Improving Time-related Predictions in Predictive Process Monitoring Through Contextual Feature Engineering

Maximilian Vouk[1,*]

[1]Vienna University of Economics and Business, Vienna, Austria

Abstract

Predictive Process Monitoring focuses on forecasting time-related properties of ongoing process instances, such as remaining time until case completion, to optimize resource scheduling and operational efficiency. However, despite recent architectural advances, state-of-the-art models still yield substantial error rates that limit their real-world utility. A fundamental limitation causing these deviations is the predominant reliance on a case-centric perspective with limited contextual factors. Consequently, existing approaches exhibit a systematic blindness to the multi-layered context that inherently shapes process velocity. To address these challenges, we investigate how multi-layered contextual knowledge can be systematically embedded into data representations via explicit feature engineering. This paper presents our doctoral research proposal, structured around three distinct, bottom-up pillars: immediate transactional friction via path coherence, internal pacing via working-time normalization, and environmental pacing via multi-scale seasonal feature selection. Our initial experimental results demonstrate that all three algorithmic artifacts significantly reduce predictive error compared to context-agnostic baselines. Finally, we outline the remaining roadmap to integrate these components into a unified design framework.

Keywords

Predictive Process Monitoring, Feature Engineering, Process Context, Machine Learning

1. Introduction

Predictive Process Monitoring (PPM) develops methods to predict the future unfolding of business process instances using data captured in event logs [1]. A specific subclass of these predictions focuses on time-related aspects, such as the remaining time until instance completion [2]. Accurately predicting these temporal properties generates significant value for businesses and customers [3] by facilitating resource scheduling [4] or providing information about waiting times [5].

Despite recent advances driven by sophisticated machine learning architectures, state-of-the-art approaches still yield substantial error rates [6]. For example, the Mean Absolute Error (MAE) in remaining time prediction frequently exceeds 10% of the true value, often translating to a variance of several days [2, 7]. These deviations severely limit the operational utility of PPM tools, as inaccurate forecasts lead to sub-optimal resource allocation and misinformed decision-making.

A fundamental contribution to those errors is a predominant case-centric perspective, which constructs predictions exclusively using the isolated instance's data payload [8]. Conversely, broader Business Process Management (BPM) research proves the benefits of considering the multi-dimensional execution *context* [9, 10]. Within PPM, context-aware models that, e.g., integrate concurrent workloads [11], resource experience [4], or environmental factors [12] achieve superior accuracy. This is vital because process instances never execute in isolation, but are deeply influenced by their shared operational environment [13].

While prior PPM literature extensively addresses design choices regarding log bucketing, trace encoding, and predictive model selection [2, 7], a critical phase in the predictive pipeline remains unexamined: systematic *feature engineering*. Feature engineering defines the explicit transformation of raw event log attributes into informative input representations for a model [1, 14]. For instance, raw timestamps from an event log can be transformed into features capturing the respective time within

Proceedings of the BPM 2026 Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Forum (BPM 2026) , Toronto, Canada, September 27th to October 2nd, 2026.
maximilian.alexander.vouk@s.wu.ac.at (M. Vouk)
0009-0007-0450-5048 (M. Vouk)

the current day or week, which matter in business contexts where resource availability is bound to typical working hours [15]. Consequently, feature engineering represents the primary mechanism for embedding complex contextual knowledge into predictive models before training.

In order to take advantage of this opportunity, this research aims to improve the predictive accuracy of time-related predictions via the means of incorporating contextual knowledge into engineered features.

2. Background

In this section, we present the background of this research. First, we outline the foundations of time-related PPM; second, we discuss the multi-dimensional layers that constitute process context; and third, we present feature engineering as a vehicle for context integration, including existing literature in the field.

2.1. Predictive Process Monitoring Foundations

PPM aims to predict the development of active business process instances (or cases) using historical event logs [1]. These logs consist of discrete events, each defined by a specific case ID, completed activity, and precise timestamp, alongside optional attributes like involved resources [16].

The *time-related prediction* problem focuses on temporal targets, primarily the remaining time until case completion, but also durations to the next event or specific milestones [2, 7]. To train predictive models, historical cases are divided into prefixes—snapshots of ongoing cases up to a specific point, which are used to estimate target durations [17]. These predictive *workflows* are operationalized through offline training and online testing phases (Figure 1). Because machine learning models require structured inputs, raw logs are converted into feature vectors [2]. However, conventional techniques are strictly case-centric. They isolate internal data from specific instances while disregarding the broader operational environment [15, 18].

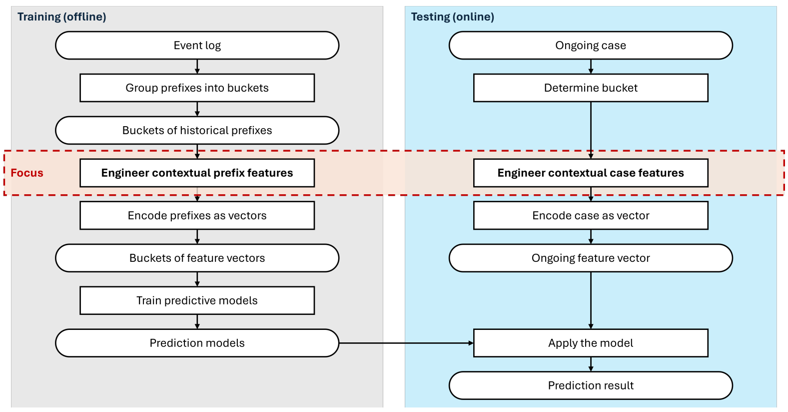

Figure 1: Focus of PhD research along workflow of predictive process monitoring (adapted from [2])

2.2. Dimensions of Context

The *context* of process executions is highly diverse and dependent on a variety of factors [10, 19]. Context can be defined as any information that characterizes the state of an entity [20], where in the setting of PPM the central entity is the business process itself [9]. To systematically structure this information, Rosemann et al. [9] propose a stratified layer model that categorizes the process environment into four concentric layers:

- **Immediate layer:** Elements directly facilitating process execution beyond the pure control flow. This includes essential input/output data, responsible organizational resources, and supporting IT systems.

- **Internal layer:** The wider organizational system exerting indirect influence. It encompasses corporate strategies, operational objectives, organizational culture, internal stakeholders, and logistical or financial infrastructure.
- **External layer:** The broader business network beyond the organization's direct control. It captures the significant impact of network stakeholders, including suppliers, competitors, investors, and customers, alongside their respective demands and limitations.
- **Environmental layer:** The overarching macro-environment entirely outside the business network. It comprises societal, natural, technological, and economic dimensions, including variables like weather conditions, holiday schedules, workforce shortages, and legislative regulations.

Categorizing the process environment into these distinct layers highlights the limitations of traditional PPM pipelines. Conventional workflows rely heavily on the *immediate context*, yet the temporal behavior of process instances also depends on the outer layers. Execution time fluctuations are rarely caused by an isolated case's data payload alone but rather, they stem from *internal* operational bottlenecks, *external* supplier failures, and *environmental* factors like holidays. Systematically integrating these multi-layered dimensions into predictive frameworks breaks the boundary of the immediate context, capturing the broader operational reality necessary to minimize prediction errors.

2.3. Feature Engineering in Predictive Process Monitoring

To effectively leverage context for time-related predictions, structural trace encoding must be distinguished from explicit feature engineering. While trace encoding formally shapes sequential logs into machine-readable vectors [2], it merely transforms the data's shape. In contrast, *feature engineering* semantically enriches this data into domain-specific attributes prior to model training. By explicitly extracting complex patterns rather than relying on predictive algorithms to discover them implicitly, feature engineering serves as the primary mechanism for integrating multi-layer context.

Transforming raw variables into domain-specific predictors is highly effective for incorporating context into PPM. While standard practices routinely engineer time-based features to capture resource availability [7, 14], some recent works expand across complex contextual boundaries. At the *internal layer*, Senderovich et al. [11] formalize cross-case competition by modeling queueing congestion, Klijn and Fahland [6] calculate operational distances to batch releases, and Kim et al. [4] capture human resource experience via task learning curves. Bridging the *internal and external layers*, Folino et al. [21] blend workload indicators with arrival times to isolate execution scenarios through predictive clustering. Finally, at the *environmental layer*, Yeshchenko et al. [12] dynamically create macro-environmental features using sentiment from online news repositories.

Despite these advancements, PPM lacks a unified framework for contextual feature creation. Because current literature is highly fragmented, effective techniques, such as seasonality extraction or resource friction modeling, operate in isolation. Without a cohesive architecture to organize and combine these multi-layered methods, context integration remains non-systematic and decoupled from an overarching taxonomy.

3. Research Agenda

The overarching goal of this dissertation is to deepen the understanding of how multi-layered contextual knowledge can be systematically embedded into data representations for time-related predictions in PPM. To guide this inquiry, this research addresses the following master research question:

RQ0: How can multi-layered contextual information be leveraged to improve time-related predictions in PPM?

To address the fragmented state of current literature and systematically resolve RQ0, this research adopts a pillar-based structure as illustrated in Figure 2. Each pillar investigates a distinct layer of the process context, progressing bottom-up from immediate transactional friction to external environmental factors. Finally, a comprehensive synthesis consolidates these findings, culminating in a unified architectural framework.

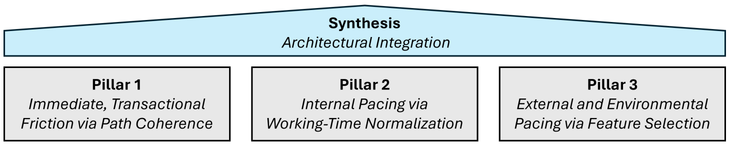

Figure 2: Structure of the PhD research agenda

3.1. Pillar 1: Immediate, Transactional Friction via Path Coherence

RQ1: How can path coherence be operationalized as a representation of immediate context, and to what extent does it improve remaining time prediction accuracy?

At the *immediate* context layer, process duration is heavily dictated by the micro-dynamics of execution, particularly the invisible coordination friction and setup costs generated when execution contexts shift between consecutive events. Traditional PPM models conventionally encode execution history by treating traces as sequences of structurally independent occurrences [7, 2], creating a fundamental methodological blindness to these transitional dynamics. This pillar [22] addresses this limitation by introducing a theory-informed approach to feature engineering that operationalizes the concept of *path coherence* [23, 24] as a predictive transition signal. By quantifying the continuity of contextual attributes, such as resource assignments or physical locations, between successive events, this algorithmic design explicitly models the social and operational friction inherent in organizational handoffs. The resulting artifacts capture the real-time continuity of the process state rather than static, isolated events.

3.2. Pillar 2: Internal Pacing via Working-Time Normalization

RQ2: How can working schedules be operationalized as a time-axis transformation representing internal context, and to what extent does it improve time until next event prediction accuracy?

At the *internal* organizational context layer, process velocity is strictly bounded by corporate policies, active shifts, and resource availability calendars. Traditional predictive models operate on standard clocktime, which introduces systemic bias by penalizing predictions during non-operational "dead time" (e.g., nights and weekends) when no resources are available to advance the case. This pillar [25] addresses this limitation by introducing an algorithmic design that warps the underlying temporal structure, transforming standard clocktime into a continuous worktime axis. By dynamically extracting active working intervals from event logs, this time-axis transformation eliminates non-working periods and is applied consistently to both input features and prediction targets.

3.3. Pillar 3: External and Environmental Pacing via Feature Selection

RQ3: How can seasonality be operationalized as a feature selection guideline representing external and environmental context, and to what extent does it improve remaining time prediction accuracy?

Reaching the outermost *external and environmental* layers, predictive accuracy is heavily constrained by temporal contingencies that exist outside the organization's direct control. While conventional encoding paradigms utilize standardized feature templates to capture these temporal dynamics, they are typically restricted to rigid, low-order cyclical indicators like basic daily and weekly rhythms. Consequently, they fail to account for more complex, long-term environmental patterns that systematically distort process execution times. This pillar addresses this limitation by developing a rule-based algorithmic design that automatically detects and tailors temporal feature sets to the multi-scale seasonal characteristics—such as monthly, quarterly, or yearly cycles—of the specific process at hand. Grounded in a conceptual analysis of the temporal feature space, this artifact systematically embeds external pacing trends into the features.

3.4. Synthesis: Architectural Integration

RQ4: How can multi-layered contextual feature engineering be integrated into a unified architecture, and to what extent do the individual components contribute to time-related prediction accuracy?

While the three individual pillars provide specialized mechanisms for isolated contextual layers, the final component of this research acts as their consolidation. Current literature remains fragmented because these disparate techniques operate in isolation. This synthesis addresses this gap by integrating the artifacts from all three pillars into a single, unified architecture. Crucially, this extensible pipeline allows for the future integration of additional approaches. To validate this architecture, an ablation study across heterogeneous event logs is employed. Selectively deactivating individual components isolates and quantifies the marginal accuracy gains attributable to each contextual layer. If feasible, a practical field-validation within an industrial case study will optionally supplement this evaluation. The ultimate outcome is a formalized set of prescriptive design principles establishing a reproducible blueprint for context-aware data representation in PPM.

On an overarching level, this research contributes to the field of *algorithm engineering* [26]. Specifically, the individual parts of the research focus on algorithmic designs, aiming to deepen the knowledge of and about designs that fulfill time-related predictive tasks in PPM. Methodologically, this objective is realized by adopting a design science research methodology [27]. Crucially, each component of this research structure introduces a novel artifact in the form of an algorithmic design. Following the process of design science research methodology, these design artifacts are operationalized via concrete software implementations and empirically evaluated against state-of-the-art baselines to demonstrate superior predictive accuracy on synthetic and real-world event data.

4. Current Progress and Outlook

This dissertation aims to demonstrate that incorporating multi-layered contextual knowledge via feature engineering significantly improves the accuracy of time-related predictions in PPM. The efficacy of all feature engineering approaches described by the three pillars has been empirically demonstrated in initial experiments based in Long Short-Term Memory (LSTM) model-based prediction workflows. When comparing against suitable, context-agnostic baselines they each yield a significant reduction of MAE. The approaches of Pillar 1 [22] and Pillar 2 [25] have subsequently been developed into papers that have been accepted for the BPM 2026 conference.

Moving forward, we will focus on extending Pillar 2 and further developing Pillar 3. For Pillar 2, the focus is on extending the currently proposed design, building upon findings of the initial paper, and strengthening the empirical evaluation. For Pillar 3, the focus is on refining the initial design and conducting the empirical evaluation. The final phase of this doctoral project will center on executing the overarching Synthesis. This will involve extensive empirical experimentation to assess different contextual feature representations in an ablation study, as well as potentially deploying the proposed solution together with an industry partner.

Declaration on Generative AI

During the preparation of this work, the author used Gemini 3.5 Flash in order to check spelling, grammar, and improve writing style. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] P. Ceravolo, M. Comuzzi, J. De Weerdt, C. Di Francescomarino, F. Maggi, Predictive process monitoring: concepts, challenges, and future research directions, Process Science 1 (2024) 2. doi:10.1007/s44311-024-00002-4.

[2] I. Verenich, M. Dumas, M. L. Rosa, F. M. Maggi, I. Teinemaa, Survey and Cross-benchmark Comparison of Remaining Time Prediction Methods in Business Process Monitoring, ACM Transactions on Intelligent Systems and Technology 10 (2019) 1–34. doi:10.1145/3331449.

[3] W. M. P. Van Der Aalst, J. Carmona (Eds.), Process Mining Handbook, volume 448 of *Lecture Notes in Business Information Processing*, Springer International Publishing, 2022. doi:10.1007/978-3-031-08848-3.

[4] J. Kim, M. Comuzzi, M. Dumas, F. M. Maggi, I. Teinemaa, Encoding resource experience for predictive process monitoring, Decision Support Systems 153 (2022) 113669. doi:10.1016/j.dss.2021.113669.

[5] W. Van Der Aalst, M. Schonenberg, M. Song, Time prediction based on process mining, Information Systems 36 (2011) 450–475. doi:10.1016/j.is.2010.09.001.

[6] E. L. Klijn, D. Fahland, Identifying and Reducing Errors in Remaining Time Prediction due to Inter-Case Dynamics, in: 2020 2nd International Conference on Process Mining (ICPM), IEEE, 2020, pp. 25–32. doi:10.1109/ICPM49681.2020.00015.

[7] E. Rama-Maneiro, J. Vidal, M. Lama, Deep Learning for Predictive Business Process Monitoring: Review and Benchmark, IEEE Transactions on Services Computing (2022) 1–1. doi:10.1109/TSC.2021.3139807.

[8] A. E. Marquez-Chamorro, M. Resinas, A. Ruiz-Cortes, Predictive Monitoring of Business Processes: A Survey, IEEE Transactions on Services Computing 11 (2018) 962–977. doi:10.1109/TSC.2017.2772256.

[9] M. Rosemann, J. Recker, C. Flender, Contextualisation of business processes, International Journal of Business Process Integration and Management 3 (2008) 47. doi:10.1504/IJBPIM.2008.019347.

[10] J. Vom Brocke, S. Zelt, T. Schmiedel, On the role of context in business process management, International Journal of Information Management 36 (2016) 486–495. doi:10.1016/j.ijinfomgt.2015.10.002.

[11] A. Senderovich, C. D. Francescomarino, F. M. Maggi, From knowledge-driven to data-driven inter-case feature encoding in predictive process monitoring, Information Systems 84 (2019) 255–264. doi:10.1016/j.is.2019.01.007.

[12] A. Yeshchenko, F. Durier, K. Revoredo, J. Mendling, F. Santoro, Context-Aware Predictive Process Monitoring: The Impact of News Sentiment, in: H. Panetto, C. Debruyne, H. A. Proper, C. A. Ardagna, D. Roman, R. Meersman (Eds.), On the Move to Meaningful Internet Systems. OTM 2018 Conferences, volume 11229, Springer International Publishing, 2018, pp. 586–603. doi:10.1007/978-3-030-02610-3_33.

[13] J. Mendling, N. Berente, S. Seidel, T. Grisold, The philosopher's corner: Pluralism and pragmatism in the information systems field: The case of research on business processes and organizational routine, ACM SIGMIS Database 52 (2021) 127–140. doi:10.1145/3462766.3462773.

[14] R. S. Oyamada, G. M. Tavares, S. B. Junior, P. Ceravolo, Enhancing Predictive Process Monitoring with Time-Related Feature Engineering, in: G. Guizzardi, F. Santoro, H. Mouratidis, P. Soffer (Eds.), Advanced Information Systems Engineering, volume 14663, Springer Nature Switzerland, 2024, pp. 71–86. doi:10.1007/978-3-031-61057-8_5.

[15] N. Tax, I. Verenich, M. La Rosa, M. Dumas, Predictive Business Process Monitoring with LSTM Neural Networks, in: E. Dubois, K. Pohl (Eds.), Advanced Information Systems Engineering, volume 10253, Springer International Publishing, 2017, pp. 477–492. doi:10.1007/978-3-319-59536-8_30.

[16] W. Van Der Aalst, Process Mining, Springer Berlin Heidelberg, 2016. doi:10.1007/978-3-662-49851-4.

[17] H. Weytjens, J. De Weerdt, Creating Unbiased Public Benchmark Datasets with Data Leakage Prevention for Predictive Process Monitoring, in: A. Marrella, B. Weber (Eds.), Business Process Management Workshops, Springer International Publishing, 2022, pp. 18–29. doi:10.1007/978-3-030-94343-1_2.

[18] N. Navarin, B. Vincenzi, M. Polato, A. Sperduti, LSTM networks for data-aware remaining time prediction of business process instances, in: 2017 IEEE Symposium Series on Computational Intelligence (SSCI), IEEE, 2017, pp. 1–7. doi:10.1109/SSCI.2017.8285184.

[19] T. Da Cunha Mattos, F. M. Santoro, K. Revoredo, V. T. Nunes, A formal representation for context-aware business processes, Computers in Industry 65 (2014) 1193–1214. URL: https://linkinghub.elsevier.com/retrieve/pii/S0166361514001407. doi:10.1016/j.compind.2014.07.005.

[20] A. K. Dey, Understanding and Using Context, Personal and Ubiquitous Computing 5 (2001) 4–7. doi:10.1007/s007790170019.

[21] F. Folino, M. Guarascio, L. Pontieri, Discovering Context-Aware Models for Predicting Business Process Performances, in: R. Meersman, H. Panetto, T. Dillon, S. Rinderle-Ma, P. Dadam, X. Zhou, S. Pearson, A. Ferscha, S. Bergamaschi, I. F. Cruz (Eds.), On the Move to Meaningful Internet Systems: OTM 2012, Springer, 2012, pp. 287–304. doi:10.1007/978-3-642-33606-5_18.

[22] M. Vouk, T. Grisold, K. Revoredo, J. Mendling, Continuity matters: Improving predictive process monitoring based on path coherence, in: Business Process Management, Springer Nature Switzerland, 2027. Accepted for publication.

[23] B. T. Pentland, J. Recker, G. Wyner, Rediscovering Handoffs, Academy of Management Discoveries 3 (2017) 284–301. doi:10.5465/amd.2016.0018.

[24] I. Kim, B. T. Pentland, K. A. Frank, J. R. Wolf, Path Coherence and Disruption in Routine Dynamics, Organization Science 36 (2025) 862–880. doi:10.1287/orsc.2022.16749.

[25] M. Vouk, K. Revoredo, J. Mendling, WARP: Worktime-aware representation for predictive process monitoring, in: Business Process Management Forum, Springer Nature Switzerland, 2027. Accepted for publication.

[26] J. Mendling, H. Leopold, H. Meyerhenke, B. Depaire, Methodology of Algorithm Engineering, ACM Computing Surveys 58 (2026) 1–38. doi:10.1145/3769071.

[27] K. Peffers, T. Tuunanen, M. A. Rothenberger, S. Chatterjee, A Design Science Research Methodology for Information Systems Research, Journal of Management Information Systems 24 (2007) 45–77. doi:10.2753/MIS0742-1222240302.