manufacturing Science and Technology*, vol. 24, pp. 6–16, 2019.
- [139] S. Pereira Detro, E. A. P. Santos, H. Panetto, E. D. Loures, M. Lezoche, and C. Cabral Moro Barra, "Applying process mining and semantic reasoning for process model customisation in healthcare," *Enterprise Information Systems*, vol. 14, no. 7, pp. 983–1009, 2020.
- [140] K.-M. Osei-Bryson, "A context-aware data mining process model based framework for supporting evaluation of data mining results," *Expert Systems with Applications*, vol. 39, no. 1, pp. 1156–1164, 2012.
- [141] A. Ivanchikj, S. Serbout, and C. Pautasso, "From text to visual bpmn process models: Design and evaluation," in *Proceedings of the 23rd ACM/IEEE international conference on model driven engineering languages and systems*, 2020, pp. 229–239.
- [142] Y. Liu and G. Salvendy, "Design and evaluation of visualization support to facilitate decision trees classification," *International journal of human-computer studies*, vol. 65, no. 2, pp. 95–110, 2007.
- [143] Q. Liu, G. Feng, N. Wang, and G. K. Tayi, "A multi-objective model for discovering high-quality knowledge based on data quality and prior knowledge," *Information Systems Frontiers*, vol. 20, pp. 401–416, 2018.
- [144] S. Yang *et al.*, "A data-driven process recommender framework," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017, pp. 2111–2120.
- [145] C. Combi, B. Oliboni, and P. Sala, "Customizing bpmn diagrams using timelines," in *26th International Symposium on Temporal Representation and Reasoning (TIME 2019)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [146] D. Raheja, J. Llinas, R. Nagi, and C. Romanowski, "Data fusion/data mining-based architecture for condition-based maintenance," *International journal of production research*, vol. 44, no. 14, pp. 2869–2887, 2006.
- [147] I. Bruha, "Evaluation of decision rules by qualities for decision-making systems," in *Encyclopedia of Data Warehousing and Mining, Second Edition*, IGI Global, 2009, pp. 795–801.
- [148] B. Al-Salim and M. Abdoli, "Data mining for decision support of the quality improvement process," *AMCIS 2005 Proceedings*, p. 115, 2005.
- [149] J. Gallardo, G. Giacaman, C. Meneses, and Ó. Marbán, "Framework for decisional business modeling and requirements modeling in data mining projects," in *Intelligent Data Engineering and Automated Learning-IDEAL 2009: 10th International Conference, Burgos, Spain, September 23-26, 2009. Proceedings 10*, Springer, 2009, pp. 268–275.
- [150] S. Ghosh and J. Maiti, "Data mining driven DMAIC framework for improving foundry quality—a case study," *Production Planning & Control*, vol. 25, no. 6, pp. 478–493, 2014.
- [151] J. Wątróbski, J. Jankowski, and Z. Piotrowski, "The selection of multicriteria method based on unstructured decision problem description," in *Computational Collective Intelligence. Technologies and Applications: 6th International Conference, ICCCI 2014, Seoul, Korea, September 24-26, 2014. Proceedings 6*, Springer, 2014, pp. 454–465.
- [152] M. Žnidaršič, M. Bohanec, and B. Zupan, "Data-driven revision of decision models," in *Encyclopedia of Data Warehousing and Mining, Second Edition*, IGI Global, 2009, pp. 617–623.
- [153] K. Tsang, H. Lau, and S. Kwok, "Development of a data mining system for continual process quality improvement," *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 221, no. 2, pp. 179–193, 2007.
- [154] D. C. Corrales, A. Ledezma, and J. C. Corrales, "A conceptual framework for data quality in knowledge discovery tasks (fdq-kdt): A proposal," *J. Comput.*, vol. 10, no. 6, pp. 396–405, 2015.

- [155] L. Gu and R. Baxter, "Decision models for record linkage," *Data Mining: Theory, Methodology, Techniques, and Applications*, pp. 146–160, 2006.
- [156] M.-A. Mortada and S. Yacout, "Cbmlad-using logical analysis of data in condition based maintenance," in *2011 3rd International Conference on Computer Research and Development*, IEEE, vol. 4, 2011, pp. 30–34.
- [157] S. Peeters, "Decision mining in prom: From log files to dmn decision logic," Ph.D. dissertation, Ph. D. dissertation, Ghent University, 2016.
- [158] F. G. Alizamini, M. M. Pedram, M. Alishahi, and K. Badie, "Data quality improvement using fuzzy association rules," in *2010 International Conference on Electronics and Information Engineering*, IEEE, vol. 1, 2010, pp. V1–468.
- [159] S. Suriadi, M. T. Wynn, C. Ouyang, A. H. ter Hofstede, and N. J. van Dijk, "Understanding process behaviours in a large insurance company in australia: A case study," in *Advanced Information Systems Engineering: 25th International Conference, CAiSE 2013, Valencia, Spain, June 17-21, 2013. Proceedings 25*, Springer, 2013, pp. 449–464.
- [160] I. van de Weerd and S. Brinkkemper, "Meta-modeling for situational analysis and design methods," in *Handbook of research on modern systems analysis and design technologies and applications*, IGI Global, 2009, pp. 35–54.
- [161] N. Martin *et al.*, "Opportunities and challenges for process mining in organizations: Results of a delphi study," *Business & Information Systems Engineering*, vol. 63, pp. 511–527, 2021.
- [162] T. Chiu and M. Jans, "Process mining of event logs: A case study evaluating internal control effectiveness," *Accounting Horizons*, vol. 33, no. 3, pp. 141–156, 2019.
- [163] C. Di Ciccio, A. Marrella, and A. Russo, "Knowledge-intensive processes: Characteristics, requirements and analysis of contemporary approaches," *Journal on Data Semantics*, vol. 4, pp. 29–57, 2015.
- [164] B. Löhr, K. Brenning, C. Bartelheimer, D. Beverungen, and O. Müller, "Process Mining of Knowledge-Intensive Processes: An Action Design Research Study in Manufacturing," in *International Conference on Business Process Management*, Springer, 2022, pp. 251–267.
- [165] E. H. Nooijen, B. F. van Dongen, and D. Fahland, "Automatic discovery of data-centric and artifact-centric processes," in *Business Process Management Workshops: BPM 2012 International Workshops, Tallinn, Estonia, September 3, 2012. Revised Papers 10*, Springer, 2013, pp. 316–327.
- [166] R. J. C. Bose, R. S. Mans, and W. M. P. Van Der Aalst, "Wanna improve process mining results?" In *2013 IEEE symposium on computational intelligence and data mining (CIDM)*, IEEE, 2013, pp. 127–134.
- [167] M. De Leoni, W. M. P. Van der Aalst, and M. Dees, "A general framework for correlating business process characteristics," in *International Conference on Business Process Management*, Springer, 2014, pp. 250–266.
- [168] W. M. P. van der Aalst, "On the pareto principle in process mining, task mining, and robotic process automation," in *DATA*, 2020, pp. 5–12.
- [169] B. Hompes, J. Buijs, W. Van der Aalst, P. Dixit, and J. Buurman, "Discovering deviating cases and process variants using trace clustering," in *27th Benelux Conference on Artificial Intelligence (BNAIC 2015)*, 2015.
- [170] J. De Weerd, M. De Backer, J. Vanthienen, and B. Baesens, "A multi-dimensional quality assessment of state-of-the-art process discovery algorithms using real-life event logs," *Information Systems*, vol. 37, no. 7, pp. 654–676, 2012.
- [171] A. Rozinat and W. M. P. Van der Aalst, "Conformance checking of processes based on monitoring real behavior," *Information Systems*, vol. 33, no. 1, pp. 64–95, 2008.
- [172] D. M. V. Sato, S. C. De Freitas, J. P. Barddal, and E. E. Scalabrin, "A survey on concept drift in process mining," *ACM Computing Surveys (CSUR)*, vol. 54, no. 9, pp. 1–38, 2021.
- [173] S. Pravalovic, A. Appice, and D. Malerba, "Process mining to forecast the future of running cases," in *New Frontiers in Mining Complex Patterns: Second International Workshop, NFMCP 2013, Held in Conjunction with ECML-PKDD 2013, Prague, Czech Republic, September 27, 2013, Revised Selected Papers 2*, Springer, 2014, pp. 67–81.

- [174] I. Teinemaa, M. Dumas, M. L. Rosa, and F. M. Maggi, "Outcome-oriented predictive process monitoring: Review and benchmark," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 13, no. 2, pp. 1–57, 2019.
- [175] M. Bozkaya, J. Gabriels, and J. M. van der Werf, "Process diagnostics: A method based on process mining," in *2009 International Conference on Information, Process, and Knowledge Management*, IEEE, 2009, pp. 22–27.
- [176] P. Harmon, *Business process change: a business process management guide for managers and process professionals*. Morgan Kaufmann, 2019.
- [177] M. Dumas *et al.*, "Process redesign," *Fundamentals of Business Process Management*, pp. 297–339, 2018.
- [178] S. Son *et al.*, "Process mining for manufacturing process analysis: A case study," in *Proceeding of 2nd Asia Pacific Conference on Business Process Management, Brisbane, Australia*, 2014.
- [179] R. Lorenz, J. Senoner, W. Sihn, and T. Netland, "Using process mining to improve productivity in make-to-stock manufacturing," *International Journal of Production Research*, vol. 59, no. 16, pp. 4869–4880, 2021.
- [180] S. Rinderle-Ma and J. Mangler, "Process automation and process mining in manufacturing," in *International conference on business process management*, Springer, 2021, pp. 3–14.
- [181] J. Cardoso, "Business process control-flow complexity: Metric, evaluation, and validation," *International Journal of Web Services Research (IJWSR)*, vol. 5, no. 2, pp. 49–76, 2008.
- [182] T. Derkink, "An insight in the planning of the maintenance process of wheel axles at nederlandse spoorwegen," M.S. thesis, University of Twente, 2021.
- [183] T. Kluyver *et al.*, "Jupyter notebooks—a publishing format for reproducible computational workflows," in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, IOS Press, 2016, pp. 87–90.
- [184] W. McKinney *et al.*, "Data Structures for Statistical Computing in Python," in *Proceedings of the 9th Python in Science Conference*, Austin, TX, vol. 445, 2010, pp. 51–56.
- [185] H. Verbeek, J. C. Buijs, B. F. Van Dongen, and W. M. P. Van Der Aalst, "XES, XESame, and ProM 6," in *Information Systems Evolution: CAiSE Forum 2010, Hammamet, Tunisia, June 7-9, 2010, Selected Extended Papers 22*, Springer, 2011, pp. 60–75.
- [186] A. Berti, S. van Zelst, and D. Schuster, "PM4Py: A process mining library for Python," *Software Impacts*, vol. 17, p. 100556, 2023, issn: 2665-9638. doi: <https://doi.org/10.1016/j.simpa.2023.100556>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2665963823000933>.
- [187] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [188] S. J. Leemans, D. Fahland, and W. M. P. van der Aalst, "Using life cycle information in process discovery," in *Business Process Management Workshops: BPM 2015, 13th International Workshops, Innsbruck, Austria, August 31–September 3, 2015, Revised Papers 13*, Springer, 2016, pp. 204–217.
- [189] W. M. P. Van Der Aalst, "Process mining: Discovering and improving spaghetti and lasagna processes," in *2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, IEEE, 2011, pp. 1–7.
- [190] C. W. Gunther and W. M. P. van der Aalst, "Fuzzy mining—adaptive process simplification based on multi-perspective metrics," in *Business Process Management: 5th International Conference, BPM 2007, Brisbane, Australia, September 24-28, 2007, Proceedings*, Springer, vol. 4714, 2007, p. 328.
- [191] A. Berti and W. M. P. van der Aalst, "A novel token-based replay technique to speed up conformance checking and process enhancement," in *Transactions on Petri Nets and Other Models of Concurrency XV*, Springer, 2021, pp. 1–26.
- [192] S. J. Leemans, D. Fahland, and W. M. P. Van Der Aalst, "Discovering block-structured process models from event logs containing infrequent behaviour," in *Business Process Management Workshops: BPM 2013 International Workshops, Beijing, China, August 26, 2013, Revised Papers 11*, Springer, 2014, pp. 66–78.

- [193] S. J. Leemans, "Robust Process Mining with Guarantees," in *BPM (Dissertation/Demos/Industry)*, Springer, 2018, pp. 46–50.
- [194] M. Saunders, P. Lewis, and A. Thornhill, *Research methods for business students*. Pearson Education, 2009.
- [195] J. J. Koorn *et al.*, "Bringing rigor to the qualitative evaluation of process mining findings: An analysis and a proposal," in *2021 3rd International Conference on Process Mining (ICPM)*, IEEE, 2021, pp. 120–127.
- [196] N. R. Haddaway, M. J. Page, C. C. Pritchard, and L. A. McGuinness, "PRISMA2020: An R package and Shiny app for producing PRISMA 2020-compliant flow diagrams, with interactivity for optimised digital transparency and Open Synthesis," *Campbell Systematic Reviews*, vol. 18, no. 2, e1230, 2022.
- [197] F. Mannhardt, M. De Leoni, H. A. Reijers, and W. M. P. Van Der Aalst, "Decision mining revisited-discovering overlapping rules," in *International Conference on Advanced Information Systems Engineering*, Springer, 2016, pp. 377–392.
- [198] E. Bazhenova and M. Weske, "Deriving decision models from process models by enhanced decision mining," in *Business Process Management Workshops: BPM 2015, 13th International Workshops, Innsbruck, Austria, August 31–September 3, 2015, Revised Papers 14*, Springer, 2016, pp. 444–457.
- [199] W. Es-Soufi, E. Yahia, and L. Roucoules, "A process mining based approach to support decision making," in *Product Lifecycle Management and the Industry of the Future: 14th IFIP WG 5.1 International Conference, PLM 2017, Seville, Spain, July 10-12, 2017, Revised Selected Papers 14*, Springer, 2017, pp. 264–274.
- [200] D. Breuker, M. Matzner, P. Delfmann, and J. Becker, "Comprehensible predictive models for business processes," *MIS Quarterly*, vol. 40, no. 4, pp. 1009–1034, 2016.
- [201] N. Mehdiyev and P. Fettke, "Explainable artificial intelligence for process mining: A general overview and application of a novel local explanation approach for predictive process monitoring," *Interpretable artificial intelligence: A perspective of granular computing*, pp. 1–28, 2021.
- [202] A. E. Márquez-Chamorro, M. Resinas, A. Ruiz-Cortés, and M. Toro, "Run-time prediction of business process indicators using evolutionary decision rules," *Expert Systems with Applications*, vol. 87, pp. 1–14, 2017.
- [203] R. Petrusel, I. Vanderfeesten, C. C. Dolean, and D. Mican, "Making decision process knowledge explicit using the decision data model," in *Business Information Systems: 14th International Conference, BIS 2011, Poznań, Poland, June 15-17, 2011. Proceedings 14*, Springer, 2011, pp. 172–184.
- [204] W. Es-Soufi, E. Yahia, and L. Roucoules, "On the use of Process Mining and Machine Learning to support decision making in systems design," in *Product Lifecycle Management for Digital Transformation of Industries: 13th IFIP WG 5.1 International Conference, PLM 2016, Columbia, SC, USA, July 11-13, 2016, Revised Selected Papers 13*, Springer, 2016, pp. 56–66.
- [205] W. Jareevongpiboon and P. Janecek, "Enhancing decision patterns discovered by process mining with semantic related data," in *ICEB 2011 Proceedings (Bangkok)*, 2011, p. 5.
- [206] E. Brzychczy and A. Trzcionkowska, "Creation of an event log from a low-level machinery monitoring system for process mining purposes," in *Intelligent Data Engineering and Automated Learning—IDEAL 2018: 19th International Conference, Madrid, Spain, November 21–23, 2018, Proceedings, Part II 19*, Springer, 2018, pp. 54–63.



SLR PRISMA 2020 report

The process of applying the inclusion / exclusion criteria and the resulting properties of the final set of articles are shown in Figure A.1.

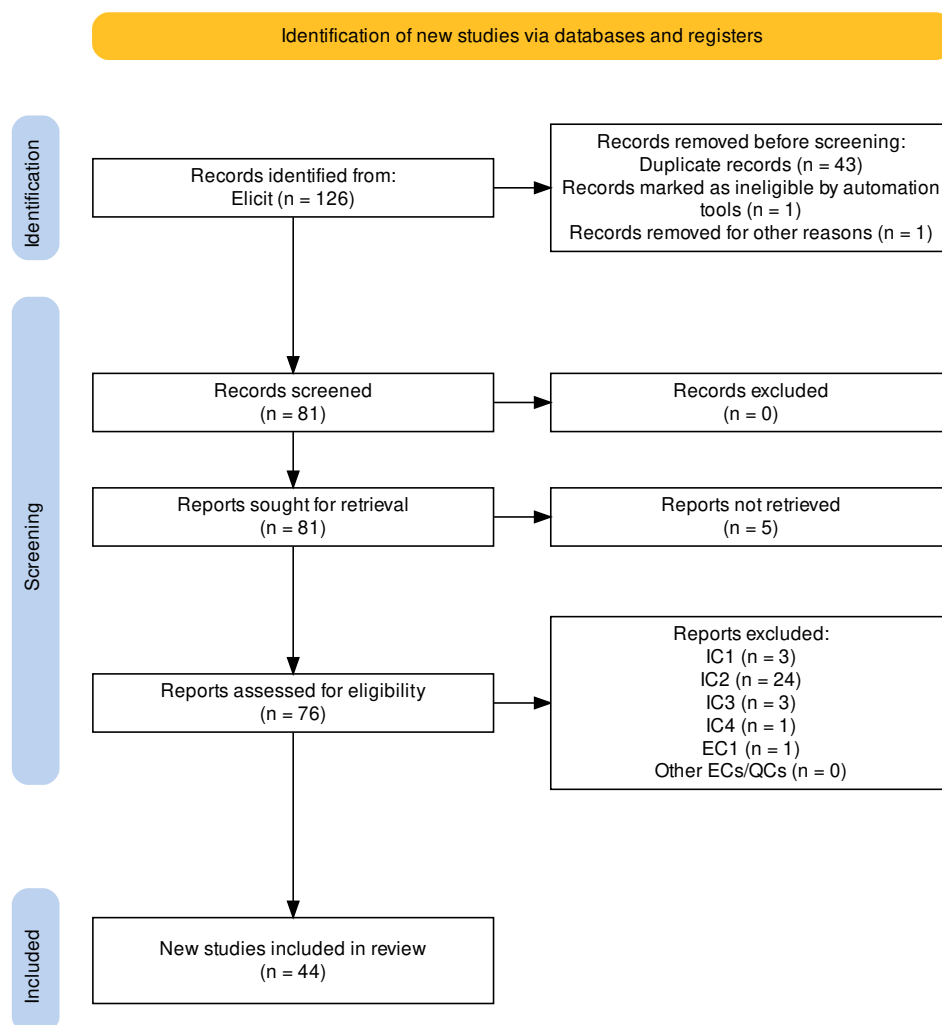


Figure A.1: Overview of the search-screen phases and the resulting amount of papers, template by Haddaway *et al.* [196].

B

Overview of decision mining approaches

Table B.1: The typology of the different approaches for decision mining.

Paper	Process domain	Analysis approach	Data type	Data structuredness	Data-case relation	Learning algorithm
[25]	Logistics and medical	Process discovery for DPA	Time series continuous	Vector of recurring measurement values	Per-case set	Decision trees
[128]	Loan applications	Separate discovery of data-flow	Numerical / Date	Single attribute per event	Attributes	Decision trees (C4.5)
[130]	ICT infrastructure ticketing	Complex business rule discovery	Text	Unstructured text message exchanges	Per-case set	Decision trees, Text mining
[83]	Patient trajectories, Hospital billing	Overlapping and infrequent paths	Numerical, Boolean	Multiple attributes per event [197]	Attributes	Decision trees (C4.5)
[126]	Loan application, Purchase order	Computing offline situation tables	Text, Numerical	Multiple attributes at case level	Attributes	Decision trees
[198]	Credit assessment	Enrich model with KPI-based decision logic	Numerical, Boolean	Multiple attributes at case level	Attributes	Decision trees
[199]	Industrial network supervision	Data-flow discovery for guard conditions	Numerical, Boolean	Decision rules related to enriched event log	Attributes	N/A
[200]	Several	Predict running case outcomes from historical data	Categorical, Dirichlet	Categorical representations of predictive Petri nets	Attributes	Expectation-Maximization, Probabilistic models
[201]	IT incident management	Model explanation for process behavior	Numerical, Categorical	Derived process metrics converted into decision rules	Metrics	Deep learning, Decision trees
[21]	Loan application procurement	Predictive models relate variables to activities	Numerical, Categorical, Boolean	Multiple attributes at case level	Variable-Activity Pairs	Decision trees

Continued on next page

Table B.1: The typology of the different approaches for decision mining. (Continued)

[22]	Insurance liability claim	Convert decision points into classification problems	Numerical, Categorical	Multiple attributes at case level	Attributes	Decision trees (J48, C4.5)
[20]	IT incident management	Uncover long-distance dependencies	Numerical, Categorical	Multiple attributes at activity level	Attributes	Decision trees
[202]	IT incident management	Predict process indicators from decision rules	Numerical, Categorical, DateTime	Multiple attributes at case level	Event classes	Evolutionary
[203]	Loan contracting	Explicitly capture decision logic from users	Numerical, Categorical	Multiple attributes at activity level	Attributes	Self-developed
[61]	Industrial manufacturing	Derive decision rules from accumulated events	Categorical	Multiple attributes at process state level	State Properties	Decision trees
[204]	Product design	Identify common patterns from attribute evolution	Categorical	Multiple attributes at activity level	Attributes	Decision trees (J48, C4.5)
[205]	Logistics	Improve complex rule discovery with semantic data	Numerical, Boolean, Text	Multiple attributes at activity level	Attributes	Decision trees (C4.5)
[132]	Housing acquisition	Explicitly capture decision logic from users	Numerical	Multiple attributes at activity level	Attributes	Decision trees
[133]	Loan application and approval	Automatic decision activity service discovery	Categorical	Single property at event level	Properties	Self-developed
[137]	Health care	Goal-oriented process-state based decision model	Boolean	Derived goal metrics that represent decision rules	Attributes	Self-developed
[135]	Clinical logistics	Context-based cost perspective and data analysis	Numerical	Multiple attributes at activity level	Attributes	Decision trees (J48)
[138]	Robotized car manufacturing	Abstracting to operation and resource-product flow	Numerical, Boolean, Text, DateTime	Multiple key-value pairs, multiple events per operation	Aggregates	Self-developed
[136]	Industrial roof bolter machine	PLC and sensor data abstracted into an event log	Numerical, Boolean, Timestamp	Derived events from sensor ranges and thresholds [206]	Attributes	Self-developed
[139]	Health care treatment procedure	Ontology-based process model customization	Numerical, Boolean	Multiple attributes at activity level	Attributes	Decision trees (J48)
[144]	Medical patient trajectories	Process variant recommendation by trace clustering	Numerical, Boolean, Nominal, Text	Multiple attributes at case level	Attributes	Logistic regression

Continued on next page

Table B.1: The typology of the different approaches for decision mining. (Continued)

[146]	Condition-based maintenance	Predict future health using sensor monitoring data	Numerical, Boolean	Derived events from sensor ranges and thresholds	Attributes	Decision trees (Fault trees), Association rules
-------	-----------------------------	--	--------------------	--	------------	---



PDDs, Activity and Concept tables

This appendix presents the PDD of the PM^2 methodology and the related activity and concept tables. In addition, the PDD of the extended $PM^2 \times DM$ methodology is presented. The additions and modifications in the latter model are indicated with a gray-colored overlay.

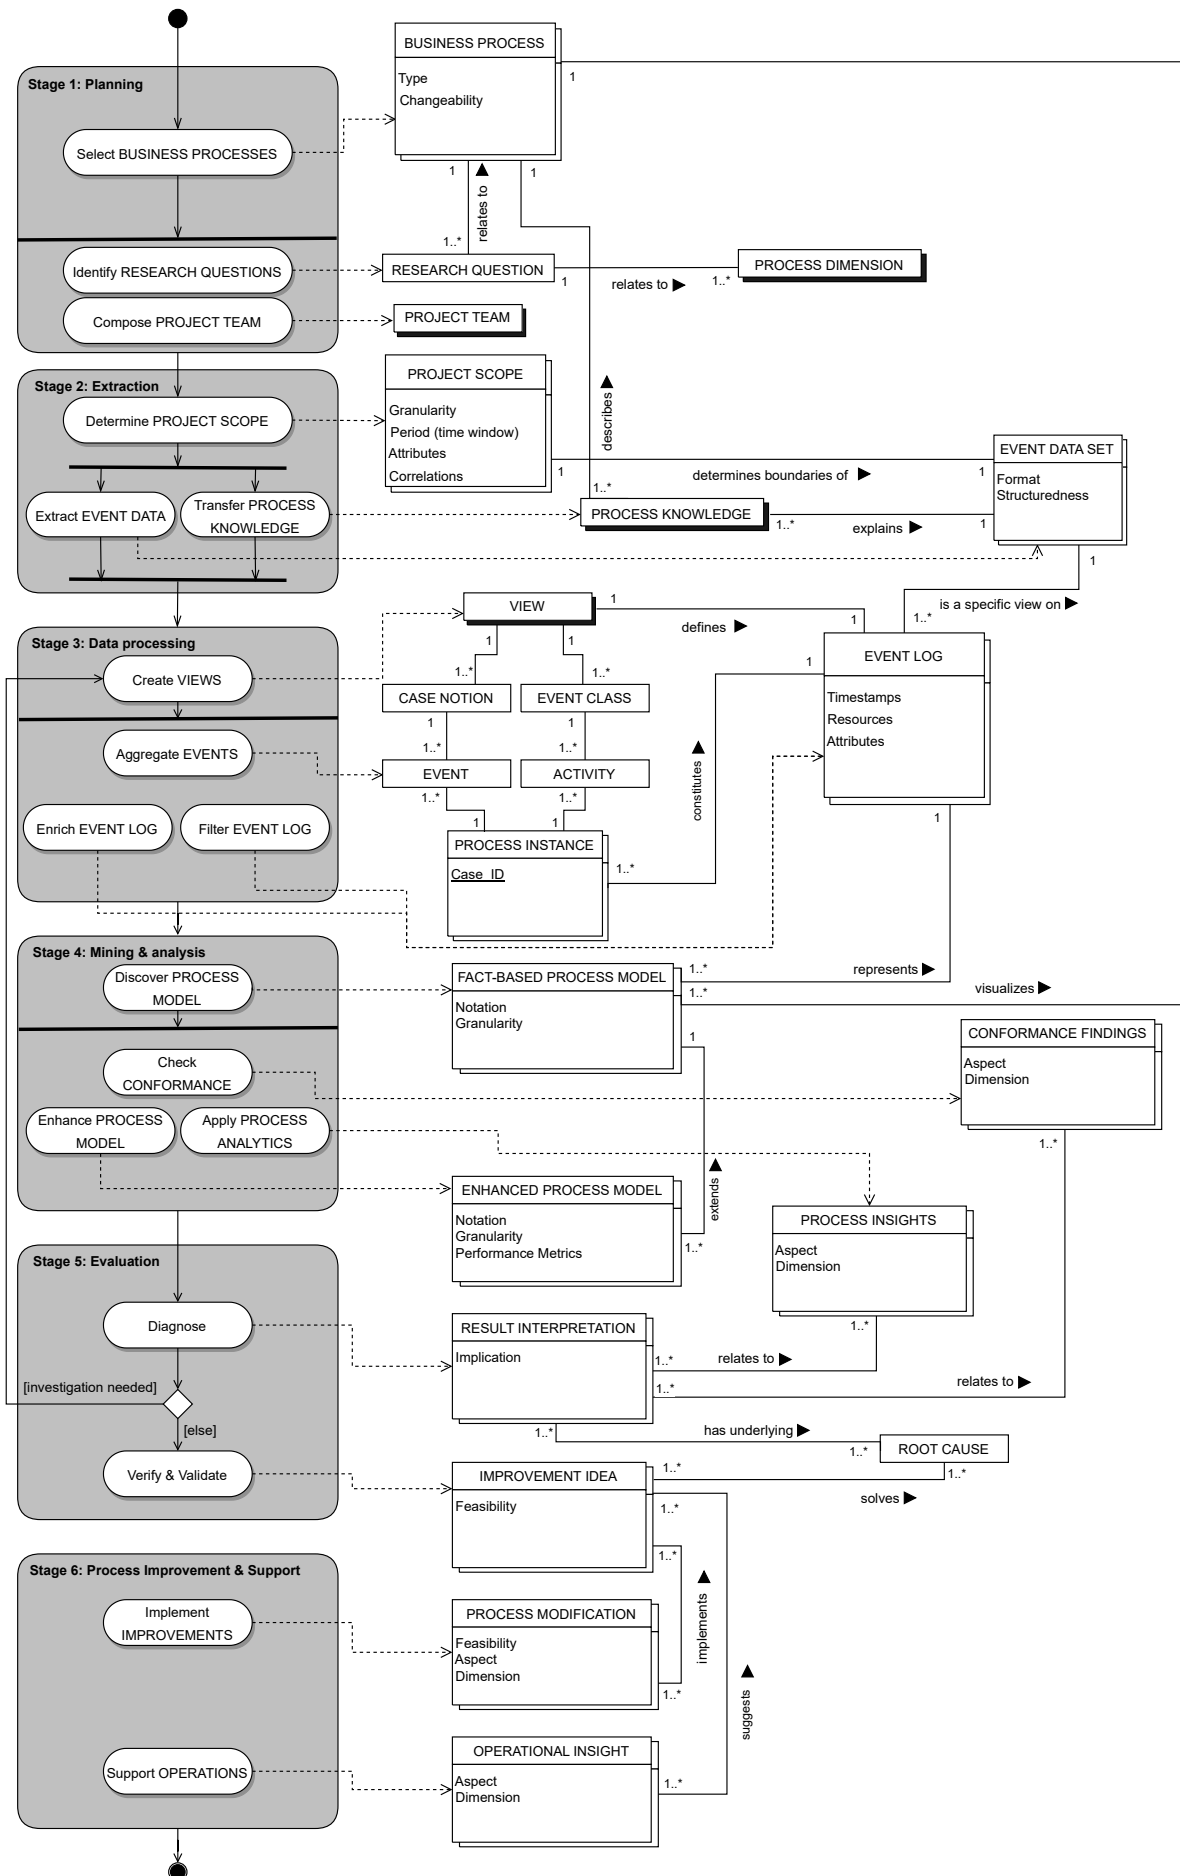


Figure C.1: The PDD model of the PM² methodology by Van Eck *et al.* [99].

Table C.1: Activity table for the PDD of PM², descriptions based on [59, 99].

Activity	Sub-activity	Description
Planning	Select BUSINESS PROCESSES	Identify and select the business processes for analysis.
	Identify RESEARCH QUESTIONS	Formulate specific questions or objectives for the process mining project.
	Compose PROJECT TEAM	Assemble a team with the necessary skills and expertise for the project.
Extraction	Determine PROJECT SCOPE	Define the boundaries and extent of the process to be analyzed.
	Extract EVENT DATA	Gather data regarding events related to the business process from various sources.
	Transfer PROCESS KNOWLEDGE	Share and utilize existing knowledge about the process within the project team.
Data processing	Create VIEWS	Develop different perspectives or views of the event data for analysis.
	Aggregate EVENTS	Combine event data from different sources or timeframes.
	Enrich EVENT LOG	Add additional information to the event log to enhance analysis.
	Filter EVENT LOG	Remove irrelevant or redundant data from the event log or focus the analysis.
Mining and analysis	Discover PROCESS MODEL	Use process mining techniques to uncover the actual process model.
	Check CONFORMANCE	Compare the discovered process model against the predefined or expected model.
	Enhance PROCESS MODEL	Refine the process model based on insights gained from analysis.
	Apply PROCESS ANALYTICS	Perform advanced analytics for deeper insights into the process.
Evaluation	Diagnose	Assess the outcomes of the process mining and analysis to identify issues or areas of improvement.
	Verify and validate	Confirm the accuracy and relevance of the process mining results.
Process improvement and support	Implement IMPROVEMENTS	Apply the insights gained to make improvements to the business process.
	Support OPERATIONS	Provide ongoing support and refinement to ensure the process continues to operate effectively.

Table C.2: Concept table for the PDD of PM², descriptions based on [59, 99].

Concept	Description
BUSINESS PROCESS	A set of activities designed to achieve a specific organizational goal.
RESEARCH QUESTION	A question that guides the focus and objectives of the process mining project.
PROJECT TEAM	Group of individuals with varied expertise collaborating on the process mining project.
PROJECT SCOPE	The extent and boundaries of the process mining project.

Continued on next page

Table C.2: Concept table for the PDD of PM², descriptions based on [59, 99]. (Continued)

PROCESS DIMENSION		Aspects of a process under analysis, such as time, cost, or quality.
EVENT DATA		Data representing occurrences in a business process, typically logged by IT systems.
PROCESS KNOWLEDGE		Understanding of the business process, often held by stakeholders or documented.
VIEW		A specific perspective or aspect of the process being analyzed.
CASE NOTION		The concept identifying unique instances in a process, such as a customer order.
EVENT CLASS		A category of events in the process, often representing a type of activity.
EVENT		A single occurrence within a process, typically recorded in an event log.
ACTIVITY		A distinct action or step within a business process.
EVENT LOG		A record of events occurring within a process, used as data for process mining.
PROCESS INSTANCE		A specific occurrence of a business process, represented by a sequence of events.
PROCESS MODEL		A formal representation of a business process, often as a flowchart or diagram.
CONFORMANCE		The degree to which a process instance adheres to the defined process model.
PROCESS ANALYTICS		Analysis of process data to gain insights into process performance and compliance.
FACT-BASED MODEL	PROCESS MODEL	A process model built directly from and validated by empirical data.
ENHANCED MODEL	PROCESS MODEL	A process model improved with additional data, annotations, or insights.
PROCESS INSIGHTS		Understandings and findings derived from process mining analysis.
CONFORMANCE FINDINGS		Results from comparing process instances with the model to identify deviations.
RESULT INTERPRETATION		The analysis and understanding of the outcomes from process mining activities.
IMPROVEMENT IDEA		Suggestions for changes or enhancements to the business process.
ROOT CAUSE		The underlying reason for a problem or deviation in the process.
PROCESS MODIFICATION		Changes made to the existing business process based on insights.
OPERATIONAL INSIGHT		Practical understanding of the day-to-day functioning and performance of a business process.

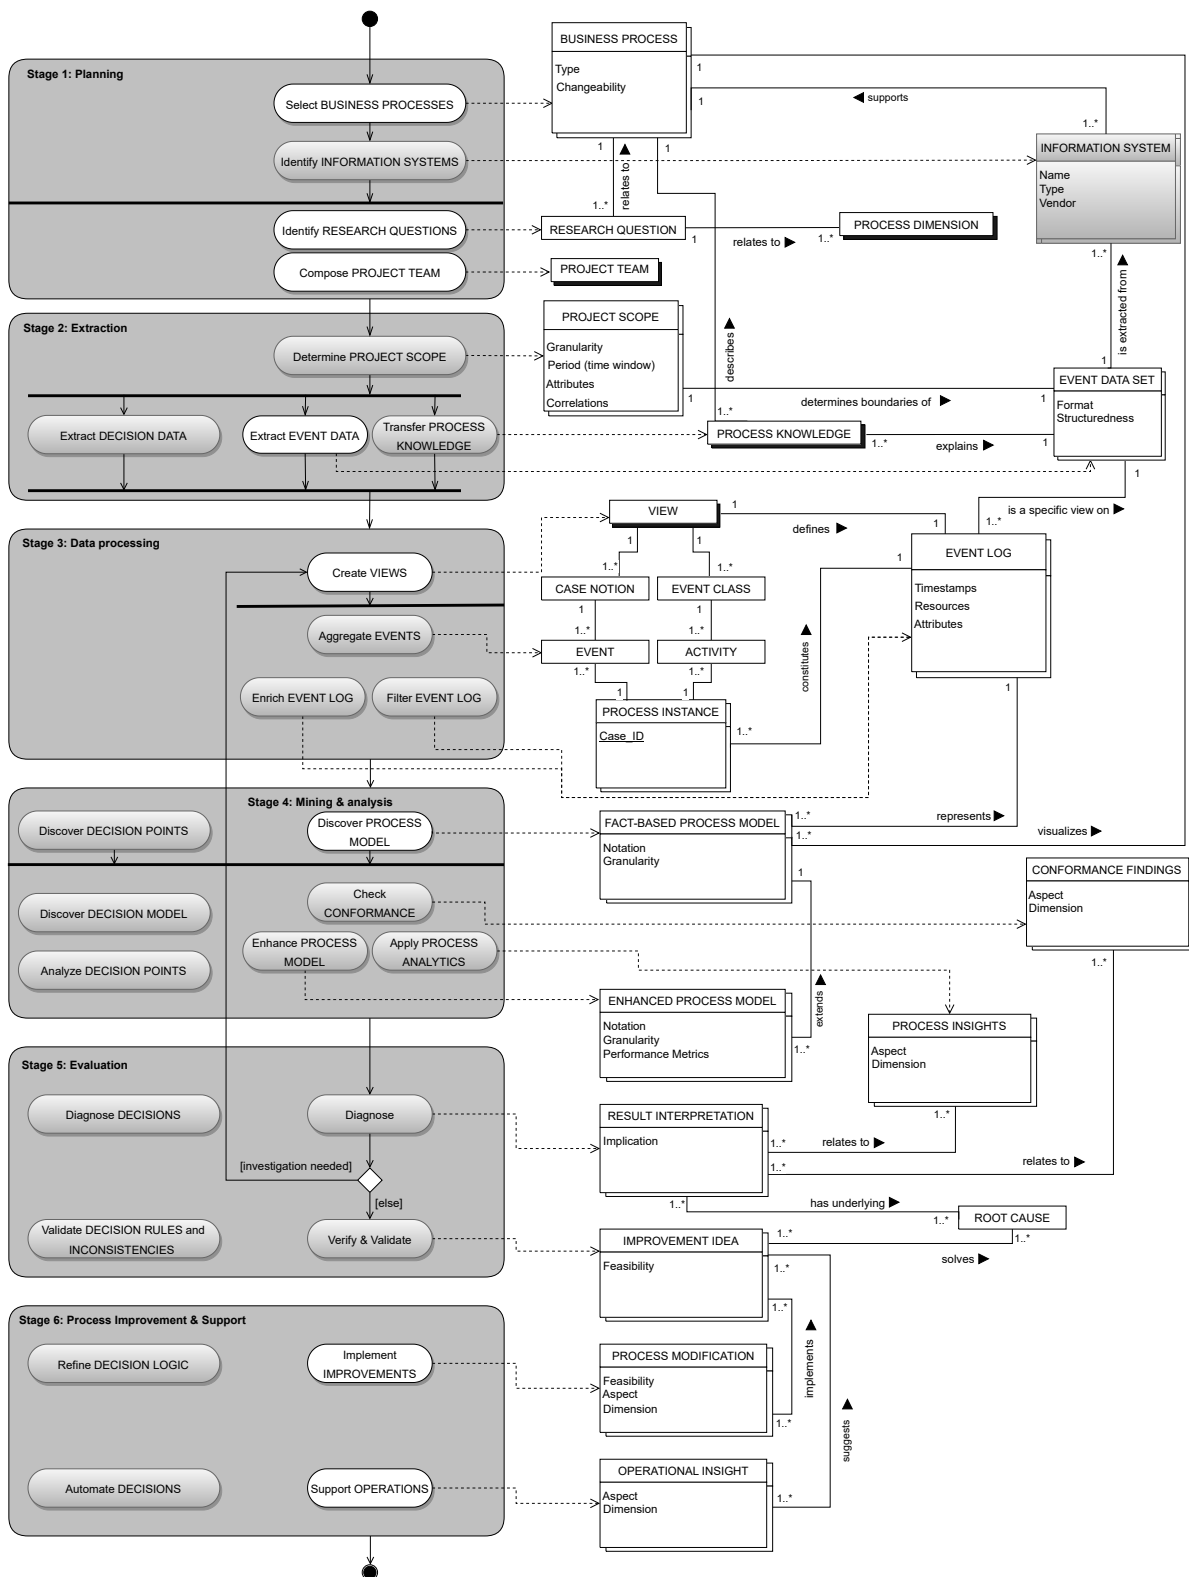



Figure C.2: The PDD model of the extended PM²xDM methodology.



Case study protocol

D.1. Preamble

- A mutual confidentiality agreement was established between the case study organization and the researcher and was signed by both.
- Data storage of collected documents, data sets, and other materials will be exclusively on the secured cloud environment of the organization, or on the laptop of the researcher that is equipped with strong password protection and full disk encryption.
- The publication of this thesis document, as well as any derivatives, should be reviewed and approved by the HR department before being made public.
- A separate version of this thesis suitable for publication will be available upon request. Any sensitive information from the case study context or about the organization will be removed.
- Documentation collected and created for the purpose of this research is saved in the case study database. This may include but is not limited to slide decks, descriptive documents, audio recordings, and interview transcripts.
- This protocol further specifies the general procedures, terms, and conditions under which the case study has been executed.

D.2. General

- The research project involves the development of a methodological framework for the integration of process and decision mining. Therefore, a process mining project is implemented using an existing methodology in the real-world context of a wheelset revision process at the Dutch railway company NS. The proposed methodology is implemented concurrently and quantitatively evaluated with a case study. Qualitative validation and evaluation are conducted using a focus group.
- The case study research method entails a single-case embedded design.

D.3. Procedures

- The organization was selected for convenience reasons as part of an internship project.
- Several internal cases have been evaluated as candidates for the case study. However, the resulting case of the wheelset revision process was left over due to sufficient availability of the process context, data, and adequate involvement of stakeholders.
- Initial contact with stakeholders of this respective case was established on June 6th, 2023. Over the course of six months, documents were collected and several collaborative sessions with different stakeholders were organized. The complete report is embedded in the results of the case study that are disseminated in Chapter 6.

D.4. Research instruments

- The quantitative research instrument is the case study of the application of the original methodology and its proposed extensions within the practical process context.
- The qualitative research instrument is the evaluation of artifacts and insights in a focus group with a predefined protocol of open-ended questions.

D.5. Data analysis guidelines

- Data analysis is performed by synthesizing the results of the implementation of the case study and the resulting insights from the focus group evaluation.
 - The convergence of data from multiple sources is obtained by thematic analysis of both quantitative and qualitative data.
 - Triangulation of perspectives is achieved by collecting the different inputs of the focus group participants.
- The within-case analysis employs three types of data, namely:
 - Descriptive data: documentation and the results of the case study implementation
 - Explanatory data: the results of the intermediate cooperative sessions with the SMEs that are used within the case study, as well as the results of the focus group.
 - Individual case report: this is disseminated in Chapter 6.
- The data schema distinguishes primary and secondary data, including but not limited to the following:
 - Primary data: raw process data, case study implementation results, process models and visualizations, focus group recordings and transcripts. These are used to empirically validate the research project itself.
 - Secondary data: descriptive documents, earlier process mining project results, work instructions. These are used to describe the case study context at hand, the specific process and other relevant characteristics.
- The SLR uses tabular representations, whereas the case study uses descriptive tables and process models. The methodological frameworks are disseminated using proprietary visualizations and models in PDD notation.

D.6. Appendix


The participation request letter was distributed in the form of an integrated information letter and consent form, which can be found in Appendix E.



Focus group and interview consent form

The consent form was distributed digitally using a Google Forms web page. It was designed so that only the name of the respective participant, their agreement, and the date of signing were recorded. This is in line with the ethics guidelines of Utrecht University, which encourage one not to register additional personal information, such as signatures. The content of the actual consent form is presented on the next page, which was designed according to the consent form template provided by Utrecht University¹. Please note that the content of the consent form is in Dutch for convenience reasons.

¹<https://www.uu.nl/en/research/institute-of-information-and-computing-sciences/ethics-and-privacy>


Universiteit Utrecht

Sharing science,
shaping tomorrow

Informatiebrieven toestemmingsformulier deelname wetenschappelijk onderzoeksproject

Lees de onderstaande informatie en stellingen aandachtig door. Geef het aan het einde aan als alles duidelijk voor u is en u daarmee toestemming verleent voor deelname aan dit onderzoek.

[Log in bij Google](#) om je voortgang op te slaan. [Meer informatie](#)

*** Verplichte veld**

Onderzoekstitel: A case study on an application of decision mining for process model enhancement and conformance checking in the context of an industrial production process

Het onderzoek met bovenstaande titel behelst het evalueren van een methodologisch raamwerk om **decision mining** in een industriële procescontext toe te passen. Het onderzoek heeft als doel te onderzoeken of deze techniek bij kan dragen aan het verbeteren van procesmodellen en het controleren of het proces zich conformeert aan de vereisten. De evaluatie vindt plaats in de vorm van een **case study** waarbij u als deelnemer input kan geven in een **focus groep** en/of **interview**. Dit formulier legt het onderzoek uit, wat deeltname precies inhoudt en wat uw rechten hierbij zijn.

Onderzoekbeschrijving

De focus groep of het interview waaraan u deelneemt bestaat primair uit zeven delen:

- Introductie en kenmerken deelnemers
- Korte uitleg van basisbegrippen en concepten: process mining en decision mining
- Analyse van de kenmerken van het proces in de context van procesanalyse en verbetering
- Uitleg over en evaluatie van verschillende soorten procesmodellen
- Analyse en evaluatie van modellen van het wiskundeviesproces zonder decision mining
- Analyse en evaluatie van modellen van het wiskundeviesproces met decision mining
- Evaluatie van diverse inzichten die met de analyse van het proces zijn verbonden

De onderzoeker zal hiervoor diverse antwoorden en inzichten aan u verzamelen en hier op een semi-gestructureerde manier vragen over stellen. Het is de bedoeling dat de sessie een zeer interactief karakter heeft en het is daarom van belang dat alle perspectieven aan bod kunnen komen. Er zijn dan ook geen goede of foute antwoorden en meningsverschillen of discussies worden gewaardeerd omdat deze illustratief van karakter kunnen zijn of voor het onderzoek relevante inzichten kunnen opleveren.

De opnames van de focus groep kunnen worden gezien als persoonlijke informatie, echter is het verzamelen persoonlijke informatie niet nodig voor het doel van het onderzoek. Alleen de audio wordt gebruikt voor transcriptie en de twee opnames inclusief video zullen uiterlijk binnen 90 dagen na transcriptie worden vernietigd. Eventuele informatie waarmee u persoonlijk kunt worden herleid is alleen toegankelijk voor de onderzoeker en supervisor.

Na afloop van de transcriptie zal deze met alle deelnemers worden gedeeld en is er de mogelijkheid om informatie te verwijderen die niet zal mee worden genomen in de vervolganalyse. Een samenvatting van de verkregen inzichten uit de focus groep of interview zal vervolgens weer met u worden gedeeld, welke u eventueel nog een laatste maal van commentaar kunt voorzien voordat deze definitief wordt verwerkt in de resultaten.

Dit onderzoek is onderdeel van het scriptieproject van Bart Hoonstra voor het afstuderen van de Master Business Informatics aan de Universiteit Utrecht en staat onder supervisie van dr. Inge van de Weerd. Voor het onderzoek is een ethische toets (Ethics Quick Scan*) uitgevoerd conform de reguleringen van de Universiteit Utrecht. Het onderzoek werd goedgekeurd als laag risico, waardoor geen aanvullende ethische toetsing of privacy beoordeling is uitgevoerd. Indien u verder nog vragen of opmerkingen over het onderzoek heeft, dan kunt u contact opnemen met Bart Hoonstra op b.a.hoonstra@studie.uu.nl of Inge van de Weerd op i.van.de.weerd@uu.nl.

Voor uw deelname aan dit onderzoek wordt geen vergoeding beschikbaar gesteld. Er is geen indicatie voor belangversterving van toepassing op de uitvoering van dit onderzoek.

Toestemmingsverklaring

- Ik bevestig dat ik 18 jaar of ouder ben.
- Ik bevestig dat het onderzoeksproject aan mij is uitgelegd. Ik heb de gelegenheid gehad om vragen te stellen over het project en deze zijn naar tevredenheid beantwoord. Ik had voldoende tijd om na te denken of ik mee zou doen.
- Ik ga ermee akkoord dat het materiaal dat ik bijdraag, wordt gebruikt om inzichten te genereren voor het onderzoeksproject. Dit is inclusief eventueel (betreft) gebruik van illustratieve citaten.
- Ik geef toestemming voor het gebruik van audio/video-opnamen in dit onderzoek, zoals uitgelegd in het informatieblad. Ik begrijp dat ik op elk moment kan verzoeken om de opnames stop te zetten.
- Ik begrijp dat als ik toestemming geef, de audio/video-opnamen vertrouwelijk worden bewaard, zodat alleen de hoofdonderzoeker en primaire supervisor zoals hierboven genoemd toegang hebben tot de opname. De opnamen zelf worden bewaard in de beveiligde Microsoft-omgeving van NS gedurende maximaal 90 dagen, waarna ze alleen nog in anoniem getranscribeerde/gecodeerde vorm blijven bestaan en het origineel veilig wordt vernietigd. Conform de Algemene Verordening Gegevensbescherming (AVG) kan ik op elk moment gedurende deze periode toegang krijgen tot mijn opnames en kan ik verzoeken deze te verwijderen.
- Ik begrijp dat mijn deelname aan dit onderzoek vrijwillig is en dat ik me op elk moment kan terugtrekken uit het onderzoek zonder opgaaf van reden, en dat als ik me terugtrek, alle gegevens die al van mij zijn verzameld, zullen worden gewist.
- Ik geef toestemming om de volledig geanonimiseerde gegevens te gebruiken in toekomstige publicaties en andere wetenschappelijke middelen om de bevindingen van het onderzoeksproject te verspreiden.
- Ik begrijp dat de verkregen gegevens veilig zullen worden opgeslagen door onderzoekers, maar dat op passende wijze geanonimiseerde gegevens in de toekomst beschikbaar kunnen worden gesteld aan andere voor onderzoeksdoeleinden. Ik begrijp dat de universiteit op passende wijze geanonimiseerde gegevens mag publiceren in geschikte gegevensopslagplaatsen voor wetenschappelijke doeleinden en om deze toegankelijk te maken voor onderzoekers en andere onderzoeksgebruikers.

Voor- en achternaam *

Ik bevestig dat ik de bovenstaande verklaringen heb gelezen en begrepen, en ga akkoord met deelname aan het onderzoek (vink het vakje aan).

Ja, ik ga akkoord

Datum *

Selecteer hier de datum van ondertekening, doorgaans is dit vandaag.

Datum

dd/mm/yyyy

Verzenden
Pagina 1 van 1
Formulier wissen

Verzend nooit wachtwoorden via Google Formulieren.
[Wetelijk aansprakelijk](#)
 Google Formulieren

Figure E.1: The consent form for the focus group or interview that was signed beforehand by the participants.



Focus group and interview protocol

Actions are highlighted in **bold** and questions are presented in *italics* with the related evaluation criteria abbreviated between parentheses, if applicable. A legend is provided at the end of this protocol in Table F.1. The duration of the focus group is 1.5-2 hours in total. The protocol is defined in English, but the focus group is conducted in Dutch for convenience reasons. The results are partially transcribed in Dutch and literally translated into English for analysis and reporting purposes in the respective results sections of this document.

F.1. Outline

1. Welcome and introduction
2. Explanation of process mining and decision mining
3. Process characteristics
4. Process modeling techniques
5. Evaluation of baseline process models
6. Evaluation of enhanced process models
7. Evaluation of derived insights and process improvement potential
8. Final remarks and closing

F.2. Welcome

Min duration: 2 minutes

Max duration: 5 minutes

- Test of recording setup: **check legibility of all participants**
- Short welcome word, introduction of the researcher
- Thank the participants for their participation
- Explanation of the procedure and focus group etiquette
- Check if all consent forms have been signed and verify that there are no remaining questions
- Notify everyone that the recording will be started next

F.3. Introduction

Min duration: 5 minutes

Max duration: 10 minutes

- **Start the recording**
- Explain the purpose of this research
- Verify background information of participants
 - Project team roles
 - Organizational roles
 - Tenure (process-related)
 - Professional experience
 - Process mining experience
- Short description of the six remaining parts of the session

F.4. Explanation of process mining and decision mining

Min duration: 5 minutes

Max duration: 10 minutes

- Shortly outline the concept of process mining
- Explain the position of decision mining in relation to process mining
- Verify if there are no remaining questions regarding these concepts

F.5. Process characteristics

Min duration: 10 minutes

Max duration: 20 minutes

- Evaluation of process changeability
 - *How do you consider the changeability of the process?*
 - *Which aspects are most easy/important/hard in this respect?*
- Evaluation of relation between decisions and changeability / configurability
 - *How would you describe the relation between changeability and the decisions taken within the process?*
 - *How influential are the decisions taken in the process on the routing of an individual case?*

F.6. Process modeling techniques

Min duration: 10 minutes

Max duration: 20 minutes

- Present exemplary process behaviors expressed as Directly-Follows Graph (DFG), Petri net and BPMN
 - Sequential
 - Optional / Exclusive (XOR)
 - Inclusive (OR)
 - Parallel (AND)
- *How interpretable are the representations of the different process behaviors? (A-UN-C)*
- *How would you rate the readability of the different models? (A-UN-R)*

F.7. Evaluation of baseline process models

Min duration: 10 minutes

Max duration: 20 minutes

- Present the Petri net models of the process variants for the three most common distinct wheelset compositions
 1. 328 SLT M active with gearbox and braking plates
 2. 327 VIRM2/3/4 L passive with three braking discs
 3. 156 FLIRT L passive with braking plates
 4. Optional, if time permits: present the two models of the axle line that include all wheelset types
- For each model, evaluate the following questions:
 - *How well does the model truthfully represent the process? (A-Q-CR)*
 - *Is the model complete? If not, what is missing? (A-Q-CP)*
 - *How do you consider the conciseness? Is the level of detail appropriate? (A-Q-CC)*
 - *Are there any contradictions present in the model? How should this be solved? (A-Q-CS)*
- For the first model, also evaluate relevance and generalizability:
 - *How relevant and valuable is this process model to your organizational division? (I-R)*
 - *How could similar models be used for other parts of the process or other organizational divisions/contexts? (I-G)*

F.8. Evaluation of enhanced process models

Min duration: 10 minutes

Max duration: 20 minutes

- Present the same models again, however, now with decision perspective
- The evaluation now focuses on the decision representations that have been added to the models
- For all models, evaluate the following questions:
 - *How understandable are the depictions of routing decisions at the decision points?* (A-UN-C-DM)
 - *Is the model complete with respect to decisions? If not, what is missing?* (A-Q-CP-DM)
 - *How do you consider the conciseness of the decisions? Is the level of detail appropriate?* (A-Q-CC-DM)
 - *Are there any contradictions present in the model? How should this be solved?* (A-Q-CS-DM)
- For the first model, also evaluate relevance and generalizability:
 - *In which ways is this enhanced process model relevant and valuable to your organizational division?* (I-R-DM)
 - *How could this be used for other parts of the process or other organizational divisions/contexts?* (I-G-DM)

F.9. Evaluation of derived insights and process improvement potential

Min duration: 10 minutes

Max duration: 20 minutes

- Present a summary of the insights derived using decision mining, such as unlogical routing of wheelsets.
- Discuss the following questions:
 - *Are the insights presented valid based on your knowledge about the process?* (I-FC-DMI)
 - *How relevant are these insights to optimization of this specific process?* (I-R-DMI)
 - *In what ways could these insights contribute to potential process improvements?* (I-R-DMI)
 - *Can you think of any other ways that these insights might contribute to process improvement?* (I-G-DMI)

F.10. Closing

Min duration: 2 minutes

Max duration: 5 minutes

- Concluding remarks
 - *Is there anything that has not been discussed and that should be added at this point?*
 - *Are there any further questions regarding the artifacts and insights that have been presented?*
- Thank all participants again for the contributions and investment of their valuable time in academic research
- **Stop the recording**

Table F.1: A mapping of the evaluation criteria with the questions in the focus group / interview protocol, based on [195].

Abbreviation	Goal	Sub-goal	Criterion	Context	Decision Mining
A-UN-C	Artifact	Understandability	Complexity	Modeling techniques	×
A-UN-R	Artifact	Understandability	Readability	Modeling techniques	×
A-Q-CR	Artifact	Quality	Correctness	Baseline models	×
A-Q-CP	Artifact	Quality	Completeness	Baseline models	×
A-Q-CC	Artifact	Quality	Conciseness	Baseline models	×
A-Q-CS	Artifact	Quality	Consistency	Baseline models	×
I-R	Insights	Relevance		Baseline model	×
I-G	Insights	Generalizability		Baseline model	×
A-UN-C-DM	Artifact	Understandability	Complexity	Enhanced models	✓
A-Q-CP-DM	Artifact	Quality	Completeness	Enhanced models	✓
A-Q-CC-DM	Artifact	Quality	Conciseness	Enhanced models	✓
A-Q-CS-DM	Artifact	Quality	Consistency	Enhanced models	✓
I-R-DM	Insights	Relevance		Enhanced model	✓
I-G-DM	Insights	Generalizability		Enhanced model	✓
I-FC-DMI	Insights	Findings confirmation		Decision mining insights	✓
I-R-DMI	Insights	Relevance		Decision mining insights	✓
I-G-DMI	Insights	Generalizability		Decision mining insights	✓



Source code

This appendix presents the source code that was used in this research project. First, it presents the code for the data preparation steps that resulted in the file used in Disco to create the actual event logs. In addition, it provides the decision mining code to generate the enhanced Petri nets with the guard conditions and frequency information.

```
1 """
2 Data preparation
3 """
4
5 ## Import libraries
6 import pandas
7 import numpy
8
9 import warnings
10 warnings.simplefilter(action='ignore', category=FutureWarning)
11
12 from tqdm.notebook import tqdm
13
14 ## Set file paths
15 FILE = './WSR_MESPODATA_Jaar_2022.csv'
16 OUTPUT_CSV = './MESPO_2022_EVENT_LOG.csv'
17
18 ## Read the CSV file
19 input = pandas.read_csv(FILE, encoding='windows-1252', delimiter=';')
20
21 ## Use underscore as separator
22 input.loc[:, 'Takt'] = input.loc[:, 'Takt'].str.replace('Takt_', '')
23
24 ## Define the default event type
25 input['Takt_subtype'] = 'MES'
26 input['Takt_desc'] = pandas.NA
27
28 ## Define separate rows for EQU and MOB events
29 equ_start_rows = input[input['Karakteristiek'] == 'StarttijdEQU']
30 equ_end_rows = input[input['Karakteristiek'] == 'EindtijdEQU']
31 mob_start_rows = input[input['Karakteristiek'] == 'StarttijdMOB']
32 mob_end_rows = input[input['Karakteristiek'] == 'EindtijdMOB']
33
34 ## Add a suffix to the respective events to distinguish them
35 equ_start_rows.loc[:, 'Takt_subtype'] = 'EQU'
36 equ_start_rows.loc[:, 'Karakteristiek'] = equ_start_rows.loc[:, 'Karakteristiek'].str.replace(
37     ('EQU', ''))
38 equ_end_rows.loc[:, 'Takt_subtype'] = 'EQU'
39 equ_end_rows.loc[:, 'Karakteristiek'] = equ_end_rows.loc[:, 'Karakteristiek'].str.replace(
40     ('EQU', ''))
41 mob_start_rows.loc[:, 'Takt_subtype'] = 'MOB'
42 mob_start_rows.loc[:, 'Karakteristiek'] = mob_start_rows.loc[:, 'Karakteristiek'].str.replace(
43     ('MOB', ''))
44 mob_end_rows.loc[:, 'Takt_subtype'] = 'MOB'
45 mob_end_rows.loc[:, 'Karakteristiek'] = mob_end_rows.loc[:, 'Karakteristiek'].str.replace(
46     ('MOB', ''))
47
```



```

44 ## Merge it in the DF
45 input = pandas.concat([input, equ_start_rows])
46 input = pandas.concat([input, equ_end_rows])
47 input = pandas.concat([input, mob_start_rows])
48 input = pandas.concat([input, mob_end_rows])
49
50 ## Enrich with the actual activity names
51 takt_desc_map = {
52 '70': 'Inslag_vuile_voorraad',
53 '180': 'Transport_naar_invoerbaan_afperspers',
54 '185': 'Invoerbaan_afperspers_1',
55 '190': 'Invoerbaan_afperspers_2',
56 '200': 'Invoerbaan_afperspers_3',
57 '205': 'Voorbereiding_revisie',
58 '210': 'Afpersen_onderdelen_van_as',
59 '215': 'Afnemen_onderdelen_van_as',
60 '221': 'Ontkoppelpunt_assenlijn',
61 '225': 'Demontage_cardan_uit-&inslede',
62 '230': 'Deconserveren_BIP',
63 '235': 'NDO-US',
64 '237': 'Demontage_TWK_uit-&inslede',
65 '240': 'ATG_NDO-MAG',
66 '245': 'Meten',
67 '247': 'Takt_247',
68 '250': 'Horizontale_draaibank',
69 '252': 'Aswissel_/Montage',
70 '260': 'Reinigen',
71 '265': 'Maskering_aanbrengen_op_de_as',
72 '267': 'Wachtstand_conserveren_(ATEX)',
73 '270': 'Primer_sputten',
74 '275': 'Uitdampen_primer',
75 '280': 'Laklaag_sputten',
76 '281': 'Uitdampen_laklaag',
77 '285': 'Drogen_conserveerstraat_1',
78 '290': 'Drogen_conserveerstraat_2',
79 '295': 'Drogen_conserveerstraat_3',
80 '300': 'Drogen_conserveerstraat_4',
81 '305': 'Drogen_conserveerstraat_5',
82 '310': 'Drogen_conserveerstraat_6',
83 '315': 'Drogen_conserveerstraat_7',
84 '320': 'Drogen_conserveerstraat_8',
85 '321': 'Drogen_conserveerstraat_9',
86 '322': 'Drogen_conserveerstraat_10',
87 '323': 'Drogen_conserveerstraat_11',
88 '325': 'Maskering_verwijderen_van_de_as',
89 '330': 'Voorbereiden_oppersen',
90 '333': 'Invoer_montagelij',
91 '335': 'Oppersen',
92 '336': 'Persbussen_en_schalen_afnemen',
93 '337': 'Toetsen',
94 '338': 'Ontkoppelbuffer_Hegenscheidt',
95 '342': 'Jack-out_tussenbuffer',
96 '340': 'Profileren_en_vlakken',
97 '345': 'Onbalans_meten',
98 '350': 'Monteren_lagers',
99 '355': 'Afkoelen_in_buffer_1',
100 '360': 'Afkoelen_in_buffer_2',
101 '365': 'Afkoelen_in_buffer_3',
102 '370': 'Monteren_aspotten_en_vullen_met_vet',
103 '375': 'Vullen_TWK_en_monteren_vlerken',
104 '380': 'Uitvoeren_eindcontrole',
105 '385': 'Ontkoppelbuffer_Eindcontrole_1',
106 '395': 'Ontkoppelbuffer_Eindcontrole_2',
107 '396': 'Wielstellen_reparatie',
108 '397': 'Takt_397',
109 '410': 'Verzendgereed_maken_Wielstellen',
110 '420': 'Verzendgereed_TWK',
111 '430': 'Verzendgereed_maken_SLT',
112 '435': 'Verzendgereed_olie_vullen',
113 '440': 'Controle_verzendgereed_maken',
114 '450': 'Wielstel_transport/inslag_naar_SV1',


```

```

115 '520': 'Afpersen_Tandwiel/Sterstuk',
116 '533': 'Reparatie_na_Oppersen',
117 '540': 'Oppersen_Tandwiel',
118 '545': 'Opkrimpen_Tandwiel/Sterstuk',
119 '590': 'Reparatie_na_eindcontrole',
120 '595': 'Klein_herstel_na_eindcontrole',
121 '821': 'Aanvoeren_herbruikbare_componenten',
122 '850': 'Borstelen,_meten_en_visuele_controle_boringen',
123 '865': 'Demonderen_remplaten',
124 '885': 'Afvoeren_niet_herbruikbare_wielen_e/o_remschijven',
125 '890': 'Monteren_remplaten',
126 '905': 'Draaien_en_meten_boringen',
127 }
128 input['Takt_desc'] = input['Takt'].map(takt_desc_map)
129
130 ## Add general case level attributes and fill unknown names
131 input['Takt'] = input['Takt'].fillna('Algemeen')
132 input['Takt_desc'] = input['Takt_desc'].fillna('Onbekend')
133
134 ## Retain possibility to export to Excel
135 input['ExcelColumn'] = input['Takt'] + "|" + input['Karakteristiek'] + "[" + input['
    Karakteristiek_Instantie'].apply(str) + "]"
136
137 ## Check the amount of resulting cases
138 print("Aantal_MESPOs:", len(input['MESPONUMMER'].unique()))
139
140 ## Melt the DF around the id columns
141 pandas.melt(input, id_vars=['MESPONUMMER', 'Takt', 'Takt_instantie', 'Karakteristiek_Instantie',
    Karakteristiek'])
142
143 ## Pivot the DF around the identity columns to turn it into a long format
144 final = input.pivot_table(index=['MESPONUMMER', 'Takt', 'Takt_subtype', 'Takt_desc', 'Takt_
    instantie', 'Karakteristiek_Instantie'], columns='Karakteristiek', values='Karakteristiek_
    Waarde', aggfunc='first')
145
146 ## Write the to-be event log as CSV file
147 final.to_csv(OUTPUT_CSV)
148
149 """
150 Decision mining
151 """
152
153 ## Import libraries
154 import os
155 import pandas as pd
156 import pm4py
157 import ipykernel
158 import pprint
159 from pm4py.visualization.petri_net import visualizer
160 from pm4py.algo.decision_mining import algorithm as decision_mining
161
162 ## Import the event log
163 log = pm4py.read_xes("NS_RLW_MES_2022_328.xes.gz")
164
165 ## Drop columns with associative values and anonymized attributes (e.g. resources)
166 log = log.drop(['UNIT', 'RapportNaam', 'Opmerking', 'DekraOpmerking', 'StarttijdMOB', '
    StarttijdMOB_IN', 'StarttijdMOB_UIT', 'StarttijdEQU', 'EindtijdMOB', 'EindtijdMOB_IN', '
    EindtijdMOB_UIT', 'EindtijdEQU', 'WielstelLocatie', 'Batchnummer', 'Kratnummer'], axis=1,
    errors='ignore')
167 log = log.drop(log.filter(regex='Attribute').columns, axis=1, errors='ignore')
168
169 ## Only retain the completion timestamps from the event log as
170 ## the inductive miner does not support interval logs or lifecycle information
171 event_log_int = log[log["lifecycle:transition"] == "complete"]
172 ## Fill empty values with zeroes
173 event_log_int = event_log_int.fillna(0)
174
175 ## Discover a Petri net on the log with 100% fitness
176 net, im, fm = pm4py.discover_petri_net_inductive(event_log_int, noise_threshold=0.00)
177
178 ## Save and view the initial Petri net

```

```
179 pn = visualizer.apply(net, im, fm)
180 pn.render("NS_RLW_MES_2022_328_th000", format="svg")
181 visualizer.view(pn)
182
183 ## Visualize Petri net without labels
184 pn_debug = visualizer.apply(net, im, fm, parameters={visualizer.Variants.WO_DECORATION.value.
    Parameters.DEBUG: True})
185 visualizer.view(pn_debug)
186
187 ## List all decision points in the Petri net in pretty ascending format
188 dps = decision_mining.get_decision_points(net, labels=True)
189 pp = pprint.PrettyPrinter()
190 dps_asc = sorted(dps.items(), key=lambda x: int(x[0][2:]))
191 pp.pprint(dps_asc)
192
193 ## Replay the event log with TBR and add the criteria to the decision points
194 net, im, fm = decision_mining.create_data_petri_nets_with_decisions(event_log_int, net, im,
    fm)
195
196 ## List all guard conditions that were discovered
197 for t in net.transitions:
198     if "guard" in t.properties:
199         print("")
200         print(t)
201         print(t.properties["guard"])
202
203 ## Visualize the enhanced Petri net and save it
204 dpn = visualizer.apply(net, im, fm)
205 visualizer.view(dpn)
206 dpn.render("NS_RLW_MES_2022_328_th000_dpn", format="svg")
207
208 ## Visualize the enhanced Petri net, add frequency information, and save it
209 dpn_frq = visualizer.apply(net, im, fm, variant=visualizer.Variants.FREQUENCY, log=
    event_log_int)
210 visualizer.view(dpn_frq)
211 dpn.render("NS_RLW_MES_2022_328_th000_dpn_frq", format="svg")
```



Petri net with data for wheelset type 328

The Petri net model with data for wheelset type 328 unfortunately is too large for integral inclusion in this PDF. This would lead to problems with zooming in and out of page width. Therefore, the actual model is available as a separate PDF file attachment. This attachment is not included with this redacted version due to the confidential nature of the model.