Process Management after ChatGPT: How Generative and Predictive AI Relate to Process Mining

Prof.dr.ir. Wil van der Aalst
Process and Data Science, RWTH Aachen University, D-52056 Aachen, Germany.
WWW: www.vdaalst.com, E-mail: wvdaalst@pads.rwth-aachen.de

This short article discusses the role of generative and predictive AI in the context of process management, focusing on process mining in particular. Generative and predictive AI will significantly impact process management, but dedicated process-mining approaches remain essential. Organizations need to master basic process mining techniques before advancing to more sophisticated AI and machine learning applications. As long as organizations struggle with finding and extracting data and have difficulties implementing obvious process improvements, it is not realistic to assume that a Large Language Model (LLM) trained on general-purpose data will be of much help. As the adage says: "Crawl before you walk, and walk before you run". Therefore, to ensure efficient, compliant, and optimized operational processes, one needs to take things step-by-step. This article aims to give directions for people interested in using these new technologies to improve their processes.

Understanding Process Mining and Its Types

Generative and predictive AI complement process mining in three different ways. Before explaining key relationships between the different technologies, we briefly summarize the essence of process mining. Process mining is a transformative approach in the field of business process management that leverages data-driven process-centric techniques to analyze and improve organizational processes [1,5]. It starts from event data that reflect the actual processes. An event has a timestamp and refers to an activity.

The main types of process mining are [1]:

- **Discover**: Process discovery involves extracting process models from event logs, providing an accurate visualization of the real process flows within an organization.
- **Check**: Conformance checking focuses on comparing the actual process, as recorded in event logs, with the pre-defined process model, highlighting deviations and compliance issues.
- **Predict**: This type of process mining aims at improving existing processes based on insights from event logs. It is part of the broader set of process (model) enhancement techniques.

The importance of process mining lies in its ability to offer real-time insights, operational efficiency, and compliance adherence.

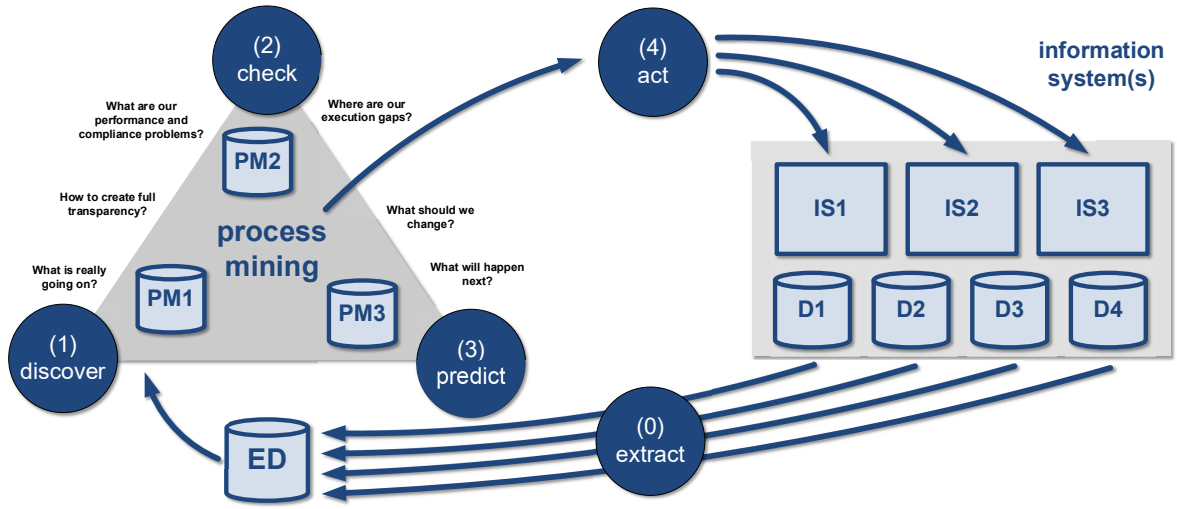

Figure 1: Overview of the different process mining steps.

The above diagram positions the three types of process mining. However, the diagram also shows two additional steps:

- **Extract**: Finding, extracting, and transforming data in the source systems (e.g., ERP and CRM systems). Data may be scattered over multiple systems, each having hundreds or even thousands of tables.
- **Act**: Translating the insights from process mining into improvements. This includes organizational change and automation. Revealing execution gaps does not automatically lead to improvements. People tend to slip back into inefficient behavioral patterns. Moreover, process mining may show the potential to automate parts of the process, but this requires dedicated implementation efforts.

These last two steps are not really specific to process mining but often turn out to be the biggest challenges in real-life applications of process mining. The technology itself works without problems. The main inhibitors are often cultural or organizational.

Object-Centric Process Mining (OCPM), a relatively new approach, expands the scope of traditional process mining by focusing on multiple objects (like orders, invoices, and deliveries) simultaneously [2,3,4]. This approach is particularly important in complex environments where traditional, case-centric process mining may fall short. The approach recognizes that real-life events and processes often involve various interconnected objects. OCPM aims to provide a more comprehensive and accurate understanding of the process by examining how these different objects interact and influence each other within the process flow.

In traditional event data, an event refers to a single case. In OCPM, an event can refer to any number of objects [2,3,4]. Objects and events are typed, and objects may refer to other objects. There is no need to pick a specific case notion, i.e., all events and objects can be stored in a system-agnostic manner. Hence, it is possible to view all operational activities from any perspective using a single source of truth. There is no need to extract the data when changing the viewpoint. This allows for flexibility, using on-demand process-mining views. Traditional distortions like divergence and convergence can be avoided. Finally, it is possible to see problems involving multiple object types. Note that performance and compliance problems often live at the intersection points of processes and organizational units. For example, delayed customer orders may be explained by production problems, outbound logistic problems, problems in procurement, or understaffing of the sales department.

Generative vs. Predictive AI

The distinction between generative and predictive Artificial Intelligence (AI) is crucial in understanding their roles in process mining. **Generative AI** focuses on creating new data or patterns, like ChatGPT, which generates human-like text. **Predictive AI**, on the other hand, analyzes existing data to predict future outcomes. Advances in AI are mostly fueled by breakthroughs in Machine Learning (ML); therefore, the terms are often used interchangeably.

Predictive AI is designed to analyze existing data and make predictions about future events or outcomes. It's extensively used in forecasting, risk assessment, and trend analysis. Most of the techniques have been around for quite some time, e.g., regression analysis, time-series analysis, and various machine learning models, ranging from decision trees to neural networks. Predictive AI relies on historical data to learn patterns and relationships. The output is usually a specific prediction, a proposed decision (recommendation), or a probability score, indicating the likelihood of a future event or outcome. In most cases, only data specific to the problem are used.

Generative AI, on the other hand, focuses on creating new content or data that are similar to but distinct from the training data. It can generate text, images, music, and more. Unlike predictive AI, the focus is not a specific phenomenon, and huge amounts of data not specific to the problem are used. The output is new, original content that did not exist before, created based on the patterns learned from the input data. Large Language Models (LLM), like ChatGPT, try to generate new, original content that mimics the learned data.

Large Language Models (LLMs) like ChatGPT have experienced a rapid uptake across various industries due to their advanced natural language processing capabilities. Their ability to generate human-like text, answer queries, and automate tasks is revolutionizing how businesses and individuals interact with technology, fostering innovation and efficiency. LLMs can also generate simple programs and database queries. This is actually unsurprising because Python and SQL can be seen as languages and many example “texts” are available. Nevertheless, generative AI has a huge potential and is more transformative than classical predictive AI because it takes over tasks formerly done by humans.

Integrating Process Mining with AI

Due to the spectacular developments in AI and ML, especially since the release of ChatGPT by OpenAI on November 30, 2022, organizations are eager to use AI and ML for process improvement. It is clear that human tasks can be partly automated. For example, Microsoft Copilot, a tool integrated into Microsoft 365 applications, helps accelerate tasks typically done by people. In applications like Excel, Copilot can analyze data, create complex formulas, and offer insights, significantly reducing the time and expertise required for data management tasks. The creation of PowerPoints and Word documents and e-mail interactions using Outlook can be accelerated. These examples have in common that content is created, or interactions with software become more human-like. However, thus far, AI (both generative and predictive) does not play a role in the core process mining algorithms for process discovery and conformance checking.

While traditional process mining techniques like process discovery and conformance checking do not benefit from mainstream AI and ML methods, they complement each other.

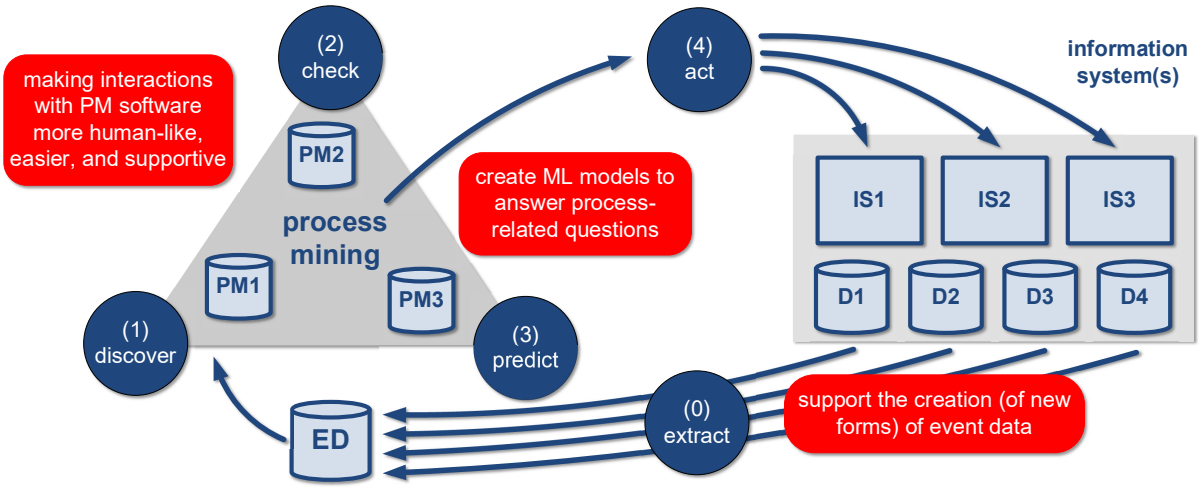

Figure 2: AI/ML techniques complement process mining in three main areas: data extraction, building models to answer process-related questions, and making interactions more human-like, easier, and supportive.

As the above diagram shows, there are three areas where AI/ML techniques complement process mining:

- **Support the creation (of new forms) of event data.** AI/ML techniques facilitate the extraction of data from source systems and can be used to create we forms of event data. Process mining requires structured event data, just like a spreadsheet needs numbers. However, it is possible to turn unstructured data into structured data using AI/ML techniques. For example, textual messages can be classified using supervised learning techniques (e.g., this message is a "service request", and the customer is "angry", or this sensor is now malfunctioning). Other examples: ML techniques can be used to decide whether two scanned invoices are duplicates to avoid double payments, and LLMs can also be used to generate SQL code to extract data from source systems. These examples show that the threshold to create event data can be substantially lowered using AI/ML.
- **Create ML models to answer process-related questions.** Between 2005 and 2010, process mining was extended to include conformance checking, decision mining, predictions, etc. [1]. Decision mining analyzes choice points in the process and builds an ML model to predict or explain a choice. For example, a decision tree can be discovered for a choice and subsequently turned into logical conditions for activities. At the same time, we developed techniques to predict or explain the remaining flow time of a case using ML models. The same ideas can be applied to conformance checking. For example, given the characteristics of a case, we can predict whether a case will deviate or not. The general approach is to create so-called "situation tables" as input for ML techniques. Often, plain tabular data suffice (e.g., one row per decision or case), and therefore, a wide range of ML techniques can be used.
- **Making interactions with PM software more human-like, easier, and supportive.** Generative AI is quickly changing the way that people interact with software, including process mining software. Users would like to pose process-related questions in natural language. Just like ChatGPT can generate SQL queries, it is possible to ask a process mining tool to create process-related queries. Example questions include: "What is the biggest bottleneck in this production process?", "Which suppliers are causing production delays?", and "What do these deviating cases have in common?".

Hence, AI can help in generating and interpreting new types of event data, which can then be utilized for more advanced process mining, and the use of LLMs like ChatGPT can simplify interactions with process mining software, making it more accessible and user-friendly.

Note that a **general-purpose pre-trained LLM** knows nothing about the process, the organization, and specific objects (e.g., customers, suppliers, etc.). However, without this knowledge, it is still possible to answer questions. There are two main approaches:

- No specific data is exchanged, but the LLM helps to generate the query based on metadata about tables, etc.
- Abstracted data is added to the prompt. One can encode trace variants, their frequencies, and flow times to the prompt. It is also possible to encode the Directly-Follows Graph (DFG).

All kinds of mixtures of the above are possible [6,7]. For example, Retrieval Augmented Generation (RAG) is a technique that combines prompt engineering with context retrieval from external data sources to improve the performance and relevance of LLMs. GPT-4 is able to load and analyze CSV, Excel, and JSON files with tabular data. Note that this requires sharing data with OpenAI. However, using the Azure OpenAI Service, it is possible to use GPT-4 without sharing data. In the future, it may even be possible to train LLMs specifically for an organization using information about suppliers, customers, products, etc. Note that ERP systems like SAP may contain hundreds of thousands of tables with information on an organization and its processes. This cannot be added to a prompt. However, creating a dedicated LLM would provide access to this untapped information.

Object-Centric Process Mining (OCPM) provides a way to structure information about objects and events in a system-agnostic manner. Source systems store data in a structured format. Information about event types and object types provides the ontological information needed to avoid well-known hallucinations. When adding up numbers, one should not “guess” the sum, but simply “compute” the sum. When data is structured, one should use the structure and not throw it into a “soup of data”, hoping that the LLM will get it right. Therefore, Object-Centric Event Data (OCED [8]) and OCPM [2,3] are key enablers of a more hybrid approach combining structured and unstructured data.

Conclusion: The Future of Process Management

Generative and predictive AI, exemplified by advancements like ChatGPT, will significantly impact process management. However, the need for dedicated process mining techniques remains the key to success. Just like ChatGPT cannot replace basic tools like calculators, AI cannot replace core process-mining techniques, such as process discovery and conformance checking. Process models annotated with performance and compliance problems remain crucial for organizations to understand and optimize their processes. As we move forward, it's essential for organizations to first master the “process and data management basics” before embracing more sophisticated forms of AI and ML. The future of process management lies in a balanced integration of Object-Centric Process Mining (OCPM) and AI, leading to more efficient, compliant, and optimized operational processes.

References

[1] W.M.P. van der Aalst. Process Mining: Data Science in Action. Springer-Verlag, Berlin, 2016.

[2] W.M.P van der Aalst. Object-Centric Process Mining: Unraveling the Fabric of Real Processes. Mathematics, 11(12):2691, 2023.

[3] W.M.P van der Aalst, Object-Centric Process Mining: The Next Frontier in Business Performance, 2023. celon.is/OCPM-Whitepaper.

[4] W.M.P. van der Aalst and A. Berti. Discovering Object-Centric Petri Nets. Fundamenta Informaticae, 175(1-4):1–40, 2020.

[5] W.M.P. van der Aalst and J. Carmona, editors. Process Mining Handbook, volume 448 of Lecture Notes in Business Information Processing. Springer-Verlag, Berlin, 2022.

[6] A. Berti, D. Schuster, and W.M.P. van der Aalst. Abstractions, Scenarios, and Prompt Definitions for Process Mining with LLMs: A Case Study, NLP4BPM'2023, 2023

[7] A. Berti and M.S. Qafari. Leveraging Large Language Models (LLMs) for Process Mining, 2023.

[8] Process and Data Science Group: Object-Centric Event Log (OCEL 2.0) Website Providing Data Sets, Software, Etc. www.ocel-standard.org, 2023.