Investigating the Value of Predictive Process Monitoring in Manufacturing Supply Chains

Thanh Phan[1,*]

[1]*Kühne Logistics University, Großer Grasbrook 17, 20457, Hamburg, Germany*

Abstract

The past fifteen years have seen continuous development in Predictive Process Monitoring (PPM) that optimizes for prediction accuracy. Yet, when this accuracy is actually valuable, and how PPM deployment affects operations and customer relationships, remain understudied. My PhD program responds to the call for building industry impact of PPM, and centers on its value in manufacturing supply chains. The first paper of the PhD proposes a contingency view linking two structural characteristics of the manufacturing process to the relative accuracy gain of PPM over state-based approaches. Using discrete event simulation, we find that the process's structural complexity accounts for a substantial part of the variation in PPM's accuracy advantage across 70 synthetic process configurations. Building on this, for the second and third papers I plan to study the consequences of PPM deployment, how it affects operational performance within the firm and customer relationships across the supply chain. I look forward to feedback on the first paper and on shaping this agenda.

Keywords

online time prediction, predictive process monitoring, manufacturing, operations management

1. Introduction

Predictive Process Monitoring (PPM) methods exploit historical process execution logs to generate real-time predictions about running instances, including remaining time and process outcomes [1]. In B2B manufacturing, PPM enables suppliers to update and improve the lead time prediction after every event completion, following real-time monitoring and intervention. The past 15 years have seen continuous development in PPM algorithms to optimize for prediction accuracy, starting from statistical, control-flow based models [2], to state-of-the-art models using LSTM and neural networks [3, 4]; as well as models for Object-Centric PPM [5, 6] and settings with resource contention [7, 8]. But to what extent are these valuable to actual manufacturing supply chains that are operating on state-based lead-time estimations? And how does the deployment of PPM impact operational performance and customer relationships for the focal firm? These are the inquiries guiding my PhD, which centers on remaining time (lead time) prediction in manufacturing.

The PhD consists of three papers. In the first paper (currently in the results analysis and writing phase), I study the predictive accuracy of PPM, which relies on the prefix of the running job to predict its remaining future, relative to classical OR approaches which utilize the aggregate state of the system (e.g. queueing models, time series), and identify in what type of process the prefix-based PPM algorithms add the most value compared to the traditional state-based approaches. The following two papers will expand on my PhD's objective of investigating and optimizing the value of PPM in manufacturing operations.

Section 2 of the proposal details the research problem, methodology, initial results as well as expected theoretical and practical implications of the first paper. Section 3 concludes with the expected contribution of my PhD to BPM literature and the PhD roadmap.

Proceedings of the BPM 2026 Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Forum (BPM 2026), Toronto, Canada, September 27th to October 2nd, 2026.

*Corresponding author.

thanh.phan@klu.org (T. Phan)

2. Paper 1: The Contingent Accuracy Advantage of PPM in Manufacturing Lead Time Prediction

2.1. Research Problem

In traditional Operations Research (OR) and Operations Management (OM), the most common approach to predict lead time is one that relies on the aggregate state of the system, including work in progress, throughput rates, and temporal momentum. Prominent examples include Queueing Theory [9, 10], and Time Series Analysis [11]. These predictors typically generate a single lead time prediction at the beginning of the case. As the process executes, this prediction is generally not revised to incorporate the case's trajectory so far.

PPM algorithms, which condition each prediction on partial traces, including the sequence of machines visited, the timestamps of each completed activity, and the elapsed times, naturally address this gap. Because a new prediction can be made after every completed activity, PPM models not only update dynamically as the job progresses but are also able to reflect on the ongoing case to predict its remaining future. We call this the case-path information advantage of PPM models against the state-based approaches grounded in OR/OM. Exploiting this case-path information from event logs is a relatively mature stream of work within PPM. Recent studies seek improvement through incremental model updating to counter concept drift [12], richer inter-case and resource-level features [13], and alternative encodings [14].

What remain understudied are the conditions under which the case-path information meaningfully improves predictions compared to the state-based approach. Because the two approaches are structurally different, they have been developed and evaluated in isolation. Even within PPM studies, most perform mass evaluation to find the most accurate modeling approaches using standard BPI Challenge datasets, without exploring how their relative performances are affected by the process itself as a central research problem. This leaves firms considering investment in PPM capabilities with no structural basis for assessing whether the accuracy gain will justify the cost. Our paper addresses this tension by studying what levels of process complexity benefit the most from the case-path information that PPM provides, and where aggregate system state information provides sufficient signal — a crucial understanding for building effective and efficient PPM pipelines. We formally frame our research question as:

Under what structural conditions of the manufacturing process does a PPM case-path approach offer the greatest accuracy gains over classical state-based approaches for remaining time prediction?

2.2. Conceptual Framework

We predict that the accuracy advantage of case-path prediction over state-based prediction depends on two structural characteristics of the manufacturing process.

The first characteristic is *Between-path cycle time variance* (D_{path}), which measures how much of total cycle time variance is explained by the choice of path — i.e., the sequence of machines a job visits:

$$D_{\text{path}} = \frac{\text{Var}_{\text{between}}}{\text{Var}_{\text{total}}}$$

When D_{path} is high, paths differ substantially in their typical completion times. Knowing which path a job is on is therefore highly informative. When D_{path} is low, all paths complete in roughly similar times and path knowledge adds little.

The second characteristic is *Path entropy* (H_{norm}), which measures the distribution of jobs across available paths, using normalized Shannon entropy:

$$H_{\text{norm}} = \frac{-\sum_{p=1}^{K} \pi_p \log \pi_p}{\log K}$$

where π_p is the proportion of jobs following path p and K is the number of distinct paths. When H_{norm} is high, jobs are spread across many paths and state-based approaches are expected to be less reliable.

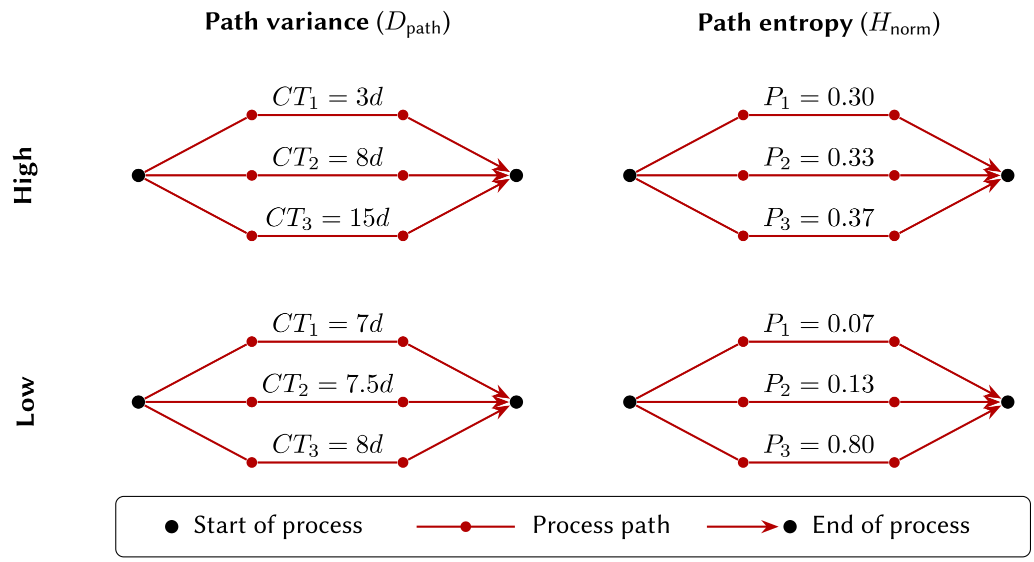

Figure 1: Illustrative examples of D_{path} and H_{norm} behavior

Together, these two dimensions form a contingency framework linking the structural properties of a manufacturing process to the predictive value of case-path information. We formalize the accuracy advantage as the Relative Accuracy Gain (RAG) of the PPM model over the best-performing state-based baseline, measured by mean absolute error (MAE):

$$\text{RAG} = \frac{MAE_{\text{state-based}} - MAE_{\text{PPM}}}{MAE_{\text{state-based}}}$$

We expect that the RAG value will be largest in processes with low H_{norm} and high D_{path}.

2.3. Methodology

Establishing the relationship between process structure and prediction accuracy requires controlled variation of structural parameters. This is difficult to achieve with empirical data where different real-world event logs come from vastly different settings. We therefore generate synthetic event logs using Discrete Event Simulation (DES) with configurable process parameters, building on the framework of [15].

The simulation models an open manufacturing line with probabilistic routing, exponential service times, Poisson arrivals, and FCFS queue discipline. Path variation arises at the machine level and through probabilistic rework loops. Four experimental factors are systematically varied: queue style, quality control pass rate, number of production stages, and machine heterogeneity, producing 81 synthetic event logs of which 70 pass stability checks and are retained for analysis.

Three models are implemented on each dataset: Flow Analysis (FA) [16] as the PPM representative, which is a model-based method that reconstructs process structure from the event log and uses the observed prefix to predict the remaining lead time; Simple Exponential Smoothing (SES) [17] and Little's Law [18] as state-based baselines. We chose FA as the PPM model because it is interpretable and well-suited for our configured simulation logs.

2.4. Initial Results

The FA model outperforms the state-based baselines across all 9 combinations of H_{norm} and D_{path} (Low, Medium, High, defined by 1/3 and 2/3 cutoff points). More importantly, the magnitude of the advantage varies substantially and predictably with process structure. Specifically, RAG correlates positively with both metrics, though the relationships are non-linear with a potential threshold. We also found a statistically significant interaction effect between the two dimensions: entropy matters more when variance is high, and vice versa. Together, D_{path} and H_{norm} explain $R^2 = 0.59$ of the variation in RAG across the 70 logs.

Notably, the effect of H_{norm} on RAG is positive — contrary to the initial expectation that low entropy would maximize PPM's advantage. Examining this result reveals a second, competing effect that we had overlooked: in low-entropy settings, where few paths dominate, state-based predictions naturally converge toward the dominant path, closing the accuracy gap with PPM. This suggests that a more comprehensive view of case-path information value should consider not only the process structure, but also the characteristics of the state-based baselines.

2.5. Theoretical and Practical Implications

Theoretical implications. We introduce a contingency perspective to PPM research, where studies typically evaluate model accuracy in isolation from the process context. By demonstrating that D_{path} and H_{norm} together account for a substantial share of the variation in PPM's accuracy advantage, we show that process structure is an important indicator that should be considered in future PPM research. This can benefit not only evaluative studies but also model engineering. For example, the two metrics can guide the decision to prioritize longer prefix lengths or richer state information. It also sheds light on PPM's relative disadvantage in low H_{norm} settings, where the dominance of one or a few paths allows state-based aggregates to reliably predict lead time without observing the case path.

Practical implications. For manufacturing firms, the contingency framework offers a diagnostic tool. Processes with high between-path variance and broadly distributed routing will benefit most from PPM adoption, while those with homogeneous, predictable routing may find that their existing OR tools already capture most of the predictive signal.

2.6. Threats to validity

The study was conducted using synthetic data to isolate the effect of the process structure, optimizing for internal validity. This inevitably means that external validity, which would require testing against empirical manufacturing data, is compromised. In terms of construct validity, one potential issue is the choice of the representative PPM method. We chose FA since its model-based approach works well with the simulated logs with cleanly predefined configurations. However, to the best of our understanding, path knowledge and inference are common characteristics of both model-based and supervised learning methods, even if the latter do not explicitly model it. Another issue is the need for a firmer theoretical grounding of H_{norm} and D_{path}, rather than the largely intuitive motivation presented here. I look forward to feedback on all of these issues at the Consortium.

3. Positioning in State of the Art BPM and PhD Roadmap

3.1. Positioning in State of the Art BPM

The BPM field has invested heavily in improving PPM algorithms. [19] observe that PPM research has been primarily driven by algorithmic engineering at the technical level, with comparatively less emphasis on organizational and ecosystem impacts. They identify improving industry impact as a major open challenge, outlining three specific directions: data scarcity, tooling, and industrial case studies. My PhD is both a direct response and an extension of this call. Paper 1 partially addresses the data scarcity issue: using synthetic data from simulation, I explore how the predictive accuracy of PPM relative to state-based baselines depends on the structure of the process, guiding industry adoption decisions. I extend the issue of industry impact further with two dimensions that form the prerequisite to meaningful industry adoption:

The first dimension is contextual. Manufacturing processes remain grossly understudied in BPM, and especially in PPM, which predominantly works with data from service processes. To better demonstrate the benefits of applying PPM techniques in any setting, one must first understand the setting itself, as it is critical to defining the role PPM should play within it. Another contextual layer is BPM's relation with OR/OM. Both disciplines share the goal of optimizing operational values for organizations, and there has been ample algorithmic work that combines elements from both [20]. However, when it comes to industry adoption, the relation can be competitive rather than collaborative. Paper 1 explores where PPM solutions stand in the industry landscape and their strengths and disadvantages in relation to traditional OR/OM tools.

The second dimension my PhD aims to address is consequential. The main motivation of PPM research is to support decision-making and improve organizational outcomes. Obviously, accuracy and explainability of the prediction must be continuously studied and improved. However, how they affect downstream operational, organizational and inter-organizational outcomes remains largely unexplored. In my second and third papers, I plan to study these aspects.

3.2. Roadmap

The tentative period for the entire PhD is September 2025 to August 2029. Paper 1 is currently in the results analysis and writing phase, with journal submission targeted by the end of 2026. Papers 2 and 3 turn from predictive accuracy to exploring the consequential aspect of PPM, namely how its deployment affects operational performance, customer relationships and business outcomes in manufacturing supply chains.

In paper 1, as presented in Section 2, we study the value of case-path monitoring in yielding more accurate, frequent and dynamic lead time estimates. This leads to another question: given that a supplier can update lead-time predictions more frequently and accurately, how should it communicate them? Drawing upon operational transparency [21] and schedule nervousness theory [22] from OM literature, in paper 2, we study how the frequency and framing of lead-time updates, independent of predictive accuracy, shape buyer trust and switching behavior, since volatile updates may erode trust even when the supplier is ultimately accurate. Paper 2 thereby connects PPM's predictive capability to its downstream outcomes across the supply chain. The precise scope for paper 3 is still being defined; here we turn to the within-firm lens, examining how PPM deployment affects operational performance within the focal firm.

I look forward to advice from senior researchers and fellow PhDs at the Consortium in shaping this agenda.

Declaration on Generative AI

The author used a generative AI tool (Claude, Anthropic) to assist with grammar checking and language polishing of the manuscript text. The tool was not used to generate research ideas, data, analysis, or conclusions.

References

[1] C. Di Francescomarino, C. Ghidini, Predictive Process Monitoring, in: W. M. P. Van Der Aalst, J. Carmona (Eds.), Process Mining Handbook, volume 448, Springer International Publishing, Cham, 2022, pp. 320–346. doi:10.1007/978-3-031-08848-3_10.

[2] 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.

[3] K. A. Elyasi, H. van der Aa, H. Stuckenschmidt, PGTNet: A Process Graph Transformer Network for Remaining Time Prediction of Business Process Instances, 2024. doi:10.48550/ARXIV.2404.06267.

[4] N. Navarin, B. Vincenzi, M. Polato, A. Sperduti, LSTM Networks for Data-Aware Remaining Time Prediction of Business Process Instances, 2017. doi:10.48550/ARXIV.1711.03822.

[5] J. N. Adams, G. Park, S. Levich, D. Schuster, W. M. P. van der Aalst, A Framework for Extracting and Encoding Features from Object-Centric Event Data, Springer Nature Switzerland, 2022, p. 36–53. URL: http://dx.doi.org/10.1007/978-3-031-20984-0_3. doi:10.1007/978-3-031-20984-0_3.

[6] T. K. Smit, H. A. Reijers, X. Lu, HOEG: A New Approach for Object-Centric Predictive Process Monitoring, Springer Nature Switzerland, 2024, p. 231–247. URL: http://dx.doi.org/10.1007/978-3-031-61057-8_14. doi:10.1007/978-3-031-61057-8_14.

[7] A. Senderovich, M. Weidlich, A. Gal, A. Mandelbaum, Queue Mining – Predicting Delays in Service Processes, in: D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, A. Kobsa, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, D. Terzopoulos, D. Tygar, G. Weikum, M. Jarke, J. Mylopoulos, C. Quix, C. Rolland, Y. Manolopoulos, H. Mouratidis, J. Horkoff (Eds.), Advanced Information Systems Engineering, volume 8484, Springer International Publishing, Cham, 2014, pp. 42–57. doi:10.1007/978-3-319-07881-6_4.

[8] A. Senderovich, M. Weidlich, A. Gal, A. Mandelbaum, Queue mining for delay prediction in multi-class service processes, Information Systems 53 (2015) 278–295. doi:10.1016/j.is.2015.03.010.

[9] S. T. Enns, Job shop leadtime requirements under conditions of controlled delivery performance, European Journal of Operational Research 77 (1994) 429–439. doi:10.1016/0377-2217(94)90408-1.

[10] R. E. Schelasin, Estimating wafer processing cycle time using an improved G/G/M queue, in: 2013 Winter Simulations Conference (WSC), IEEE, Washington, DC, USA, 2013, pp. 3789–3795. doi:10.1109/WSC.2013.6721738.

[11] I. Tirkel, Forecasting flow time in semiconductor manufacturing using knowledge discovery in databases, International Journal of Production Research 51 (2013) 5536–5548. doi:10.1080/00207543.2013.787168.

[12] 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 Trans. Serv. Comput. 18 (2025) 1308–1320.

[13] R. Aalikhani, M. Fathian, M. R. Rasouli, Enhancing remaining time prediction in business processes by considering system-level and resource-level inter-case features, Softw. Syst. Model. 24 (2025) 1515–1547.

[14] C. Liu, W. Liu, N. Guo, R. Song, Y. Gu, L. Cheng, Q. Zeng, Comparative evaluation of encoding techniques for workflow process remaining time prediction for cloud applications, J. Cloud Comput. Adv. Syst. Appl. 14 (2025).

[15] K. Busch, H. Leopold, Efficient allocation of shared resources across multiple processes, Information Systems 138 (2026) 102663. URL: http://dx.doi.org/10.1016/j.is.2025.102663. doi:10.1016/j.is.2025.102663.

[16] I. Verenich, H. Nguyen, M. La Rosa, M. Dumas, White-box prediction of process performance indicators via flow analysis, in: Proceedings of the 2017 International Conference on Software and System Process, ACM, Paris France, 2017, pp. 85–94. doi:10.1145/3084100.3084110.

[17] R. G. Brown, R. F. Meyer, The fundamental theorem of exponential smoothing, Operations Research 9 (1961) 673–685. URL: http://dx.doi.org/10.1287/opre.9.5.673. doi:10.1287/opre.9.5.673.

[18] J. D. C. Little, A Proof for the Queuing Formula: $L = \lambda W$, Operations Research 9 (1961) 383–387. doi:10.1287/opre.9.3.383.

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

[20] M. Kunkler, F. Schumann, S. Rinderle-Ma, A systematic review of business process improvement: Achievements and potentials in combining concepts from operations research and business process management, 2024. URL: https://arxiv.org/abs/2409.01276. doi:10.48550/ARXIV.2409.01276.

[21] R. W. Buell, M. I. Norton, The labor illusion: How operational transparency increases perceived

value, Manage. Sci. 57 (2011) 1564–1579.
[22] I. N. Pujawan, A. U. Smart, Factors affecting schedule instability in manufacturing companies, International Journal of Production Research 50 (2012) 2252–2266. URL: https://doi.org/10.1080/00207543.2011.575095. doi:10.1080/00207543.2011.575095. arXiv:https://doi.org/10.1080/00207543.2011.575095.