From Theory to Practice: Real-World Use Cases on Trustworthy LLM-Driven Process Modeling, Prediction and Automation

Peter Pfeiffer

peter.pfeiffer@dfki.de German Research Center for Artificial Intelligence (DFKI) and Saarland University 66123 Saarbrücken, Saarland, Germany

Maxim Majlatow maxim.majlatow@dfki.de German Research Center for Artificial Intelligence (DFKI) and Saarland University 66123 Saarbrücken, Saarland, Germany

Abstract

Traditional Business Process Management struggles with rigidity, opacity, and scalability in dynamic environments, while emerging Large Language Models (LLMs) present transformative opportunities alongside risks. This paper explores four real-world use cases that demonstrate how LLMs, augmented with trustworthy process intelligence, redefine process modeling, prediction, and automation. Grounded in early-stage research projects with industrial partners, the work spans manufacturing, modeling, life-science, and design processes, addressing domain-specific challenges through human-AI collaboration. In manufacturing, an LLM-driven framework integrates uncertainty-aware explainable Machine Learning with interactive dialogues, transforming opaque predictions into auditable workflows. For process modeling, conversational interfaces democratize BPMN design. Pharmacovigilance agents automate drug safety monitoring via knowledge-graph-augmented LLMs. Finally, sustainable textile design employs multi-agent systems to navigate regulatory and environmental trade-offs. We intend to examine tensions between transparency and efficiency, generalization and specialization, and human agency versus automation. By mapping these trade-offs, we advocate for context-sensitive integration - prioritizing domain needs, stakeholder values, and iterative human-in-the-loop workflows over universal solutions. This work provides actionable insights for researchers and practitioners aiming to operationalize LLMs in critical BPM environments.

CCS Concepts

• Information systems \rightarrow *Extraction, transformation and loading; Expert systems; Data analytics; Process control systems.*

2025 ACM SIGMOD/PODS, Berlin, GER

Alexander Rombach

alexander_michael.rombach@dfki.de German Research Center for Artificial Intelligence (DFKI) and Saarland University 66123 Saarbrücken, Saarland, Germany

Nijat Mehdiyev nijat.mehdiyev@dfki.de German Research Center for Artificial Intelligence (DFKI) and Saarland University 66123 Saarbrücken, Saarland, Germany

Keywords

Business Process Management, Process Prediction, Process Modeling, Large Language Models, Explainable AI, Human Computer Interaction

ACM Reference Format:

1 Introduction

The integration of Artificial Intelligence (AI) into Business Process Management (BPM) has long been recognized as a pathway to operational efficiency [40]. Nevertheless, conventional approaches – reliant on manual process modeling, rigid statistical methods, and non-transparent Predictive Process Monitoring (PPM) systems – struggle to address the dynamism and complexity of modern industrial workflows [5, 23]. While Process Mining (PM) has enabled organizations to visualize and analyze event logs, critical gaps persist: static models fail to adapt to real-time operational shifts, domain experts face technical barriers [17, 34], and "black-box" AI systems weaken trust in high-stakes decision-making [2, 13, 35]. These limitations are particularly evident in industries such as manufacturing, healthcare, and sustainable design, where urgent customization, regulatory compliance, and multi-stakeholder collaboration demand flexible, transparent solutions [31].

Recent advances in Large Language Models (LLMs) and Explainable AI (XAI) signal a transformative shift, framing AI not just as an automation tool but as a collaborative partner in process management [36]. LLMs' ability to interpret natural language, generate contextual insights, and interact with diverse data sources aligns with the evolving demands of BPM, where adaptability and human-AI collaboration are essential [38]. However, realizing this potential requires addressing practical challenges: How can LLMs bridge the expertise gap between process modeling experts and domain specialists? Can AI-driven systems balance automation with accountability in regulated environments? What technical innovations are needed to scale these solutions across industries?

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

^{© 2025} Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN XXX-X-XXXX-XXXX-X/XXXX/XX https://doi.org/XXXXXXXXXXXXXXXX

This paper addresses these questions through four real-world use cases, each stemming from recently funded research initiatives that explore distinct applications of LLMs in BPM along the BPM lifecycle [6]. Spanning manufacturing, cross-industry process modeling, drug safety monitoring, and sustainable textile design, these projects demonstrate how LLMs can support activities along the entire BPM lifecycle — from discovery and execution to monitoring and optimization — while tackling domain-specific challenges:

- In LLM-mediated interaction with trustworthy process predictions, a framework integrating uncertainty-aware Machine Learning (ML), PM, and multi-agent LLMs transforms opaque predictions into auditable, interactive workflows. By grounding explanations in predictions from Manufacturing Execution System (MES) event logs and enabling natural language dialogues, the system empowers users to validate and adapt AI recommendations in real-time.
- For **conversational process modeling**, a chat-based interface allows users to model and refine Business Process Model and Notation (BPMN) models in an interactive manner, democratizing access to process design and optimization.
- In **pharmacovigilance**, knowledge graph (KG)-augmented LLM agents support regulatory experts in drug safety monitoring by synthesizing unstructured reports and structured medical domain knowledge.
- For **sustainable textile design**, a multi-agent LLM system navigates regulatory, environmental, and social trade-offs, empowering designers to align creative choices with circular economy principles.

Though distinct in scope, these cases share three transformative themes for BPM. First, they demonstrate human-centric AI, where LLMs act as bridges between technical systems and nonexpert users, whether through process visualizations or natural language queries. Second, they emphasize trustworthy automation through XAI techniques - such as uncertainty quantification (UQ) and Retrieval-Augmented Generation (RAG) - that anchor LLM outputs in verifiable data sources. Third, they expose scalability and performance challenges in integrating LLMs into legacy BPM infrastructures, from reconciling XML-based process models with generative AI to reducing errors in KG queries. By presenting these diverse perspectives, this paper contributes to next-generation process management in three ways. Practically, it shows how LLMs address industry pain points like manual process modeling bottlenecks. Methodologically, it highlights hybrid architectures that merge LLMs with traditional ML and participatory design. Theoretically, it challenges the view of LLMs as standalone solutions, advocating instead for their role in human-in-the-loop workflows.

The remainder of this paper is structured as follows: Sections 2–5 detail the four use cases, emphasizing their technical architectures, LLM integration strategies, and (initial) empirical findings. Section 6 synthesizes challenges and opportunities, framing a research agenda for AI-driven BPM in the LLM era. Finally, section 7 summarizes and concludes the paper.

2 Use Case 1 - LLM-Mediated Interaction with Trustworthy Process Predictions

2.1 Motivation

Modern manufacturing systems generate vast volumes of event data through MES, capturing granular details of production machine states, operator actions, material flows, and quality checks [25]. PM leverages these event logs to reconstruct and analyze production processes, offering insights into bottlenecks, inefficiencies, and deviations from planned schedules. While traditional PM excels at descriptive analytics, its capacity for predictive and prescriptive optimization remains limited, particularly in dynamic environments requiring real-time adaptation to urgent orders, material variability, or machine downtime [39].

ML has emerged as a promising tool to augment PM in production scheduling and planning [25]. PPM models, trained on historical event logs, can forecast production durations, optimize resource allocation, and simulate alternative workflows. However, their adoption introduces a critical challenge: the opacity of ML decision-making undermines trust in high-stakes scenarios. Production planners, tasked with validating schedules against contractual deadlines and quality standards, face a dilemma. They must reconcile ML's predictive accuracy with the need for intuitive, auditable reasoning — a requirement unmet by "black-box" ML models that obscure the rationale behind their outputs.

Current XAI methodologies fail to resolve these barriers [30]. Most approaches generate static explanations that do not align with the dynamic priorities of domain experts [21]. This creates a critical bottleneck: machine-generated explanations lack the flexibility to accommodate human intuition, limiting users' ability to interrogate, refine, or contextualize model outputs. Consequently, even accurate predictions risk rejection, perpetuating reliance on suboptimal traditional methods like expert estimations. The consequences of this gap are severe. Inaccurate or mistrusted predictions risk costly delays, contractual penalties, and resource underutilization. For example, overly conservative scheduling to hedge against opaque ML recommendations may idle machines or operators, while overly optimistic forecasts jeopardize on-time delivery.

2.2 Approach

The challenges outlined above, namely non-transparent ML-driven decision-making, misalignment between static explanations and human intuition, and the operational risks of mistrusted predictions, underscore a pressing need for frameworks that harmonize technical explainability with domain-specific interpretability. Traditional approaches treat explainability as a post-hoc justification layer, divorced from the dynamic, context-rich realities of production planning. To bridge this gap, we propose a process-driven human-AI collaboration framework that reimagines explainability as an iterative, process-aware dialogue. The conceptual framework, consisting of three modules, is shown in Figure 1.

Central to this approach is the integration of PM, which grounds ML predictions in the temporal and causal dependencies captured by MES event logs. Module 1 collects and prepares event logs from various manufacturing steps, such as pressing, bending, or grinding. Supervised learning is applied to forecast processing duration



Figure 1: Three-module AI solution that integrates human-centered ML, LLM-based interaction, and an intelligent AI interface

of production activities. Predictions are contextualized within historical workflows, linking delays to prior bottlenecks or resource shortages. XAI techniques - including Shapley-based approaches [22], Individual Conditional Expectation [11], and Partial Dependence Plots [8] - produce multiple layers of interpretability [28]. However, grounding alone is insufficient. To foster trust, the framework embeds UQ at its core, transforming probabilistic ML outputs into risk-aware decision boundaries. UQ methods (such as Monte Carlo Dropout [9] or Split Conformal Prediction [41]) quantify the reliability of these predictions. Planners no longer face a binary choice between accepting or rejecting a prediction; instead, they assess predictions or credible intervals (e.g., "Cutting time: 8.2h ± 1.1h") alongside PM insights to make informed trade-offs [24, 26].

Since transparency without interactivity risks perpetuating the status quo, our framework introduces Module 2: a multi-agent LLM architecture that transforms static explanations into dynamic conversations. By decomposing user queries into specialized tasks - explanation generation, scenario testing, validation - the system mirrors human problem-solving: probing assumptions, testing hypotheses, and refining conclusions. This agentic design ensures explanations adapt not just to the data, but to the planner's evolving priorities, whether optimizing for speed, cost, or resource fairness. The LLM module is refined through prompt engineering [10], ensuring domain-specific accuracy and mitigating bias to provide context-aware recommendations. In addition to its prompt-driven capabilities, the LLM layer employs a RAG [20] mechanism to store and retrieve previously generated or curated XAI explanations: Through vector-based indexing, the LLM ranks potentially relevant explanations according to semantic similarity, quantifies their relevance scores, and selectively incorporates them into its responses. This approach not only promotes consistency and coherence in the

explanations but also ensures that domain experts receive the most pertinent information for their context.

To translate these enriched conversations into actionable insights, the framework adds Module 3. This layer presents an adaptive, context-aware user interface that integrates dashboards, conversational widgets, and visual analytics. It renders predictions, uncertainty intervals, and process-mining traces in an intuitive format, and guides planners through what-if analyses. By surfacing the right information at the right granularity, the interface closes the loop between analytical depth and operational usability.

Ultimately, the framework positions PM as the backbone of a collaborative loop, where human expertise and AI predictions iteratively refine one another. Planners interrogate models through natural language, which in turn respond with process-contextualized insights, and shared understanding emerges from this dialogue. In doing so, the work moves beyond technical explainability to operational trust — a prerequisite for transforming predictive analytics from a reactive tool into a proactive partner in high-stakes manufacturing.

2.3 Experiences

Early experiments suggest that combining robust data pipelines with iterative user interfaces significantly enhances trust among domain experts [24–28]. Providing a dialogue-centric LLM interface encourages users from various domains to probe model outputs, interpret results more effectively, and adapt strategies on-the-fly. Moreover, from a data protection and compliance perspective, the LLM module supports the use of tools as well as systematic filtering and revision of outputs to ensure adherence to regulatory requirements. While the agentic LLM architecture addresses many domainspecific queries with high accuracy, ensuring model reliability under real-time production constraints remains a challenge. The complexity of industrial processes can lead to irregularities in data quality, requiring continual data curation and pipeline updates. Additionally, harmonizing automated explanation outputs with diverse user expectations demands a careful balance between technical detail and domain-level interpretability. Finally, overcoming skepticism about "black-box" AI systems calls for ongoing transparency measures, where the interplay of XAI and LLM-driven explanations must be demonstrably fair, unbiased, and consistent with operational requirements.

3 Use Case 2 - Conversational BPMN Modeling

3.1 Motivation

Business process models are a highly valuable source of information in business contexts [19]. They are necessary for the adequate design of information systems, requirements engineering, and for the means of process communication. However, creating these models requires significant manual effort and knowledge about modeling notations such as BPMN. At the same time, process modeling involves various stakeholders with different levels of expertise [6]. For example, a domain expert has knowledge about the underlying business workflows; however, they might not be able to create or understand a process model on their own [29]. Therefore, it would be beneficial to enable process modeling through natural language while eliminating the requirement of understanding the corresponding modeling notation. Furthermore, being able to ask questions about a certain process model could also improve understanding for all stakeholders.

For many years, researchers have explored automated text-tomodel approaches, ranging from rule-based methods to more advanced systems using Natural Language Processing (NLP) techniques. LLMs have emerged as a promising technology for this task, given their strong reasoning capabilities over natural language and ability to solve tasks not explicitly trained on. Thus, LLMs can also support conversational process modeling, allowing for interactive, collaborative process modeling with users in a chat-based environment [15].

A key challenge in this context is ensuring that LLMs can accurately interpret and generate BPMNs. Also, one needs to decide how the process models can be represented in a machine-readable format. We noticed that most open-source LLMs struggle to generate valid BPMN models, even though it's a frequently used modeling notation [18]. This calls for specializing LLMs for this task, e.g., by fine-tuning (preferably lightweight) LLMs. In particular, Supervised Fine-Tuning on the text-to-BPMN task can enhance their ability to generate valid and accurate BPMN models.

3.2 Approach

We implemented a prototypical chat environment that enables users to interact with an LLM in order to generate BPMN models of business processes. Users can prompt the LLM to generate a BPMN, modify an existing BPMN, or ask clarifying questions. This interactive system requires the LLM to understand different user intents, such as initial modeling, refinement, and updates. To guide Pfeiffer, Rombach, Majlatow and Mehdiyev



Figure 2: Architecture for Conversational Process Modeling

the LLM, a system prompt first defines the general setting and task at hand. It also defines the LLM's target output — in our case, a JSON object that can easily be parsed. It contains the textual response as well as the actual BPMN model. Additionally, one-shot prompting is used to provide an example output, which tends to improve response quality. The LLM then processes user requests, generates a new BPMN, updates an existing one, or answers questions about a BPMN model as needed.

In our scenario, the exchange format for BPMNs between the application and LLM was defined as the official BPMN-XML standard¹. To optimize efficiency, the layout of BPMN elements (i.e., precise coordinates) is handled algorithmically. This approach speeds up inference by eliminating the need for the LLM to generate layoutrelated tokens. At the same time, it prevents placement errors during the layouting procedure resulting from incorrect LLM outputs. Finally, a refinement loop was added to ensure BPMN validity, allowing the system to correct any issues before finalizing the model and presenting it to the user. The overall architecture and user collaboration is shown in Figure 2.

In preliminary experiments, we tested various open-source LLMs with sizes between 1B and 70B on publicly available text-model pairs [18]. To ensure the validity of the generated BPMN-XML models, we used an XML validator². We assess the quality of generated BPMN models based on different evaluation metrics, covering different modeling aspects, such as syntactic correctness or semantic expressiveness, compared to the ground truth models.

3.3 Experiences

At first, the quality of the LLM-generated BPMNs seems convincing. The LLMs are able to generate BPMN models that depict the description, with bigger LLMs generating BPMNs with more details compared to smaller ones. However, we observe various limitations throughout our experiments (see [18] for more details).

Our findings reveal that LLMs struggle with BPMN-XML as exchange format. The XML standard requires many characters to represent a process model, such as specific XML tags that define the XML schema. Therefore, the majority of generated tokens are spent on generating content that does not address the process model itself, but rather necessary structuring elements. Generating BPMN-XML is also error-prone and results in many invalid models, even though

¹https://www.bpmn.org/

²https://github.com/bpmn-io/bpmn-moddle/tree/main/resources/bpmn/xsd

2025 ACM SIGMOD/PODS, June 22-27, 2025, Berlin, GER

a refinement loop was added. Particularly, smaller LLMs struggled to produce valid and syntactically correct BPMNs. In contrast, larger LLMs performed noticeably better, as expected.

Further, the BPMNs have semantic deficits. Often, the LLMs struggle to identify all the relevant behaviors of a business process and thus generate BPMNs that do not contain all paths. This means that either a manual refinement or additional instructions on how to update specific details of the BPMN are necessary. We also observed that when prompted to alter an existing BPMN, LLMs often failed to incorporate all necessary changes, especially in terms of updating, removing, or adding the right sequence flows.

When asked to generate complex BPMNs, LLMs tend to produce faulty code or hallucinate non-existent XML tags. Therefore, it is advisable to introduce a more compact process model representation. In literature, different approaches are used, e.g., utilizing bullet point lists [12] or partially ordered workflow language, which can be converted to process models [16]. Further, Chain-of-Thought should be used to increase the accuracy of the generated models and to make the modeling intention of the LLM transparent to the user, fostering trust in the LLM and understandability of the BPMN.

Additionally, inference times were quite long, making larger LLMs less suitable for real-time, chat-based interactions. This highlights the need for lightweight LLMs for this use case to overcome scalability issues. Fine-tuning via techniques such as LoRA [14] can be an efficient way to enhance LLM performance in conversational process modeling. However, creating high-quality fine-tuning datasets can be challenging and time-consuming. So-called Preference Tuning can also be used to let the LLM obtain a deeper understanding of process modeling. For example, one could provide a textual process description and different corresponding process models - in this scenario BPMNs - with different levels of quality and let the LLM decide which process model fits best. Coming from the other direction, one could provide a process model as well as different textual descriptions that describe this process more or less well, and let the LLM decide which description is the best fitting one. In real-world settings, it's possible to model a certain process in different correct ways. It is precisely this aspect that could be trained with such an approach, where the LLM learns nuances in business process modeling. However, this also requires high-quality preference datasets as well as additional fine-tuning stages (e.g., leveraging Direct Preference Optimization [32]).

4 Use Case 3 - Monitoring Drug Safety with KGs and LLM-Agents

4.1 Motivation

The safety of medicines has to be monitored continuously throughout their use. Pharmacovigilance (PV) describes the activities carried out to ensure the safety of medicines in the market, i.e., the detection, assessment, understanding, and prevention of adverse effects with medicines. The PV lifecycle consists of the phases "collect, manage and analyze", "review and assess", "decide and act", "communicate", and "monitor" [7]. Information on medicines in the market has to be collected and analyzed, reviewed, and assessed by regulatory experts to decide whether and which actions are required. These actions need to be communicated and their consequences monitored.



Figure 3: Architecture of PV KG and LLM-Agents

These tasks require screening large amounts of textual reports and their assessment with domain knowledge. For instance, they involve querying and scanning relevant documents in public databases such as PubMed (scientific reports) and MAUDE (reports from medical experts and patients). Further, internal company case reports or social media posts can also be of interest. Domain knowledge in the form of publicly³ or internally available KGs is another relevant source of information, such as drug-drug interactions or terminologies. This requires querying graph databases. In order to analyze and decide about following actions in PV, experts need to interpret information with vast domain knowledge, e.g., to assess whether a reported issue is known or unknown, or even an adverse event.

LLMs' abilities can assist regulatory experts with PV. LLMs with tool access can retrieve, collect, and analyze reports from different data sources. Graph-structured databases can also be queried by LLMs [37], and the information linked with reports. This supports regulatory experts in the initial phases of the PV lifecycle. Additionally, LLM agents can support other lifecycle phases, like "communicate" and "monitor" by writing reports and using agents to perform the data collection phase continuously. Thereby, they support experts by relieving them from repetitive and time-consuming tasks such as information retrieval and initial analysis. However, data is distributed across various databases and stored in different formats such as text and graphs. Further, strict regulations apply when integrating AI systems into highly regulated environments. Therefore, a custom approach has been designed.

4.2 Approach

The main parts of the approach are outlined in Figure 3, consisting of database extractors (top row), NLP preprocessing, the PV KG, LLM agents, and a frontend to interact with users. Central is a custom PV KG containing entities and relations relevant to PV. Based on this, LLM agents, which are able to utilize the information in the PV KG and other KGs, are used to answer user queries and PV questions.

³e.g., https://bioportal.bioontology.org/

Database Extractor. Different database extractors retrieve relevant reports and other information from diverse data sources, e.g., scientific publications from PubMed, reports from MAUDE, or social media posts. Using language models, these texts are annotated and classified, and relevant entities are extracted, which allows information to be stored in a KG. If data is readily available as a graph, e.g., publicly available KGs, we link or integrate this data directly into the PV KG (e.g., terminologies like SNOMEDCT or RXNORM).

PV KG. A dedicated KG for PV is created, consisting of relevant information in a structured format, such as drug-drug interactions, adverse effects, populations, risks, and subsets of terminologies. Further, it should link to the source document, e.g., the PubMed report, social media post, or other KG, if detailed information is required. As such, it serves as a unified knowledge base from which to retrieve information.

LLM Agents. LLMs with tool access are used to retrieve information from the PV KG for analytical tasks. For this, LLM-based question answering over KG is employed [37], where the LLM queries the KG given a question from the user and answers that question based on the results from the KG query. By using structured knowledge in KG, resembling domain knowledge, the answers from the LLM should adhere to factual knowledge instead of only learned knowledge, reducing the risk of hallucinations [42]. LLM agents will be used for more complex tasks that require multiple steps, e.g., collecting data from databases that cannot be queried or more complex queries that require analyzing multiple reports and documents in depth. They are trained from recorded process executions by users to perform the same task similarly. Individual steps can be performed by LLMs with tool access. Thereby, they support humans in information retrieval and analysis.

Frontend. The frontend is used for displaying results and interacting with the databases and agents. It allows the user to manually query the databases, in case they want to analyze the data by hand. It also allows to trigger agentic workflows.

4.3 Experiences

Preliminary results indicate the overall approach's feasibility. For instance, it is feasible to retrieve, annotate, and extract entities of interest from PubMed and MAUDE and store them in the PV KG using language models. Small-sized LLMs are sufficient for most tasks, while larger LLMs are beneficial for complex reports. Additionally, LLMs are capable of formulating queries to the PV KG given the respective KG scheme, as well as interpreting the output. This should allow multi-step pipelines where the LLM decides to query the KG and/or other documents to extract relevant information. Next steps include validating the quality of the generated answers as well as the performance and accuracy of the results of the agentic workflows. Multi-step LLM agent workflows will be conceptualized and trained using the recorded process executions to validate the feasibility of instructing agents with click stream data. We expect the agents to take over a substantial amount of repetitive work with some applications on the analysis part, too. Further, a framework for evaluating the compliance of the approach has to be developed.

5 Use Case 4 - Multi-Agent Assistance System for the Textile Design Process

5.1 Motivation

The textile industry is one of the most significant global economic sectors and faces various challenges such as environmental concerns, resource inefficiency, waste generation, and the growing pressure to meet regulatory and sustainability standards [4, 33]. Increasing consumer demand for environmentally responsible products and more stringent regulations have created a shift in the textile industry towards sustainability and circular economy principles. However, this transformation is complex, requiring innovative solutions from the sourcing of raw materials to managing the end-of-life of textiles [1, 33].

In response to these challenges, the EU has launched a strategy for sustainable and circular textiles [3]. The strategy aims to promote sustainable practices, reduce environmental impact, and increase the efficiency of material usage throughout the textile supply chain. For instance, the Digital Product Passport (DPP) aims to enhance transparency throughout a fashion product's lifecycle. However, operationalizing this strategy is challenging. It requires industry-wide collaboration, significant technological advances, and the development of reliable decision-making frameworks based on data about materials, production processes, and supply chains. As a result, there is a pressing need for more intelligent, data-driven approaches to help manufacturers meet these challenges.

5.2 Approach

The goal is to implement a trustworthy AI system that processes heterogeneous data such as DIN Standards, EU guidelines, or material characteristics, to support textile designers during the design process. The system shall facilitate adherence to regulatory requirements and actively allow for sustainable textile designs. It will assist designers by providing actionable recommendations for sustainable practices across multiple design criteria.

For instance, the longevity of textiles can be enhanced by using durable fabrics. However, sometimes, a tradeoff has to be made between durability and recyclability, since not all durable fabrics can be recycled easily or infinitely. Similarly, tradeoffs between social and economic factors have to be made when choosing suppliers. The proposed system should support the designer in making informed decisions that align with sustainability goals. By analyzing the specific context of each design decision, the system will also generate explanations to ensure transparency and reliability. Further, DPPs should be created automatically by the system, thus reducing the designers' need for manual work.

The core of the solution will be a multi-agent LLM system to facilitate the interaction between designers and sustainability criteria. The first key component consists of specialized LLM-agents, as shown in Figure 4, that each perform a specialized task. For example, the regulation agent searches the knowledge base for relevant data on regulations and laws in order to integrate them into the design process. The material agent specializes in identifying materials and suppliers following the requirements requested by the designer. The underlying knowledge base includes various internal and external data sources. The user can upload company-owned



Figure 4: Architecture of Multi-Agent Textile Design Assistance System

data, such as historic orders or internal sustainability criteria. This shall allow the system to access relevant data at all times and ensure that the final results are tailored to the end-user. All agents are controlled by a central orchestrator, which controls the information flow between the agents to ensure an adequate execution.

Given the results of the specialized agents, the LLM-core is used to generate corresponding outputs. To build confidence in the system's recommendations, a key component will be trustworthy AI. The idea is to leverage XAI techniques such as counterfactual explanations, or UQ approaches like Bayesian approaches or Conformal Prediction. These methods are built on top of the LLMs and are used to analyze their outputs, which would otherwise be opaque from a human point of view. The final outputs and explanations will then be forwarded to the designer. This approach will allow for an easy integration of textual data and contextual inputs from various sources. Designers will be able to interact with the system to quickly generate ideas, validate choices, and access data-driven insights.

5.3 Experiences

We started implementing this use case by developing a recommender system for fabrics based on user preferences. Unfortunately, there is very little data on fabric combinations and their longevity publicly available, which is why expert interviews have been conducted. Thereby, we noticed that several trade-offs have to be made, which requires extensive explanation and communication with the designer about the consequences of substituting fabrics, other implications, e.g., social aspects, and alternatives. Soon, the impact of suppliers and their certificates became evident, which adds another dimension to the decision problem. Using LLMs as well as Trustworthy AI techniques should allow the unaware end-user to understand this multi-criteria decision problem and balance the effects on each dimension. For example, the system could identify that selecting a specific supplier for a specific material could reduce the production cost by a considerable amount while still adhering to the same sustainability criteria as before. There is an ongoing challenge on how to adequately integrate corresponding knowledge into the LLM system and how to make those trade-offs explainable to the end-user by means of XAI and related methods.

6 Discussion

6.1 Overview

Table 1 summarizes and compares the four use cases in terms of input modality, BPM lifecycle phases, industry impact, LLM Role, XAI perspectives, challenges, and future directions. As demonstrated throughout the paper, LLMs can assist various BPM use cases in different industries - from production to consulting and life-science to sustainable textiles. All BPM lifecycle phases can be supported with LLM techniques, impacting different business goals such as efficiency, compliance, and sustainability. To enforce trust in AIdriven BPM solutions, various trust mechanisms such as Shapley values, Chain-of-Thought, or KG grounding can be used. In all use cases, human interaction is enforced, and no decisions are made without human supervision.

6.2 Challenges

Each use case has its own challenges. In the first use case, generating real-time UQ and XAI outcomes in industrial PPM is a key challenge. Explanations have to balance technical details and domain-level interpretability to be trusted by the users.

In collaborative BPMN modeling, fine-tuning small LLMs to generate high-quality BPMN models in near real-time is of particular interest. The lack of high-quality datasets for this task makes training LLMs particularly challenging. Therefore, alternative training strategies have to be developed to enhance the modeling competencies of open-source LLMs.

Use Case	LLM-based Interaction with Predictions	Conversational BPMN Modeling	Pharmacovigilance with AI Agents	Sustainable Textile De- sign Process
Input Modality	Event logs	Text	Text, KG	Text, Tabular, Images
BPM Lifecycle Phases	Implementation, Execu- tion, Monitoring	Discovery, Redesign	Implementation, Moni- toring	Execution
Industry Impact	Transparent, process- driven production planning	Democratized process de- sign (cross-industry)	Supporting regulatory experts in medicine safety	Sustainable textile design
LLM Role	Multi-agent dialogues with Trustworthy AI- based PPM	Conversational BPMN translation	KG-QA agents for phar- macovigilance	Multi-agent orches- tration for trade-off resolution
Trust Mechanism	UQ, XAI, RAG	SFT, Chain-of-Thought	KG grounding	XAI, Algorithmic Fair- ness, RAG
Ethical Risks	Accountability in high- stakes decisions	Over-reliance on auto- mated modeling	Hallucinations in medical compliance	Bias in sustainability trade-offs
Human-AI Interaction	Multi-agent dia- logues with process- contextualized feedback	Natural language dia- logue	Agentic workflows with human oversight	Committee-based multi- agent negotiation
Key Challenges	Scalability	Improving BPMN quality	LLM-KG integration	Quantifying multidimen- sional trade-offs
Data Challenge	Event log complexity	Efficient BPMN notation modeling	Multisource integration (text + KG)	Multimodal integration (materials, regulations)
Future Directions	Bilateral integration of UQ and XAI; Object-centric PM	Lightweight LLMs for real-time BPMN model- ing	Hybrid KG-LLM architec- tures	Multi-objective optimiza- tion frameworks

Table 1: Comparative Analysis of LLM-Driven BPM Use Cases

To support regulatory experts in PV, challenges remain in combining LLMs with KGs. Especially when multiple KGs are available or multiple steps are required to reach conclusions. Instructing agents to perform business processes following human demonstrations will be another challenge to face.

For sustainable recommendations in textile designs, trade-offs in multiple dimensions and domains, which are difficult to quantify, have to be made. Whether LLMs can balance these trade-offs remains to be seen.

Spanning across all use cases, we see that more research is needed for process-aware LLMs. While LLMs can generate text that is difficult to distinguish from expert writing, LLM-generated BPMN models are often not valid XML files and contain syntactical errors. Further, agentic workflows need instructions, showing that such systems are not process-aware by design. However, we see the potential to transfer knowledge between the use cases to enhance the processual thinking of LLMs.

Challenges also remain when it comes to ethical risks. Overreliance on predictions or generated BPMN models, which can potentially be faulty, as well as biases in LLMs, poses the risk of leading to bad consequences. For regulatory use cases, hallucinations have to be controlled, and the LLMs have to be enforced to only use factual knowledge available in KGs.

7 Conclusion

This paper introduced four LLM-based use cases, demonstrating that all phases of the BPM lifecycle can be supported. These cases span different industries, such as manufacturing, consulting, lifescience, and sustainable textiles. They support companies by increasing efficiency, enhancing business processes, ensuring compliance, and adhering to diverse regulations. Various ways on how to apply LLMs in productive settings have been shown - from explaining process predictions, generating BPMN models, querying knowledge graphs to balancing trade-offs. In all cases, it emerged that human-machine interactions are essential for successful process executions. In addition, AI methods - and LLM applications in particular - should be supplemented by XAI techniques in order to make their predictions auditable for decision-making in real-world scenarios. Key challenges across the use cases remain in improving LLMs' capabilities to think in processes rather than support single tasks. Further, trustworthiness should be enforced by design, and ethical constraints should be considered from early on.

Acknowledgments

This research was funded in part by the German Federal Ministry of Education and Research under grant number 01IS24048C (EIN-HORN), 01IS24053C (KICoPro), and 01IS24046D (PVRadar).

2025 ACM SIGMOD/PODS, June 22-27, 2025, Berlin, GER

References

- Stefano Abbate, Piera Centobelli, Roberto Cerchione, Simon Peter Nadeem, and Emanuela Riccio. 2023. Sustainability trends and gaps in the textile, apparel and fashion industries. *Environment, Development and Sustainability* 26, 2 (feb 2023), 2837–2864. doi:10.1007/s10668-022-02887-2
- [2] Kamala Aliyeva and Nijat Mehdiyev. 2024. Uncertainty-aware multi-criteria decision analysis for evaluation of explainable artificial intelligence methods: A use case from the healthcare domain. *Information sciences* 657 (2024), 119987. doi:10.1016/j.ins.2023.119987
- [3] EUROPEAN COMMISSION. 2022. EU Strategy for Sustainable and Circular Textiles. https://ec.europa.eu/eurostat
- [4] European Commission. 2022. Questions and Answers on EU Strategy for Sustainable and Circular Textiles. Technical Report. EU.
- [5] Chiara Di Francescomarino, Chiara Ghidini, Fabrizio Maria Maggi, and Fredrik Milani. 2018. Predictive Process Monitoring Methods: Which One Suits Me Best?. In Business Process Management, Mathias Weske, Marco Montali, Ingo Weber, and Jan vom Brocke (Eds.). Springer International Publishing, Cham, 462–479. doi:10.1007/978-3-319-98648-7_27
- [6] Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. 2018. Introduction to Business Process Management. Fundamentals of Business Process Management (2018), 1-33. doi:10.1007/978-3-662-56509-4_1
- [7] EMA. 2021. EMA Pharmacovigilance System Manual (Version 1.3). (2021). https: //www.ema.europa.eu/en/human-regulatory-overview/pharmacovigilanceoverview
- [8] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. Annals of statistics (2001), 1189–1232. doi:10.1214/aos/1013203451
- [9] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48 (New York, NY, USA) (ICML'16). JMLR.org, 1050–1059.
- [10] Louie Giray. 2023. Prompt engineering with ChatGPT: a guide for academic writers. Annals of biomedical engineering 51, 12 (2023), 2629–2633.
- [11] Alex Goldstein, Adam Kapelner, Justin Bleich, and Emil Pitkin. 2015. Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics* 24, 1 (2015), 44–65. doi:10.1080/10618600.2014.907095
- [12] Michael Grohs, Luka Abb, Nourhan Elsayed, and Jana-Rebecca Rehse. 2024. Large Language Models Can Accomplish Business Process Management Tasks. In Business Process Management Workshops, Jochen De Weerdt and Luise Pufahl (Eds.). Springer Nature Switzerland, Cham, 453–465. doi:10.1007/978-3-031-50974-2 34
- [13] Vikas Hassija, Vinay Chamola, Atmesh Mahapatra, Abhinandan Singal, Divyansh Goel, Kaizhu Huang, Simone Scardapane, Indro Spinelli, Mufti Mahmud, and Amir Hussain. 2024. Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence. *Cognitive Computation* 16, 1 (jan 2024), 45–74. doi:10.1007/s12559-023-10179-8
- [14] Edward J Hu, Yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In International Conference on Learning Representations. https: //openreview.net/forum?id=nZeVKeeFYf9
- [15] Nataliia Klievtsova, Janik-Vasily Benzin, Timotheus Kampik, Juergen Mangler, and Stefanie Rinderle-Ma. 2023. Conversational Process Modelling: State of the Art, Applications, and Implications in Practice. In Business Process Management Forum, Chiara Di Francescomarino, Andrea Burattin, Christian Janiesch, and Shazia Sadiq (Eds.). Springer Nature Switzerland, Cham, 319–336. doi:10.1007/ 978-3-031-41623-1_19
- [16] Humam Kourani, Alessandro Berti, Daniel Schuster, and Wil M P van der Aalst. 2024. Process Modeling with Large Language Models. In Enterprise, Business-Process and Information Systems Modeling, Han van der Aa, Dominik Bork, Rainer Schmidt, and Arnon Sturm (Eds.). Springer Nature Switzerland, Cham, 229–244. doi:10.1007/978-3-031-61007-3_18
- [17] Aleksandra Kuzior, Mariya Sira, and Paulina Brożek. 2023. Use of Artificial Intelligence in Terms of Open Innovation Process and Management. *Sustainability* 15, 9 (apr 2023), 7205. doi:10.3390/su15097205
- [18] Chantale Lauer, Peter Pfeiffer, Alexander Rombach, and Nijat Mehdiyev. 2025. Conversational Business Process Modeling using LLMs: Initial Results and Challenges. In *EMISA 2025*. Gesellschaft für Informatik e.V., Bonn. doi:10.18420/ EMISA2025_03
- [19] Henrik Leopold. 2013. Generation of Natural Language Texts from Process Models. In Natural Language in Business Process Models. Vol. 168. 103–136. doi:10.1007/978-3-319-04175-9_5
- [20] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20). Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.

- [21] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/331831.3376590
- [22] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. Advances in neural information processing systems 30 (2017).
- [23] Fabrizio Maria Maggi, Chiara Di Francescomarino, Marlon Dumas, and Chiara Ghidini. 2014. Predictive Monitoring of Business Processes. In Advanced Information Systems Engineering, Matthias Jarke, John Mylopoulos, Christoph Quix, Colette Rolland, Yannis Manolopoulos, Haralambos Mouratidis, and Jennifer Horkoff (Eds.). Springer International Publishing, Cham, 457–472. doi:10.1007/978-3-319-07881-6_31
- [24] Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. 2024. Communicating uncertainty in machine learning explanations: a visualization analytics approach for predictive process monitoring. In *World Conference on Explainable Artificial Intelligence*. Springer, 420–438.
- [25] Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. 2024. Counterfactual Explanations in the Big Picture: An Approach for Process Prediction-Driven Job-Shop Scheduling Optimization. Cognitive Computation 16, 5 (2024), 2674–2700. doi:10.1007/s12559-024-10294-0
- [26] Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. 2024. Quantifying and explaining machine learning uncertainty in predictive process monitoring: an operations research perspective. Annals of Operations Research (2024), 1–40. doi:10.1007/s10479-024-05943-4
- [27] Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. 2025. Augmenting post-hoc explanations for predictive process monitoring with uncertainty quantification via conformalized Monte Carlo dropout. *Data & Knowledge Engineering* 156 (2025), 102402. doi:10.1016/j.datak.2024.102402
- [28] Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. 2025. Integrating permutation feature importance with conformal prediction for robust Explainable Artificial Intelligence in predictive process monitoring. *Engineering Applications of Artificial Intelligence* 149 (2025), 110363. doi:10.1016/j.engappai.2025.110363
- [29] Martin Meitz, Henrik Leopold, and Jan Mendling. 2013. An Approach to Support Process Model Validation based on Text Generation. *EMISA FORUM* 33, 2 (aug 2013), 7–20. doi:10.1007/BF03345896
- [30] Tim Miller. 2023. Explainable AI is Dead, Long Live Explainable AI! Hypothesisdriven Decision Support using Evaluative AI. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 333–342. doi:10. 1145/3593013.3594001
- [31] Sina Mohseni, Niloofar Zarei, and Eric D. Ragan. 2021. A Multidisciplinary Survey and Framework for Design and Evaluation of Explainable AI Systems. ACM Trans. Interact. Intell. Syst. 11, 3–4, Article 24 (Sept. 2021), 45 pages. doi:10.1145/3387166
- [32] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. Direct preference optimization: your language model is secretly a reward model. In Proceedings of the 37th International Conference on Neural Information Processing Systems (NIPS '23). Curran Associates Inc., Red Hook, NY, USA.
- [33] Héctor Guadalupe Ramírez-Escamilla, María Concepción Martínez-Rodríguez, Alejandro Padilla-Rivera, Diego Domínguez-Solís, and Lorena Elizabeth Campos-Villegas. 2024. Advancing Toward Sustainability: A Systematic Review of Circular Economy Strategies in the Textile Industry. *Recycling* 9, 5 (oct 2024), 95. doi:10. 3390/recycling9050095
- [34] Williams Rizzi, Chiara Di Francescomarino, and Fabrizio Maria Maggi. 2020. Explainability in Predictive Process Monitoring: When Understanding Helps Improving. In Business Process Management Forum, Dirk Fahland, Chiara Ghidini, Jörg Becker, and Marlon Dumas (Eds.). Springer International Publishing, Cham, 141–158. doi:10.1007/978-3-030-58638-6_9
- [35] Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. 2022. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys* 16 (2022), 1 – 85. doi:10.1214/21-SS133
- [36] Busra L Sebin, Nazim Ali Taskin, and Nijat Mehdiyev. 2024. Exploring the Intersection of Large Language Models (LLMs) and Explainable AI (XAI): A Systematic Literature Review (Research-in-Progress). CONF-IRM 2024 Proceedings. 25 (2024).
- [37] Juan Sequeda, Dean Allemang, and Bryon Jacob. 2024. A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model's Accuracy for Question Answering on Enterprise SQL Databases. In Proceedings of the 7th Joint Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA) (Santiago, AA, Chile) (GRADES-NDA '24). Association for Computing Machinery, New York, NY, USA, Article 5, 12 pages. doi:10.1145/3661304.3661901
- [38] Dylan Slack, Satyapriya Krishna, Himabindu Lakkaraju, and Sameer Singh. 2023. Explaining machine learning models with interactive natural language conversations using TalkToModel. *Nature Machine Intelligence* 5, 8 (2023), 873–883. doi:10.1038/s42256-023-00692-8

2025 ACM SIGMOD/PODS, June 22-27, 2025, Berlin, GER

- [39] Irene Teinemaa, Marlon Dumas, Marcello La Rosa, and Fabrizio Maria Maggi. 2019. Outcome-Oriented Predictive Process Monitoring: Review and Benchmark. ACM Trans. Knowl. Discov. Data 13, 2, Article 17 (March 2019), 57 pages. doi:10. 1145/3301300
- [40] Wil Van Der Aalst. 2012. Process mining. Commun. ACM 55, 8 (Aug. 2012), 76-83. doi:10.1145/2240236.2240257
 [41] Vladimir Vovk, Ivan Petej, Ilia Nouretdinov, Valery Manokhin, and Alexander
- Gammerman. 2020. Computationally efficient versions of conformal predictive

Pfeiffer, Rombach, Majlatow and Mehdiyev

distributions. Neurocomputing 397 (2020), 292-308. doi:10.1016/j.neucom.2019. 10.110

Wenjie Wu, Yongcheng Jing, Yingjie Wang, Wenbin Hu, and Dacheng Tao. 2025. Graph-augmented reasoning: Evolving step-by-step knowledge graph retrieval for llm reasoning. *arXiv preprint arXiv:2503.01642* (2025). [42]

Received 15.05.2025