Explainable Artificial Intelligence Meets Uncertainty Quantification for Predictive Process Monitoring

Nijat Mehdiyev[1,2,*], Maxim Majlatow[1,2] and Peter Fettke[1,2]

[1]German Research Center for Artificial Intelligence (DFKI), Campus D3.2. Saarbrücken, Germany
[2]Saarland University, Campus D3.2. Saarbrücken, Germany

Abstract
The increasing reliance on Machine Learning (ML) models for critical decision-making necessitates the development of trustworthy systems that foster confidence among all relevant stakeholders. To address this need, Explainable Artificial Intelligence (XAI) aims to improve the transparency of ML models, enabling a better understanding of their decision-making processes. Concurrently, Uncertainty Quantification (UQ) has emerged as a crucial ML research area, emphasizing the estimation and communication of uncertainties inherent to model predictions. Despite the importance of both XAI and UQ in facilitating informed decision-making, there is a noticeable gap in integrating these techniques effectively. This paper highlights our recent research endeavors to explore the synergy between XAI and UQ for predictive process monitoring. Our first contribution, submitted to the Decision Support Systems journal, involves leveraging UQ to communicate the uncertainty present in ML explanations, ultimately promoting trust in the generated course of actions. Our second contribution, submitted to the Annals of Operations Research, employs XAI techniques to elucidate the factors contributing to ML model uncertainty, providing valuable insights for refining and enhancing the models. These insights are intended to shape future ML research in predictive process monitoring, fostering the development of more transparent, robust, and reliable systems.

Keywords
Predictive Process Monitoring, Explainable Artificial Intelligence, Uncertainty Quantification, Trustworthy AI

1. Introduction

Artificial Intelligence (AI) provides a promising avenue for corporations to transform their business and operational processes [1]. Recently, we have observed a surge in successful implementations of data-driven process analytics in different application domains, including but not limited to insurance, healthcare, manufacturing, public administration, and service management [2]. However, upon further examination, it becomes evident that the prevalent analytical methodologies - those of a descriptive and diagnostic nature - dominate industrial applications and commercial toolkits [3]. While significant progress has been made in machine learning-aided business process analytics and monitoring, a noticeable gap exists between academic advancements and their practical implementation. This lag can be attributed to users' trust and reliance on such AI-based systems [4].

The acceptance of AI as a credible source of guidance is still a hurdle to be overcome, signifying the need for transparency. In response, two promising areas of ML research have risen to prominence: Explainable Artificial Intelligence (XAI) and Uncertainty Quantification (UQ). XAI aims to demystify complex, opaque AI algorithms, making them understandable to human users. Various concepts, methods, and frameworks have been proposed recently to enhance this collaboration between AI systems and their human counterparts [5]. These strategies seek to increase the transparency and interpretability of the AI models, enabling stakeholders to gain a deeper understanding of the systems they interact with. Parallel to this, UQ is another intensively explored research field. Its focus lies in estimating and effectively communicating the uncertainty associated with an AI model's predictions [6]. It can be viewed as a complementary form of transparency, augmenting the explainability of solutions proposed for various decision-making tasks.

Despite the significance of both research areas in promoting better-informed decision-making, a gap exists in developing integrated, trustworthy solutions that combine XAI and UQ. This deficiency is particularly noticeable in the context of predictive process monitoring problems. While there is an increasing trend in XAI research to make predictive process monitoring methods more explainable [7, 8, 9, 10], and a few studies have begun to explore how to address uncertainty in this field [11, 12], the integration of these two aspects remains largely unexplored.

To address this gap, we have recently focused on the bidirectional integration of these research fields (see Figure 1). Our findings and progress in this endeavor are presented in this short paper. Section 2 discusses our first contribution, where we employ the chosen UQ approach

2nd International Workshop on Process Management in the AI era (PMAI), at 32nd International Joint Conference On Artificial Intelligence, August 19, 2023, Macao, S.A.R
nijat.mehdiyev@dfki.de (N. Mehdiyev); maxim.majlatow@dfki.de (M. Majlatow); peter.fettke@dfki.de (P. Fettke)

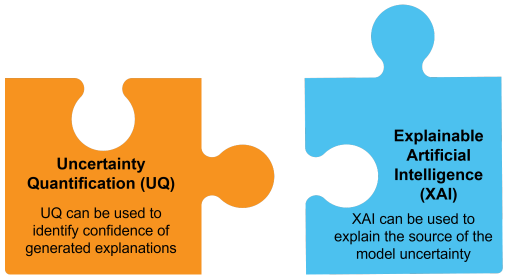

Figure 1: Bilateral Integration of UQ and XAI

to quantify and communicate the uncertainty inherent in ML explanations. This approach enhances trust and reliability in AI-powered predictive process monitoring solutions. Our second contribution, introduced in Section 3, takes a different direction. Here, we utilize XAI techniques to explain the factors contributing to uncertainty in ML models. This analysis provides valuable insights for refining and improving these models. Together, these findings and propositions have the potential to lay the groundwork for the development of more transparent, robust, and reliable AI systems for predictive process monitoring.

2. Communicating Uncertainty in Machine Learning Explanations

This section summarizes our recently proposed methodology for assessing and conveying uncertainty in ML explanations [13]. The focus of the manuscript[1] lies in examining the effective communication of model uncertainty in the explanations generated through global and local post-hoc explanation techniques, namely Individual Conditional Expectation (ICE) plots and Partial Dependence Plots (PDP).

Our proposed approach quantifies the model uncertainty using the Monte Carlo dropout technique, a well-established method in the deep learning domain. Utilizing this method allows us to not only generate point estimates from the posterior distributions but also to calculate corresponding credible intervals for assessing predictive uncertainties. However, a notable drawback of this UQ method is the absence of formal guarantees. To address this limitation, we implement Conformal Prediction, a post-processing technique designed to tackle this specific challenge [14]. Through conformalization, we can offer the theoretical guarantees necessary for ensuring the model's trustworthiness.

The pivotal step in integrating UQ with XAI in this study lies in constructing uncertainty profiles. These profiles are generated by processing conformalized credible intervals. Initially, the widths of these intervals are sorted in ascending order. Subsequently, the chosen thresholds are employed to categorize them, often determined through percentile-based estimations such as the 25th and 75th percentiles. In our specific use case, this methodology yields three distinct uncertainty profiles—low, medium, and high—that indicate the model's varying confidence levels. These profiles enhance the interpretability of PDP and ICE plots. Each profile is represented by a different color, offering a more nuanced understanding of the model's confidence in its predictions.

For ICE plots, the approach is based on a local explanation method that focuses on a single observation from the underlying dataset and a single predictor of interest. For each unique value of the predictor of interest, a copy of the selected observation is created, and the original value of the predictor in the copy is replaced with the unique value. The model predictions for this modified observation are then computed using the adopted UQ approach. The calculated conformalized credible interval for the model prediction is used to identify the uncertainty profile to which the observation belongs. This process is repeated for all unique values of the predictor. The pairs of predictor values and model predictions are then plotted to create the ICE plot, with each pair colored according to the uncertainty profile to which it belongs (see Figure 2). Additionally, conformalized credible intervals are incorporated in the same fashion and color-coding.

For uncertainty-aware PDPs, we consider a set of selected predictor variables. For each unique value of these selected predictors, a copy of the training dataset is generated with the original values replaced by this unique value. The average model prediction for this adjusted dataset is then computed. Concurrently, for each data point corresponding to a specific predictor value, a series of stochastic forward passes is performed to calculate credible intervals for model predictions. After conformalization, these credible intervals serve as the basis for identifying uncertainty profiles in PDP. Subsequently, we count the number of allocations to each uncertainty profile and place the majority profile for the examined predictor value, which is used to define the color in the PDP. This process is repeated for all unique values of the predictors. The pairs of predictor values and average model predictions are plotted to create the PDP, with each point being colored according to the majority uncertainty profile. Additionally, averaged conformalized credible intervals are integrated similarly to ICE plots and color-coded according to the predominant uncertainty profile for each point. Further visualizations, such as unconformalized credible intervals or pie charts representing uncertainty profile memberships for each predictor value,

[1]submitted to Decision Support Systems

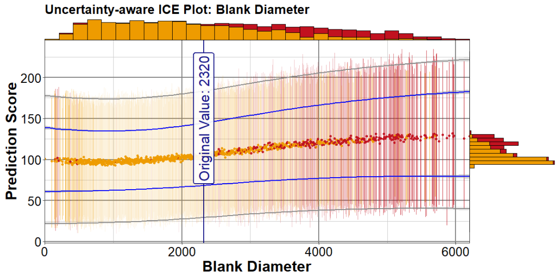

Figure 2: Uncertainty-aware ICE plot[13]

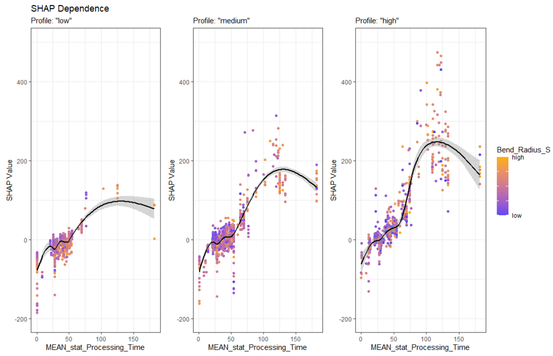

Figure 3: SHAP Dependence Plots for different Uncertainty Profiles [15]

are introduced for a more comprehensive understanding of the uncertainty distribution.

The study includes expert interviews to assess the suitability of the proposed approach and designed interface for a predictive process monitoring problem in the manufacturing domain.

3. Explaining the Machine Learning Uncertainty

Our second manuscript[2] employs an XAI technique to elucidate the factors contributing to ML model uncertainty [15]. The study revolves around a comprehensive, multi-stage ML methodology that interconnects information systems and AI to enhance decision-making within operations research (OR). The framework addresses common limitations of existing solutions, such as the lack of data-driven estimation for crucial production parameters, the generation of point forecasts without considering model uncertainty, and the absence of explanations for such uncertainties. The approach integrates various key technical elements to address these issues.

The study uses supervised learning to probabilistically estimate a production-related parameter, specifically the processing time of production events. This aspect of the method involves collecting and preparing process event data from Manufacturing Execution Systems (MES), which coordinate and track operational processes. The problem is treated as a predictive process monitoring problem, necessitating specific preprocessing, encoding, and feature engineering techniques to align with business and operational process data requirements.

To tackle model uncertainty, the study proposes using Quantile Regression Forests (QRF), an extension of the traditional Random Forests technique. QRF is designed for estimating conditional quantiles for high-dimensional predictor variables, offering a non-parametric and precise approach for estimating prediction intervals. An interval-based representation provides valuable insights into the uncertainty of model outcomes, facilitating a deeper comprehension of the model's predictive capabilities and inherent limitations.

To ensure model explainability, the study utilizes SHapley Additive Explanations (SHAP), which provide local and global post-hoc explanations of the model's uncertainties (see Figure 3). Our methodology represents a departure from traditional practices. Instead of relying on point predictions, we use prediction intervals as the output. This shift enables a more comprehensive analysis of how feature values influence prediction intervals. As a result, we can directly attribute features to the model's uncertainty. The explanations are refined to a granular level, offering a focused understanding of the factors contributing to uncertainty.

The proposed approach is demonstrated to be effective through a real-world production planning case study, highlighting the use of prescriptive analytics in refining decision-making procedures. The paper emphasizes the importance of fully leveraging the extensive data resources available for informed decision-making.

4. Conclusion

This short paper presents our two recent research contributions concerning the bidirectional integration of UQ and XAI. To the best of our knowledge, our study is the first to merge UQ and XAI within the context of predictive process monitoring problems. We are confident that this work lays the foundation for developing responsible AI solutions for this domain. Future research can further enhance these studies by incorporating insights from relevant related research areas, such as privacy-preserving AI, algorithmic fairness, reliability, and safety, as well as human-centered design. By amalgamating these diverse perspectives, we hope to contribute to the development

[2] submitted to Annals of Operations Research

of more robust, ethically sound, and user-oriented AI systems. This integrated approach aims to address not just the technical aspects of AI development but also the ethical and societal implications, leading to more holistic and beneficial AI solutions.

References

[1] J. Evermann, J.-R. Rehse, P. Fettke, Predicting process behaviour using deep learning, Decision Support Systems 100 (2017) 129–140.

[2] W. Van Der Aalst, A. Adriansyah, A. K. A. De Medeiros, F. Arcieri, T. Baier, T. Blickle, J. C. Bose, P. Van Den Brand, R. Brandtjen, J. Buijs, et al., Process mining manifesto, in: Business Process Management Workshops: BPM 2011 International Workshops, Clermont-Ferrand, France, August 29, 2011, Revised Selected Papers, Part I 9, Springer, 2012, pp. 169–194.

[3] N. Mehdiyev, P. Fettke, Explainable artificial intelligence for process mining: A general overview and application of a novel local explanation approach for predictive process monitoring, Interpretable Artificial Intelligence: A Perspective of Granular Computing (2021) 1–28.

[4] N. R. Sanders, K. B. Manrodt, The efficacy of using judgmental versus quantitative forecasting methods in practice, Omega 31 (2003) 511–522.

[5] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, D. Pedreschi, A survey of methods for explaining black box models, ACM Computing Surveys (CSUR) 51 (2018) 1–42.

[6] U. Bhatt, J. Antorán, Y. Zhang, Q. V. Liao, P. Sattigeri, R. Fogliato, G. Melançon, R. Krishnan, J. Stanley, O. Tickoo, et al., Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty, in: Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, 2021, pp. 401–413.

[7] M. Harl, S. Weinzierl, M. Stierle, M. Matzner, Explainable predictive business process monitoring using gated graph neural networks, Journal of Decision Systems 29 (2020) 312–327.

[8] A. Stevens, J. De Smedt, J. Peeperkorn, Quantifying explainability in outcome-oriented predictive process monitoring, in: Process Mining Workshops: ICPM 2021 International Workshops, Eindhoven, The Netherlands, October 31–November 4, 2021, Revised Selected Papers, Springer, 2022, pp. 194–206.

[9] M. Velmurugan, C. Ouyang, C. Moreira, R. Sindhgatta, Evaluating fidelity of explainable methods for predictive process analytics, in: Intelligent Information Systems: CAiSE Forum 2021, Melbourne, VIC, Australia, June 28–July 2, 2021, Proceedings, Springer, 2021, pp. 64–72.

[10] N. Mehdiyev, P. Fettke, Prescriptive process analytics with deep learning and explainable artificial intelligence, in: 28th European Conference on Information Systems (ECIS), An Online AIS Conference, 2020.

[11] H. Weytjens, J. De Weerdt, Learning uncertainty with artificial neural networks for predictive process monitoring, Applied Soft Computing 125 (2022) 109134.

[12] M. Shoush, M. Dumas, When to intervene? prescriptive process monitoring under uncertainty and resource constraints, in: Business Process Management Forum: BPM 2022 Forum, Münster, Germany, September 11–16, 2022, Proceedings, Springer, 2022, pp. 207–223.

[13] N. Mehdiyev, M. Majlatow, P. Fettke, Communicating uncertainty in machine learning explanations: A visualization analytics approach for predictive process monitoring, arXiv preprint arXiv:2304.05736 (2023).

[14] G. Shafer, V. Vovk, A tutorial on conformal prediction., Journal of Machine Learning Research 9 (2008).

[15] N. Mehdiyev, M. Majlatow, P. Fettke, Quantifying and explaining machine learning uncertainty in predictive process monitoring: An operations research perspective, arXiv preprint arXiv:2304.06412 (2023).