ESCP BUSINESS SCHOOL

IT ALL STARTS HERE

BERLIN I LONDON I MADRID I PARIS I TURIN I WARSAW

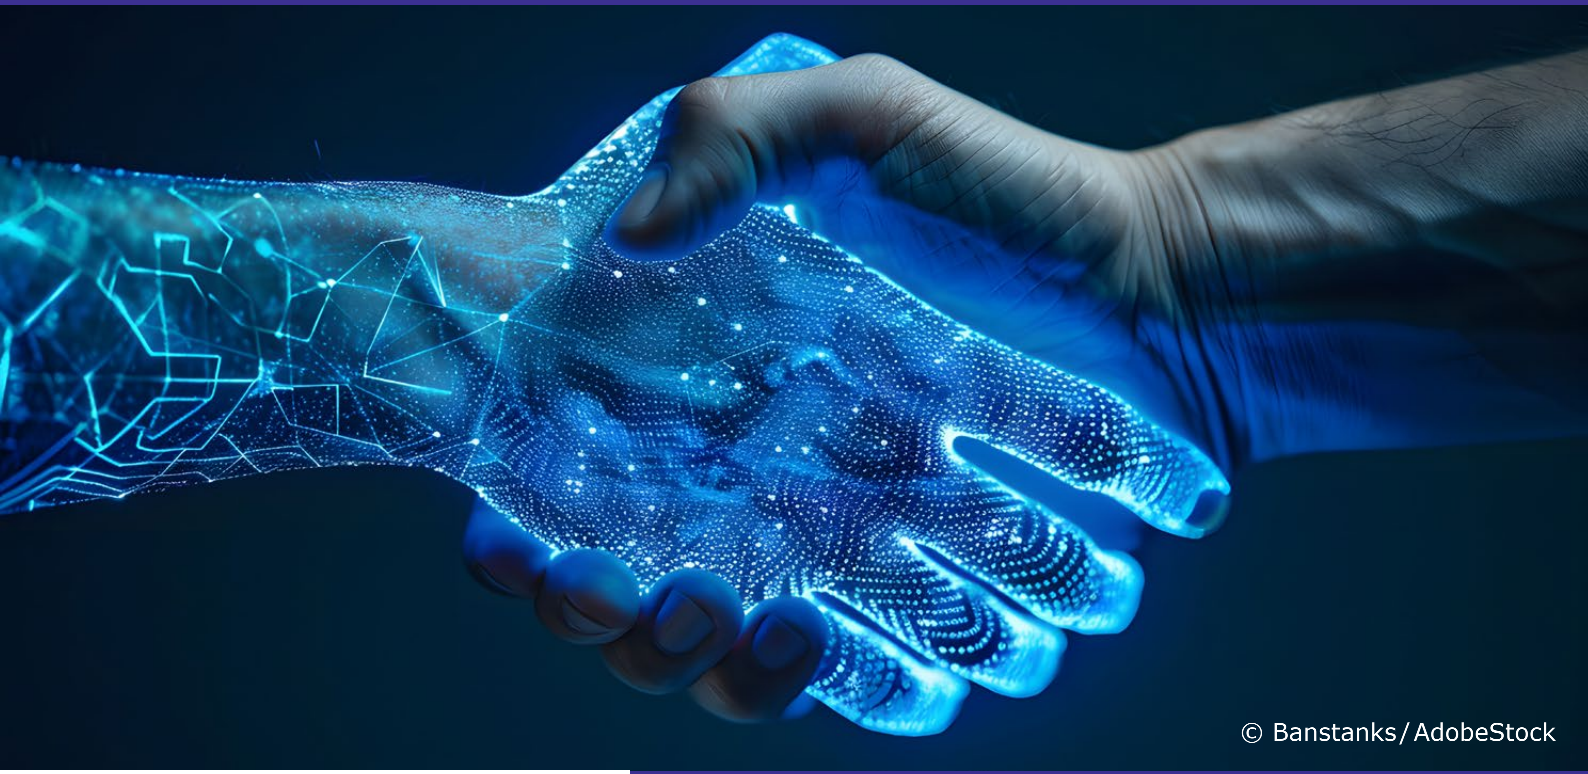

© Banstanks / AdobeStock

LIGHTS - Technology

From data to meaning: Contextual judgment in AI-enabled process mining

ESCP Impact Paper No.2026-41-ENG

Antolin CLAUDE & Markus BICK
ESCP Business School

ESCP RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ESCP Impact Paper No.2026-41-ENG

From data to meaning:
Contextual judgment in AI-enabled process mining

Antolin Claude*
Markus Bick**

ESCP Business School

Abstract

Organizations use process mining to extract insights from digital trace data. Generative AI is now being integrated into the workflow to make process analytics more accessible. Research in process mining, business analytics, and organizational sense-making suggests that analytical outputs do not carry self-evident meaning: their interpretation requires contextual knowledge distributed across people, documents, and institutional routines. This paper draws on recent contributions in these fields to argue that the value organizations derive from AI-assisted process analysis depends on whether AI is positioned as a contextual resource that supports human interpretation or as an endpoint that substitutes for it. It examines how context frameworks, sense-giving cycles, and emerging AI-assisted designs relate to one another, and draws practical implications for decision-makers governing AI adoption in their analytical workflows.

Keywords: process mining, sense-making, generative AI, organization, context

*PhD Candidate, ESCP Business School
**Professor, ESCP Business School

From data to meaning:
Contextual judgment in AI-enabled process mining

Introduction

A growing number of organizations use process mining to understand how their business processes actually unfold, based on digital traces left in information systems. Process mining "starts from event data and uses process models in various ways, e.g., process models are discovered from event data, serve as reference models, or are used to project bottlenecks on" (van der Aalst, 2016). With the recent integration of large language models (LLMs) into process mining tools, analysts can now query process data in natural language and receive AI-generated interpretations of the results (Brützke et al., 2025). This development suggests a potential broadening of process analytics accessibility, beyond the circle of specialists who traditionally use it.

However, accessibility does not guarantee understanding. This paper argues that the value organizations derive from process analysis depends on whether the people interpreting the results can connect them to the organizational context in which the processes operate. Research in process mining, business analytics, and organizational sense-making suggests that analytical outputs do not carry self-evident meaning. The interpretation of results requires contextual knowledge that is distributed across people, documents, and institutional routines. When the interpretive work is partially or fully delegated to an AI system, one potential risk is to lose "epistemic oversight": nobody in the organization is in a position that allows assessing the veracity or origin of the claim at hand. Furthermore, the data itself often had quality issues. For example, if an LLM queried for the root cause analysis of a bottleneck returns that the cause is a delay in payment, one would need to verify that such a claim actually matches the real context behind the data and the LLM's interpretation. To keep the argument concrete, one process example is given throughout the paper: a procurement-to-pay process in which a purchase order should be matched against a goods-receipt note and an invoice (a three-way match) before payment is approved. The example returns as the discussion moves from raw outputs to interpretation, sense-making, and governance.

The paper develops this argument in four steps. It first describes why context is constitutive of meaning in process analysis. It then examines how sense-making research characterizes the interpretive work that connects analytical outputs to organizational decisions. Third, it discusses how current AI-assisted approaches relate to this interpretive work. Finally, it draws implications for decision-makers considering how to govern AI adoption in their analytical workflows.

Process mining outputs require interpretation, not just computation

Process mining algorithms produce structured representations of organizational processes: discovered models, conformance reports, performance dashboards. The meaning of such representations of process data inside the organization is not always self-explanatory. For example, a process deviation flagged by a conformance check may indicate a compliance failure, an informal workaround that has become standard practice, or a legitimate adaptation to a regulatory change. The event log alone does not distinguish between these possibilities.

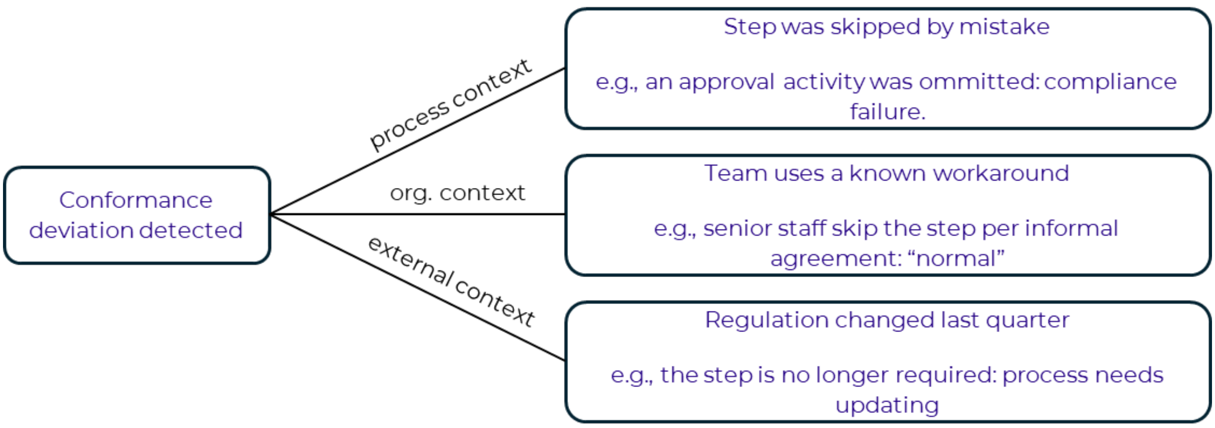

Figure 1: Different meanings of a deviation, based on context (own figure)

Grisold et al. (2024) formalize this dependency by distinguishing three layers of context that analysts need to mobilize when interpreting process mining results: process-immediate context (changes in activities, resources, or case attributes), organization-internal context (role structures, hierarchies, learning dynamics), and organization-external context (regulations, market conditions, industry norms). Their framework explicitly outlines what experienced process analysts know implicitly: interpreting process data requires organizational knowledge that sits outside the data itself. Research extended this work into a taxonomy that specifies how each contextual layer shapes the sense-making paths available to the analyst. They found that while some contextual changes can be detected computationally (for instance, shifts in resource behavior or concept drifts), understanding what these changes mean for the organization requires a different kind of work: interpretive, dialogic, and inherently situated.

Figure 1 illustrates this dependency concretely. A conformance check flags a deviation because an expected activity was not performed. Without further context, the analyst cannot determine whether this represents a compliance failure (a step was mistakenly omitted), a known organizational practice (senior staff skip the step per informal agreement), or an adaptation to changed circumstances (a regulatory update rendered the step unnecessary). Each interpretation calls for a different organizational response. While the event log contains the same data in all three cases, only contextual knowledge allows to distinguish between the cases. In the procurement example, a mismatch between, say, the invoice and the receipt may have various causes. The management may have committed fraud risk, but it could also be due to a tolerated practice for trusted suppliers below a value threshold, or a deliberate fast-track introduced during a supply shortage. In a customer service process, repeated case reopenings may reflect poor first-contact resolution or a newly added quality-assurance step. In healthcare, a diagnostic test performed out of the expected order may indicate a protocol breach or appropriate clinical prioritization for an urgent case. In financial operations, a payment released before reconciliation may be an due to an internal-control issue or an approved exception for time-critical settlements. In each case, the event log is identical across readings, and only organizational context can inform which one applies.

Research reports on a seminar at the intersection of visual analytics and process mining. It stresses they state that “it is important for analysts to thoroughly explore, investigate, and understand the data before and during the implementation of interactive analysis methods, such as PM algorithms, to determine which methods are appropriate and to be able to make sense of the output generated by the selected methods”. Even the most sophisticated algorithm can produce a process model that is too complex to read or too simplified to capture the full picture. The analyst must navigate between these extremes, and doing so requires domain knowledge, exploratory capacity, and an understanding of what questions matter to the organization.

Sense-making as the bridge between data and organizational understanding

The organizational literature has studied the interpretation of data or situations under the term “sense-making”. Organizational theory conceptualizes sense-making as a social and cognitive process through which individuals construct plausible accounts of ambiguous situations. Research proposes the Data-Frame model, where analysts match data against interpretive frames and revise those frames when discrepancies appear. The crucial insight, based on a review of 120 case studies, is that genuine analytical insight tends to emerge from frame revision rather than from the accumulation of confirming evidence. In other words, insight requires questioning one’s initial interpretation, not just finding more data that fits it. Namvar and Intezari (2021) translate these ideas into the specific context of business analytics. Based on 32 interviews with data analysts and consultants, they identified a cycle of sense-giving activities, which was adapted and depicted in Figure 2. Data integration (assembling relevant sources), trustworthiness analysis (assessing whether the output can be relied upon), appropriateness analysis (checking whether the output fits the business context), and alternative selection (comparing different analytical approaches) are the four activities that enable seven sensemaking properties explaining how the identified sensegiving activities assist end users in gaining a better understanding of the business environment. Their key finding is that the quality of decision-makers’ understanding depends on whether these steps are actually performed. When time pressure or resource constraints lead analysts to skip some steps, the final analysis outputs that support the decision may not have been adequately examined. Returning to the procurement example, these activities are what would allow the analyst to make sense of the flagged mismatch: integrating the supplier master data and approval policy (data integration), checking whether the event log reliably records the match step (trustworthiness analysis), asking whether skipping is permissible for this supplier and value band (appropriateness analysis), and weighing fraud against tolerated exception (alternative selection).

Ultimately, these steps might help addressing the issue of volume and velocity of contemporary data, which can overwhelm interpretive capacity. Calvard (2016) labelled this issue the “big data hubris”: treating algorithmic outputs as reliable without further contextual scrutiny. The research states that organizations continually alternate between exploratory learning (generating new interpretive possibilities) and sense-making (reducing those possibilities into actionable conclusions).

Bani-Hani et al. (2022) add an institutional layer to the picture, showing that whether employees engage in analytical questioning at all depends on organizational norms, personal dispositions, and the social relationships that surround data use.

AI-assisted process analysis : supplying context or substituting for judgment?

Brützke et al. (2025) develop a process mining AI assistant using a design science approach. Their system retrieves contextual information from organizational documents (process descriptions, hierarchies, regulatory texts) and combines it with process mining outputs via an LLM-based conversational interface. The architecture is deliberately designed to position the AI as a contextual intermediary rather than a final interpreter: the prompt template instructs the LLM to rely only on retrieved data, and an agent selection heuristic prevents indiscriminate tool deployment. Evaluation with 14 process mining experts found that the artifact supports both novice and experienced users in conducting context-sensitive analyses.

By automatically surfacing context alongside results, the system reduces the cognitive burden on the analyst without removing the analyst from the interpretive loop. It operationalizes, in technical form, the context framework proposed by Grisold et al. (2024): the AI retrieves process-immediate, organization-internal, and organization-external context from available documentation and presents it as part of the analytical response. Research points in a similar direction when suggesting that LLMs and agentic systems could process unstructured organizational data sources to provide richer contextual understanding.

The same technological capabilities can also be deployed differently. Krishnan and Bhat (2025), describing hyperautomation frameworks in financial services, document how generative AI can automate data interpretation, report generation, and compliance verification at scale. The efficiency gains are real and documented. From the perspective of the sense-giving cycle described by Namvar and Intezari (2021), several of these automated tasks overlap with the very activities that mediate organizational understanding: trustworthiness analysis, appropriateness checking, and dialogue between analysts and decision-makers. When such steps are automated, a question arises: does the organization retain the capacity to evaluate whether the AI's interpretation fits its specific situation?

Brützke et al. (2025) acknowledge that LLM hallucination remains a problem that prompting strategies alone do not fully resolve. Ferrario et al. (2025), developing an ontological framework for classifying bias, identify what they call deployment bias: a situation in which a system trained on one class of objects is applied to an incompatible class. Applied to LLM-assisted process analysis, this concept highlights that a model producing generally plausible process interpretations may still be unsuitable for the specific organizational context at hand, and that the analyst may lack a clear signal of this mismatch. Research also documents a parallel problem in human reasoning: confirmation bias can lead analysts to accept interpretations that fit their expectations rather than revising their interpretive frame. When an AI system produces a fluent, confident interpretation, the conditions for this kind of bias may be reinforced rather than mitigated. One solution could be to automatically infer reasoning from low-level user interaction logs.

At a more fundamental level, Leonelli (2025) and Chang et al. (2025) argue that the rapid development of AI has outpaced the construction of the semantic and ontological infrastructure needed to ensure that AI outputs are interpretable in context. Chang et al. (2025) describe this as a growing misalignment between AI capabilities and human-understandable intent, arguing that structured knowledge and formal ontologies should be treated as integral design components rather than afterthoughts. Leonelli (2025) proposes a process-sensitive approach to data semantics that foregrounds the role of human judgment and environmental interactions in giving data their meaning.

Implications for decision-makers

What does this mean for organizations considering how to govern AI in their analytical workflows? The research reviewed cannot provide a simple "yes or no" recommendation for or against AI adoption in process analysis. Instead, it suggests that the relevant question is how to position and integrate AI within the organizations' interpretive process.

In terms of design, two orientations appear. In the first, AI serves as a contextual resource: it retrieves relevant organizational knowledge, surfaces data quality issues, and presents its outputs as inputs to human judgment. The system developed by Brützke et al. (2025) exemplifies this orientation. In the second, AI serves as an endpoint: it produces finished interpretations, automates report generation, and delivers conclusions ready for action. Hyperautomation frameworks (Krishnan and Bhat, 2025) tend toward this orientation, at least for routine tasks.

Both orientations have legitimate uses. Routine compliance checks, standardized reporting, and high-volume data processing may benefit from end-to-end automation. But for analytical tasks where the meaning of the output depends on organizational context, where the stakes of misinterpretation are significant, and where the situation may differ from what the AI has been trained on, the first orientation appears better aligned with what the sense-making literature describes as the conditions for sound organizational judgment.

Three practical considerations follow. First, organizations should distinguish between tasks where AI can produce reliable outputs autonomously and tasks where its outputs require contextual evaluation by someone with relevant domain knowledge. The context framework of Grisold et al. (2024) offers a useful diagnostic: the more the interpretation depends on organization-internal or organization-external context, the more human judgment is likely to matter.

Second, the sense-giving cycle (Namvar and Intezari, 2021) provides a template for embedding interpretive checkpoints into AI-assisted workflows: prompting analysts to assess trustworthiness, appropriateness, and alternative readings before conclusions are forwarded to decision-makers.

Third, organizations should attend to the institutional conditions that support or inhibit sense-making. As Bani-Hani et al. (2022) and Calvard (2016) document, whether people actually engage in critical analytical thinking depends on organizational culture, time, and institutional incentives. In sum, AI adoption that increases output volume may well produce more data-driven decisions, but decision makers, together with the AI integration team, should ensure the corresponding investment in interpretive capacity .

These considerations can be condensed into a simple screening question for any AI-generated interpretation: can the organization say why a certain interpretation, rather than a plausible alternative, is the right one for its situation? The checklist below aims to support practitioners in answering that question using the context layers of Grisold et al. (2024) and the sense-giving activities of Namvar and Intezari (2021). If the answers fall mostly in the left column, an AI output may be trusted more confidently. If the output falls in the right column, contextual review by a domain expert may be essential.

Question to ask of the AI output	AI output can stand alone if...	Human contextual review needed if...
Which context layer does the reading rely on? (Grisold et al., 2024)	Mostly process-immediate context that is visible in the event data itself.	It depends on organization-internal or organization-external context not present in the log.
Are alternative readings plausible? (alternative selection; Namvar and Intezari, 2021)	One reading is clearly dominant and the others are easily excluded.	Compliance failure, tolerated practice, and legitimate adaptation remain competing explanations.
What are the stakes of a wrong reading?	Low: routine, reversible, high-volume reporting.	High: legal, financial, safety, or reputational consequences.
Does the situation match what the model was built for? (deployment bias; Ferrario et al., 2025)	The process resembles the cases the system was designed and evaluated on.	The process differs from the training or evaluation setting in ways the analyst cannot rule out.

Seen this way, the question is not specific to process mining. It is one instance of a broader challenge in human-AI collaboration: how to combine the analytical scale and speed of AI with the contextual understanding, judgment, and critical reflection that remain distinctively human. The orientations and checkpoints above describe one division of labor in which AI widens what an organization can examine while humans retain responsibility for what its outputs are taken to mean. Treating that division of labor as a design choice, rather than a by-product of adoption, is what allows the analytical power of AI to be used without ceding the interpretive authority on which sound decisions depend.

Conclusion

Process mining and generative AI promise more efficient and accessible analytics, but they also reconfigure where and how interpretation occurs. This paper has shown that when contextual judgment is weakened or bypassed, organizations risk replacing understanding with plausibility. The implication for organizations may well be an efficiency gain, but that gain will not be sustained without an implemented routine assessment of AI-mediated analytical outputs. Returning to the procurement example, the difference is that organizations that internalized the issue described can confidently interpret why a mismatch between two customer order documents was acceptable, whereas organizations satisfied with plausible outcomes might only become aware of the fact once auditors find out about it.

The takeaway is that the challenge organizations are faced with is primarily organizational, not technical. Practitioners can address the issue of contextual validity by designing AI-assisted workflows that preserve epistemic responsibility. Without such safeguards, the proliferation of AI-generated insights may lead to more confident mistakes instead of better decisions.

References

Bani-Hani, I., Chowdhury, S., & Kurti, A. (2022). The Triadic Relationship of Sense-Making, Analytics, and Institutional Influences. Informatics, 9(1), 3. https://doi.org/10.3390/informatics9010003

Brützke, P., Killewald, R., Franzoi, S., & Brocke, J. vom. (2025). AI-Assisted Process Mining for Context-Sensitive Analysis Support. ECIS 2025 Proceedings. https://aisel.aisnet.org/ecis2025/bpm/bpm/7

Calvard, T. S. (2016). Big data, organizational learning, and sensemaking: Theorizing interpretive challenges under conditions of dynamic complexity. Management Learning, 47(1), 65–82. https://doi.org/10.1177/1350507615592113

Chang, C., Chen, Y.-T., & Chen, Y. (2025). Reinvigorating Structured Knowledge and Ontologies for Trustworthy and Beneficial AI and Robotics. In Formal Ontology in Information Systems (pp. 125–136). IOS Press. https://doi.org/10.3233/FAIA250489

Ferrario, R., Porello, D., Bottazzi, E., Florio, C. D., & Fumagalli, M. (2026). Taming the Sea of Errors: An Ontological Study of Biases in DOLCE. In Formal Ontology in Information Systems FOIS 2025 (pp. 64–78). IOS Press. https://doi.org/10.3233/faia250484

Grisold, T., Aa, H. van der, Franzoi, S., Hartl, S., Mendling, J., & Brocke, J. vom. (2024). A Context Framework for Sense-making of Process Mining Results. 6th International Conference on Process Mining, ICPM 2024, Kgs. Lyngby, Denmark, October 14-18, 2024, 57–64. https://doi.org/10.1109/ICPM63005.2024.10680665

Krishnan, G., & Bhat, A. K. (2025). Empower Financial Workflows: Hyper Automation Framework Utilizing Generative Artificial Intelligence and Process Mining. 2025 3rd International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI), 2041–2047. 2025 3rd International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI). https://doi.org/10.1109/ICoICI65217.2025.11254280

LEONELLI, S. (2025). Bringing Humans into Ontologies. In T. P. Sales, C. Masolo, & C. M. Keet (Eds.), Formal Ontology in Information Systems. SAGE Publications. https://doi.org/10.3233/FAIA250478

Namvar, M., Intezari, A., & Im, G. (2021). Sensegiving in organizations via the use of business analytics. Information Technology & People, 34(6), 1615–1638. https://doi.org/10.1108/ITP-10-2020-0735

van der Aalst, W. (2016). "Process Mining: Data Science in Action". Springer Publishing Company. https://doi.org/10.1007/978-3-662-49851-4