
*Developing the Organizational and Technical Foundations of Process Mining:
On The Role of Governance and Data Quality*

*Dissertation
zur Erlangung des Grades eines Doktors der Wirtschaftswissenschaft
der Rechts- und Wirtschaftswissenschaftlichen Fakultät
der Universität Bayreuth*

*Vorgelegt von
Sebastian Johannes Schmid
aus
Augsburg*

Dekan: Prof. Dr. Claas Christian Germelmann

Erstberichterstatter: Prof. Dr. Maximilian Röglinger

Zweitberichterstatter: Prof. Dr. Jan Mendling

Tag der mündlichen Prüfung: 03.06.2025

*But you know that when the truth is told,
that you can get what you want, or you can just get old?*

William Martin "Billy" Joel

Was war das nicht für eine Reise, die hier hinter mir liegt – es überrascht kaum, dass es sich nun fast surreal anfühlt, die letzten Zeilen dieser Dissertation zu Papier zu bringen. Dieser Meilenstein markiert nicht nur das Ende meines akademischen Werdegangs – und damit auch zehn Jahre Bayreuth – sondern vor allem den Abschluss dreier intensiver wie auch lehrreicher Jahre. Drei Jahre voller Höhen und Tiefen, Rückschläge und Fortschritte – vor allem aber auch: voller persönlicher Entwicklung. Wie viele scheinbar unüberwindbare Hürden es waren, kann ich rückblickend kaum mehr zählen. Trotzdem hat es am Ende irgendwie doch alles geklappt. Und somit bleibt unterm Strich ein klares Fazit: das war es alles wert. Auch wenn heute mein Name auf dem Titelblatt dieser Dissertation steht, war diese Reise bei Weitem keine Individualleistung. Ich bin unendlich dankbar für all jene, die mich auf diesem Weg begleitet haben. Worte können kaum ausdrücken, wie viel Rückhalt, Unterstützung und Ermutigung ich in dieser Zeit erfahren durfte - und doch will ich es zumindest versuchen.

Mein besonderer Dank gilt zunächst meinem Doktorvater, Max Röglinger. Du hast mir nicht nur die Chance gegeben, Teil deines Teams zu werden, sondern mir von Beginn an großes Vertrauen entgegengebracht. Durch dich habe ich gelernt, wie wertvoll Ambition ist und was es heißt, Exzellenz mit Menschlichkeit zu verbinden. Als Mentor hast du mich fachlich wie persönlich geprägt, und gerade im Endspurt der Promotion war deine Unterstützung unbezahlbar. Für diese Erfahrung bin ich dir zutiefst dankbar und werde sie stets in großer Wertschätzung bewahren.

Ebenso danke ich allen Freunden, Kollegen und Co-Autoren, die ich im Umfeld des FIM Forschungsinstituts, der Universität Bayreuth sowie am Fraunhofer FIT kennenlernen durfte. Die inhaltlichen Herausforderungen einer Promotion sind das eine, das tägliche Miteinander mit großartigen Menschen das andere. Es war ein Privileg, mit euch zusammenzuarbeiten, Ideen zu teilen und lachen zu können. Ihr habt diese Reise leichter und deutlich schöner gemacht.

Ein weiterer großer Dank gilt all meinen Freunden aus der Heimat. Ich weiß, dass ich in dieser Zeit nicht immer durch Anwesenheit gegläntzt habe und doch wart ihr da: ohne Vorwurf, mit ehrlichem Interesse und echter Freude über meine Erfolge. Diese besondere Art der Freundschaft ist alles andere als selbstverständlich und ich bin tief dankbar, immer auf euch zählen zu können.

Meinen Eltern, Elfriede und Hermann, möchte ich von Herzen danken. Ihr kommt aus einer anderen Welt – der Pflege und dem Handwerk – und die akademische Sphäre war euch nie sonderlich vertraut. Trotzdem habt ihr mir von klein auf vermittelt, wie wichtig Bildung ist und was es bedeutet, sich im Leben anzustrengen, Verantwortung zu übernehmen und dabei stets menschlich wie auch bodenständig zu bleiben. Durch eure harte Arbeit, euren Verzicht und eure unermüdliche Unterstützung habt ihr mir das Studium überhaupt erst ermöglicht. Woran ich hier eigentlich genau geforscht habe, blieb wohl manchmal ein Rätsel - was euch aber keine Sekunde davon abgehalten hat, mir euer bedingungsloses Vertrauen zu schenken. Oder um es in euren Worten zu sagen: „Der Bub wird das schon richtig machen“. Ich bin unendlich stolz, auf so tolle Eltern wie euch beide zählen zu dürfen.

Das Beste kommt zum Schluss: Liebe Karo, die vergangenen Jahre haben uns beide gefordert – jeden für sich, ganz individuell, aber auch gemeinsam. Die eigenen Bedürfnisse in solch einer Zeit zurückstellen zu können, um den anderen aufzufangen, kann nicht jeder. Aber du konntest es – und wie du es konntest! Du hast mich aufgebaut, wenn ich meine Zweifel hatte und Freiräume für uns geschaffen, wenn die Promotion mal wieder alles andere in den Hintergrund rückte. Du hast nie geklagt, sondern mitgetragen. Jeden Abend, wenn ich nach Hause kam, wusste ich: Da wartet jemand auf mich, der über all dem steht, jemand der an mich glaubt, mit ganz viel Wärme und Liebe. Dass wir nun diesen Meilenstein gemeinsam feiern können, macht mich unendlich glücklich. Noch glücklicher macht es mich zu wissen, dass ich all das, was noch vor mir liegt, mit dir an meiner Seite bestreiten darf.

Copyright Statement

The following sections are partly comprised of content from the research papers included in this thesis. To improve the readability of the text, I omit the standard labelling of these citations.

Abstract

Process mining is a data-driven technique that generates insights into business processes by analyzing event log data. By using process mining, organizations gain an evidence-based understanding of how processes are actually executed, enabling the identification of inefficiencies, deviations, and opportunities for improvement. While research has predominantly focused on algorithm engineering for analysis tasks, many organizations face difficulties in translating these technical capabilities into sustained business value. In particular, persistent challenges in process mining governance and process data quality management limit the effective use of process mining in practice. To overcome these challenges, the overarching purpose of this dissertation is to contribute to advancing the organizational and technical prerequisites of process mining. Following the Design Science Research paradigm, the dissertation presents five research papers that introduce and evaluate novel artifacts with both theoretical and practical relevance.

Process mining governance emerges as a critical area of concern as organizations seek to embed process mining beyond isolated projects and into enterprise-wide practice. However, existing research has largely overlooked the organizational and managerial prerequisites for successful adoption. Key challenges include the absence of structured guidance for designing process mining setups, limited understanding of how individuals respond to increased process transparency, and a lack of strategies for embedding process mining as a continuous capability rather than a one-off initiative. To address these gaps, research paper P1 develops a taxonomy of organizational process mining setups, offering a structured framework to support governance design across diverse organizational contexts. Research paper P2 complements this by introducing a capability framework for managing process-based behavioral visibility, outlining how organizations can leverage transparency for business value while mitigating potential negative side effects. Together, these contributions advance the field of process mining governance by equipping practitioners and researchers with conceptual and practical tools for sustained and value-generating adoption.

Process data quality management is a critical research area, as the accuracy and reliability of process mining insights fundamentally depend on the quality of the underlying event logs. In reality, however, event logs are often affected by multiple imperfections such as incorrect timestamps, missing case identifiers, or mislabeled activities which undermine the validity of process analyses and complicate data preparation. Addressing these issues is particularly challenging when imperfections co-occur and interact, making sequential repair approaches

unfit. Both researchers and practitioners have identified poor event log quality and the resulting complex data preparation as highly relevant challenges that hinder the effective application of process mining in real-world settings. In response to these challenges, the dissertation proposes three generative AI-based artifacts for event log repair. Research paper P3 proposes a method based on Generative Adversarial Networks for repairing identical timestamp errors, achieving state-of-the-art accuracy in timestamp estimation. Research paper P4 extends this work by introducing a hybrid method for repairing missing case identifiers, combining rule-based logic with Transformer models and human-in-the-loop elements. Building on both contributions, research paper P5 presents a fine-tuned Large Language Model for multi-imperfection event log repair, offering a unified and adaptable approach that overcomes limitations of the toolchain paradigm of event log repair. Together, these contributions advance both single- and multi-imperfection event log repair, addressing the challenge of poor event log quality. By offering methods that automate parts of the repair process, they also help reduce the effort required for complex data preparation.

Collectively, the five research papers included in this dissertation contribute to advancing the process mining discipline by strengthening both its organizational and technical foundations. By addressing critical challenges in process mining governance and process data quality management, this dissertation advances the socio-technical understanding of process mining and supports its sustained, value-generating use in organizations.

Table of Contents

I	Introduction	1
I.1	Motivation	1
I.2	Research Objectives	5
I.3	Structure of the Thesis and Embedding of the Research Papers	6
II	Process Mining Governance	9
II.1	A Taxonomy of Organizational Process Mining Setups	9
II.2	Capabilities for Managing Process-Based Behavioral Visibility.....	13
III	Process Data Quality Management.....	17
III.1	Repairing Identical Timestamp Errors with Generative Adversarial Networks	17
III.2	Repairing Elusive Cases in Event Logs with Transformer Networks	21
III.3	Multi-Imperfection Event Log Repair with LLMs.....	26
IV	Conclusion.....	31
IV.1	Summary	31
IV.2	Limitations and Future Research.....	32
V	References	36
VI	Appendix	42
VI.1	Index of Research Articles.....	42
VI.2	Individual Contribution to the Included Research Articles	43
VI.3	Research Paper 1: Navigating the Landscape of Organizational Process Mining Setups: A Taxonomy Approach	46
VI.4	Research Paper 2: Capabilities for Building and Managing Process-Based Behavioral Visibility in Organizations.....	47
VI.5	Research Paper 3: Everything at the Proper Time: Repairing Identical Timestamp Errors in Event Logs with Generative Adversarial Networks	50
VI.6	Research Paper 4: Case ID Revealed HERE: Hybrid Elusive Case Repair Method for Transformer-Driven Business Process Event Log Enhancement.....	51
VI.7	Research Paper 5: One to Rule Them All: Large Language Models for Multi-Imperfection Business Process Event Log Repair	53

List of Figures

Figure 1: Process mining challenges addressed by this dissertation (vom Brocke et al. 2021; Martin et al. 2021).....	3
Figure 2: The relationship between data, behavioral visibility, and business value as well as the necessary capabilities.....	14
Figure 3 Overview of the proposed method for repairing identical timestamp errors	18
Figure 4: Architecture of the generator	19
Figure 5: Architecture of the discriminator	19
Figure 6: Comparison of effectiveness in estimating times between events	20
Figure 7: Overview of HERE	23
Figure 8: Repair routine of HERE.....	23
Figure 9: Effectiveness of HERE in repairing the hospital billing event log	25
Figure 10: Overview of the proposed method for multi-imperfection event log repair	28
Figure 11: Effectiveness of repair with accuracy metrics.....	29

List of Tables

Table 1: Structure of this thesis and embedding of the research papers	7
Table 2: Taxonomy of organizational process mining setups.....	11
Table 3: Overview of foundational capabilities.....	15
Table 4: Overview of transformational capabilities	15
Table 5: Overview of continual capabilities	16
Table 6: Research articles included in the dissertation.....	42
Table 7: Further research articles not included in the dissertation	43

Abbreviations

AI	Artificial Intelligence
BPM	Business Process Management
C2E	Conceptual-to-Empirical
cGAN	Conditional Generative Adversarial Network
CRediT	Contributor Role Taxonomy
DSR	Design Science Research
E2C	Empirical-to-Conceptual
FEDS	Framework for Evaluation in Design Science
GAN	Generative Adversarial Network
HERE	Hybrid Elusive Case Repair Engine
LLM	Large Language Model
LoRA	Low-Rank Adaptation
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
RMSE	Root Mean Squared Error
XES	eXtensible Event Stream

I Introduction

I.1 Motivation

Process mining is a data-driven technique that generates valuable insights into business processes by analyzing process executions recorded in event log data (vom Brocke et al. 2021). By leveraging these digital footprints, process mining has advanced traditional business process management (BPM) approaches – typically based on idealized “should-be” models – by providing an “as-is” perspective of how processes are actually performed (van der Aalst 2016). This new perspective on business processes has enabled a variety of novel use cases. For instance, conformance checking and performance analysis reveal deviations, bottlenecks, and inefficiencies, thereby providing actionable insights for data-driven process improvement (van der Aalst 2022). Furthermore, technologies such as artificial intelligence (AI) can be integrated in process mining to enable a forward-looking “to-be” perspective on business processes (Di Francescomarino and Ghidini 2022). This allows for use cases such as predictive monitoring, which forecasts outcomes like remaining processing times (Ceravolo et al. 2024), or prescriptive monitoring, which provides stakeholders with recommendations to prevent potential issues in process executions (Weinzierl et al. 2020).

These use cases allow process mining to realize a range of economic benefits, including reduced process costs, optimized working capital, and enhanced customer satisfaction (Badakhshan et al. 2022). Hence, it comes as no surprise that process mining has experienced significant uptake in practice. For example, a study by Deloitte found that 74% of the companies surveyed had either already implemented process mining or were planning to do so (Deloitte 2025). Moreover, recent market analysis indicates that global process mining software revenue grew by 40% in 2023 and is projected to reach \$1.5 billion by 2025, reflecting a compound annual growth rate of 33% from 2020 to 2025 (Kerremans et al. 2024). Notably, process mining was initially driven by research, with the first process discovery algorithms emerging in the early 2000s (van der Aalst et al. 2004), while industry adoption only began to accelerate around 2015 (van der Aalst 2020). Yet, academic interest in process mining has not waned; on the contrary, the research discipline is experiencing a considerable surge in publications (van der Aalst 2020).

Process mining research has historically focused heavily on the development and improvement of algorithms for analysis tasks (vom Brocke et al. 2021). For instance, early contributions in process discovery, such as the Alpha Miner, laid the groundwork for automatically extracting process models from event log data (van der Aalst et al. 2004). Subsequent innovations like the

Inductive Miner further advanced these capabilities by extracting more sound and fit process models within a finite timeframe (Leemans et al. 2014). Similar progress has been made with algorithms beyond process discovery, leading to continuous improvements in the discipline's strong technological foundations (van der Aalst 2022).

While this strong foundation has undoubtedly contributed to the success and widespread adoption of process mining, recent studies indicate that many of today's challenges extend well beyond technological limitations. In practice, issues such as strategic alignment, governance, and cultural aspects are emerging as the primary challenges that must be addressed to fully realize process mining's potential (Martin et al. 2021). This shift in focus reflects a broader pattern recognized in information systems research. Socio-technical systems theory, for instance, states that organizational effectiveness depends on the joint optimization of technical and social subsystems, each influencing and reinforcing the other (Cherns 1976). Viewed through this lens, it becomes evident that addressing the social context is essential for advancing the process mining discipline (vom Brocke et al. 2021). In response, scholars have increasingly called for a more holistic research perspective – one that integrates technological progress with managerial and organizational considerations so as to better address the challenges practitioners face in the application of process mining (Grisold et al. 2021; Martin et al. 2021).

To guide research within this broader perspective, the five-level framework for research on process mining has been proposed by vom Brocke et al. (2021). This framework structures research across five interrelated dimensions, as partly illustrated in Figure 1. First, research on the *technical* level consists of design and knowledge contributions with a historic focus on algorithm engineering for analysis tasks. Second, research on the *individual* level examines how stakeholders engage with process mining to perform certain tasks, emphasizing usability, acceptance, and outcomes. Third, research on the *group* level investigates how groups of actors collaborate to interpret and act upon process mining insights. Fourth, research on the *organizational* level examines how the integration and use of process mining influences strategic alignment, governance structures, and organizational culture. Fifth, research on the *ecosystem* level explores the inter-organizational dynamics and network effects arising from process mining.




Figure 1: Process mining challenges addressed by this dissertation (vom Brocke et al. 2021; Martin et al. 2021)

Although research in process mining has been predominantly technical, a critical aspect remains underexplored at this level: process data quality management (Hofstede et al. 2023). As a fundamentally data-driven approach, process mining relies on high-quality input data to produce reliable and useful insights (Suriadi et al. 2017). Real-world event logs are, however, frequently affected by various imperfections such as identical timestamp errors, missing case identifiers, and mislabeled activities (Andrews et al. 2018) – any of which can significantly undermine the accuracy of process mining insights (Mannhardt 2022). For instance, incorrect timestamps cause events to be misordered, resulting in unrepresentative process models and inaccurate performance metrics (Fischer et al. 2022). Unsurprisingly, up to 80% of the effort spent in process mining is dedicated to data preparation (Wynn et al. 2022). This stage is especially challenging when multiple quality issues coexist. In such scenarios, current research typically suggests applying single-imperfection repair methods in succession within a toolchain (Andrews et al. 2020). However, the sequential application of interdependent repair methods may introduce unintended side effects or leave residual errors that cascade through the toolchain (Suriadi et al. 2017; Hofstede et al. 2023). Consequently, process data quality management remains an important area of research at the technical level, with practitioners and researchers rating the persistent challenges of poor event log quality (cf. *Challenge 7*, Martin et al. 2021) and the resulting complex data preparation (cf. *Challenge 12*, Martin et al. 2021) as extremely relevant (Martin et al. 2021).

While process data quality management fills a critical gap in advancing a more holistic research

perspective on the technical level, its value remains limited if the organizational and managerial implications of process mining are not equally well understood and addressed (vom Brocke et al. 2021). As process mining becomes more embedded in operational and strategic contexts, research must increasingly consider how its insights affect people and decision-making within organizations. For instance, the increasing visibility into human behavior also introduces new organizational opportunities and challenges (Martin et al. 2021). In recent years, the ubiquity of digital technologies results in trace data being increasingly generated by professional activities (Franzoi et al. 2023). By constructing event logs from such trace data and applying process mining (Weber et al. 2024), the behavior of individuals becomes increasingly visible. This phenomenon known as (process-based) behavioral visibility (Leonardi and Treem 2020) offers various benefits, such as allowing managers to optimize resource allocation and improve performance measures (Badakhshan et al. 2022; Zorina et al. 2021). Furthermore, process mining has the potential to become the technological backbone of a new management paradigm based on behavioral visibility. Thereby, real-time insights of process mining allow managers to intervene proactively during work processes, shifting from traditional retrospective evaluations towards a continuous management paradigm (Bernstein and Li 2017; Leonardi and Treem 2020). This is particularly important as organizations struggle to continuously integrate process mining beyond the stage of a one-off initiative (Martin et al. 2021).

Alongside these benefits, however, considerable challenges are associated with process-based behavioral visibility. For instance, organizations can exploit behavioral visibility to implement intrusive control practices and impose unrealistic performance expectations (e.g., Vaujany et al. 2021). This, in turn, can lead to emotional stress, resistance, and even the gaming of managerial systems by workers (Aaltonen and Stelmaszak 2024; Benlian et al. 2022; Newlands 2021; Spicer 2017; Zorina et al. 2021). Consequently, research has yet to develop strategies for managing process-based behavioral visibility, capitalizing on its benefits while mitigating negative side effects. In particular, existing studies fall short in addressing employees' defensive reactions to transparency (cf. *Challenge 23*, Martin et al. 2021) and concerns over intrusive monitoring (cf. *Challenge 26*, Martin et al. 2021). Furthermore, approaches to embed process mining as the backbone of a continuous management paradigm instead of a one-off initiative, remain underdeveloped (cf. *Challenge 27*, Martin et al. 2021).

Effectively managing process-based behavioral visibility and capitalizing on the benefits of process mining, not only requires capabilities at the individual and group level but must also be embedded within broader organizational structures and governance mechanisms (Eggers et al. 2021). Without adequate governance, efforts to institutionalize process mining and manage its

socio-technical implications may remain fragmented or unsustainable (Martin et al. 2021). Although research into BPM governance offers initial insights (Kerpedzhiev et al. 2021; Rosemann and vom Brocke 2015), it falls short in addressing process mining’s unique, data-driven requirements (van der Aalst et al. 2012). Moreover, process mining setups vary widely (Reinkemeyer et al. 2022), likely due to the diverse contextual factors that influence their design. So far, research has not described the various factors that need to be considered when integrating process mining into an organization. Unsurprisingly, practitioners lack guidance in implementing an organizational setup that ensures effective use of process mining across varied contexts (cf. *Challenge 3 and 6*, Martin et al. 2021). For instance, the optimal organizational anchoring of process mining expertise remains unclear (cf. *Challenge 10*, Martin et al. 2021).

I.2 Research Objectives

Given the previously introduced research needs, this dissertation contributes to two key areas. First, in the research area of *process mining governance*, scholars have highlighted a lack of organizational and managerial perspective (vom Brocke et al. 2021), resulting in persistent challenges for establishing effective governance structures (Martin et al. 2021). To address this gap, the dissertation provides two research contributions. The first contribution introduces a taxonomy of organizational process mining setups, offering a comprehensive description of their unique characteristics. By doing so, light is shed on the governance requirements of process mining while guiding practitioners in designing and refining their process mining setups. Furthermore, research lacks strategies for effectively managing process-based behavioral visibility. While this phenomenon offers vast potential for generating business value (Badakhshan et al. 2022; Zorina et al. 2021), it is unclear what is necessary to effectively address some of its undesired consequences such as defensive reactions in employees due to digital surveillance-induced stress (Grisold et al. 2024; Martin et al. 2021). Hence, the second contribution focuses on identifying the capabilities needed to convert process-based behavioral visibility into sustained business value. This contribution not only outlines capabilities for generating business value from process-based behavioral visibility but also proposes measures to mitigate undesired side effects, such as resistance to transparency. Together, these contributions advance the understanding of process mining governance, thus bridging the gap between its robust technological foundations and the complex organizational environments in which it is deployed. Consequently, the proposed frameworks address multiple challenges at the individual, group, and organizational level of process mining research.

In the research area of *process data quality management*, the accuracy and reliability of process

mining outcomes is fundamentally dependent on the quality of the input event logs (Hofstede et al. 2023). Real-world logs frequently suffer from various imperfections and can be incomplete, noisy, or imprecise (Fischer et al. 2022). If such poor quality event logs are used for process mining, the insights will be unrepresentative and misleading, ultimately diminishing business value (Suriadi et al. 2017). Hence, the majority of effort in process mining is currently spent on data preparation (Wynn et al. 2022). To address this research need, the dissertation makes three key contributions based on generative AI, as it has been successfully used to repair data quality issues in other domains such as audio and image data (Hofmann et al. 2021). First, it introduces a method for automatically repairing identical timestamp errors by combining established error detection and reordering techniques with a novel Generative Adversarial Network (GAN) based approach for timestamp estimation. Second, it proposes a hybrid method that repairs missing case identifiers in event logs by integrating a rule-based approach with a Transformer-based architecture, enhanced through human-in-the-loop elements. Third, it explores the use of fine-tuned Large Language Models (LLMs) to simultaneously repair multiple event log imperfections, offering a single, adaptable framework that advances the current toolchain paradigm in process data quality management. Together, these contributions enhance process data quality management by proposing effective repair strategies that range from single to multi-imperfection event log repair. By doing so, the dissertation not only provides design improvements and exaptations, but also knowledge contributions in the form of performance, sensitivity, and explanatory propositions (vom Brocke et al. 2021).

By integrating both research areas of process mining governance and process data quality management, the overall purpose of this dissertation is to advance the *technical and organizational prerequisites of process mining*. In particular, the dissertation contributes by making a step towards establishing an initial balance between the technical and social subsystems of process mining. To do so, the dissertation follows the Design Science Research (DSR) paradigm to address its research objectives. Thereby, multiple artifacts in the form of novel methods and instantiations are designed and evaluated, ensuring scientific rigor with practical relevance (Hevner et al. 2004). Ultimately, the dissertation advances process mining research on the technical, individual, group, and organizational level (vom Brocke et al. 2021) while addressing the prerequisites needed to generate business value from its sustained use.

I.3 Structure of the Thesis and Embedding of the Research Papers

The cumulative dissertation comprises five research papers that collectively address the research objectives outlined in Section I.2. An overview of the overall structure and embedding

of the research contributions is provided in Table 1.

I	Introduction
II	Process Mining Governance
P1	Navigating the Landscape of Organizational Process Mining Setups: A Taxonomy Approach <i>Marcus L, Schmid SJ, Friedrich F, Röglinger M, Grindemann P</i>
P2	Capabilities for Building and Managing Process-Based Behavioral Visibility in Organizations <i>Franzoi S, Kipping G, Marcus L, Schmid SJ, vom Brocke J, Grisold T, Mendling J, Röglinger M</i>
III	Process Data Quality Management
P3	Everything at the Proper Time: Repairing Identical Timestamp Errors in Event Logs with Generative Adversarial Networks <i>Schmid SJ, Moder L, Hofmann P, Röglinger M</i>
P4	Case ID Revealed HERE: Hybrid Elusive Case Repair Method for Transformer-Driven Business Process Event Log Enhancement <i>Zetsche F, Andrews R, ter Hofstede AHM, Röglinger M, Schmid SJ, Wynn MT</i>
P5	One to Rule Them All: Large Language Models for Multi-Imperfection Business Process Event Log Repair <i>Schmid SJ, Zetsche F, Röglinger M</i>
IV	Conclusion
V	References
VI	Appendix

Table 1: Structure of this thesis and embedding of the research papers

The dissertation is structured as follows. In Section I, the overarching research field is motivated, and the research objectives are derived. Section II presents the contributions in the area of process mining governance, featuring two papers that explore the organizational and managerial dimensions of process mining. Research paper P1 proposes a taxonomy of organizational process mining setups, offering a comprehensive understanding of the decision-making factors and configurations underlying process mining governance. Research paper P2 introduces a capability framework for managing process-based behavioral visibility, suggesting strategies that enable organizations to generate business value while mitigating drawbacks.

Section III presents the contributions in the research area of process data quality management, featuring three papers that design generative AI-based methods for enhancing event log quality. Research paper P3 presents a method for automatically repairing identical timestamp errors by integrating established error detection techniques with a GAN-based approach for timestamp estimation. Research paper P4 develops a hybrid approach for repairing missing case identifiers,

combining a rule-based approach with a Transformer-based architecture, complemented by human-in-the-loop elements. Research paper P5 explores the application of fine-tuned LLMs for multi-imperfection event log repair, offering a unified and adaptable framework that advances the current toolchain paradigm in process data quality management.

Section IV concludes the dissertation by summarizing the contributions, discussing limitations, and outlining directions for future research. Section V lists all references used, and Section VI provides an index of the research papers, details of the author's individual contributions as well as the full research papers.

II Process Mining Governance

As previously motivated, research in process mining has long focused on technological advancements, particularly algorithm engineering for analysis tasks (vom Brocke et al. 2021). However, with growing adoption in practice, the organizational and managerial challenges surrounding process mining have become increasingly apparent (Martin et al. 2021). For instance, a lack of comprehensive understanding and guidance prevents many organizations from designing context-appropriate and effective process mining setups (Martin et al. 2021). Hence, research paper P1 (Section II.1) proposes a taxonomy of organizational process mining setups, systematically describing the key dimensions and decision points that shape how process mining is governed across varying contexts. Furthermore, process mining offers powerful opportunities to generate transparency regarding process work from digital trace data (Eggers et al. 2021). Yet, it remains unclear how this process-based behavioral visibility can be translated into sustained business value – especially while mitigating the potential negative effects associated with digital surveillance and organizational resistance. Thus, research paper P2 (Section 0) proposes a capability framework for holistically managing process-based behavioral visibility – from establishing its socio-technical foundations to ensuring sustained business value generation.

II.1 A Taxonomy of Organizational Process Mining Setups

Despite significant technological advancements in process mining, various organizational challenges in establishing effective governance structures remain a critical barrier to its successful adoption (vom Brocke et al. 2021; Martin et al. 2021). As noted by Martin et al. (2021), the unclear organizational anchoring of process mining poses a significant challenge, since integrating process mining capabilities within the organizational structure requires both strategic alignment and operational effectiveness. While previous studies have examined process mining in the context of individual use cases (e.g., Yang and Su 2014), single organizations (e.g., Reinkemeyer 2020), or specific industries such as healthcare (Rojas et al. 2016), a holistic understanding of process mining governance across diverse organizational contexts is still missing. This lack of a comprehensive perspective leaves many organizations struggling to determine an appropriate process mining setup, ultimately preventing them from fully capitalizing on the technology’s potential. In light of this gap, the central research question addressed by research paper P1 is: *What are the characteristics of organizational process mining setups?*

To answer the research question, the study's goal is to identify and systematically categorize the key dimensions and factors that define effective process mining governance, providing both a descriptive and practical framework. Taxonomies are particularly well-suited for this purpose as they enable the systematic classification of complex domains, support the organization of conceptual knowledge, and serve as a foundation for theory building and structured decision-making (Nickerson et al. 2013). Hence, the study follows the taxonomy development method by Nickerson et al. (2013) and Kundisch et al. (2021). This iterative process combines Empirical-to-Conceptual (E2C) and Conceptual-to-Empirical (C2E) iterations. In E2C iterations, characteristics and dimensions are derived inductively from empirical observations of real-world objects, while in C2E iterations, they are derived deductively based on existing theoretical knowledge and are then validated against empirical data (Kundisch et al. 2021). Initially, the taxonomy was developed using data gathered from a survey of 214 process mining adopters, which provided a broad view of organizational practices (E2C). This was followed by 15 semi-structured interviews with process mining experts from various industries, designed to generate deeper insights into specific dimensions of process mining setups than the preceding survey could offer (E2C). A subsequent literature review further enriched the taxonomy and ensured that the naming of its dimensions, characteristics, and layers aligned with established terminology, thereby standardizing its presentation and enhancing conceptual clarity (C2E). Additional evaluation interviews and an online survey were conducted to assess the taxonomy's understandability, completeness, and usefulness (E2C). The final taxonomy was then applied in three exemplary cases, representing organizations with beginner, intermediate, and experienced levels of process mining adoption.

The central outcome of research paper P1 is a multilayer taxonomy (depicted in Table 2) that categorizes organizational process mining setups into 12 distinct dimensions across four layers:

- **Governance and Structure:** Defining how the process mining unit is embedded within the broader organizational context.
- **Operationalization and Scope:** Capturing the strategic focus and operational mode of process mining activities.
- **Funding and Planning:** Outlining the financial approach to process mining within the organization.
- **Roles and Responsibilities:** Detailing how responsibilities, decision-making authority, and support structures are assigned.

Layer	Dimension	Characteristic				E/N*	Guiding questions			
Governance and structure	<i>Degree of centralization</i>	Centralized		Hybrid		Decentralized		E	What is your process mining unit's degree of centralization?	
	<i>Anchoring</i>	IT	Business	Shared services	Executive level	N	Where in the organization is your process mining unit anchored?			
	<i>Institutionalization</i>	Integrated in a (business) department		Integrated in a CoE		Cross-functional organization		Standalone department / CoE	N	How is your process mining unit institutionalized?
	Operationalization and scope	<i>Key activities</i>	Demand generation and assessment	Data science and engineering	Project management	Governance and steering	Change and community management	Value management and scaling	N	Which activities are part of your process mining unit's value proposition?
	<i>Prioritization of projects</i>	Long-term roadmap		Mid-term pipeline		Short-term ad hoc		N	How are incoming projects prioritized by your process mining unit?	
	Funding and Planning	<i>Budgeting</i>	Global	Project-based	Process-based	Per department		N	Where does the financial budget for process mining activities originate?	
Roles and responsibilities	<i>Internal cost management</i>	Profit center		Hybrid		Cost center		E	What is your process mining unit's financial setup?	
	<i>Role allocation</i>	Based on (business) department	Based on key activities	Based on end-to-end processes	Flexible		N	How are the roles in/of the process mining unit allocated?		
	<i>Internal leadership</i>	Process mining lead		Executive sponsor		Champion		N	Which process mining leadership roles exist in your organization?	
	<i>External support**</i>	Vendor		Consultancy		None		N	Which external parties provide services for your process mining activities?	
	<i>Data ownership</i>	IT		Business		Process mining unit		N	Who has primary ownership of the source data used in process mining activities?	
	<i>Tool ownership</i>	IT		Business		Process mining unit		N	Who has primary ownership of the tools used in process mining activities?	

Notes: * E = exclusive, N = non-exclusive.

** Selecting both "None" and another option simultaneously is not applicable.

Table 2: Taxonomy of organizational process mining setups

Each dimension is marked as either exclusive or non-exclusive and is accompanied by guiding questions to facilitate its practical application. Rigorous evaluation through twelve expert interviews and a follow-up online survey confirmed the taxonomy's completeness, understandability, and practical usefulness. Interview participants praised the taxonomy for offering a clear, comprehensive view of organizational process mining setups, describing it as a helpful "bird's-eye perspective" to assess and compare their own or clients' configurations. Several participants reported successfully classifying existing projects with the taxonomy and highlighted its value in guiding stakeholders unfamiliar with the topic. Suggestions for improvement focused primarily on clarifying specific terminology and exploring the rationale behind setup choices, many of which were incorporated into the final version. To complement the interviews, an anonymous online survey was conducted using a five-point Likert scale and three evaluation criteria: completeness, understandability, and usefulness. Results were strongly positive: 92% of respondents rated both completeness and understandability as "strongly agree," while usefulness received either "strongly agree" (58%) or "agree" (42%) ratings. No participant rated any criterion lower than "agree." Minor critique regarding generic terminology in some characteristics was addressed through selective refinements.

Finally, the taxonomy's practical relevance was demonstrated through its application in three real-world organizations – each representing a different maturity level in process mining adoption – where it successfully described diverse setups and supported structured reflection on design decisions. The cases revealed notable differences in organizational configurations, shaped by factors such as company size, industry, and available resources, emphasizing that there is no one-size-fits-all approach. Importantly, the taxonomy also captured how setups evolve over time, particularly as organizations scale their use of process mining across additional processes and subsidiaries. Finally, the cases highlighted recurring success factors, such as strong executive sponsorship and the presence of internal champions, which appeared beneficial regardless of the organization's maturity.

Research paper P1 contributes by not only advancing the descriptive knowledge of process mining governance but also laying a robust theoretical foundation for future research in this underexplored area. Developed with active involvement from process mining experts, the taxonomy systematically categorizes the key dimensions that define effective process mining integration within diverse organizational contexts. The taxonomy reveals significant variability in configurations, underscoring that conventional BPM governance approaches often fail to capture the complexities of process mining. Practically, it provides a holistic framework that guides practitioners in the initial setup and subsequent refining of their process mining setups

by clarifying critical decision-making factors. The exemplary cases further illustrate how organizational characteristics – such as size, industry, and maturity level – influence process mining setup choices, suggesting that future research should investigate the interrelationships among these dimensions to derive process mining archetypes and develop higher-level theories. Moreover, while the focus is on process mining, the taxonomy’s comprehensive structure may also be applicable to related fields such as data analytics broadening the impact of contributions.

In doing so, P1 responds to key challenges of process mining such as the lack of guidance in implementing an organizational setup that ensures effective use of process mining across varied contexts (cf. *Challenge 3 and 6*, Martin et al. 2021). Furthermore, the taxonomy provides an indication of how process mining expertise can be anchored within organizations (cf. *Challenge 10*, Martin et al. 2021), while the exemplary cases illustrate the underlying rationales for different anchoring choices across organizations at varying levels of process mining maturity. Research paper P1 advances process mining research primarily at the organizational and group levels and lays the foundation for embedding process mining into enterprise-wide practice. However, while it clarifies the general conditions under which process mining can thrive, it does not specify how to act within this framework. Specifically, it lacks information on the capabilities required to continuously generate business value from process mining. Hence, research paper P2 addresses this gap by exploring process mining as a source of process-based behavioral visibility that, if managed effectively, can be transformed into business value. Building on the structural foundation of P1, it identifies the capabilities needed to act within this framework and continuously generate value from process mining.

II.2 Capabilities for Managing Process-Based Behavioral Visibility

Almost all professional activities generate digital trace data (Leonardi and Treem 2020), which, when converted into event logs (Weber et al. 2024), yield process-based behavioral visibility. When effectively leveraged, process-based behavioral visibility offers new opportunities to manage organizations and create business value (Badakhshan et al. 2022; Vaujany et al. 2021; Leonardi and Treem 2020). However, to realize this potential, organizations must first develop capabilities to govern, interpret, and act upon the behavioral visibility derived from digital traces (vom Brocke et al. 2021). Current research on behavioral visibility has predominantly focused on its negative consequences such as digital surveillance, intrusive control practices, and employee resistance (Benlian et al. 2022; Newlands 2021; Spicer 2017; Zorina et al. 2021), neglecting the potential for business value if these challenges can be effectively managed (Badakhshan et al. 2022). Specifically, it remains unclear how organizations use and manage

the behavioral visibility created by process mining to generate business value. In other words, the strategic approaches to establishing and managing process-based behavioral visibility are still underexplored. In light of this gap, the central research question addressed by research paper P2 is: *How do organizations implement and manage process-based behavioral visibility to generate value?*

To answer this question, the study employs a grounded theory-based approach. Thirty in-depth, semi-structured interviews were conducted with process mining experts, including process analysts, senior managers, and unit heads, from a variety of industries and organizations. Purposive sampling ensured a diverse range of perspectives. All interviews were transcribed and analyzed, following an inductive, iterative coding process as outlined by Gioia et al. (2013) and Corbin and Strauss (2008). In this iterative process, first-order concepts were initially identified and later aggregated into second-order themes, which then formed aggregate capabilities and overarching themes. The resulting analysis led to the identification of three distinct groups of capabilities – foundational, transformational, and continual – that enable organizations to continuously convert data into behavioral visibility and business value. These relationships are illustrated in the overarching model in Figure 2.




Figure 2: The relationship between data, behavioral visibility, and business value as well as the necessary capabilities

Thereby, foundational capabilities serve as the starting point, as recorded trace data on work-related activities forms the basis for creating process-based behavioral visibility. These capabilities, as detailed in Table 3, encompass both technical elements (such as infrastructure) and people-related factors (such as organizational culture), emphasizing the need for a socio-technical approach. For example, “behavioral data modeling” involves ensuring that the data

accurately reflects real-world behavior – a task made challenging by fragmented systems and the presence of blind spots in the event log. Similarly, “behavioral data integration” highlights the need to stitch together disparate data sources into a cohesive view of behavior. Finally, “organizational structuring” includes not just governance mechanisms but also fostering a culture that encourages transparency while addressing privacy concerns.

Capability	Definition
Behavioral Data Modeling	<i>...refers to the capability of defining and using real-time data sources that represent relevant behavior in reliable, complete, and secure ways.</i>
Behavioral Data Integration	<i>...refers to the capability of integrating and centralizing behavioral fragments for the subsequent comprehensive analysis of work performances.</i>
Organizational Structuring	<i>...refers to the capability of designing and implementing a framework that seamlessly integrates socio-technical knowledge, promotes an empowering culture, and ensures alignment between managerial logic and behavioral visibility-based management.</i>

Table 3: Overview of foundational capabilities

Once a solid socio-technical foundation is in place, transformational capabilities are required to convert the established behavioral visibility into tangible business value. These capabilities, depicted in Table 4, involve aligning visible behavior with its real-world implications through “behavioral correspondence,” which includes helping teams reconcile system-based insights with their lived process knowledge. “Evidence-based management” refers to acting on insights, moving beyond awareness to continuous improvement loops and real-time interventions. Lastly, “strategic behavior mapping” ensures that behavioral KPIs are meaningfully linked to organizational goals while avoiding misaligned metrics that can unintentionally distort performance.

Capability	Definition
Behavioral Correspondence	<i>...refers to the capability of mapping and contextualizing visible behavior to corresponding instances in the physical world.</i>
Evidence-Based Management	<i>...refers to the capability of leveraging evidence-based insights for managerial actions.</i>
Strategic Behavior Mapping	<i>...refers to the capability of meaningfully translating strategic goals into behavioral visibility-based KPIs.</i>

Table 4: Overview of transformational capabilities

Finally, with foundational and transformational capabilities established, continual capabilities are essential to sustain and extend the benefits of behavioral visibility-based management efforts. Given that behavior is dynamic and evolves over time due to both internal shifts and managerial interventions, continual capabilities are crucial for ongoing adaptation and improvement. As outlined in Table 5, these include “opportunity recognition,” which ensures

organizations continually discover new areas where behavioral visibility can add value. “Dynamic behavioral mindset” reflects a culture of ongoing learning and adjustment, driven by new process insights. Finally, “ongoing commitment” emphasizes the long-term nature of value realization from behavioral visibility – requiring persistent stakeholder support, cultural acceptance, and continuous leadership buy-in.

Capability	Definition
Opportunity Recognition	<i>...refers to the capability of continuously perceiving opportunities for scaling and extending behavioral visibility-based management.</i>
Dynamic Behavioral Mindset	<i>...refers to the capability of continuously updating the organizational understanding of work performances.</i>
Ongoing Commitment	<i>...refers to the capability of using behavioral visibility as a sustained management effort.</i>

Table 5: Overview of continual capabilities

Research paper P2 contributes by advancing the understanding of how organizations can leverage process-based behavioral visibility to generate business value. By identifying and categorizing three essential sets of capabilities, it demonstrates that effective value realization from digital trace data requires not only technical infrastructure and data integration (foundational capabilities) but also the ability to interpret and align these insights with strategic objectives (transformational capabilities) and, importantly, to sustain and adapt these efforts over time (continual capabilities). This research contribution further shifts the narrative from the predominantly negative discourse on behavioral visibility – centered on surveillance and employee resistance – to one that highlights its strategic potential when managed properly. In doing so, the study provides a theoretical foundation for future research on process mining and the evolving role of managerial decision-making in a digitally enabled work environment.

Beyond these specific contributions, research paper P2 advances the broader research area of process mining governance by identifying the capabilities needed to act within the structural framework outlined by research paper P1 in a value-adding way. It complements the structural “where” of process mining integration with the organizational “how”: by explaining which capabilities are needed so organizations can translate digital trace data into behavioral visibility and ultimately into business value. Thereby, it contributes to process mining research at the individual, group, and organizational levels and responds to the challenges of addressing employees’ defensive reactions to transparency (cf. *Challenge 23*, Martin et al. 2021) and concerns over intrusive monitoring (cf. *Challenge 26*, Martin et al. 2021). Furthermore, it positions process mining as the backbone of a new management paradigm, allowing for its continued integration instead of being a one-off initiative (cf. *Challenge 27*, Martin et al. 2021).

III Process Data Quality Management

Process mining is a data-driven analytical approach, which is why the accuracy and reliability of its outcomes is inherently dependent on the quality of the input event logs (Hofstede et al. 2023). In practice, however, real-life event logs are rarely free of imperfections (Suriadi et al. 2017); they regularly suffer from various data quality issues, making data preparation a complex and tedious task that can consume up to 80% of process mining efforts (Wynn et al. 2022). If not properly addressed, these imperfections not only diminish the business value generated by process mining but also challenge the very feasibility of applying process mining techniques (Hofstede et al. 2023). Hence, it comes as no surprise that practitioners and researchers unanimously identify poor event log quality and complex data preparation as two of the three challenges both groups consider extremely relevant in process mining (Martin et al. 2021).

Among the most critical examples are timestamp-related imperfections as accurate timestamps are a prerequisite for many process mining use cases (Fischer et al. 2022). Hence, research Paper P3 (Section III.1) proposes a method for repairing identical timestamp errors with GANs. Similarly, many process mining techniques rely on accurately mapping events to specific process instances using high-quality case identifiers (van der Aalst 2022). Thus, research paper P4 (Section III.2) presents a method for repairing missing case identifiers combining a rule-based approach with a Transformer-based architecture, complemented by human-in-the-loop elements. In event logs where multiple imperfections are present, artifacts like those proposed in research papers P3 and P4 can be applied consecutively as part of a toolchain. While effective in some scenarios, this toolchain paradigm comes with its limitations, for instance, when real-world event logs feature interdependent imperfections where the order of repairs can introduce unintended side effects (Hofstede et al. 2023). Consequently, research paper P5 (Section III.3) proposes an LLM-based method for adaptable multi-imperfection event log repair.

III.1 Repairing Identical Timestamp Errors with Generative Adversarial Networks

Process mining depends on high-quality timestamps to determine the true sequence and temporal relationships of recorded activities (Fischer et al. 2022). However, various data quality issues can affect timestamp accuracy (Suriadi et al. 2017). For instance, identical timestamp errors occur when multiple events are erroneously associated with the same timestamp. This can result from form-based input systems that batch-submit multiple events at once, limited system time granularity that cannot distinguish events occurring in rapid succession, or system overloads that delay event logging (Conforti et al. 2020; Suriadi et al. 2017). If not properly

addressed, identical timestamp errors have severe consequences as consecutive events are misinterpreted as concurrent, leading to distorted process models, inaccurate performance measures, and ultimately, misleading insights for decision-making (Suriadi et al. 2017). While existing methods for identical timestamp repair (e.g., Conforti et al. 2020) have demonstrated effectiveness, they fall short in fully capitalizing on the long-term contextual information available in event logs. In that context, generative AI offers a promising alternative by leveraging advanced neural network architectures to capture complex temporal dependencies (Hochreiter and Schmidhuber 1997; Vaswani et al. 2017). More specifically, GANs create convincing samples of fake data based on random noise (Goodfellow et al. 2014), which allows them to repair data quality issues in other domains such as audio and image data (Hofmann et al. 2021). By adapting them for event log repair, it is possible to model the underlying distribution of correct timestamps more effectively. Consequently, the research question addressed by research paper P3 is: *How can identical timestamp errors in event logs be repaired based on Generative Adversarial Networks?*

The study follows the DSR paradigm as proposed by Peffers et al. (2007) to develop a method that integrates established techniques for error detection and event reordering with a novel GAN-based approach for timestamp repair. Thereby, the design and development phase was iterative, involving multiple cycles of design, evaluation, and refinement. Furthermore, the method was instantiated as a software prototype and used to repair six real-life and artificial event logs, each injected with identical timestamp errors by increments of 10%, ranging from 10% to 50%, encompassing a total of 30 variations. The effectiveness of repair was measured using metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), sequence-level accuracy, and Levenshtein distance.



Figure 3 Overview of the proposed method for repairing identical timestamp errors

An overview of the method is depicted in Figure 3. It begins with a preprocessing phase where the event logs are filtered to remove noise and to flag events affected by identical timestamp errors. Subsequently, key variables such as the time intervals between events are computed and normalized. Finally, the event reordering technique by Conforti et al. (2020) is applied. Having reordered and preprocessed the erroneous event log, the correct timestamps need to be estimated. Here, the core contribution lies in the application of a conditional GAN (cGAN)

architecture, which employs Long Short-Term Memory (LSTM) layers to capture long-term dependencies in the sequential data (Taymouri et al. 2020; Hochreiter and Schmidhuber 1997). In this framework, the generator produces plausible timestamps for the affected events, while the discriminator assesses the quality of these estimates by comparing them with correct data.




Figure 4: Architecture of the generator

As illustrated in Figure 4, the generator integrates three input layers, which are concatenated and passed through a stack of LSTM layers, ultimately producing a timestamp estimation for each erroneous input event.




Figure 5: Architecture of the discriminator

The discriminator, shown in Figure 5, mirrors this structure with two input layers and multiple LSTM and fully connected layers, resulting in a binary classification output that distinguishes ground truth, correct events from synthetic, repaired events generated by the generator. During

training, the generator and discriminator are optimized in alternating steps: the discriminator learns to improve its classification accuracy, while the generator adapts to produce outputs increasingly indistinguishable from correct events – meaning repaired events with plausible timestamps. Once training is complete, the discriminator is discarded, and the generator is used independently to estimate repaired timestamps for the affected events.

The method was evaluated along three criteria: effectiveness, generality, and ease of use, addressing both the reordering of erroneous events and the subsequent estimation of timestamps. In addition to the proposed method, the evaluation included several benchmark methods: the method by Conforti et al. (2020), a basic (and non-adversarial) LSTM model, a random sorting strategy, and a median-based timestamp estimation. In terms of effectiveness, the proposed method outperformed all benchmarks in estimating timestamps (cf. Figure 6) while for event reordering, it fell short of the performance achieved by the Conforti et al. (2020) method. As a result, the final method integrates the strengths of both approaches – leveraging the Conforti et al. (2020) technique for reordering and the GAN for timestamp estimation.




Figure 6: Comparison of effectiveness in estimating times between events

In terms of generality, the method was tested on a variety of artificial and real-world event logs with different error rates. Across all datasets, the method consistently delivered robust performance, demonstrating its effectiveness in diverse settings. Nonetheless, its applicability is constrained by its reliance on a specific event log format – requiring case identifiers, activity labels, and timestamps. Despite this limitation, the method is compatible with the widely adopted eXtensible Event Stream (XES) standard (Ghahfarokhi et al. 2021), ensuring relevance to most practical use cases. Regarding ease of use, the tool was designed with user-friendliness in mind. It features a high level of code abstraction, concealing technical complexities and

offering default parameters that require minimal user input. No specialized programming skills are needed beyond the use of standard libraries. Furthermore, the prototype is not very resource-hungry and can be run on low-end workstations. While certain limitations remain – such as typical GAN training challenges like mode collapse where the generator only outputs a single timestamp estimation over and over again – the method provides a practical and accessible solution for both research and practice.

In sum, research paper P3 makes a contribution to process data quality management by advancing both design and knowledge dimensions (Mendling et al. 2021). From a design perspective, the method constitutes a design improvement as it achieves state-of-the-art performance in timestamp estimation, thereby enhancing the accuracy of event log repairs compared to existing approaches. At the same time, it represents a design exaptation by adapting GANs – techniques previously successful in audio and image processing (Hofmann et al. 2021) – to the domain of event log repair. Furthermore, research paper P3 delivers explanatory knowledge contributions by rigorously documenting the iterative design activities, evaluation results, and the impact of specific design decisions on the evaluation criteria. Thus, through detailed interpretation of these findings, it contributes knowledge regarding the underlying mechanisms that affect repair performance. Finally, the evaluation also leads to performance propositions that describe the algorithm’s ability to meet its task requirements, showing that it outperforms state-of-the-art approaches in repairing identical timestamp errors. Additionally, the study contributes sensitivity propositions by analyzing how performance varies across event logs with different error rates, sizes, and complexities, offering insights into the robustness and applicability of the method across diverse real-world scenarios. Collectively, these contributions provide a robust foundation for future research on applying generative AI approaches to enhance process data quality, while simultaneously offering a software prototype for both practitioners and researchers.

Research paper P3 represents an essential first step toward a solution to the challenges of poor event log quality and complex data preparation (Martin et al. 2021). Thereby, it addresses erroneous timestamps as one of many attributes in event logs. Building on the knowledge this work contributed, subsequent research papers will focus on repairing additional event log attributes with generative AI.

III.2 Repairing Elusive Cases in Event Logs with Transformer Networks

A key requirement for process mining analysis is the accurate mapping of events to process instances using case identifiers (van der Aalst 2022). This mapping is essential for

reconstructing process flows and understanding the sequence of activities executed in each case (van der Aalst 2016). However, real-life event logs frequently suffer from missing case IDs, a problem known as the elusive case imperfection pattern (Suriadi et al. 2017). When an event log is affected by this imperfection, the inability to group events accurately into process instances leads to incorrect model discovery, inaccurate performance metrics, and ultimately, misleading conclusions (Suriadi et al. 2017). The impact of this issue can be severe, rendering process mining applications ineffective and preventing organizations from realizing the full value of their process data (Fischer et al. 2022; Tajima et al. 2023). Existing methods for elusive case repair (e.g., Bayomie et al. 2023) often depend on supplementary, well-structured data beyond the event log itself – data that is not always available in practical settings. Furthermore, they are designed to regroup the entire event log, preventing selective repair where only a subset of events is affected by the imperfection. This is particularly problematic in scenarios where the majority of events are correctly assigned, and only a small percentage require correction. In response to these challenges, generative AI offers a promising solution. Its ability to model complex data patterns and reconstruct missing or erroneous information (Hofmann et al. 2021) makes it well-suited for repairing elusive cases in event logs. Recent applications of generative AI in process data quality management – such as using Transformers for activity label repair – have demonstrated its capability to enhance data quality and support more effective analyses (Wu et al. 2024). Furthermore, process mining vendors and market research organizations anticipate that generative AI will play a significant role in future process data quality management (Reinkemeyer et al. 2023; Kerremans and Kerremans 2023). Against this background, research paper P4 explores the following research question: *How can generative AI be used to repair the elusive case imperfection pattern?*

The study follows the DSR paradigm as proposed by Peffers et al. (2007) to develop a method for repairing the elusive case imperfection pattern. The design process was iterative, consisting of multiple cycles of development and evaluation. Initially, the Transformer was selected as a foundational architecture due to its strengths in handling sequential data and long-range dependencies (Vaswani et al. 2017). In subsequent iterations, the method was enhanced by integrating domain-specific rules and a human-in-the-loop approach to capture contextual knowledge. Final refinements focused on optimizing performance through hyperparameter tuning. The resulting method, instantiated as a Python prototype, was demonstrated on three publicly available event logs, including both synthetic and real-life data. To evaluate the artifact, the Framework for Evaluation in Design Science (FEDS) by Venable et al. (2016) following the Technical Risk & Efficacy strategy was adopted. Thereby, multiple evaluation

episodes including expert interviews, controlled repair experiments with injected errors, and real-world applicability feedback based on a second round of expert interaction were conducted. The proposed artifact, HERE (Hybrid Elusive Case Repair Engine), consists of three components as depicted in Figure 7. These include data preprocessing, Transformer training, and event log repair.




Figure 7: Overview of HERE

During preprocessing, event logs are standardized, normalized, and structured into sequences of discrete and continuous attributes. Domain knowledge is incorporated as expert attributes (e.g., start activities, directly-follows relationships), which are encoded and merged with event data. All inputs are tokenized, whereby event attributes are converted into numerical IDs. In the subsequent training phase, the Transformer learns to predict case IDs from sequences of input attributes. Using multi-head attention, positional encoding, and embedding layers, the model captures dependencies across long-range event sequences (Vaswani et al. 2017). Once trained, the Transformer is used in an iterative repair process as depicted in Figure 8.




Figure 8: Repair routine of HERE

Thereby, each cycle combines model predictions with ex-ante and ex-post rule checks, applying domain-specific logic to validate or adjust the case assignments. Stakeholders can input rules

and review prediction confidence using probability thresholds. The repair process continues until all missing case identifiers are resolved or stakeholder-defined stopping criteria are met.

The evaluation of HERE was conducted in three episodes as recommended by FEDS' Technical Risk & Efficacy Strategy (Venable et al. 2016): artificial formative, artificial summative, and naturalistic summative evaluation. In the artificial formative evaluation – focused on generating insights and guiding design improvements – 11 expert interviews were conducted. The method was rated highly for understandability (6.5) and novelty (6.0), with particular appreciation for the combination of Transformer architecture and human-in-the-loop rule checks. Completeness (5.9) and applicability (5.3) were also positively evaluated, although interviewees suggested improvements such as clearer probability outputs and enhanced integration of domain knowledge. Concerns were noted regarding complexity for non-technical users and dependence on input data quality.

In the artificial summative evaluation – focused on assessing the technical performance – the method was benchmarked against a basic LSTM network, a random repair based on observed frequencies, and a fully random repair. Thereby, three event logs were injected with varying degrees of elusiveness (10–90%), resulting in the repair of 27 erroneous event logs in total. The effectiveness of repair was evaluated based on ten imperfection-specific metrics. The results, as depicted in Figure 9, demonstrate high effectiveness when sufficient training data is available. In scenarios with few training data, being equivalent to scenarios with high elusiveness levels, the integration of rule-checking mechanisms enhanced robustness and accuracy. Compared to existing methods, the approach is especially suitable when preserving existing case ID logic is important. For higher error rates, alternative approaches such as Bayomie et al. (2023) show better effectiveness. In the final naturalistic summative evaluation, 10 experts interacted with the prototype. Feedback highlighted usefulness and ease of use, though applicability scored lower due to concerns about runtime, integration, and output validation. Participants emphasized the need for visual validation features and tailored interfaces for different user roles. Generality was widely acknowledged, assuming XES-style input, although potential variability in performance across different process types was noted.




Figure 9: Effectiveness of HERE in repairing the hospital billing event log

Research paper P4 contributes to process data quality management by introducing a method for repairing the elusive case imperfection pattern. Thereby, several design and knowledge contributions are made (Mendling et al. 2021). First, this study represents a design improvement, enhancing the effectiveness of elusive case repair in scenarios with low error rates. Second, it can be understood as a design exaptation as the Transformer architecture was adapted to the domain of event log repair for the first time. Third, the study offers performance propositions by demonstrating that the method outperforms existing approaches in scenarios with large datasets and low levels of elusiveness. Fourth, it contributes sensitivity propositions by showing how internal parameter settings (e.g., confidence thresholds) and task assumptions (e.g., data volume and quality) affect the method’s robustness. Fifth, it provides explanatory propositions by offering insights into how design decisions – such as the integration of human expertise and fallback rules – influence the accuracy and completeness of the repair process.

Research paper P4 builds on the foundation established in research paper P3 by repairing another critical event log attribute with generative AI. With this contribution, two individual methods have now been developed that are highly effective in repairing timestamp and case identifier imperfections respectively. Hence, significant progress has been made towards

addressing the challenge of poor event log quality and complex data preparation (Martin et al. 2021). However, both methods rely on each other’s outputs – P3 assumes accurate case IDs, while P4 assumes correct timestamps. As a result, when both imperfections occur simultaneously, applying these methods sequentially in a toolchain leads to significantly reduced effectiveness. To address this limitation, the subsequent research focuses on multi-imperfection event log repair.

III.3 Multi-Imperfection Event Log Repair with LLMs

Real-life event logs suffer from various data quality issues, sometimes multiple imperfections simultaneously (Fischer et al. 2022; Hofstede et al. 2023). Yet, research in process data quality management has addressed event log repair by focusing on single, isolated imperfections, designing specialized methods for issues such as timestamp errors (e.g., Conforti et al. 2020), missing case identifiers (e.g., Bayomie et al. 2023), or incorrect activity labels (e.g., Sadeghianasl et al. 2024). These methods are usually highly effective within their narrow scope and can operate as standalone components in a sequential toolchain of event log preprocessing (Andrews et al. 2020).

However, approaching multi-imperfection event log repair with a toolchain entails several key limitations. First, interdependencies between imperfections and their respective repair methods can cause unintended side effects, as the sequence in which repairs are applied within the toolchain can significantly impact the quality of the results (Suriadi et al. 2017; Hofstede et al. 2023). Second, residual errors may propagate through the toolchain if not all imperfections are detected or resolved (Suriadi et al. 2017; Hofstede et al. 2023). Finally, most existing methods are designed for static environments. Evolving log formats such as object-centric event logs (Berti et al. 2024; Adams et al. 2022) or emerging data quality issues require extensive redesign of methods (Baier et al. 2020; Sato et al. 2021). Although some recent studies (e.g., Nguyen et al. 2019; Sim et al. 2021) have explored multi-attribute repair strategies, they typically address only missing attributes, overlooking more complex imperfection patterns like duplicated timestamps or inconsistent labels (Suriadi et al. 2017). Consequently, process data quality management lacks a comprehensive solution that can simultaneously detect and repair a broad range of interdependent event log imperfections in a unified, adaptable manner.

To overcome these challenges, LLMs emerge as a promising solution being increasingly recognized as powerful tools for data transformation (Zhang et al. 2024). Furthermore, they have a profound understanding of process mining and event logs being already used for process modelling (Kourani et al. 2024a, 2024b), analysis and optimization (Barbieri et al. 2025; Buss

et al. 2025; Lashkevich et al. 2024) as well as event abstraction for log enhancement (Brzychczy et al. 2025). Yet, the application of LLMs to event log repair remains underexplored. Consequently, the research question addressed by research paper P5 is: *How can LLMs be designed to support multi-imperfection event log repair?*

Following the DSR paradigm as proposed by Peffers et al. (2007), the artifact was developed through an iterative design process comprising four iterations, each aimed at improving a specific aspect: feasibility, effectiveness, robustness, and efficiency. In the first design iteration, feasibility was demonstrated by applying in-context learning with the DeepSeek-R1 LLM (DeepSeek-AI et al. 2025). While a transformed event log was generated, the repair quality was limited, often worsening imperfections. Hence, the second iteration progressed from in-context learning to fine-tuning to enhance repair quality. This included compiling a base log collection, an instructional dataset as well as the fine-tuning routine itself. While effectiveness partially improved, challenges arose as the model frequently engaged in infinite internal reasoning instead of producing a repaired event log. Hence, to improve robustness, the third iteration switched the foundational LLM to the LLaMA-3.1 instruct model and adapted its prompt format accordingly. This change significantly increased robustness as well as accurate repairs. Finally, in the fourth iteration, the artifact's efficiency was enhanced through hyperparameter tuning and optimized data preparation, cutting processing time per repair by half while improving overall repair effectiveness. For demonstration purposes, the artifact was instantiated as a Python prototype and used to repair eight publicly available event logs from various real-world domains. Each of these event logs was injected with up to five different imperfection patterns affecting case identifier, timestamp, and activity label attributes. Each imperfection pattern was injected with an error rate ranging from 10% to 100% increasing in 10% increments. For evaluation, the FEDS framework by Venable et al. (2016) was employed. Following the purely technical evaluation strategy, the effectiveness of diagnosing and repairing event logs was assessed using three diagnostic and four imperfection-specific repair metrics.

The proposed method consists of three core components: data preparation, LLM fine-tuning, and a repair routine. An overview of the method is depicted in Figure 10.




Figure 10: Overview of the proposed method for multi-imperfection event log repair

The data preparation phase ensures that event logs are structured and encoded in a way that enables effective LLM training and inference. First, a base log collection of real-life, publicly available event logs is curated. Logs are selected based on their diversity and absence of data quality issues, enabling controlled artificial injection of imperfections. To ensure generalizability, selected logs span multiple industries and processes. Next, the logs are reduced to three essential attributes: case ID, activity label, and timestamp. While omitting additional context attributes may limit repair effectiveness, this step is necessary to reduce the input size and ensure compatibility with the limited context windows of current LLMs. Standardization steps follow, including timestamp normalization and column renaming. Event logs are then partitioned into batches of up to 150 events, balancing context completeness with hardware constraints. Subsequently, batches are injected with random combinations of imperfections and error rates. From all imperfection patterns by Suriadi et al. (2017), five are selected – two affecting timestamps, two affecting activity labels, and one affecting case identifiers. Patterns unsuitable for artificial injection (e.g., those requiring domain knowledge) are excluded. This yields a comprehensive dataset for fine-tuning, with paired examples of erroneous input and correct output. Finally, the dataset is split into training and test data.

During fine-tuning the foundational LLaMA-3.1 8B instruct model is trained using LoRA (Low-Rank Adaptation) to enable efficient fine-tuning. Training prompts follow a structured three-role format: the system prompt sets the domain context of event log repair, the user prompt provides specific instructions and the erroneous log in CSV format, and the assistant prompt presents the desired response – organized into diagnosis, mitigation, and repaired log sections. Each batch is then tokenized with special markers to differentiate roles and sections, enabling the model to detect, explain, and repair multiple simultaneous imperfections.

In the final repair routine, the erroneous event logs are preprocessed and partitioned to match the data format used during training. Each batch is converted into a structured prompt with a

system and user message. Upon execution, the fine-tuned LLM generates a structured response comprising diagnosis, mitigation strategy, and repaired log segments. Finally, the repaired log is extracted and used to overwrite the corresponding erroneous batch. This repair routine is repeated for all batches, resulting in a fully repaired event log.

The evaluation focused on two aspects: the effectiveness of imperfection detection and the effectiveness of imperfection repair, using both accuracy-based and deviation-based measures. With imperfection detection, the fine-tuned LLM achieved almost perfect scores across all detection metrics – precision, recall, and F1-score – on all eight evaluation logs. Most importantly, this includes a log which was completely excluded from training, thereby validating the model’s ability to generalize to previously unseen event logs. The results indicate that the method is highly effective at detecting multiple types of imperfections simultaneously. However, a closer qualitative analysis revealed occasional discrepancies in the model’s estimates of error rates and the number of affected events. Thus, while the model reliably flags the presence of imperfections, its quantitative assessment should be treated with caution.

For repair effectiveness, the results for all accuracy-based metrics are displayed in Figure 11.




Figure 11: Effectiveness of repair with accuracy metrics

The method performed exceptionally well in repairing activity labels, reaching near-perfect accuracy even at high error rates. Repairing missing case identifiers was also highly accurate at

low and moderate error rates but declined sharply when the error rate exceeded 60 to 80%. At 100% error rate, the model failed to assign correct case IDs since there was no longer any indication in the prompt as to which case IDs can be assigned in the current 150 events. Timestamp repair showed mixed results. Some logs achieved high accuracy across all error rates, while others saw much lower performance. Accuracy remained stable within each log but varied between logs, indicating that effectiveness depends on log-specific characteristics. In the BPI Challenge 2020 Permit Log, for example, timestamp deviations were reduced by 400 hours on average. Despite this improvement, absolute deviation remained relatively high. Furthermore, in six of the eighty log error rate combinations, the LLM worsened the timestamp deviation compared to no-repair, reflecting the difficulty LLMs face with numerical precision.

In sum, research paper P5 contributes to process data quality management by advancing both the design and knowledge dimensions (Mendling et al. 2021). From a design perspective, the method constitutes a design improvement by enabling multi-imperfection event log repair within a single, adaptable framework, hence advancing the toolchain paradigm of event log repair. At the same time, it represents a design exaptation by fine-tuning LLMs – originally developed for natural language processing – to repair imperfections in event logs. The study also delivers explanatory knowledge contributions by documenting the design iterations, detailing evaluation outcomes, and analyzing how specific design decisions affected repair performance. Furthermore, the evaluation yields performance propositions by demonstrating the method’s ability to detect and repair a wide range of imperfections effectively. Sensitivity propositions complement this by revealing how repair performance varies with imperfections, error rates, and dataset characteristics, offering guidance for method deployment across diverse real-world scenarios. Together, these contributions advance the understanding of how LLMs can be adapted for effective and scalable event log repair.

While research paper P5 focuses on repairing multiple interdependent imperfections in event logs, it builds on and complements the prior contributions of research papers P3 and P4. Collectively, these three studies represent a significant step towards addressing the challenges of poor event log quality and complex data preparation, which are identified as relevant by both practitioners and scholars (Martin et al. 2021).

IV Conclusion

IV.1 Summary

Process mining enables organizations to derive valuable insights from event logs, transforming their understanding of business processes from idealized process models to evidence-based “as-is” and “to-be” perspectives (van der Aalst 2016; van der Aalst 2022). As organizations increasingly seek to optimize their processes, predict outcomes, and intervene in real-time, process mining is becoming foundational to modern BPM (Dumas et al. 2018; van der Aalst 2016). This shift is not only evident in the academic momentum behind the discipline (van der Aalst 2020), but also in its widespread industry adoption and rapidly growing software market (Kerremans et al. 2024; Deloitte 2025).

Historically, research in process mining has had a predominantly technical focus, concentrating on algorithm engineering for analysis tasks (van der Aalst 2022; vom Brocke et al. 2021). While this has led to a strong technological foundation, the implications of process mining at the individual, group, organizational, and ecosystem levels remain comparatively underexplored (vom Brocke et al. 2021). Furthermore, even within the technical level, research remains uneven: while analysis algorithms are comparatively mature, foundational issues such as event log quality have not received adequate attention (Suriadi et al. 2017; Hofstede et al. 2023). This imbalance is also reflected in practice where organizations report multiple challenges in realizing the full potential of process mining (Martin et al. 2021). In particular, both research and practice attest to the areas of *process mining governance* and *process data quality management* being critical yet underrepresented in current research (Martin et al. 2021; vom Brocke et al. 2021; Hofstede et al. 2023).

Against this backdrop, the purpose of this dissertation is to advance the *organizational and technical prerequisites of process mining*. To advance the field of process mining governance, research paper P1 introduces a taxonomy of organizational process mining setups, offering a framework to guide the initial setup and subsequent refinement of governance configurations. In order to create business value within this organizational framework, research paper P2 develops a capability framework for managing process-based behavioral visibility. Thereby, it identifies the capabilities required to generate business value from process transparency while mitigating potential drawbacks. To advance the field of process data quality management, research paper P3 presents a GAN-based method for repairing identical timestamp errors, achieving state-of-the-art performance in timestamp estimation. Building on this foundation and extending the use of generative AI for event log repair, research paper P4 proposes a hybrid

approach for addressing missing case identifiers by combining rule-based logic with a Transformer network and human-in-the-loop elements. Finally, research paper P5 brings these two strands together by introducing an LLM-based method for multi-imperfection event log repair. In doing so, it proposes a unified and adaptable framework that advances the toolchain paradigm of event log repair in process data quality management.

On the technical level, the dissertation advances process data quality management through three novel artifacts for diagnosing and repairing event log imperfections. These contributions directly address the persistent challenge of poor event log quality (cf. *Challenge 7*, Martin et al. 2021) and, through their human-in-the-loop elements, also span the technical and individual levels and address the challenge of complex data preparation (cf. *Challenge 12*, Martin et al. 2021). At the individual, group, and organizational levels, the dissertation introduces a capability framework for managing process-based behavioral visibility. Thereby, it contributes strategies to realize business value from behavioral visibility while mitigating unintended side effects such as defensive reactions to transparency (cf. *Challenge 23*, Martin et al. 2021) and concerns over intrusive monitoring (cf. *Challenge 26*, Martin et al. 2021). By positioning process-based behavioral visibility as a new management paradigm with process mining as its technological backbone, the dissertation contributes to moving process mining beyond one-off initiatives towards continuous organizational integration (cf. *Challenge 27*, Martin et al. 2021). Finally, at the organizational level, the proposed taxonomy of process mining setups provides guidance in key governance decisions. This contribution responds to the challenges of missing implementation guidance (cf. *Challenge 6*, Martin et al. 2021) and the unclear organizational anchoring of process mining expertise (cf. *Challenge 10*, Martin et al. 2021). Collectively, these five papers contribute to establishing the organizational and technical prerequisites of process mining and represent a step toward balancing the discipline’s robust technological foundation with its social and managerial dimensions.

IV.2 Limitations and Future Research

Despite its contributions spanning technical and organizational dimensions, this dissertation has overarching limitations that go beyond those discussed in the individual research papers. At the same time, these limitations point to promising avenues for future research.

First, the dissertation does not address the ecosystem level of the five-level framework for research on process mining (vom Brocke et al. 2021). While contributions were made at the technical, individual, group, and organizational levels, the inter-organizational dynamics – such as data sharing, cross-company governance, and ecosystem-wide coordination – remain

unexplored. As process mining adoption grows across supply chains and industry networks (Reinkemeyer 2020), understanding the ecosystem level becomes increasingly relevant. Therefore, future research could extend the proposed frameworks to the ecosystem level, examining how interoperability, trust, and collaborative value creation can be enabled through shared process insights and distributed governance mechanisms.

Second, while this dissertation addresses several key challenges in process mining, it does not offer exhaustive coverage of the field's open issues. Specifically, it contributes to 7 of the 32 challenges proposed by Martin et al. (2021), yet the remaining challenges are equally important and require further attention. Moreover, the addressed challenges benefit from further investigation through different theoretical lenses, methodological approaches, or stakeholder perspectives. For example, the challenges surrounding behavioral visibility may manifest differently in international and cross-cultural environments. Hence, future research should aim to both broaden the scope of inquiry and deepen the existing understanding by incorporating more contextual and interdisciplinary perspectives.

Third, this dissertation follows the DSR paradigm, which emphasizes the design and iterative evaluation of purposeful artifacts. While this approach ensures scientific rigor and practical relevance (Hevner et al. 2004), it also introduces certain methodological limitations. In particular, most evaluations were conducted in controlled or semi-controlled settings (e.g., synthetic data, expert interviews). When viewed through the lens of the FEDS framework (Venable et al. 2016), these evaluation episodes primarily fall under the category of artificial evaluation. As such, the dissertation provides limited insight into how the artifacts behave under the complex, unpredictable conditions of real-world organizations. Consequently, further naturalistic evaluation episodes are needed to better understand how the contributions hold in real-world environments. Future research should focus on how the proposed artifacts are adopted by real users, embedded in real systems, under real organizational constraints.

Fourth, while this dissertation demonstrates that generative AI-based methods are highly effective in event log repair, it does not provide insight into how these models arrive at their predictions or what potential implications this may have in real-world applications. The black-box nature of models like GANs, Transformers, and LLMs limits our understanding of their internal reasoning, which is especially critical in contexts where repaired event logs directly influence business decisions. For example, if an LLM-based repair model exhibits bias towards certain user groups – such as systematically favoring male employees – this could lead to a structural disadvantage in performance analysis. Similarly, the risk of reinforcing inaccurate or

unfair process interpretations through automated repair raises concerns around accountability and auditability. Future research should therefore focus on enhancing the transparency, explainability, and fairness of generative AI methods in process mining. This could include applying bias detection to event log repair or designing explainable-by-design architectures that enable users to understand and validate AI-driven repair – especially in business-critical or regulated environments where the integrity of process data is paramount.

Beyond the research avenues arising directly from the identified limitations, this dissertation also encourages broader reflection on the future development of the two research areas it addresses. For instance, as process mining continues to mature and become embedded in enterprise operations, governance emerges as a critical enabler of long-term value creation. Accordingly, current efforts to establish process mining governance are essential for securing its role as a foundational technology in organizations that strive for continuous improvement and evidence-based management practices. This dissertation has shown that existing governance approaches – such as BPM governance – are insufficient to address the unique requirements of process mining. Similar governance discourse is unfolding around other data-driven technologies such as AI (Mäntymäki et al. 2022) or RPA (Smeets et al. 2021), where it becomes evident that traditional IT governance frameworks fall short in managing the specific challenges posed by these technologies. For now, establishing technology-specific governance remains a pragmatic approach to rapidly scale and consolidate these emerging technologies within organizations. Yet, these technologies increasingly converge in practice, sharing common requirements such as high-quality data or translator functions that bridge business and technology. Thus, maintaining isolated governance models beyond the consolidation of these technologies risks redundancy and misses opportunities for synergy. Consequently, future research should explore what lies beyond technology-specific frameworks by investigating integrated advanced analytics governance. Hub-and-spoke models, for instance, could centralize core services like data management or solution portfolios, while satellite units provide domain-specific expertise. This shift towards integrated governance may represent the next evolutionary stage of process mining governance – enabling organizations to balance standardization with flexibility while orchestrating cohesive, efficient, and enterprise-wide analytics capabilities.

Similar to dedicated governance models, event log repair currently represents a pragmatic response to the challenges of process data quality management, ensuring the feasibility and effectiveness of process mining applications. However, from a long-term perspective, repairing imperfections after they occur should be seen as just one component of a broader data quality

strategy. Ideally, process data quality management should address the root causes of imperfections and prevent quality issues from arising in the first place. To that end, future research should move toward proactive and embedded approaches to data quality management. This includes investigating the design of monitoring systems that continuously assess event log quality in real time, as well as self-healing data pipelines that integrate validation and repair mechanisms during data generation. Building on the idea of advanced analytics governance, future work should also explore how centralized data infrastructures can be designed to ensure consistent, high-quality process data across the enterprise. Together, these efforts point toward a more comprehensive approach to process data quality management – one in which event log repair becomes a supporting element of a scalable, trustworthy, and value-generating analytics architecture.

V References

- Aaltonen A, Stelmaszak M (2024) The Performative Production of Trace Data in Knowledge Work. *Information Systems Research* 35(3):1448–1462. doi:10.1287/isre.2019.0357.
- Adams JN, Schuster D, Schmitz S, Schuh G, van der Aalst WM (2022) Defining Cases and Variants for Object-Centric Event Data. In: 2022 4th International Conference on Process Mining (ICPM). IEEE, pp 128–135.
- Allen L, O’Connell A, Kiermer V (2019) How can we ensure visibility and diversity in research contributions? How the Contributor Role Taxonomy (CRediT) is helping the shift from authorship to contributorship. *Learned Publishing* 32(1):71–74. doi:10.1002/leap.1210.
- Andrews R, Suriadi S, Ouyang C, Poppe E (2018) Towards Event Log Querying for Data Quality: Let’s Start with Detecting Log Imperfections. In: Panetto H, Debruyne C, Proper HA, Ardagna CA, Roman D, Meersman R (eds) *On the Move to Meaningful Internet Systems. OTM 2018 Conferences. Confederated International Conferences: CoopIS, C&TC, and ODBASE 2018*, Valletta, Malta, October 22–26, 2018, Proceedings, Part I. Springer, Cham, pp 116–134.
- Andrews R, van Dun C, Wynn MT, Kratsch W, Röglinger M, Hofstede A ter (2020) Quality-informed semi-automated event log generation for process mining. *Decision Support Systems* 132:113265. doi:10.1016/j.dss.2020.113265.
- Badakhshan P, Wurm B, Grisold T, Geyer-Klingenberg J, Mendling J, vom Brocke J (2022) Creating business value with process mining. *The Journal of Strategic Information Systems* 31(4):101745. doi:10.1016/j.jsis.2022.101745.
- Baier L, Reimold J, Kühl N (2020) Handling Concept Drift for Predictions in Business Process Mining. In: Aier S, Gordijn J, Proper HA, Verelst J (eds) *2020 IEEE 22nd Conference on Business Informatics. CBI 2020*, pp 76–83.
- Barbieri L, Stroeh K, Madeira E, van der Aalst WMP (2025) An LLM-based Q&A Natural Language Interface to Process Mining. In: Delgado A, Slaats T (eds) *Process Mining Workshops. ICPM 2024 International Workshops*, Lyngby, Denmark, October 14–18, 2024, Revised Selected Papers.
- Bayomie D, Di Ciccio C, Mendling J (2023) Event-case correlation for process mining using probabilistic optimization. *Information Systems* 114:102167. doi:10.1016/j.is.2023.102167.
- Benlian A, Wiener M, Cram WA, Krasnova H, Maedche A, Möhlmann M, Recker J, Remus U (2022) Algorithmic Management. *Business & Information Systems Engineering* 64(6):825–839. doi:10.1007/s12599-022-00764-w.
- Bernstein ES, Li S (2017) Seeing Where You Stand: From Performance Feedback to Performance Transparency. *Academy of Management Proceedings* 2017(1):14752. doi:10.5465/AMBPP.2017.14752abstract.
- Berti A, Koren I, Adams JN, Park G, Knopp B, Graves N, Rafiei M, Liß L, Unterberg LTG, Zhang Y, Schwanen C, Pegoraro M, van der Aalst WMP (2024) OCEL (Object-Centric Event Log) 2.0 Specification.
- Brzychczy E, Kluza K, Szała L (2025) Enhancement of Low-Level Event Abstraction with Large Language Models (LLMs). In: Gdowska K, Gómez-López MT, Rehse J-R (eds) *Business Process Management Workshops. BPM 2024 International Workshops*, Krakow, Poland, September 1–6, 2024, Revised Selected Papers, pp 209–220.
- Buss A, Kratsch W, Schmid SJ, Wang H (2025) ProcessLLM: A Large Language Model Specialized in the Interpretation, Analysis, and Optimization of Business Processes. In: Gdowska K, Gómez-López MT, Rehse J-R (eds) *Business Process Management Workshops. BPM 2024 International Workshops*, Krakow, Poland, September 1–6, 2024, Revised Selected Papers, pp 221–232.

- Ceravolo P, Comuzzi M, De Weerd J, Di Francescomarino C, Maggi FM (2024) Predictive process monitoring: concepts, challenges, and future research directions. *Process Science* 1(1). doi:10.1007/s44311-024-00002-4.
- Cherns A (1976) The Principles of Sociotechnical Design. *Human Relations* 29(8):783–792. doi:10.1177/001872677602900806.
- Conforti R, Rosa M, Hofstede AT, Augusto A (2020) Automatic Repair of Same-Timestamp Errors in Business Process Event Logs. *Business Process Management: 18th International Conference*. <https://www.semanticscholar.org/paper/Automatic-Repair-of-Same-Timestamp-Errors-in-Event-Conforti-Rosa/3c0f50273bbbfd7ac8eb9ada86cd0ddc8a95f2a>.
- Corbin J, Strauss A (2008) *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Sage Publications, Los Angeles, London, New Delhi.
- DeepSeek-AI, Guo D, Yang D, Zhang H, Song J, Zhang R, Xu R, Zhu Q, Ma S, Wang P, Bi X, Zhang X, Yu X, Wu Y, Wu ZF, Gou Z, Shao Z, Li Z, Gao Z, Liu A, Xue B, Wang B, Wu B, Feng B, Lu C, Zhao C, Deng C, Zhang C, Ruan C, Dai D, Chen D, Ji D, Li E, Lin F, Dai F, Luo F, Hao G, Chen G, Li G, Zhang H, Bao H, Xu H, Wang H, Ding H, Xin H, Gao H, Qu H, Li H, Guo J, Li J, Wang J, Chen J, Yuan J, Qiu J, Li J, Cai JL, Ni J, Liang J, Chen J, Dong K, Hu K, Gao K, Guan K, Huang K, Yu K, Wang L, Zhang L, Zhao L, Wang L, Zhang L, Xu L, Xia L, Zhang M, Zhang M, Tang M, Li M, Wang M, Li M, Tian N, Huang P, Zhang P, Wang Q, Chen Q, Du Qiushi, Ge R, Zhang R, Pan R, Wang R, Chen RJ, Jin RL, Chen R, Lu S, Zhou S, Chen S, Ye S, Wang S, Yu S, Zhou S, Pan S, Li SS, Zhou S, Wu S, Yun T, Pei T, Sun T, Wang T, Zeng W, Zhao W, Liu W, Liang W, Gao W, Yu W, Zhang W, Xiao WL, an Wei, Liu X, Wang X, Chen X, Nie X, Cheng X, Liu X, Xie X, Liu X, Yang X, Li X, Su X, Lin X, Li XQ, Jin X, Shen X, Chen X, Sun X, Wang X, Song X, Zhou X, Wang X, Shan X, Li YK, Wang YQ, Wei YX, Zhang Y, Xu Y, Li Y, Zhao Y, Sun Y, Wang Y, Yu Y, Zhang Y, Shi Y, Xiong Y, He Y, Piao Y, Wang Y, Tan Y, Ma Y, Liu Y, Guo Y, Ou Y, Wang Y, Gong Y, Zou Y, He Y, Xiong Y, Luo Y, You Y, Liu Y, Zhou Y, Zhu YX, Huang Y, Li Y, Zheng Y, Zhu Y, Ma Y, Tang Y, Zha Y, Yan Y, Ren ZZ, Ren Z, Sha Z, Fu Z, Xu Z, Xie Z, Zhang Z, Hao Z, Ma Z, Yan Z, Wu Z, Gu Z, Zhu Z, Liu Z, Li Z, Xie Z, Song Z, Pan Z, Huang Z, Xu Z, Zhang Z, Zhang Z (2025) DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning.
- Deloitte (2025) *Deloitte Global Process Mining Survey 2025*.
- Di Francescomarino C, Ghidini C (2022) Predictive Process Monitoring. In: van der Aalst WMP, Carmona J (eds) *Process Mining Handbook*. Springer International Publishing, Cham, pp 320–346.
- Dumas M, La Rosa M, Mendling J, Reijers HA (2018) *Fundamentals of Business Process Management*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Eggers J, Hein A, Böhm M, Krcmar H (2021) No Longer Out of Sight, No Longer Out of Mind? How Organizations Engage with Process Mining-Induced Transparency to Achieve Increased Process Awareness. *Business & Information Systems Engineering* 63(5):491–510. doi:10.1007/s12599-021-00715-x.
- Fischer DA, Goel K, Andrews R, van Dun C, Wynn MT, Röglinger M (2022) Towards interactive event log forensics: Detecting and quantifying timestamp imperfections. *Information Systems* 109:102039. doi:10.1016/j.is.2022.102039.
- Franzoi S, Grisold T, vom Brocke J (2023) STUDYING DYNAMICS AND CHANGE WITH DIGITAL TRACE DATA: A SYSTEMATIC LITERATURE REVIEW. In: *ECIS 2023 Proceedings*.
- Ghahfarokhi AF, Park G, Berti A, van der Aalst WMP (2021) OCEL: A Standard for Object-Centric Event Logs. In: *European Conference on Advances in Databases and Information Systems*, pp 169–175.
- Gioia DA, Corley KG, Hamilton AL (2013) Seeking Qualitative Rigor in Inductive Research. *Organizational Research Methods* 16(1):15–31. doi:10.1177/1094428112452151.

-
- Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative Adversarial Networks. *Advances in Neural Information Processing Systems* 27 (NIPS 2014). <https://arxiv.org/pdf/1406.2661>.
- Grisold T, Mendling J, Otto M, vom Brocke J (2021) Adoption, use and management of process mining in practice. *Business Process Management Journal* 27(2):369–387. doi:10.1108/bpmj-03-2020-0112.
- Grisold T, Seidel S, Heck M, Berente N (2024) Digital Surveillance in Organizations. *Business & Information Systems Engineering* 66(3):401–410. doi:10.1007/s12599-024-00866-7.
- Hevner, March, Park, Ram (2004) Design Science in Information Systems Research. *MIS Quarterly* 28(1):75. doi:10.2307/25148625.
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Computation* 9(8):1735–1780. doi:10.1162/neco.1997.9.8.1735.
- Hofmann P, Rückel T, Urbach N (2021) Innovating with Artificial Intelligence: Capturing the Constructive Functional Capabilities of Deep Generative Learning. In: Bui T (ed) *Proceedings of the 54th Hawaii International Conference on System Sciences*. Hawaii International Conference on System Sciences.
- Hofstede AHM ter, Koschmider A, Marrella A, Andrews R, Fischer DA, Sadeghianasl S, Wynn MT, Comuzzi M, Weerdt J de, Goel K, Martin N, Soffer P (2023) Process-Data Quality: The True Frontier of Process Mining. *Journal of Data and Information Quality* 15(3):1–21. doi:10.1145/3613247.
- Kerpedzhiev GD, König UM, Röglinger M, Rosemann M (2021) An Exploration into Future Business Process Management Capabilities in View of Digitalization. *Business & Information Systems Engineering* 63(2):83–96.
- Kerremans I, Kerremans M (2023) The Impact of Generative AI on Process Mining.
- Kerremans M, Sugden D, Duffy N (2024) Gartner Magic Quadrant for Process Mining Platforms.
- Kourani H, Berti A, Schuster D, van der Aalst WMP (2024a) Process Modeling with Large Language Models. In: van der Aa H, Bork D, Schmidt R, Sturm A (eds) *Enterprise, Business-Process and Information Systems Modeling. 25th International Conference, BPMDS 2024, and 29th International Conference, EMMSAD 2024, Limassol, Cyprus, June 3–4, 2024, Proceedings*, pp 229–244.
- Kourani H, Berti A, Schuster D, van der Aalst WMP (2024b) ProMoAI: Process Modeling with Generative AI. In: Larson K (ed) *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pp 8708–8712.
- Kundisch D, Muntermann J, Oberländer AM, Rau D, Röglinger M, Schoormann T, Szopinski D (2021) An Update for Taxonomy Designers. *Business & Information Systems Engineering* 64(4):421–439.
- Lashkevich K, Milani F, Avramenko M, Dumas M (2024) LLM-Assisted Optimization of Waiting Time in Business Processes: A Prompting Method. In: Marrella A, Resinas M, Jans M, Rosemann M (eds) *Business Process Management. 22nd International Conference, BPM 2024, Krakow, Poland, September 1–6, 2024, Proceedings*, pp 474–492.
- Leemans SJJ, Fahland D, van der Aalst WMP (2014) Discovering Block-Structured Process Models from Event Logs Containing Infrequent Behaviour. In: Lohmann N, Song M, Wohe P (eds) *Business Process Management Workshops*. Springer International Publishing, Cham, pp 66–78.

- Leonardi PM, Treem JW (2020) Behavioral Visibility: A new paradigm for organization studies in the age of digitization, digitalization, and datafication. *Organization Studies* 41(12):1601–1625. doi:10.1177/0170840620970728.
- Mannhardt F (2022) Responsible Process Mining. In: van der Aalst WMP, Carmona J (eds) *Process Mining Handbook*. Springer International Publishing, Cham, pp 373–401.
- Mäntymäki M, Minkkinen M, Birkstedt T, Viljanen M (2022) Defining organizational AI governance. *AI and Ethics* 2(4):603–609. doi:10.1007/s43681-022-00143-x.
- Martin N, Fischer DA, Kerpedzhiev GD, Goel K, Leemans SJJ, Röglinger M, van der Aalst WMP, Dumas M, La Rosa M, Wynn MT (2021) Opportunities and Challenges for Process Mining in Organizations: Results of a Delphi Study. *Business & Information Systems Engineering* 63(5):511–527. doi:10.1007/s12599-021-00720-0.
- Mendling J, Depaire B, Leopold H (2021) *Theory and Practice of Algorithm Engineering*.
- Newlands G (2021) Algorithmic Surveillance in the Gig Economy: The Organization of Work through Lefebvrian Conceived Space. *Organization Studies* 42(5):719–737. doi:10.1177/0170840620937900.
- Nguyen HTC, Lee S, Kim J, Ko J, Comuzzi M (2019) Autoencoders for improving quality of process event logs. *Expert Systems with Applications* 131:132–147. doi:10.1016/j.eswa.2019.04.052.
- Nickerson RC, Varshney U, Muntermann J (2013) A method for taxonomy development and its application in information systems. *European Journal of Information Systems* 22(3):336–359.
- Peffer K, Tuunanen T, Gengler CE, Rossi M, Hui W, Virtanen V, Bragge J (2007) Design Science Research Process: A Model for Producing and Presenting Information Systems Research. in the *Proceedings of the First International Conference on Design Science Research in Information Systems and Technology*, pp. 83-16. <https://arxiv.org/pdf/2006.02763>.
- Reinkemeyer L (ed) (2020) *Process Mining in Action*. Springer International Publishing, Cham.
- Reinkemeyer L, Grindemann P, Egli V, Röglinger M, Marcus L, Fabri L (2022) *Accelerating Business Transformation With Process Mining Centers of Excellence (CoEs)*.
- Reinkemeyer L, Röglinger M, Kratsch W, Fabri L, Schmid S, Wittmann J (2023) *Exploring the Interplay of Process Mining and Generative AI: Research and Recommendations for CoEs*.
- Rojas E, Munoz-Gama J, Sepúlveda M, Capurro D (2016) Process mining in healthcare: A literature review. *Journal of Biomedical Informatics* 61:224–236.
- Rosemann M, vom Brocke J (2015) The Six Core Elements of Business Process Management. In: vom Brocke J, Rosemann M (eds) *Handbook on Business Process Management 1*. Springer, pp 105–122.
- Sadeghianasl S, Hofstede AH ter, Wynn MT, Türkay S (2024) Humans-in-the-loop: Gamifying activity label repair in process event logs. *Engineering Applications of Artificial Intelligence* 132:107875. doi:10.1016/j.engappai.2024.107875.
- Sato DMV, Freitas SC de, Barddal JP, Scalabrin EE (2021) A Survey on Concept Drift in Process Mining. *ACM Computing Surveys* 54(9):1–38. doi:10.1145/3472752.
- Sim S, Bae H, Liu L (2021) Bagging Recurrent Event Imputation for Repair of Imperfect Event Log with Missing Categorical Events. *IEEE Transactions on Services Computing* 16(1):108–121. doi:10.1109/TSC.2021.3118381.
- Smeets M, Erhard R, Kaußler T (2021) Introduction of RPA Governance. In: Smeets M, Erhard R, Kaußler T (eds) *Robotic Process Automation (RPA) in the Financial Sector*. Springer Fachmedien Wiesbaden, Wiesbaden, pp 99–117.

- Spicer A (2017) Surveillance used to be a bad thing. Now, we happily let our employers spy on us. <https://www.theguardian.com/commentisfree/2017/aug/04/surveillance-employers-spy-implanted-chipped>. Accessed 2025-03-22.
- Suriadi S, Andrews R, Hofstede A ter, Wynn MT (2017) Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs. *Information Systems* 64:132–150. doi:10.1016/j.is.2016.07.011.
- Tajima K, Du B, Narusue Y, Saito S, Iimura Y, Morikawa H (2023) Step-by-Step Case ID Identification Based on Activity Connection for Cross-Organizational Process Mining. *IEEE Access* 11:60578–60589. doi:10.1109/ACCESS.2023.3284053.
- Taymouri F, La Rosa M, Erfani S, Bozorgi ZD, Verenich I (2020) Predictive Business Process Monitoring via Generative Adversarial Nets: The Case of Next Event Prediction. In: Fahland D, Ghidini C, Becker J, Dumas M (eds) *Business Process Management. 18th International Conference, BPM 2020, Seville, Spain, September 13–18, 2020, Proceedings*, pp 237–256.
- van der Aalst W (2016) *Process Mining: Data Science in Action*. Springer, Berlin, Heidelberg.
- van der Aalst W (2020) Academic View: Development of the Process Mining Discipline. In: Reinkemeyer L (ed) *Process Mining in Action*. Springer International Publishing, Cham, pp 181–196.
- van der Aalst W, Adriansyah A, Medeiros AKA de, Arcieri F, Baier T, Blickle T, Bose JC, van den Brand P, Brandtjen R, Buijs J, Burattin A, Carmona J, Castellanos M, Claes J, Cook J, Costantini N, Curbera F, Damiani E, Leoni M de, Delias P, van Dongen BF, Dumas M, Dustdar S, Fahland D, Ferreira DR, Gaaloul W, van Geffen F, Goel S, Günther C, Guzzo A, Harmon P, Hofstede A ter, Hoogland J, Ingvaldsen JE, Kato K, Kuhn R, Kumar A, La Rosa M, Maggi F, Malerba D, Mans RS, Manuel A, McCreesh M, Mello P, Mendling J, Montali M, Motahari-Nezhad HR, zur Muehlen M, Munoz-Gama J, Pontieri L, Ribeiro J, Rozinat A, Seguel Pérez H, Seguel Pérez R, Sepúlveda M, Sinur J, Soffer P, Song M, Sperduti A, Stilo G, Stoel C, Swenson K, Talamo M, Tan W, Turner C, Vanthienen J, Varvaressos G, Verbeek E, Verdonk M, Vigo R, Wang J, Weber B, Weidlich M, Weijters T, Wen L, Westergaard M, Wynn M (2012) *Process Mining Manifesto*. In: Daniel F, Barkaoui K, Dustdar S (eds) *Business Process Management Workshops. BPM 2011 International Workshops, Clermont-Ferrand, France, August 29, 2011, Revised Selected Papers, Part I*. Springer, Berlin, Heidelberg, pp 169–194.
- van der Aalst W, Weijters T, Maruster L (2004) Workflow mining: discovering process models from event logs. *IEEE Transactions on Knowledge and Data Engineering* 16(9):1128–1142. doi:10.1109/TKDE.2004.47.
- van der Aalst WMP (2022) *Process Mining: A 360 Degree Overview*. In: van der Aalst WMP, Carmona J (eds) *Process Mining Handbook*. Springer International Publishing, Cham, pp 3–34.
- van Dongen B (2012) *BPI Challenge 2012*. Eindhoven University of Technology. https://data.4tu.nl/articles/_/12689204/1.
- van Dongen B (2019) *BPI Challenge 2019*. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12715853/1.
- van Dongen B (2020a) *BPI Challenge 2020: Domestic Declarations*. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12692543/1.
- van Dongen B (2020b) *BPI Challenge 2020: International Declarations*. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12687374/1.
- van Dongen B (2020c) *BPI Challenge 2020: Prepaid Travel Costs*. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12696722/1.

-
- van Dongen B (2020d) BPI Challenge 2020: Request For Payment. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12706886/1.
- van Dongen B (2020e) BPI Challenge 2020: Travel Permit Data. 4TU.Centre for Research Data. https://data.4tu.nl/articles/_/12718178/1.
- van Dongen B, Borchert F (2018) BPI Challenge 2018. Eindhoven University of Technology. https://data.4tu.nl/articles/_/12688355/1.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I (2017) Attention Is All You Need.
- Vaujany F-X de, Leclercq-Vandelannoitte A, Munro I, Nama Y, Holt R (2021) Control and Surveillance in Work Practice: Cultivating Paradox in ‘New’ Modes of Organizing. *Organization Studies* 42(5):675–695. doi:10.1177/01708406211010988.
- Venable J, Pries-Heje J, Baskerville R (2016) FEDS: a Framework for Evaluation in Design Science Research. *European Journal of Information Systems* 25(1):77–89. doi:10.1057/ejis.2014.36.
- vom Brocke J, Jans M, Mendling J, Reijers HA (2021) A Five-Level Framework for Research on Process Mining. *Business & Information Systems Engineering* 63(5):483–490.
- Weber B, Abbad-Andaloussi A, Franceschetti M, Seiger R, Völzer H, Zerbato F (2024) Leveraging Digital Trace Data to Investigate and Support Human-Centered Work Processes. In: Kaindl H, Mannion M, Maciaszek LA (eds) *Evaluation of Novel Approaches to Software Engineering*. Springer Nature Switzerland, Cham, pp 1–23.
- Weinzierl S, Dunzer S, Zilker S, Matzner M (2020) Prescriptive Business Process Monitoring for Recommending Next Best Actions. In: Fahland D, Ghidini C, Becker J, Dumas M (eds) *Business Process Management Forum*. Springer International Publishing, Cham, pp 193–209.
- Wu P, Fang X, Fang H, Gong Z, Kan D (2024) An Event Log Repair Method Based on Masked Transformer Model. *Applied Artificial Intelligence* 38(1). doi:10.1080/08839514.2024.2346059.
- Wynn MT, Leberherz J, van der Aalst WMP, Accorsi R, Di Ciccio C, Jayarathna L, Verbeek HMW (2022) Rethinking the Input for Process Mining: Insights from the XES Survey and Workshop. In: Munoz-Gama J, Lu X (eds) *Process Mining Workshops*. Springer International Publishing, Cham, pp 3–16.
- Yang W, Su Q (2014) Process mining for clinical pathway: Literature review and future directions. In: *Proceedings of ICSSSM ‘14*. June 25-27, 2014, Beijing, China. IEEE, [Piscataway, N.J.], pp 1–5.
- Zhang H, Dong Y, Xiao C, Oyamada M (2024) Jellyfish: Instruction-Tuning Local Large Language Models for Data Preprocessing. In: Al-Onaizan Y, Bansal M, Chen Y-N (eds) *The 2024 Conference on Empirical Methods in Natural Language Processing*. *Proceedings of the Conference*, pp 8754–8782.
- Zorina A, Bélanger F, Kumar N, Clegg S (2021) Watchers, Watched, and Watching in the Digital Age: Reconceptualization of Information Technology Monitoring as Complex Action Nets. *Organization Science* 32(6):1571–1596. doi:10.1287/orsc.2021.1435.

VI Appendix

VI.1 Index of Research Articles

Research Paper P1	Navigating the Landscape of Organizational Process Mining Setups: A Taxonomy Approach Marcus L, Schmid SJ, Friedrich F, Röglinger M, Grindemann P (2024) Published in: <i>Business & Information Systems Engineering</i> DOI: https://doi.org/10.1007/s12599-024-00908-0 (VHB-24 ¹ : B, VHB-JQ3 ² : B, SJR ³ : Q1, IF ⁴ : 7.9)
Research Paper P2	Capabilities for Building and Managing Process-Based Behavioral Visibility in Organizations Franzoi S, Kipping G, Marcus L, Schmid SJ, vom Brocke J, Grisold T, Mendling J, Röglinger M (2024) Submitted to: <i>Outlet hidden due to the double-blind review process of the journal</i>
Research Paper P3	Everything at the Proper Time: Repairing Identical Timestamp Errors in Event Logs with Generative Adversarial Networks Schmid SJ, Moder L, Hofmann P, Röglinger M (2023) Published in: <i>Information Systems</i> DOI: https://doi.org/10.1016/j.is.2023.102246 (VHB-24 ¹ : B, VHB-JQ3 ² : B, SJR ³ : Q1, IF ⁴ : 3.0)
Research Paper P4	Case ID Revealed HERE: Hybrid Elusive Case Repair Method for Transformer-Driven Business Process Event Log Enhancement Zetzsche F, Andrews R, ter Hofstede AHM, Röglinger M, Schmid SJ, Wynn MT (2025) Published in: <i>Business & Information Systems Engineering</i> DOI: https://doi.org/10.1007/s12599-025-00935-5 (VHB-24 ¹ : B, VHB-JQ3 ² : B, SJR ³ : Q1, IF ⁴ : 7.9)
Research Paper P5	One to Rule Them All: Large Language Models for Multi-Imperfection Business Process Event Log Repair Schmid SJ, Zetzsche F, Röglinger M (2025) Submitted to: <i>Outlet hidden due to the double-blind review process of the journal</i>

Table 6: Research articles included in the dissertation

Furthermore, I have co-authored the following research papers, which are not included in this dissertation.

¹ VHB Publication Media Rating 2024

² VHB-JOURQUAL3

³ Scimago Journal & Country Rank 2023

⁴ Impact Factor

How to Leverage Process Mining in Organizations – Towards Process Mining Capabilities

Kipping G, Djurica D, Franzoi S, Grisold T, Marcus L, Schmid SJ, vom Brocke J, Mendling J, Röglinger M (2022)

Published in: *International Conference on Business Process Management 2022*

DOI: https://doi.org/10.1007/978-3-031-16103-2_5

Exploring the Interplay of Process Mining and Generative AI: Research and Recommendations for CoEs

Reinkemeyer L, Röglinger M, Kratsch W, Fabri L, Schmid SJ, Wittmann J (2023)

DOI: <https://doi.org/10.24406/publica-2358>

ProcessLLM: A Large Language Model Specialized in the Interpretation, Analysis, and Optimization of Business Processes

Buss A, Kratsch W, Schmid SJ, Wang H (2025)

Published in: *International Conference on Business Process Management 2024*

DOI: https://doi.org/10.1007/978-3-031-78666-2_17

Table 7: Further research articles not included in the dissertation

VI.2 Individual Contribution to the Included Research Articles

This dissertation is cumulative in nature and comprises five research papers. All included papers were co-authored with other researchers in collaborative team settings. This section provides a transparent account of my individual contributions to each research paper. The descriptions follow the Contributor Roles Taxonomy (CRediT) by Allen et al. (2019), which allows for a differentiated view of contributions across roles such as conceptualization, methodology, investigation, writing, and more. Unless explicitly stated otherwise, all research papers were the result of shared efforts, with author roles and responsibilities agreed upon jointly by the respective research teams.

Research paper P1 entitled “*Navigating the Landscape of Organizational Process Mining Setups: A Taxonomy Approach*” (Section VI.3) was written in equal authorship. I was actively involved in conceptualization, helping to shape the overall direction and research objectives of the study. In terms of methodology, I co-developed the research design and played a central role in data curation and investigation, including conducting interviews, preparing transcripts, and organizing the data for analysis. I also contributed to formal analysis, particularly in coding the interview data, which supported the development of the taxonomy. Additionally, I contributed to the visualization of the taxonomy by shaping its structure, selecting appropriate wording, and presenting it in a coherent and communicable format. Finally, I was responsible for writing sections of the original draft and contributed significantly to review and editing throughout the revision process.

Research paper P2 entitled “*Capabilities for Building and Managing Process-Based Behavioral Visibility in Organizations*” (Section VI.4) was written in equal authorship. I was actively involved in conceptualization, helping to shape the overall direction and research

objectives of the study. In terms of methodology, I co-developed the research design and played a central role in data curation and investigation, including conducting interviews, preparing transcripts, and organizing the data for analysis. I also contributed to formal analysis, particularly in coding the interview data and applying the Gioia et al. (2013) methodology to develop the capability framework. Additionally, I contributed to visualization by designing the central result figure. I was responsible for writing sections of the original draft, including the results section, and contributed significantly to reviewing and editing the entire paper.

Research paper P3 entitled “*Everything at the Proper Time: Repairing Identical Timestamp Errors in Event Logs with Generative Adversarial Networks*” (Section VI.5) was written in equal authorship. In line with my role as the first author, I was responsible for project administration, ensuring coordination and timely submission. Furthermore, I was responsible for conceptualization and methodology, helping to define the research objectives and design of the study. I was actively developing the artifact and software prototype used for demonstration and evaluation purposes. I was also responsible for data curation, preparing the evaluation data by generating synthetic logs through error injection, and conducted the formal analysis to assess the artifact’s utility. Furthermore, I created key visualizations to present the research method, artifact design, and evaluation results. I authored multiple sections of the original draft and was actively involved in reviewing and editing the entire paper, including all revisions.

Research paper P4 entitled “*Case ID Revealed HERE: Hybrid Elusive Case Repair Method for Transformer-Driven Business Process Event Log Enhancement*” (Section VI.6) was written in equal authorship. I was actively involved in conceptualization and methodology, contributing to the design of both the artifact and the evaluation strategy. I also took on a supervision role, mentoring another author throughout the project. My work further contributed to formal analysis by conducting and analyzing the artificial formative evaluation interviews as well as leading the artificial summative evaluation, where I defined evaluation metrics and developed the evaluation pipeline. I additionally participated in the naturalistic evaluation interviews and implemented the software interface used for interacting with the artifact. I created key visualizations to illustrate evaluation results and wrote parts of the original draft, including the evaluation section. I was also actively involved in reviewing and editing all sections of the original draft. Finally, I contributed significantly to reviewing and editing throughout the revision process.

Research paper P5 entitled “*One to Rule Them All: Large Language Models for Multi-Imperfection Business Process Event Log Repair*” (Section VI.7) was written in lead

authorship. In line with this leading role, I was responsible for project administration, ensuring coordination and execution of the research activities. I contributed significantly to conceptualization and methodology, helping to shape the research design and the resulting artifact. Furthermore, I was responsible for developing the software prototype as an instantiation of the artifact. My work further contributed to data curation by compiling the base log collection and generating the evaluation data through controlled error injection. I carried out the formal analysis by evaluating the artifact's utility and calculating key metrics, which I subsequently illustrated through visualizations. Finally, I authored multiple sections of the original draft and was actively involved in reviewing and editing the entire manuscript.

VI.3 Research Paper 1: Navigating the Landscape of Organizational Process Mining Setups: A Taxonomy Approach

Authors:

Laura Marcus, Sebastian Johannes Schmid, Franziska Friedrich, Maximilian Röglinger, Philipp Grindemann

Published in:

Business & Information Systems Engineering (November 2024)

ISSN: 1867-0202

DOI: <https://doi.org/10.1007/s12599-024-00908-0>

Abstract:

Process mining (PM) technology evolves around the analysis, design, implementation, and ongoing improvement of business processes. While it has experienced a lot of attention and significant technological advancements, contributions to the field have mostly revolved around technical matters, neglecting managerial and organizational aspects. Thus, researchers have called for a more holistic view of the application and adoption of PM in enterprises. To address this gap, this paper presents a taxonomy for organizational PM setups. Its applicability and usefulness are shown in three exemplary cases. This study extends the descriptive knowledge at the intersection of PM and business process management governance, highlighting the unique governance requirements associated with PM that cannot be effectively addressed through traditional governance approaches. The taxonomy provides practitioners with orientation when developing an effective PM setup and helps to characterize existing setups.

Keywords:

Process Mining; Organizational Setup; BPM Governance; Center of Excellence; Taxonomy Development

VI.4 Research Paper 2: Capabilities for Building and Managing Process-Based Behavioral Visibility in Organizations

Authors:

Sandro Franzoi, Gregor Kipping, Laura Marcus, Sebastian Johannes Schmid, Jan vom Brocke, Thomas Grisold, Jan Mendling, Maximilian Röglinger

Submitted to:

Outlet hidden due to the double-blind review process of the journal

Extended Abstract:

The increasing digitization of organizational processes is producing a wealth of behavioral data, offering new possibilities for managing work in digital environments and creating business value (Badakhshan et al. 2022; Vaujany et al. 2021; Leonardi and Treem 2020). This paper explores how organizations can develop capabilities to generate and manage process-based behavioral visibility. Rather than focusing solely on the potential surveillance risks or ethical implications (Benlian et al. 2022; Newlands 2021; Spicer 2017; Zorina et al. 2021), this study adopts a managerial perspective, investigating how organizations create business value from behavioral visibility.

The research employs a qualitative, grounded theory methodology (Gioia et al., 2013; Glaser & Strauss, 2017; Strauss & Corbin, 1998). Data was collected through 30 expert interviews with practitioners who have implemented or managed process mining initiatives in their organizations. The interviews spanned various industries, functions, and organizational roles, providing a rich empirical foundation. The iterative coding and analysis process led to the development of a conceptual framework that identifies three interrelated types of organizational capabilities essential for using process-based behavioral visibility:

Foundational capabilities establish the socio-technical groundwork necessary for behavioral visibility to emerge. They address both technical and organizational challenges, such as capturing reliable behavioral data, integrating fragmented systems, and fostering a culture that balances transparency with trust. These capabilities ensure that visibility efforts are built on robust, ethically sound, and strategically aligned infrastructure.

Transformational capabilities enable organizations to convert behavioral visibility into actionable insights and business value. They focus on interpreting digital traces in context,

aligning behavioral data with real-world work practices, and embedding visibility into decision-making. These capabilities are essential for closing the gap between insight and value, ensuring that visibility leads to informed action rather than mere observation.

Continual capabilities support the long-term adaptability and evolution of behavioral visibility initiatives. They promote a mindset of continuous learning and refinement, helping organizations stay responsive to changing behaviors, new technologies, and shifting strategic priorities. These capabilities safeguard the relevance and resilience of behavioral visibility over time, transforming it from a one-off project into an enduring management practice.

The study emphasizes that behavioral visibility is not simply a technological artifact but a dynamic organizational phenomenon that emerges from the interplay between digital trace data, technological infrastructure, and managerial practices. The capabilities identified are not only technical but also social in nature. They enable organizations to make informed decisions, foster transparency, and align business operations with strategic goals. By focusing on the capability perspective, this paper advances both theoretical understanding and practical implementation of process-based behavioral visibility in organizational contexts. It contributes to the literature on digital trace data, information systems capabilities, and data-driven management by showing how behavioral visibility can be purposefully used to drive organizational learning, alignment, and performance.

Keywords:

Behavioral Visibility; Capability Framework; Process Mining; Digital Trace Data

References:

- Badakhshan P, Wurm B, Grisold T, Geyer-Klingeberg J, Mendling J, vom Brocke J (2022) Creating business value with process mining. *The Journal of Strategic Information Systems* 31(4):101745. doi:10.1016/j.jsis.2022.101745.
- Benlian A, Wiener M, Cram WA, Krasnova H, Maedche A, Möhlmann M, Recker J, Remus U (2022) Algorithmic Management. *Business & Information Systems Engineering* 64(6):825–839. doi:10.1007/s12599-022-00764-w.
- Gioia DA, Corley KG, Hamilton AL (2013) Seeking Qualitative Rigor in Inductive Research. *Organizational Research Methods* 16(1):15–31. doi:10.1177/1094428112452151.
- Glaser BG, Strauss AL (2017) *The Discovery of Grounded Theory*. Routledge, New York.
- Leonardi PM, Treem JW (2020) Behavioral Visibility: A new paradigm for organization studies in the age of digitization, digitalization, and datafication. *Organization Studies* 41(12):1601–1625. doi:10.1177/0170840620970728.

- Newlands G (2021) Algorithmic Surveillance in the Gig Economy: The Organization of Work through Lefebvrian Conceived Space. *Organization Studies* 42(5):719–737. doi:10.1177/0170840620937900.
- Spicer A (2017) Surveillance used to be a bad thing. Now, we happily let our employers spy on us. <https://www.theguardian.com/commentisfree/2017/aug/04/surveillance-employers-spy-implanted-chipped>.
- Strauss AL, Corbin JM (1998) *Basics of qualitative research: Techniques and procedures for developing grounded theory*. SAGE Publications, Thousand Oaks, CA.
- Vaujany F-X de, Leclercq-Vandelannoitte A, Munro I, Nama Y, Holt R (2021) Control and Surveillance in Work Practice: Cultivating Paradox in ‘New’ Modes of Organizing. *Organization Studies* 42(5):675–695. doi:10.1177/01708406211010988.
- Zorina A, Bélanger F, Kumar N, Clegg S (2021) Watchers, Watched, and Watching in the Digital Age: Reconceptualization of Information Technology Monitoring as Complex Action Nets. *Organization Science* 32(6):1571–1596. doi:10.1287/orsc.2021.1435.

VI.5 Research Paper 3: Everything at the Proper Time: Repairing Identical Timestamp Errors in Event Logs with Generative Adversarial Networks

Authors:

Sebastian Johannes Schmid, Linda Moder, Peter Hofmann, Maximilian Röglinger

Published in:

Information Systems 118, 102246 (September 2023).

ISSN: 0306-4379

DOI: <https://doi.org/10.1016/j.is.2023.102246>

Abstract:

Process mining generates valuable insights into business processes through the analysis of event logs. However, event logs are commonly subject to various data quality issues which hinder the success of process mining initiatives in organizations. Identical timestamp errors, for example, occur when multiple events of a process instance mistakenly share the same timestamp. This error causes discovered process models to be unrepresentative and process performance analysis results to be misleading. To address this problem, we propose a method for automatically repairing identical timestamp errors in event logs. To that end, we combine existing method components for error detection and reordering of erroneous events with a novel approach for repairing timestamps based on Generative Adversarial Networks. To allow for a rigorous evaluation, we instantiate our approach as a software prototype, and use it to repair a total of six real-life and artificial event logs with overall 30 variations. Thereby, we show that the proposed method shows improved results compared to alternative approaches for repairing identical timestamp errors in event logs.

Keywords:

Process Mining; Process Data Quality; Generative Adversarial Networks; Event Log Repair; Business Process Management; Machine Learning

VI.6 Research Paper 4: Case ID Revealed *HERE*: Hybrid Elusive Case Repair Method for Transformer-Driven Business Process Event Log Enhancement

Authors:

Felix Zetzsche, Robert Andrews, Arthur H. M. ter Hofstede, Maximilian Röglinger, Sebastian Johannes Schmid, Moe Thandar Wynn

Published in:

Business & Information Systems Engineering (March 2025).

ISSN: 1867-0202

DOI: <https://doi.org/10.1007/s12599-025-00935-5>

Abstract:

Process mining is a data-driven technique that leverages event logs to analyze, visualize, and improve business processes. However, data quality is often low in real-world settings due to various event log imperfections, which, in turn, degrade the accuracy and reliability of process mining insights. One notable example is the elusive case imperfection pattern, describing the absence of case identifiers responsible for linking events to a specific process instance. Elusive cases are particularly problematic, as process mining techniques rely heavily on the accurate mapping of events to instances to provide meaningful and actionable insights into business processes. To address this issue, the study follows the Design Science Research paradigm to iteratively develop a method for repairing the elusive case imperfection pattern in event logs. The proposed Hybrid Elusive Case Repair Method (*HERE*) combines a traditional, rule-based approach with generative artificial intelligence, specifically the Transformer architecture. By integrating domain knowledge, *HERE* constitutes a comprehensive human-in-the-loop approach, enhancing its ability to accurately repair elusive cases in event logs. The method is evaluated by instantiating it as a software prototype, applying it to repair three publicly accessible event logs, and seeking expert feedback in a total of 21 interviews conducted at different points during the design and development phase. The results demonstrate that *HERE* makes significant progress in addressing the elusive case imperfection pattern, particularly when provided with sufficient data volume, laying the groundwork for resolving further data quality issues in process mining.

Keywords:

Process Mining; Event Log Quality; Event Log Repair; Generative Artificial Intelligence; Transformer; Business Process Management

VI.7 Research Paper 5: One to Rule Them All: Large Language Models for Multi-Imperfection Business Process Event Log Repair

Authors:

Sebastian Johannes Schmid, Felix Zetsche, Maximilian Röglinger

Submitted to:

Outlet hidden due to the double-blind review process of the journal

Extended Abstract:

Process mining is a data-driven technique that analyzes event logs of historical process executions to uncover critical insights into business processes (Rott et al. 2024). Yet, in real-world applications, such event logs often suffer from multiple, co-occurring data quality issues, including missing case identifiers, inconsistent activity labels, and erroneous or duplicate timestamps (Fischer et al. 2022; Suriadi et al. 2017). These imperfections undermine the accuracy and usefulness of process mining insights (Goel et al. 2022). Conventional repair approaches typically focus on individual imperfection types in isolation and are combined as modular toolchains in multi-imperfection settings (Andrews et al. 2020). While highly effective in certain scenarios, such toolchains suffer from cascading errors when imperfections are interdependent (Hofstede et al. 2023; Suriadi et al. 2017), limited adaptability to new data formats, and high maintenance costs (Baier et al. 2020; Sato et al. 2021).

This study explores whether Large Language Models (LLMs), known for their conceptual understanding of process mining and event log quality, can provide a unified solution to simultaneously diagnose and repair multiple types of event log imperfections. Building on the transformer-based LLaMA-3.1 8B instruct model, we fine-tune it using parameter-efficient techniques to reduce hardware requirements while maintaining task-specific performance (Lingam et al. 2024; Hu et al. 2021). The result is an open-source artifact designed to detect and repair combinations of five exemplary imperfection patterns.

Following the Design Science Research paradigm by Peffers et al. (2007), the artifact is developed and refined over four design iterations ensuring feasibility, robustness, effectiveness, and computational efficiency. The evaluation focused on repairing eight real-life event logs representing diverse domains. Each log was artificially injected with 17 valid combinations of imperfection patterns at error rates ranging from 10% to 100%. Detection performance was

measured using precision, recall, and F1-score, while repair effectiveness was evaluated with imperfection-specific metrics: timestamp deviation, timestamp accuracy, case ID accuracy, and label accuracy.

Results show that the fine-tuned LLM achieves promising diagnostic performance with average precision, recall, and F1-scores close to 97% across all logs, including one log that was unseen during training. For repair, the model excels in correcting activity labels ($\geq 99\%$ accuracy) and performs strongly on timestamp normalization and moderate levels of case ID reconstruction. However, accuracy drops sharply at approximately 70% case ID loss, where not enough contextual cues remain. Timestamp deviation is significantly reduced in most cases, though performance is limited by the LLM's inherent difficulty with precise arithmetic operations.

Beyond overcoming key limitations of the traditional toolchain paradigm of event log repair, this work also lays the foundation for LLM-enhanced process mining lifecycles. Future applications may extend toward automated data preprocessing or process improvement. Finally, the prototype released as open-source software provides a valuable resource for practitioners and researchers seeking to operationalize or extend this approach.

Keywords:

Process Mining; Process Data Quality; Generative Artificial Intelligence; Event Log Repair; Large Language Models; Machine Learning; Business Process Management

References:

- Andrews R, Emamjome F, Hofstede AHM ter, Reijers HA (2022) Root-cause analysis of process-data quality problems. *Journal of Business Analytics* 5(1):51–75. doi:10.1080/2573234X.2021.1947751.
- Baier L, Reimold J, Kühl N (2020) Handling Concept Drift for Predictions in Business Process Mining. In: Aier S, Gordijn J, Proper HA, Verelst J (eds) 2020 IEEE 22nd Conference on Business Informatics. CBI 2020, pp 76–83.
- Fischer DA, Goel K, Andrews R, van Dun CGJ, Wynn MT, Röglinger M (2022) Towards interactive event log forensics: Detecting and quantifying timestamp imperfections. *Information Systems* 109. doi:10.1016/j.is.2022.102039.
- Goel K, Leemans SJJ, Martin N, Wynn MT (2022) Quality-Informed Process Mining: A Case for Standardised Data Quality Annotations. *ACM Transactions on Knowledge Discovery from Data* 16(5):1–47. doi:10.1145/3511707.
- Hofstede AHM ter, Koschmider A, Marrella A, Andrews R, Fischer DA, Sadeghianasl S, Wynn MT, Comuzzi M, Weerdt J de, Goel K, Martin N, Soffer P (2023) Process-Data Quality: The True Frontier of Process Mining. *Journal of Data and Information Quality* 15(3):1–21. doi:10.1145/3613247.
- Hu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L, Chen W (2021) LoRA: Low-Rank Adaptation of Large Language Models.

- Lingam V, Tejaswi A, Vavre A, Shetty A, Gudur GK, Ghosh J, Dimakis A, Choi E, Bojchevski A, Sanghavi S (2024) SVFT: Parameter-Efficient Fine-Tuning with Singular Vectors.
- Peffer K, Tuunanen T, Rothenberger MA, Chatterjee S (2007) A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems* 24(3):45–77. doi:10.2753/MIS0742-1222240302.
- Rott J, Böhm M, Krcmar H (2024) Laying the ground for future cross-organizational process mining research and application: a literature review. *Business Process Management Journal* 30(8):144–206. doi:10.1108/BPMJ-04-2023-0296.
- Sato DMV, Freitas SC de, Barddal JP, Scalabrin EE (2021) A Survey on Concept Drift in Process Mining. *ACM Computing Surveys* 54(9):1–38. doi:10.1145/3472752.
- Suriadi S, Andrews R, Hofstede AHM ter, Wynn MT (2017) Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs. *Information Systems* 64:132–150. doi:10.1016/j.is.2016.07.011.