Analyzing Continuous Process Dynamics for Comprehensive Process Mining

Sana Dodangeh[1,2]

[1]Faculty of Computer Science, University of Vienna, Nordbergstraße 1, 1090, Vienna, Austria
[2]UniVie Doctoral School Computer Science DoCS, University of Vienna, Währinger Str. 29, 1090 Vienna, Austria

Abstract

Business processes rarely remain stable over the period captured in an event log. Instead, they continuously fluctuate over time. However, many process mining analyses either assume process behavior to be stable or focus on distinguishing different process versions, thereby overlooking systematic fluctuations within process versions. This doctoral project addresses this limitation by developing algorithmic approaches for the data-driven analysis of continuous process dynamics from event logs. The goal is to characterize such dynamics, explain their drivers, and use the resulting insights to support downstream process mining tasks. Together, these contributions aim to provide fine-grained and reliable insights into how processes evolve over time.

Keywords

process dynamics, continuous fluctuations, temporal analysis

Supervision

This doctoral research is conducted at the Faculty of Computer Science, University of Vienna, under the supervision of Univ.-Prof. Han van der Aa.

1. Introduction

Process mining is a discipline used to analyze, monitor, and improve organizational processes based on event logs recorded by information systems. It bridges data mining and process management by offering analytical tools to visualize and understand process execution, identify inefficiencies, and optimize process performance [1, 2]. Hence, process mining helps analysts understand how processes function in practice.

Many process mining techniques assume static behavior over the timeline covered by an event log. However, in reality, processes continuously fluctuate [3, 4]. These fluctuations can stem from intentional or unintentional adaptations, environmental changes, technological progress, and seasonal trends [5, 6].

For example, consider a hospital check-in process whose procedure remains unchanged while patient arrivals, staff availability, and workload vary over time. If these dynamics are ignored, gradual trends, recurring cycles, and interrelations between process perspectives such as performance and resource use may remain hidden. Their omission can also mislead common process mining analyses: aggregated measures may hide recurring delays, and predictions may underestimate remaining times during peaks. Analysts may therefore dismiss systematic patterns as noise or target average behavior rather than problematic periods [7]. Consequently, process mining may yield incomplete or inaccurate insights and lead to suboptimal decisions.

Although process dynamics have been studied in the literature, existing works are limited in their scope, since they generally focus on identifying transitions between distinct process versions, commonly referred to as concept drift detection [3, 5, 8]. In this proposal, we take a broader view and use the term *continuous process dynamics* for time-dependent variations in process behavior throughout process

Proceedings of the BPM 2026 Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Forum (BPM 2026), Toronto, Canada, September 27th to October 2nd, 2026.
sana.dodangeh@univie.ac.at (S. Dodangeh)

execution. We use *continuous* because we analyze process behavior across the full timeline, covering variations from changes between process versions, as studied under concept drift, to gradual trends and recurring cycles within a version. For example, waiting times may rise every Monday and fall by Friday without creating a new process version; only the measurable characteristics fluctuate cyclically. Considering only concept drift may overlook these variations and their drivers, despite their effects on performance and decision-making. This provides an incomplete picture of how processes evolve over time. In the remainder of this paper, *process dynamics* refers to continuous process dynamics.

In light of these limitations, the main goal of this doctoral project is to *develop algorithmic approaches for the data-driven analysis of process dynamics to provide insights into how processes evolve over time.* Inspired by time series analysis, this project examines how relevant characteristics extracted from event logs can be represented over time to observe inherent continuous dynamics. More specifically, it identifies and characterizes different forms of continuous process fluctuations, investigates the factors that drive these fluctuations, and uses this knowledge to support downstream process mining tasks, such as predictive process monitoring.

2. Related Work

This section briefly reviews related work that forms the basis of our project. A comprehensive understanding of process dynamics requires insights from multiple research streams, in particular, concept drift detection, resilience assessment, and routine dynamics.

Concept drift detection. Concept drift detection aims to identify changes in process behavior that indicate a transition between different process versions. Existing work has studied different forms of such transitions, including sudden, gradual, incremental, and recurring drifts, from perspectives such as control flow, data, time, and resources [5, 8, 9, 10, 11, 12, 13]. However, this view mainly captures changes that separate one version from another. By contrast, this doctoral project recognizes that process dynamics may also unfold within an ongoing process, where behavior continuously fluctuates over time without necessarily marking a transition to a new state.

Resilience assessment. Organizations often face disruptions that affect process execution. Process resilience refers to a process's ability to handle such disruptions and recover from them. Therefore, resilience assessment measures how well a process can overcome adverse events. Existing approaches focus on the impact of disruptions and subsequent recovery behavior [14], frame resilience as a life cycle with quantitative metrics for different response stages [15], or identify parameter configurations under which a process maintains acceptable performance [16]. However, these studies primarily focus on individual adverse events, whereas continuous process dynamics may arise as ongoing fluctuations in process behavior, beyond a single disruption or recovery phase.

Routine dynamics. A third related stream is routine dynamics, which studies how organizational routines evolve through repeated performances by different actors. Process mining has supported diachronic analyses of process behavior in this context [17, 18, 19, 20]. While these studies highlight variability and adaptation as intrinsic properties of organizational processes, they mainly provide theoretical or empirical insights from the Information Systems and organizational science literature. In contrast, this doctoral project develops algorithmic approaches for data-driven analysis of process dynamics from event logs.

Overall, although several research streams recognize the changing nature of organizational processes, existing works do not yet provide data-driven solutions for comprehensively analyzing continuous process dynamics. This doctoral project addresses this gap by developing approaches that characterize such dynamics, explain their drivers, and use the insights in downstream process mining tasks.

3. Current State and Research Plans

To achieve the main goal stated in Section 1, this doctoral project is organized into five research projects. The first project is a systematic literature review that, following the Kitchenham methodology [21], maps existing uses of time series analysis in process mining and helps position the contributions. The following four projects directly address the doctoral research goal by developing algorithmic approaches, following the algorithm engineering methodology [22].

In the following, we summarize the current progress and planned next steps for each project, covering the addressed problem, objective and contribution, evaluation strategy, current state, and main risks.

3.1. Project 1: A Systematic Literature Review on Time Series in Process Mining

Time series representations are increasingly used in process mining to analyze process behavior over time. They support a wide range of process mining tasks, such as prediction, concept drift detection, and process dynamics analysis [14, 23, 24, 25]. Understanding how time series techniques are applied across such tasks provides a broader context for this doctoral project and helps identify opportunities for analyzing process dynamics.

Problem. Although time series representations and techniques are increasingly used across process mining tasks, their application remains fragmented. It is still unclear which techniques support which tasks, which process aspects are commonly represented as time series, and which methods from the broader time series literature remain underexplored for process mining.

Research goal, contribution, and evaluation. This project maps process mining and time series analysis to identify underexplored techniques and research opportunities. The review asks three questions: (i) which time series techniques have been applied in process mining, (ii) which techniques exist in the broader time series literature, and (iii) which research opportunities emerge at their intersection regarding underexplored techniques and process mining tasks. The review is expected to reveal time series techniques not used in process mining and process mining tasks not studied with such methods, thereby informing later projects on characterizing dynamics, analyzing the causes of these dynamics, and using the knowledge in predictive process monitoring.

Current state, future plan, and risks. The review is currently in the study selection phase, with inclusion and exclusion criteria and the search query defined. For the first research question mentioned above, database searches returned 805 records, 477 of which remain after title screening. The second and third research questions remain to be addressed. As the review progresses, a key challenge is covering both the process mining and broader time series literature while keeping the scope manageable.

3.2. Project 2: Separating Systematic and Random Fluctuations in Process Dynamics

This project addresses the identification and characterization of continuous process fluctuations. It analyzes temporal changes in processes to distinguish systematic patterns from random fluctuations and to reveal how different process characteristics evolve together over time. Understanding these dynamics is key to making informed adjustments and decisions in process management.

Problem. Business processes fluctuate due to environmental changes, seasonal effects, and internal adaptations. Existing analyses of continuous dynamics are mainly visualization-based, plotting time-dependent behaviors such as case arrivals or activity counts to reveal drifts or recurring patterns. Their analytical support is limited because visual approaches rely on the analyst's eye rather than objective criteria: (i) without statistical testing, apparent rises or falls cannot be reliably distinguished from random noise; and (ii) separately plotted characteristics do not quantify their interrelations.

Research goal, contribution, and evaluation. This project develops a systematic, data-driven approach beyond visual analysis of continuous process dynamics. It transforms event logs into time series, decomposes them into trend, seasonal, and residual components, uses the Ljung–Box test to check for remaining structure in the residuals, and analyzes interrelations through correlation to separate systematic from random fluctuations and reveal how process aspects evolve together. Evaluation combines two analyses: a backward-looking analysis that assesses how it supports understanding systematic process dynamics and interrelations between process characteristics, and a forward-looking analysis of whether separating these dynamics improves prediction of future process behavior.

Current state and risks. Earlier versions of this work received feedback at ICPM 2025 and ICSOC 2025. Based on this, we strengthened the evaluation design, and the revised version was accepted at the BPM 2026 Forum [26]. Nevertheless, some threats to validity remain, concerning generalizability across process settings, sensitivity to design choices such as process characteristics and parameters, dependence on selected prediction models, and lack of user-based validation.

3.3. Project 3: Discovering the Causal Drivers of Performance Dynamics

Building on Project 2, this project moves from characterizing continuous process fluctuations to explaining them. While the previous project reveals systematic patterns and interrelations between process characteristics, such associations do not establish causality. This distinction is important for process management, since interventions require understanding which factors actually drive performance changes. Therefore, this project identifies the causal drivers of process performance dynamics, where process performance indicators (PPIs) may fluctuate as process characteristics evolve over time.

Problem. Although several studies have investigated causal drivers of PPIs, existing causal analyses in process mining have mainly relied on simplifying assumptions, such as linear relations or lagged effects only [23, 24]. Yet, performance dynamics may involve more complex relations between process characteristics and PPIs. Temporal causal discovery methods from the PCMCI family [27] can address such settings, but their suitability for process mining data remains unclear. Addressing this gap can reduce simplifying assumptions and better reflect real process dynamics.

Research goal, contribution, and evaluation. This work advances the causal analysis of process performance dynamics using PCMCI-based methods. It asks: (i) when these methods are suitable for process mining data, (ii) how to select a variant based on process-data properties, and (iii) how accurately they identify causal drivers of PPI dynamics compared with existing approaches. Accordingly, the project assesses the alignment between PCMCI assumptions and process-related time series, introduces a method-selection procedure, and evaluates the accuracy of these methods on synthetic event logs. Because real logs lack causal ground truth, we simulate process dynamics by configuring characteristics such as arrival rate and workload to affect the PPIs. These known relations provide ground truth, while a real-world event log demonstrates practical use.

Current state, next steps, and risks. An earlier version of this work was submitted to CAiSE 2026 and received useful feedback. We are currently revising the work for submission to the BISE journal by completing the method-selection procedure, strengthening the conceptual framing, and refining the evaluation design. Even with such planned improvements, the main open challenges concern parameter sensitivity and the design of realistic synthetic scenarios with known causal structures.

3.4. Project 4: Analyzing Changes in Causal Relations Across Process Versions

This project extends Project 3 by studying whether the causal structure explaining continuous process dynamics differs across process versions.

Problem. Existing studies on causal influences in process performance typically assume that the underlying causal structure remains stable over time [13, 24]. However, changes that lead to a new process version may also alter the causal structure explaining performance dynamics. For instance, if limited resources initially affect throughput time, increasing resource capacity may change this relation in a later version. If such changes are ignored, causal relations from different versions may be mixed, leading to misleading explanations of what drives performance changes.

Research goal, contribution, and evaluation. This project investigates how causal structures underlying performance dynamics evolve across process versions. It develops a process-specific approach that applies RPCMCI [28], a causal discovery method for temporal regimes, to version-specific log subsets. Evaluation uses synthetic event logs with known causal structures. To this end, we simulate multiple process versions with distinct configured relations and merge them into one log, providing known version-specific ground-truth structures.

Current state, next steps, and risks. This project is currently in an early conceptual stage, with the real-world problem, implementation, and evaluation planned at a high level. Next steps include formalizing the algorithmic task, implementing the approach, and designing simulation scenarios with distinct process versions and known version-specific causal structures. The main risks in this project concern the design of realistic simulation scenarios, with obvious changes in process versions.

3.5. Project 5: Dynamics-Aware Remaining Time Prediction

This project uses knowledge about process dynamics in downstream process mining tasks, focusing on remaining time prediction. Building on the previous projects, it investigates how features derived from process dynamics can be integrated into predictive process monitoring models.

Problem. Existing predictive process monitoring approaches either capture process changes implicitly or address them as concept drift, but they do not explicitly exploit structural temporal patterns for remaining time prediction [29, 30]. Moreover, while causality has been explored in predictive process monitoring [31], its use for remaining time prediction remains limited. Thus, it remains unclear how temporal patterns and causal insights can improve prediction accuracy and efficiency.

Research goal, contribution, and evaluation. The goal of this project is to examine whether outputs from the previous projects, namely decomposed process dynamics and causal drivers of performance fluctuations, can improve remaining time prediction. To this end, it develops a dynamics-aware framework that combines time series decomposition, causal information, and prediction techniques such as incremental learning [29]. The framework will be evaluated on real-life event logs in terms of prediction accuracy and computational efficiency against state-of-the-art approaches.

Current state, next steps, and risks. This project is still at an early conceptual stage, with its problem, core idea, and evaluation goal specified at a high level. At this point, the main uncertainty concerns how decomposed dynamics and causal insights can be integrated into instance-level prediction models.

4. Outlook and Roadmap

This doctoral project aims to advance process mining by developing algorithmic approaches for analyzing, explaining, and exploiting continuous process dynamics.

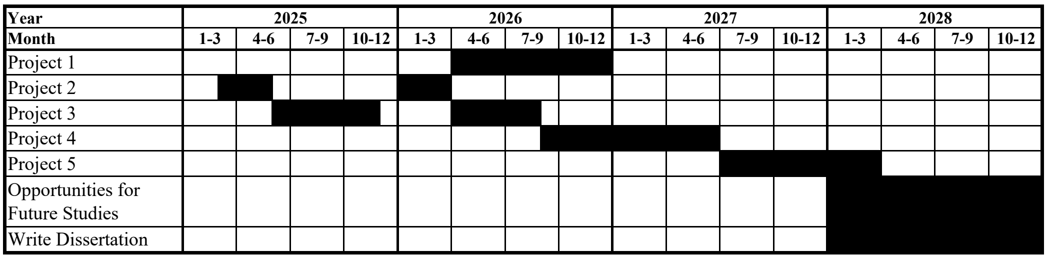

Figure 1: The doctoral project Gantt chart.

Figure 1 shows the planned roadmap. Work continues on Project (1) (systematic literature review) and Project (3) (causal drivers of process dynamics), followed by Project (4) (changing causal relations across process versions) and Project (5) (dynamics-aware remaining time prediction); 2028 is reserved for writing the dissertation. The schedule is approximate, as revisions, rejections, or journal-extension invitations may require adjustments. We target venues such as BPM, ICPM, and CAiSE, and journals such as BISE, Information Systems, and Process Science.

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT (OpenAI) for grammar checking and improving the clarity and quality of the writing.

References

[1] W. Van Der Aalst, Process mining: Overview and opportunities, ACM Transactions on Management Information Systems (TMIS) 3 (2012) 1–17.

[2] W. Van Der Aalst, W. van der Aalst, Data science in action, Springer, 2016.

[3] A. Kraus, H. van der Aa, Looking for change: a computer vision approach for concept drift detection in process mining, in: BPM, Springer, 2024, pp. 273–290.

[4] T. Grisold, C. Janiesch, M. Röglinger, M. T. Wynn, Managing dynamics in and around business processes, Business & Information Systems Engineering 66 (2024) 533–540.

[5] R. J. C. Bose, W. M. Van Der Aalst, I. Žliobaitė, M. Pechenizkiy, Dealing with concept drifts in process mining, IEEE transactions on neural networks and learning systems 25 (2013) 154–171.

[6] A. Yeshchenko, D. Bayomie, S. Gross, J. Mendling, Visualizing business process evolution, in: BPMDS, EMMSAD, Held at CAiSE, Springer, 2020, pp. 185–192.

[7] C. Browne, A change dynamic model for the online detection of gradual change, arXiv preprint arXiv:2205.01054 (2022).

[8] E. L. Klijn, F. Mannhardt, D. Fahland, Multi-perspective concept drift detection: Including the actor perspective, in: CAiSE, Springer, 2024, pp. 141–157.

[9] F. Richter, T. Seidl, Tesseract: time-drifts in event streams using series of evolving rolling averages of completion times, in: BPM, Springer, 2017, pp. 289–305.

[10] A. Kraus, H. van der Aa, Comprehensive characterization of concept drifts in process mining, 2025. Under submission to Information Systems.

[11] L. Zellner, F. Richter, J. Sontheim, A. Maldonado, T. Seidl, Concept drift detection on streaming data with dynamic outlier aggregation, in: Process Mining Workshops: ICPM, Springer, 2021, pp. 206–217.

[12] A. Yeshchenko, C. Di Ciccio, J. Mendling, A. Polyvyanyy, Comprehensive process drift analysis with the visual drift detection tool, in: Proceedings of the ER Forum and Poster & Demos Session 2019, CEUR Workshop Proceedings, 2019, pp. 108–112.

[13] B. Hompes, J. C. Buijs, W. M. van der Aalst, P. M. Dixit, J. Buurman, Detecting changes in process behavior using comparative case clustering, in: Data-Driven Process Discovery and Analysis, Springer, 2017, pp. 54–75.

[14] A. Kraus, J.-R. Rehse, H. van der Aa, Data-driven assessment of business process resilience, Process Science 1 (2024) 4.

[15] J. Stocker, N. Herda, J. Jürjens, Life cycle and metrics to measure the resilience of business processes by considering resources, Business Process Management Journal 28 (2022) 1164–1182.

[16] R. Bümel, A. Kraus, T. Kampik, H. van der Aa, Simulation-based assessment of business process robustness, in: BPM Forum, 2026.

[17] B. Pentland, E. Vaast, J. R. Wolf, Theorizing process dynamics with directed graphs: A diachronic analysis of digital trace data, MIS Quarterly 45 (2021).

[18] T. Grisold, B. Wurm, J. Mendling, J. Vom Brocke, Using process mining to support theorizing about change in organizations (2020).

[19] P. Budner, B. Wurm, C. Rosenkranz, J. Mendling, Towards routines mining–designing and implementing the argos miner, a design science artifact for studying routine dynamics with process mining (2022).

[20] G. Di Federico, C. Fernández-Llatas, Z. Ahmadi, M. Shirali, A. Burattin, Identifying variation in personal daily routine through process mining: a case study, in: ICPM, Springer, 2023, pp. 223–234.

[21] S. Keele, et al., Guidelines for performing systematic literature reviews in software engineering, Technical Report, Technical report, ver. 2.3 ebse technical report. ebse, 2007.

[22] J. Mendling, H. Leopold, H. Meyerhenke, B. Depaire, Methodology of algorithm engineering, arXiv preprint 2310.18979 (2023). URL: https://arxiv.org/abs/2310.18979. arXiv:2310.18979.

[23] B. F. Hompes, A. Maaradji, M. La Rosa, M. Dumas, J. C. Buijs, W. M. van der Aalst, Discovering causal factors explaining business process performance variation, in: CAiSE, Springer, 2017, pp. 177–192.

[24] J. N. Adams, S. J. van Zelst, L. Quack, K. Hausmann, W. M. van der Aalst, T. Rose, A framework for explainable concept drift detection in process mining, in: BPM, Springer, 2021, pp. 400–416.
[25] A. Kraus, K. A. Elyasi, H. van der Aa, On the use of steady-state detection for process mining: Achieving more accurate insights, in: CAiSE, Springer, 2025, pp. 204–220.
[26] S. Dodangeh, M. Kabierski, H. van der Aa, Making sense of continuous process dynamics, in: BPM Forum, 2026.
[27] J. Runge, A. Gerhardus, G. Varando, V. Eyring, G. Camps-Valls, Causal inference for time series, Nature Reviews Earth & Environment 4 (2023) 487–505.
[28] E. Saggioro, J. de Wiljes, M. Kretschmer, J. Runge, Reconstructing regime-dependent causal relationships from observational time series, Chaos: An Interdisciplinary Journal of Nonlinear Science 30 (2020).
[29] N. Guo, C. Liu, Q. Mo, J. Cao, C. Ouyang, X. Lu, Q. Zeng, Business process remaining time prediction based on incremental event logs, IEEE Transactions on Services Computing (2025).
[30] M. Maisenbacher, M. Weidlich, Handling concept drift in predictive process monitoring, in: SCC, IEEE, 2017, pp. 1–8.
[31] A. E. Márquez-Chamorro, A. Romero-Flores, M. Resinas, A. Ruiz-Cortés, Causally-informed predictive process monitoring, in: International Conference on Service-Oriented Computing, Springer, 2024, pp. 96–108.