KI - Künstliche Intelligenz
https://doi.org/10.1007/s13218-025-00897-6

PROJECT REPORTS

Explainable Knowledge-Aware Process Intelligence

PINPOINT Final Project Report

Eva Piccirilli[1] · Claudio Di Ciccio[2] · Marco Montali[3] · Rafael Peñaloza[1] · Luigi Pontieri[4] · Francesco Ricca[5]

Received: 4 August 2025 / Accepted: 11 September 2025

Abstract
The PINPOINT project rose from the need to combine process intelligence and domain knowledge, with logic-based formalisms allowing for effective interpretability and explainability. This project report summarises the main contributions made by the research units during the three-year period, and the next steps.

Keywords Process mining · Explainability · PRIN

1 Introduction

Contemporary organisations, from public-sector institutions to private enterprises, operate in systemically interconnected socio-technical environments. Business operational process analysis has hence shifted from indirect assumption-driven methodologies based on managerial reports, qualitative interviews and field studies, to evidence-based process intelligence techniques. Lying at the intersection of model-driven engineering and data science, *process mining* drives this transition building process-centric knowledge from event data like logs collected by enterprise systems [47]. Leveraging the fine-grained insights offered by event data, process mining techniques integrate model-based and data-driven analysis to support operational process execution refinement in alignment with factual compliance. While effective, process mining techniques remain constrained by the *garbage-in, garbage-out* factor which may compromise the reliability of its results. Significant limitations persist due to the employment of opaque (black-box) algorithms and insufficient integration of domain-specific organisational knowledge into process analysis. To address them—building on recent advancements in explainable AI and multi-perspective declarative languages and techniques—the PINPOINT project (exPlaInable kNowledge-aware PrOcess INTelligence) was conceived, building on the integrated expertise of five partner research units.[1]

This report summarises the main results achieved during the three-year project through a tight collaboration between the units. The sections follow the work package (WP) structure from the original project proposal (see Fig. 1).

Specifically, Sect. 2 describes work on *transparent, end-to-end data processing*, led by the unit at the Free University of Bozen-Bolzano; Sect. 3 focuses on *Process Knowledge Representation and Discovery*, coordinated by the unit at the University of Milano-Bicocca; Sect. 4 presents WP3 *Explainable, Knowledge-Aware Predictive Monitoring*, with leadership by the unit at ICAR-CNR; Sect. 5 consolidates the contributions to WP4 *Explainable, Knowledge-Aware Conformance Checking*, led by the unit at Sapienza University of Rome, and WP5 *Application of Explainable Process-Aware Intelligence*, coordinated by the research unit at the University of Calabria. This work lies within the scope of WP6: coordination and dissemination.

Work supported by the Italian Ministry of University and Research (MUR) under the PRIN 2020 project "exPlaInable kNowledge-aware PrOcess INTelligence" (PINPOINT) Prot. 2020FNEB27.

✉ Rafael Peñaloza
Rafael.Penaloza@unimib.it

1 University of Milano-Bicocca, Milan, Italy
2 Utrecht University, Utrecht, Netherlands
3 Free University of Bozen-Bolzano, Bolzano, Italy
4 ICAR-CNR, Rende, Italy
5 University of Calabria, Rende, Italy

1 Official PINPOINT site: https://pinpoint.unibz.it.

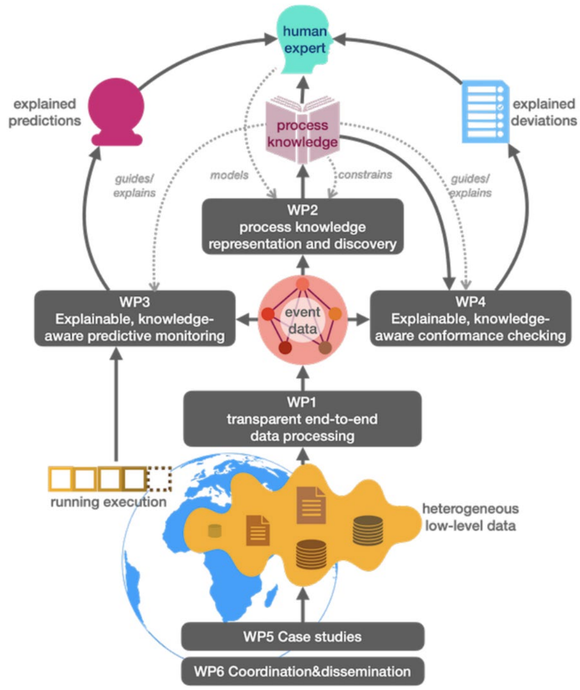

Fig. 1 Conceptual overview of PINPOINT

2 Transparent Data Processing

Formal process intelligence requires handling process data and logs, which are often part of the private industrial information and may hide user data that should not be exposed. Hence the need to develop techniques to construct transparent, explainable, data-processing pipelines, while guaranteeing that the data is not compromised, even when shared between organisations.

For transparent, end-to-end data processing, a particular focus was placed on event-case correlation enhancement for process mining. Knowledge-aware techniques for explainable event data mapping and multi-perspective event data extraction based on simulated annealing were devised. The EC-SA-RM technique [6] combines simulated annealing (SA) with iterative association rule mining (RM) to infer domain knowledge for formal specifications. The EC-SA-Data correlation engine aligns the control-flow with data perspectives through probabilistic optimisation [7].

To further extend transparent data processing for distributed, inter-organisational settings, we considered several secure processing architectures. In particular, the CONFINE toolset [32] was developed with the explicit intent of enabling process mining on process event data from multiple providers, while preserving both, the confidentiality and the integrity of the original records. CONFINE ensures that event data can be securely processed without exposure to external agents by leveraging a decentralized architecture based on Trusted Execution Environments (TEEs). Specifically, TEEs provide hardware-secured confidential computing enclaves to run verifiable software. CONFINE utilizes TEEs to deploy process mining algorithms in the form of trusted applications within those enclaves.

Other contributions, which build on programmable blockchains and distributed hash-table storage, ensure the confidentiality of data exchanged during decentralized process execution with the help of distributed systems. Two solutions were developed to enforce transparent, auditable collaboration and fine-grained control over data access and sharing, securing process data flows through applied cryptography: CAKE [41] in a centralized setting, and MARTSIA [40] for multi-authority scenarios. To enforce traceable data handling across decentralized infrastructures, a blockchain-driven usage control architecture, also built on TEEs, was proposed to ensure that once data are shared, its usage remains compliant with usage control policies [5]. The research on transparent data processing was complemented with a visual analytics framework called Tiramisù, designed to allow users to interactively visualize multi-faceted process information, thus helping them carry out complex explorative process analysis tasks [2].

A last key problem tackled pertains the extraction and processing of relational event data from general information systems (e.g., ERP and CRM), with a threefold goal: (i) semi-automate the creation of event logs from legacy information systems; (ii) provide the basis for provenance indications (linking information system records with events); (iii) support *object-centric process mining*, where event data may refer to multiple, inter-related objects.

Within the project, we have actively worked on the definition of (meta)models for supporting these three tasks [22], as well as concrete extraction pipelines, starting from the experience gained in [48].

3 Knowledge Representation and Discovery

Moving beyond the processing of data, process knowledge needs to be distilled from logs and other information sources in a general discovery process. To be usable, it needs to be stored using a formal language that allows for reasoning and derivation guarantees.

The *de-facto* standard for modelling business processes is the linear temporal logic over finite traces, $\mathrm{LTL_f}$, which provides the formal ground for specification languages like Declare [17, 45]. Declare specifications are often mined from observed behaviour and logs [36]; yet, classical mining techniques have an associated uncertainty which, left unchecked, leads to inconsistencies and further problems through the pipeline. Hence, the project studied ways to deal with uncertain specifications for standard reasoning

[39], alignment [38], that is, a correspondence between a log trace and a process model run, and monitoring [1], among others.

A new approach for satisfiability checking in bounded LTL_f [27] led to an ASP representation of Declare [15], which set the basis for enumerating minimal unsatisfiable subsets of LTL_f formulas (MUSes)—also known as unsatisfiable *cores*—using optimised methods developed for this language [3]. This MUS enumerator, along with other ASP-centric optimisations, was shown to be effective also for other kinds of logical formalisms, yielding a new efficient method for enumerating MUSes (known as *justifications* in description logics) for inexpressive description logics [35, 44]. First steps towards generalising from plain MUS enumeration to full semiring provenance were made in [43]. All these approaches aim to provide information necessary to explain, measure, and correct inconsistencies and errors in specifications.

LTL_f/Declare process specifications are centred on the process control-flow. However, it is often important to couple control-flow dependencies with data conditions, to contextualize and scope the resulting constraints. Data-aware declarative process specifications have thus been studied within the project, extending LTL_f with different types of data and corresponding conditions. Although in general such interplay is too expressive, several well-behaved fragments have been identified, bringing forward automated reasoning techniques obtained by pairing automata with SMT solving (see [31] for a summary of the main results).

Acknowledging the existence of other process modelling languages and mining approaches, the project also developed techniques to declaratively define hybrid process models using multiple formalisms [26]. In addition, the ideas of model repair were extended to also logically handle Petri net-based specifications [14].

4 Explainable Predictive Process Monitoring

Predictive monitoring aims to estimate unknown properties of ongoing process instances based on partial traces, past executions, and specifications where available. State-of-the-art Machine Learning (ML) approaches to this problem rely on training opaque ensembles or deep neural-network models, which means that explainability is typically only available *post-hoc*. Moreover, knowledge-aware modelling, where the knowledge is readily available and usable, remains in its early stages.

A novel version of the *Nirdizati Light* open-source tool [8] was used in the project as modular and flexible platform for evaluating and comparing different ML-based predictive models and post-hoc explanation methods, in diverse tasks and contexts.

An explainable-by-design alternative to post-hoc explanation, based on a sparse and shallow Mixture-of-Experts (MoE) neural-network model was devised. In it, the gate (router) and expert modules simply implement easy-to-interpret logistic regressors, trained in an end-to-end fashion [28] to model complex data distributions that go beyond the representation power of a single linear model. A MoE-based framework for clinical decision-making support was developed [16], combining locally specialized logistic regressors with *ad-hoc Gumbel-softmax* relaxations to enforce gate sparsity and per-expert feature selection seamlessly and differentiably during the training process. Notably, the modular nature of this ensemble-like predictive model allows for integrating any existing predictive model [16], defined/validated by human experts possibly using a symbolic representation. The advantages of using this framework for explainable and knowledge-aware predictions in clinical decision tasks were showcased in [16].

A different compositional approach to integrating domain knowledge in predictive monitoring was proposed in [23] for the challenging case where log events are at a lower level of abstraction than the activities to be monitored; the approach combines a neural network trained on labelled example traces and a symbolic (AAF-based) reasoner provided with prior process knowledge, in the spirit of neural-symbolic AI.

To support explainable knowledge-aware monitoring, a conditional generative model based on a Conditional Variational Autoencoder (CVAE) was developed to serve as a knowledge-aware log data generator [33]. A follow-up extension [34] broadened the scope to generate complete multi-perspective trace executions, including control-flow, temporal, and resource attributes, and condition generation on specific temporal constraints. The approach proved to be more effective than existing ML-based log generators in condition-specific trace generation tasks, supporting what-if and other causal analyses. A method for generating counterfactual explanations under temporal constraints, expressed in a variant of LTL_f and representing background knowledge, was introduced in [9]. To fit predictive monitoring settings, enabling the construction of monitors expressing predictions aligned with the given constraints, a fuzzy version of LTL_f was introduced [19] and used as a basis for infusing such temporal knowledge into deep learning architectures [4].

To discover deviance-oriented predictive models in real-life contexts where explicit data labels and domain knowledge are unavailable or scarce, the project introduced methods for grasping user knowledge interactively via active learning [29] and exploiting auxiliary ML models as a supplemental source of supervision [30].

5 Conformance Checking

Advances in process mining have introduced novel techniques for conformance checking—the task of verifying whether a trace or an event log complies with a declarative process specification. Among these, important contributions from the project include a probabilistic, event-level framework for assessing LTL_f declarative process specifications, quantifying their satisfaction over multi-sets of execution traces [12]. The presence of uncertainty makes the problem more challenging, as a deviation from a pure specification could just signal an outlying execution, rather than a specification violation. Association-rule-inspired measures were introduced to assess the quality of constraint-based specifications [12]. The measurement framework quantifies the degree to which specifications composed of LTL_f-based rules expressed as "if–then" statements are satisfied within process execution traces [11]. A further extension estimates the satisfaction of declarative specifications as a whole, thus overcoming the limitations of approaches that evaluate constraints in isolation [13]. A different but related problem is the alignment of temporal knowledge bases (TKB). In this respect, existing methods for aligning propositional LTL_f formulas were extended to produce cost-optimal alignments in highly expressive temporal description logics [25].

Seeing conformance checking as an alignment problem, the project extended traditional alignment and cost functions to account for uncertainty, e.g., about activities, timestamps as well as other data attributes, along control-flow, time, and data perspectives, leveraging techniques originally developed for Satisfiability Modulo Theories (SMT) [24]. Concurrently, alignment-based techniques based on A^* search for control-flow, and those based on SMT for dealing with data-aware processes, were combined to tackle, for the first time, alignment-based conformance checking of data-aware Declare specifications dealing with rich datatypes and corresponding conditions [10]. The feasibility of the approach was demonstrated through implementation and experimental evaluation on synthetic and real-life logs.

A large part of the research work focused on foundational aspects of Answer Set Programming (ASP) and their application to the project. First, effective algorithms for the enumeration of ASP minimal unsatisfiable subprograms [3] were developed. These algorithms and their implementations were used to provide the first approach to enumerate unsatisfiable cores (Sec. 3), which correspond to minimal reasons for inconsistencies of temporal specifications [37]. A further contribution introduced four algorithms for the extraction of unsatisfiable cores from LTL_f specifications, adapted from existing satisfiability-checking techniques [46]. This yields a first step towards basic reasoning services for explanation tasks in declarative process mining, in an attempt to provide explanatory services for process intelligence. However, in general, it is known that full explainability tasks require solving problems which go beyond NP [21]. This becomes particularly evident in LTL_f, where deciding satisfiability is PSpace-complete. Thus, the research has also focused on developing, implementing and extending tooling for the Answer Set Programming with Quantifiers ASP(Q) formalism, a quantified extension of the ASP language, that enables tackling beyond-NP reasoning tasks in a more comfortable manner than the traditional saturation technique [20]. The ASP(Q) formalism has been significantly developed within the project, with important contributions in [27, 42].

Another ASP optimization technique is based on program compilation, which attempts to avoid the pitfalls of ground-and-solve methods employed by modern systems. Compilation instead tries to minimise the need for grounding, thus reducing the overall time and space necessary for exploring solutions. This alleviates the scalability and state explosion problems which are commonly observed in process mining situations [18].

At the end of the project, a short case study was presented based on data concerning the management of applications submitted to *development contract* grants through the national development agency of a European nation. The whole data process and analysis followed the security techniques described in Sect. 2, and was managed and evaluated through techniques developed specifically for Declare-specifications and implemented directly in ASP.

6 Conclusions and Lessons Learned

The objective of the PINPOINT project was to develop techniques to yield novel process intelligence models that break the *black-box*; that is, which are explainable and interpretable, and allow for explicit knowledge management along with and beyond standard models. The techniques follow the full data and knowledge ecosystem, from methods to preserve data security and integrity, to modelling languages with their associated reasoning tasks, to process monitoring and conformance checking. Implementation of the techniques, particularly for reasoning, exploited the highly optimised tools that exist for ASP, following modern trends based on reductions, rather than full *ad-hoc* implementations. A case study was implemented using data from a national agency in a European country.

The project was funded by the Italian Ministry of University and Research (MUR) under the three-year scheme of research projects of national interest (PRIN). The partners come from research institutes (CNR) and universities covering most of the Italian territory from the South (University of Calabria) to the North (Free University of Bozen-Bolzano and University of Milano-Bicocca) through the capital city (Sapienza University of Rome). Two synchronisation events were organised, which took place in Bolzano (South Tyrol) and in Roccella Jonica (Calabria). While managing a project of this size with partners so far apart was not easy, some simple strategies contributed to its success from the beginning. One was to have a visual representation of the tasks, milestones, and contributors, readily available for easy consultation. Another was to keep a centralised, continuously updated repository enumerating all the products contributed to the project, in order to keep track of the development, milestones, and roadblocks.

We believe that the collaboration between units was a success, which may translate into other ambitious research projects in the future.

References

1. Alman A, Maggi FM, Montali M, Peñaloza R (2022) Probabilistic declarative process mining. Inf Syst 109:102033. https://doi.org/10.1016/J.IS.2022.102033
2. Alman A, Arleo A, Beerepoot I, et al. (2023) Tiramisù: a recipe for visual sensemaking of multi-faceted process information. In: ICPM 2023 workshops. Springer, pp 19–31. https://doi.org/10.1007/978-3-031-56107-8_2
3. Alviano M, Dodaro C, Fiorentino S, Previti A, Ricca F (2023) Asp and subset minimality: enumeration, cautious reasoning and muses. AIJ 320:103931. https://doi.org/10.1016/J.ARTINT.2023.103931
4. Andreaoni R, Buliga A, Daniele A, et al. (2025) T-ILR: a neurosymbolic integration for LTLf. In: Proceeding on NeSy 2025. Springer, LNCS, to appear
5. Basile D, Di Ciccio C, Goretti V, Kirrane S (2023) A blockchain-driven architecture for usage control in solid. In: Proceedings on ICDCSW 2023. IEEE, pp 19–24. https://doi.org/10.1109/ICDCSW60045.2023.00009
6. Bayomie D, Revoredo K, Di Ciccio C, Mendling J (2022) Improving accuracy and explainability in event-case correlation via rule mining. In: Proceedings on ICPM 2022. IEEE, pp 24–31. https://doi.org/10.1109/ICPM57379.2022.9980684
7. Bayomie D, Ciccio C, Mendling J (2023) Event-case correlation for process mining using probabilistic optimization. Inf Syst 114:102167. https://doi.org/10.1016/j.is.2023.102167
8. Buliga A, Graziosi R, Di Francescomarino C, et al. (2024) Nirdizati light: a modular framework for explainable predictive process monitoring. In: Proceedings on BPM 2024 D &R, CEUR-WS.org, vol 3758
9. Buliga A, Di Francescomarino C, Ghidini C, Montali M, Ronzani M (2025) Generating counterfactual explanations under temporal constraints. In: Proceedings on AAAI 2025, vol 39, pp 15622–15631. https://doi.org/10.1609/AAAI.V39I15.33715
10. Casas-Ramos J, Winkler S, Gianola A, et al. (2025) Efficient conformance checking of rich data-aware declare specifications. In: Proceedings on BPM 2025. Springer, to appear
11. Cecconi A, De Giacomo G, Di Ciccio C, Maggi FM, Mendling J (2021) Measuring the interestingness of temporal logic behavioral specifications in process mining. Inf Sys https://doi.org/10.1016/j.is.2021.101920
12. Cecconi A, Di Ciccio C, Senderovich A (2022) Measurement of rule-based LTLf declarative process specifications. In: Proceedings on ICPM 2022. IEEE, pp 96–103. https://doi.org/10.1109/ICPM57379.2022.9980690
13. Cecconi A, Barbaro L, Di Ciccio C, Senderovich A (2024) Measuring rule-based LTLf process specifications: a probabilistic data-driven approach. Inf Sys 120:102312. https://doi.org/10.1016/j.is.2023.102312
14. Chiariello F, Ielo A, Tarzariol A (2024) An ILASP-based approach to repair Petri nets. In: Proceedings on LPNMR 2024. Springer, LNCS, vol 15245, pp 85–97. https://doi.org/10.1007/978-3-031-74209-5_7
15. Chiariello F, Fionda V, Ielo A, Ricca F (2025) Direct encoding of declare constraints in asp. TPLP 25(1):92–131. https://doi.org/10.1017/S1471068424000486
16. Cuzzocrea A, Folino F, Samami M, Pontieri L, Sabatino P (2025) Towards trustworthy and sustainable clinical decision support by training ensembles of specialized logistic regressors. Neural Comp App https://doi.org/10.1007/s00521-025-11360-w
17. Di Ciccio C, Montali M (2022) Declarative process specifications: reasoning, discovery, monitoring. In: Process mining handbook. Springer, pp 108–152. https://doi.org/10.1007/978-3-031-08848-3_4
18. Dodaro C, Mazzotta G, Ricca F (2024) Blending grounding and compilation for efficient ASP solving. In: Proceedings on KR 2024. https://doi.org/10.24963/KR.2024/30
19. Donadello I, Felli P, Innes C, Maggi FM, Montali M (2025) LTL-based conformance checking of fuzzy event logs. Process Sci 2(1):14. https://doi.org/10.1007/s44311-025-00020-w
20. Faber W, Mazzotta G, Ricca F (2023) An efficient solver for ASP (Q). TPLP 23(4):948–964. https://doi.org/10.1017/S1471068423000121
21. Faber W, Mazzotta G, Ricca F (2023b) Enhancing ASP(Q) evaluation. In: AI*IA 2023 Disc. Papers, CEUR workshop proceedings, vol 3537, pp 38–46
22. Fahland D, Montali M, Lebherz J, et al. (2024) Towards a simple and extensible standard for object-centric event data (OCED). Core model, design space, and lessons learned. CoRR arXiv:2410.14495
23. Fazzinga B, Flesca S, Furfaro F, Pontieri L, Scala F (2025) Combining abstract argumentation and machine learning for efficiently analyzing low-level process event streams. arXiv:2505.05880
24. Felli P, Gianola A, Montali M, Rivkin A, Winkler S (2023) Multi-perspective conformance checking of uncertain process traces: an SMT-based approach. Eng Appl AI 126:106895. https://doi.org/10.1016/j.engappai.2023.106895

25. Fernandez-Gil O, Patrizi F, Perelli G, Turhan A (2023) Optimal alignment of temporal knowledge bases. In: Proceedings on ECAI 2023, IOS Press, vol 372, pp 708–715. https://doi.org/10.3233/FAIA230335
26. Fionda V, Ielo A, Ricca F (2023) Logic-based composition of business process models. In: Proceedings on KR 2023, pp 272–281. https://doi.org/10.24963/KR.2023/27
27. Fionda V, Ielo A, Ricca F (2025) LTLf2ASP: LTLf bounded satisfiability in ASP. In: Proceedings on LPNMR 2024. Springer, LNCS, vol 15245, pp 373–386. https://doi.org/10.1007/978-3-031-74209-5_28
28. Folino F, Pontieri L, Sabatino P (2023) Sparse mixtures of shallow linear experts for interpretable and fast outcome prediction. In: ICPM 2023 Workshops, LNBIP, vol 503. Springer, https://doi.org/10.1007/978-3-031-56107-8_11
29. Folino F, Folino G, Guarascio M, Pontieri L (2024) Data- & compute-efficient deviance mining via active learning and fast ensembles. JIIS 62:995–1019. https://doi.org/10.1007/s10844-024-00841-4
30. Folino F, Folino G, Guarascio M, Pontieri L (2025) The force of few: boosting deviance detection in data scarcity scenarios through self-supervised learning and pattern-based encoding. Soft Comput 29:3675–3690. https://doi.org/10.1007/s00500-025-10646-4
31. Gianola A, Montali M, Winkler SM (2025) Smt techniques for data-aware process mining. KI-Künstliche Intell. https://doi.org/10.1007/s13218-025-00890-z
32. Goretti V, Basile D, Barbaro L, Di Ciccio C (2024) CONFINE: Preserving data secrecy in decentralized process mining. In: ICPM 2024 DC &Demos
33. Graziosi R, Ronzani M, Buliga A et al (2024) Generating the traces you need: a conditional generative model for process mining data. ICPM 2024:25–32. https://doi.org/10.1109/ICPM63005.2024.10680621
34. Graziosi R, Ronzani M, Buliga A et al (2025) Generating multi-perspective process traces using conditional variational autoencoders. Process Sci 2:8. https://doi.org/10.1007/s44311-025-00017-5
35. Huitzil I, Mazzotta G, Peñaloza R, Ricca F (2023) ASP-based axiom pinpointing for description logics. In: Proceedings on DL 2023, CEUR-WS.org, vol 3515
36. Ielo A, Law M, Fionda V, et al. (2023) Towards ILP-based LTLf passive learning. In: Proceedings on ILP'23. LNCS, vol 14363. Springer, pp 30–45. https://doi.org/10.1007/978-3-031-49299-0_3
37. Ielo A, Mazzotta G, Ricca F, Peñaloza R (2024) Towards ASP-based minimal unsatisfiable cores enumeration for LTLf. In: Short papers OVERLAY 2024, CEUR-WS.org, vol 3904, pp 49–55
38. Ko J, Maggi FM, Montali M, Peñaloza R, Pereira RF (2023) Plan recognition as probabilistic trace alignment. In: Proceedings on ICPM 2023. IEEE, pp 33–40. https://doi.org/10.1109/ICPM60904.2023.10271943
39. Maggi FM, Montali M, Peñaloza R (2025) Probabilistic temporal reasoning using superposition semantics. ACM Trans Comput Log 26(2):1–26. https://doi.org/10.1145/3714427
40. Marangone E, Di Ciccio C, Friolo D, et al. (2023) MARTSIA: Enabling data confidentiality for blockchain-based process execution. In: Proceedings EDOC 2023. Springer, pp 58–76. https://doi.org/10.1007/978-3-031-46587-1_4
41. Marangone E, Spina M, Di Ciccio C, Weber I (2024) CAKE: Sharing slices of confidential data on blockchain. In: Proceedings on CAiSE Forum, pp 138–147
42. Mazzotta G, Ricca F, Truszczynski M (2024) Quantifying over optimum answer sets. TPLP 24(4):716–736. https://doi.org/10.1017/S1471068424000395
43. Peñaloza R (2023) Semiring provenance in expressive description logics. In: Proceedings on DL 2023, CEUR-WS.org, vol 3515
44. Peñaloza R, Ricca F (2022) Pinpointing axioms in ontologies via ASP. In: Proceedings on LPNMR 2022. Springer, LNCS, vol 13416, pp 315–321. https://doi.org/10.1007/978-3-031-15707-3_24
45. Pesic M, Schonenberg H, van der Aalst WMP (2007) DECLARE: full support for loosely-structured processes. In: Proceedings on EDOC 2007. IEEE, pp 287–300. https://doi.org/10.1109/EDOC.2007.14
46. Roveri M, Ciccio C, Francescomarino C, Ghidini C (2024) Computing unsatisfiable cores for LTLf specifications. JAIR 80:517–558. https://doi.org/10.1613/jair.1.15313
47. van der Aalst W (2016) Process mining: data science in action. Springer, Berlin. https://doi.org/10.1007/978-3-662-49851-4
48. Xiong J, Xiao G, Kalayci TE et al (2022) A virtual knowledge graph based approach for object-centric event logs extraction. In: ICPM WS, Springer LNBIP, vol 468, pp 466–478. https://doi.org/10.1007/978-3-031-27815-0_34