# Tiramisù: A Recipe for Visual Sensemaking of Multi-Faceted Process Information

Anti Alman[1][0000−0002−5647−6249], Alessio Arleo[2][0000−0003−2008−3651], Iris Beerepoot[3][0000−0002−6301−9329], Andrea Burattin[4][0000−0002−0837−0183], Claudio Di Ciccio[5][0000−0001−5570−0475], and Manuel Resinas[6][0000−0003−1575−406X]

[1] University of Tartu, Estonia, anti.alman@ut.ee
[2] TU Wien, Austria, alessio.arleo@tuwien.ac.at
[3] Utrecht University, Netherlands, i.m.beerepoot@uu.nl
[4] Technical University of Denmark, Kgs. Lyngby, Denmark, andbur@dtu.dk
[5] Sapienza University of Rome, Italy, claudio.diciccio@uniroma1.it
[6] Universidad de Sevilla, Spain, resinas@us.es

**Abstract.** Knowledge-intensive processes represent a particularly challenging scenario for process mining. The flexibility that such processes allow constitutes a hurdle as it is hard to capture in a single model. To tackle this problem, multiple visual representations of the same processes could be beneficial, each addressing different information dimensions according to the specific needs and background knowledge of the concrete process workers and stakeholders. In this idea paper, we propose a novel framework leveraging visual analytics for the interactive visualization of multi-faceted process information, aimed at easing the investigation tasks of users in their process analysis tasks. This is primarily achieved by an interconnection of multiple visual layers, which allow our framework to display process information under different perspectives and to project these perspectives onto a domain-friendly representation of the context in which the process unfolds. We demonstrate the feasibility of the framework through its application in two use-case scenarios in the context of healthcare and personal information management.

**Keywords:** Visual analytics · Process mining · Knowledge-intensive processes.

## 1 Introduction and background

Process mining (PM) is the discipline aimed at extracting information from events recorded by information systems executing processes [3]. The operational context of process mining is multi-variate since data potentially stems from multiple sources and pertains to diverse dimensions (control flow, time, resources, object lifecycles). As argued by the pivotal work of Beerepoot et al. [8], a critical, yet largely unaddressed, issue of process mining is the fixed-granularity level of process analysis, namely the inability to navigate through distinct, and possibly domain-specific, dimensions. This problem limits the effectiveness of process improvement investigation, as studied by Kubrak et al. [23], as it entails a partial view over an inherently complex search space, which is typical of knowledge-intensive processes [13, 14]. Visual Analytics (VA) is the discipline of supporting users' analytical reasoning through the use of interactive visual interfaces [22]. VA is intended to keep the user *within* the analysis loop, so that human comprehension and perception actively contribute to generating new insights and increasing confidence with the analysis results. VA has extensively studied the problem of representing complex, multi-faceted phenomena [20, 28], which is also the type of phenomena often encountered within real-life business processes.

Thus far, there has been limited cooperation between PM and VA [25] despite the remarkable achievements their interplay could bring about [19]. In this paper, we argue that VA can play a pivotal role in addressing the above limitation of fixed-granularity in process mining. And, with a cross-pollinating effect, PM can equip VA with a collection of established algorithms and techniques for the automated generation of process-oriented representations of system dynamics.

In this idea paper, we theorize a novel framework, which we call the `Tiramisù` approach. It relies on VA for the interactive investigation of factual evidence of process insights revolving around information mined from data sources that go under the name of *sequence event data*, to use VA's terminology, or *event logs*, as per the PM nomenclature. Our framework follows a multi-layered approach, providing end users with a context-aware visualization integrating classical PM representation elements (such as workflow nets [32], directly-follows graphs [2], or declarative process maps [4]) with additional diagrams and visual cues tailored for context-variable and metadata representations (e.g., timelines and calendars for time, geographical or building maps for space, etc.). We categorize the different data facets, discuss how and where they fit our framework, and provide representation examples, with the overarching goal of describing a design process that enables practitioners to integrate different data sources into an organic representation fitting a PM workflow. The interconnection of the different layers and their anchorage to a *backdrop* (i.e., a representation of the context, such as calendars or maps), onto which the representations are projected, provides the user with a means to foster explainability for process analysis and, eventually, enhancement. This overcomes a limitation of the few proposals that combine PM and VA [17, 31], which do not consider the representation of domain-specific dimensions, thus limiting their utility for knowledge-intensive processes.

In the remainder of this paper, Sec. 2 describes our framework, incl. its key concepts and their interplay. Sec. 3 showcases two scenarios in the area of knowledge-intensive processes. Sec. 4 acknowledges some known limitations of our work, paving the way for future work discussed in the final section Sec. 5.

## 2 The Tiramisù framework

`Tiramisù` is a VA framework designed to support sensemaking when dealing with complex PM event sequences. Our goal is to augment existing process models

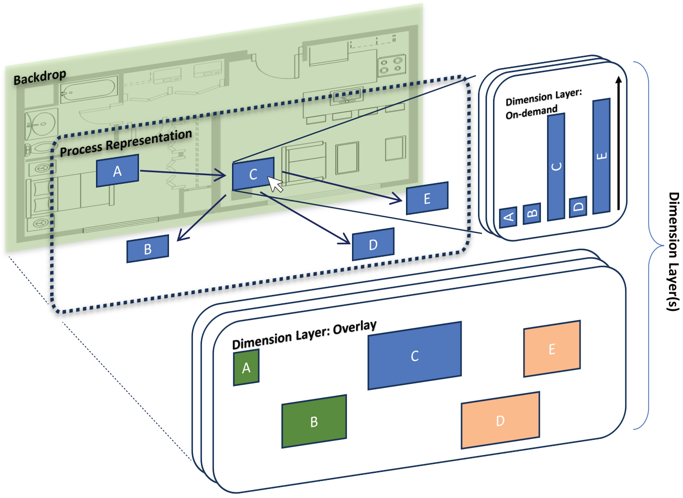


**Fig. 1.** An exemplification of the `Tiramisù` framework. Several dimension layers can be visualized concurrently, either overlaid on top of other layers or as an on-demand information. In this example, the backdrop resembles the one shown in Scenario 1 (see Section 3.1), wherein it provides the process representation within a spatial context.

and upgrade process maps with additional dimensions that can provide further context during the analysis, easing the generation of insights and generally providing a more comprehensive understanding of the process and phenomenon under investigation trough visualizations that better align with the mental model of the users. With this goal in mind, we design `Tiramisù` as a multi-level framework (see Fig. 1), reminiscent of the popular *tiramisù* dessert. The framework is structured with a "backdrop" providing context (see Section 2.1), a main layer with the process representation (Section 2.2), and one or more dimension layers (Section 2.3). We describe the individual parts of the framework in the following.

### 2.1 Backdrop

The backdrop (also *base*) layer acts as a common context for the process and for all the further dimensions that we include in the final interactive visualization. In the metaphor with the *tiramisù*, the backdrop would be the cream that permeates all the dessert's layers. It is designed to reflect the context of the current application, presenting the user with a familiar environment that acts as a "framing" device for the other process data and information. In PM, the idea of framing the process representation into a context that is not oriented to a specific task [12] is rather novel. From a visualization perspective, it provides a common 2D context to all the other layers of our framework, e.g., a spatial reference (a geographic location or the layout of a building, see Section 3.1), or a temporal reference (a calendar, see Section 3.2). The positioning of the process models' nodes encodes up to two of the available dimensions in process information, selected by the domain expert and depending on the analysis task (see also [31]). This differs from the majority of PM models, where no information is encoded in the positioning of the nodes in the plane.

### 2.2 Process representation

This section of the framework visualizes the process model. This is considered as a core layer, since it encapsulates the gist of the process behavior under study. In the tiramisù metaphor, it would consist of the coffee-imbued *Savoiardi* biscuits, which give structure and texture to the dessert. Without loss of generality, we focus on models that are graph-based. This entails the presence of two main elements: *nodes* and *edges*, where nodes commonly represent activities of the business process and edges represent (usually temporal) relations between these activities. In our framework, nodes are anchored to the backdrop, meaning that their positioning is consistent with the context expressed in the base layer. Nodes can also encode further attributes in other visual channels, such as size, color, and transparency. Node appearance can also be encoded as glyphs, as we will see in Section 3.1 with Scenario 1. Similarly, visual channels related to edges can be manipulated (e.g., width, transparency, and color). Moreover, edge visualization can be explicit (i.e., visible), or implicit. In this second case, the connections between the nodes can either be extrapolated from the context without being shown, or placed in one of the dimension layers of the framework.

### 2.3 Dimension layers

Our framework supports further layers to be superimposed on the first two, in a details-on-demand fashion [29]. These layers can serve different purposes, including: (i) provide contextual information on individual or arbitrarily small groups of nodes and edges in the process model (as in the scenario presented in Section 3); (ii) augment the backdrop visualization, such as enabling the visualization of a further dimension on the backdrop or on top of the process representation. In the tiramisù metaphor, the layers represent the different toppings, providing a distinct "character" to the dessert. In the framework, the layers are meant to expose hidden correlations between the dimensions within the auxiliary process information and the model itself. Our framework does not put any limit in the number of dimension layers. To avoid excessive visual clutter, the user can toggle the visibility of specific layers (following the details-on-demand principle). In turn, this option requires that the visualization designer knows or collaborates with a domain expert, in order to be able to rank and assign each dimension either to the backdrop or the dimension layers.

## 3 Use-case scenarios

In this section, we illustrate two scenarios that motivate our work through six user stories. We focus on two classes of knowledge-intensive processes, pertaining to healthcare [26] and personal information management [11]. Both scenarios demonstrate the capability of our solution to pinpoint deviations from the expected outcome, and investigate the cause for it in a multi-perspective fashion. The former, based on the work on [15], is described in Section 3.1 and illustrates the application of our solution with a particular focus on the spatial dimension. The latter, inspired by the investigation in [7], is discussed in Section 3.2 and involves the temporal dimension.

### 3.1 Behavioral deviation analysis in healthcare

Healthcare is one of the most difficult, but at the same time, one of the most promising domains to tackle both in the field of PM [26] and in the field of VA [10]. Indeed, the interplay of the two research fields in this context has the potential to yield remarkable outcomes, as shown in [17]. Thus, healthcare is a natural domain for applying the `Tiramisù` approach. As a demonstration, we have chosen the specific task of analyzing the sleeping routine of patients suffering from dementia or other similar diseases. Naively applying the already existing process mining approaches to this task can, for example, lead to relying on simple process maps (models) where the nodes represent various activities performed by the patients, and the arcs represent dependencies or other temporal constraints among these activities [15]. While the above models may be sufficient for a process mining expert to extrapolate meaningful process insights, it can quickly become challenging for domain experts, who may lack training in interpreting these formalisms: As discussed in [15, 16], the most likely users of such analysis tools would be the doctors and the nurses responsible for the care of the specific patient, and not PM experts.

When investigating this scenario in more depth, we also need to consider the types of analysis that would be relevant in this domain. To better structure our contribution, we identified the following user stories:

– **US1:** As a nurse, I want to inspect the typical sleeping routine of a patient;
– **US2:** As a nurse, I want to quickly spot the presence of deviations from the typical behavior and what these refer to;
– **US3:** As a nurse, I want to dive deeper into the context of a deviation.

As a backdrop for the above, we introduce the floor plan of the room of a specific patient, on top of which it is possible to project the activities performed by the patient in a way that more closely matches what doctors and nurses are familiar with, i.e., the environment that the patient interacts with. This backdrop can be seen on the left-hand side of Fig. 2.

As discussed in Section 2, the backdrop serves as a common context onto which it is possible to project process-related information. The right-hand side of Fig. 2 depicts one such floor plan, overlaid with the sleeping routine of a patient. Instead of representing activities as "meaningless" boxes (i.e., the shape and location of the boxes does not encode any dimension), we refer directly to the parts of the room, or specific appliances/furniture, related to the corresponding


**Fig. 2.** The backdrop used for Scenario 1, with the floor map where patients live (left) and the backdrop overlaid with the sleeping routine of a patient (right)

activities. This approach brings visualization closer to what the analysts (i.e., nurses and doctors) are already familiar with. It also introduces an intuitive spatial dimension to representing patient behavior. This visualization, in essence, allows us to address user story **US1**. The visual encoding of nodes (i.e., the pictograms/glyphs) refers to the frequency of the activities performed by the patient (delivered via different shades of green and transparency of nodes) while the visual encoding of edges refers to the frequency of transitioning from one activity to the other (via the thickness of the nodes).

The addition of overlays will allow for the reuse of the same backdrop also for other types of analysis. For example, the overlay shown on the left-hand side of Fig. 3 would be specifically designed in a manner that allows nurses and doctors to easily spot deviations from the sleeping routine of the patient, thus answering **US2**. In this case, we use a specific color for highlighting deviations from the reference behavior of the patient. This encoding (i.e., the color) is applied to both nodes and edges.

Furthermore, as shown in Fig. 3 (right), by interacting with the visualization, doctors and nurses will have the opportunity to select specific activities in order to dive deeper into the details related to a specific deviation via a pop-up (i.e, a window that opens on top of the glyph). Following the *details on demand* principle [29], such pop-up will report details regarding the underlying process execution, such as the actual Workflow net model (possibly with colors suggesting where deviations happened) [32], a DCR model [21], and statistical aspects such as average duration, etc. This, in turn, provides a possible solution for **US3**.

### 3.2 Personal productivity analysis in work processes

Personal information management (PIM) pertains to the organization of one's own activities, contacts, etc., through the use of software on laptops and smart devices. Similarly, personal informatics systems resort to one individual's own


**Fig. 3.** The backdrop overlaid with deviations from the typical routine. The figure on the left is meant for quickly identifying the deviation, while the figure on the right represents an interaction to dive deeper into the details of a deviation.

information to pursue the objective of aiding people to collect and reflect on their personal information [24]. Several techniques can be used to collect personal information such as non-participant observations, screen recordings, and timesheet techniques, each with their own advantages and disadvantages [30]. Regardless of the technique used, collected personal information can be seen as an event log, which can be analyzed using process mining techniques to discover personal work processes [9].

Depending on the characteristics of the work, personal work processes can be knowledge-intensive and significantly unstructured, which means they present the aforementioned challenges. In this scenario, we focus on the personal work processes performed by an academic during her daily work, which involves conducting research, preparing lessons, grading students or reviewing research papers, among many other activities. Specifically, our focus is on the retrospective analysis of the influence of the personal work processes a person has followed during a certain period of time on the positive or negative outcome of some specific task (e.g., missing the deadline for submitting a review). In particular, we identify the following user stories as representative of this scenario:

- **US4:** As a researcher, I want to investigate when I was working on the reviews;
- **US5:** As a researcher, I want to shed light on the time aspects that are relevant in the context of my reviews;
- **US6:** As a researcher, I want to dive deeper into my reviewing activity on Friday afternoon to understand what caused the delay.

To this end, we use a calendar with a user-specified time period as the backdrop of the visual representation. This backdrop is overlaid with a summary of the activities performed during the time period, as exemplified in Fig. 4. The activities are represented as boxes with the color referring to the activity type

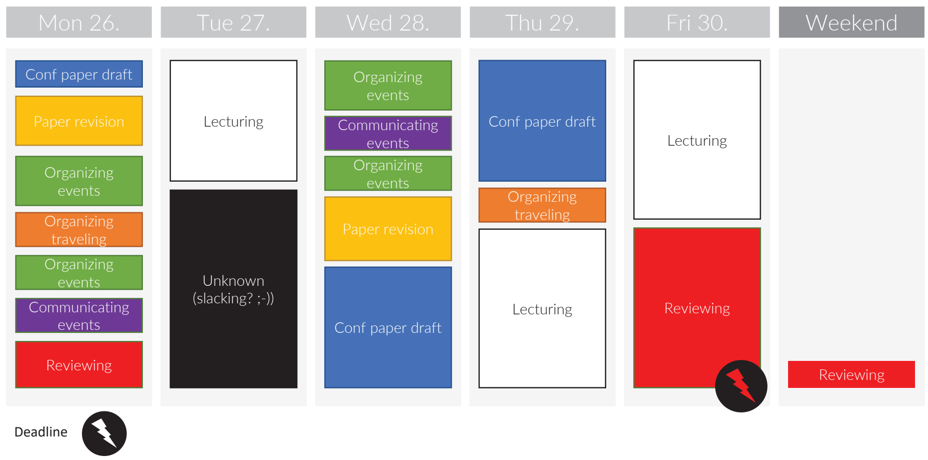


**Fig. 4.** The backdrop of Scenario 2 overlaid with the weekly summary of the activities and the deadlines, represented as lightning bolts

(e.g., all lecturing activities in white, despite possibly being lectures on different topics/courses). The position and size of each box represent when the activity started and its duration, respectively. Thereby, the researcher can spot when she is reviewing the assigned papers, thus addressing **US4**. Notice the lightning bolt icon in the figure. This additional overlay signals the reviewing deadline through that glyph. The icon uses the same color code as the activities to indicate the activity the deadline refers to. Its presence helps the researcher put in a relationship the duration and allocation of work with the expected consignment time, thus catering for **US5**.

Similarly to Section 3.1, we use as a backdrop a representation that is intuitive to the user and specific to the scenario under analysis. However, unlike in the previous scenario, the representation refers to the time dimension (a calendar) instead of the space dimension (a floor map). This aspect illustrates that the concept of backdrop can (and must) be adapted to the dimension that best fits the scenario being addressed.

Now, as the academic wants to understand what is the cause of the delay for delivering the review (**US6**), she might want to dive deeper into the details by clicking on any of the activities of the weekly summary. For instance, let us say that the academic wants to know more details about what happened on Friday evening while performing the reviewing activity that prevented her from finishing the review on time. To this end, the academic can click on the reviewing activity of Friday evening, which will open a popup with multi-faceted details related to the activity at hand (again following the *details on demand* principle [29]), as illustrated in Fig. 5. In this case, the details include a directly-follows graph representing the sub-activities performed as part of the reviewing activities and a summary of the interruptions that happened during that time interval.

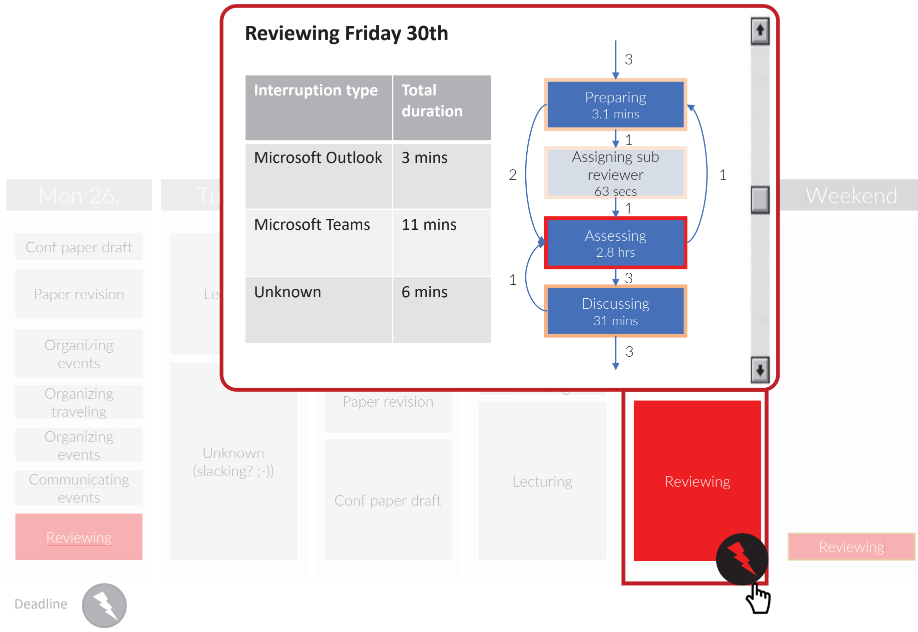


**Fig. 5.** The backdrop of Scenario 2 showing the details of the review activity on Friday evening.

## 4 Limitations

Evidence-based process discovery approaches may, in some cases, suffer from reliability issues, given that only some parts of the processes (especially in the case of knowledge-intensive processes) are recorded by computerized systems [18]. Our framework is in this sense no different. The injection of additional facts stemming from domain experts and users is complementary to this investigation and paves the path for future work.

The design of the visualizations in backdrop, and in the other layers is a challenge of its own, and should be tackled by visualization experts in close cooperation with domain experts and target users. Investigating this aspect goes beyond the scope of this paper. However, the interested reader is referred to the work of Sirgmets et al. [31], which presents a framework aimed to guide design choices for the effective visualization of analytical data, and to the paper by Munzner et al. [27], where a nested model for visualization design is presented.

Data-based process analysis tasks often involve a significant amount of time for extracting, reformatting, and filtering event logs from information systems [1]. Our framework requires these preliminary operations too, but with the addition that knowledge-intensive processes oftentimes record executions over a heterogeneous set of applications and devices in partially structured or unstructured formats [14]. Furthermore, the notion of a process instance (case) tends to be less defined in such contexts, thus requiring a prior customizable reconciliation of shared events [5, 6]. The emergence of novel and event data meta-models that are less centered around the concept of a case could be beneficial to the information processing we envision in `Tiramisù` [33].

Finally, our framework covers the visualization aspects related to VA in PM. However, the *interaction* aspects are only partly discussed, leaving a gap which is to be addressed by further iterations of this framework. This effort would be made lighter and more effective if a *task taxonomy* about VA in PM were created, opening up a novel research direction with the potential to further bring together the two disciplines.

## 5 Conclusions and future remarks

In this idea paper, we have presented `Tiramisù`, a framework based on the multi-layered representation of mined process information helping the user navigate the multi-faceted information at hand while keeping that information consistently linked and navigable across multiple different dimensions. We demonstrated our solution through its application to scenarios illustrating its suitability in the context of knowledge-intensive processes.

We foresee the following research endeavors in the future originating from our current idea. An evaluation based on an empirical study involving users is in our plans to assess the efficacy and effectiveness of the `Tiramisù` framework. The implementation of a working prototype is crucial to this extent and is also part of our agenda. This goes in tandem with the future work of identifying the different user types and their specific visualization needs within the context of the `Tiramisù` framework. In connection to the previous, we envision that the customization and refactoring of layers, their aggregation rules, and backdrops will be key challenges to be addressed in the future. Finally, we observe that stepping from a multi-layered single-backdrop design to a tree-like hierarchical structure of backdrops could initiate a new path toward a significant extension of the expressive richness guaranteed by our framework.

**Acknowledgments.** This work was started at Schloss Dagstuhl (Leibniz-Zentrum für Informatik), seminar 23271 "Human in the (Process) Mines", and we thank the institute and the organizers for giving us the opportunity to participate. The work of A. Alman was supported by the European Social Fund via "ICT programme" measure and by the Estonian Research Council grant PRG1226. The work of C. Di Ciccio was supported by the Italian Ministry of University and Research (MUR) under PRIN grant B87G22000450001 (PINPOINT) and by project SERICS (PE00000014) under the NRRP MUR program funded by the EU-NGEU. The work of M. Resinas was partially supported by projects PID2021-126227NB-C21/ AEI/10.13039/501100011033/ FEDER, UE and TED2021-131023B-C22/ AEI/10.13039/501100011033/ Unión Europea NextGenerationEU/PRTR.

## References


1. van der Aalst, W.M.P.: Process Mining - Data Science in Action, Second Edition. Springer (2016). https://doi.org/10.1007/978-3-662-49851-4
2. van der Aalst, W.M.P.: Foundations of process discovery. In: van der Aalst and Carmona [3], pp. 37–75. https://doi.org/10.1007/978-3-031-08848-3_2
3. van der Aalst, W.M.P., Carmona, J. (eds.): Process Mining Handbook, Lecture Notes in Business Information Processing, vol. 448. Springer (2022)
4. Alman, A., Di Ciccio, C., Maggi, F.M., Montali, M., van der Aa, H.: RuM: Declarative process mining, distilled. In: BPM. pp. 23–29. Springer (2021). https://doi.org/10.1007/978-3-030-85469-0_3
5. Baier, T., Mendling, J., Weske, M.: Bridging abstraction layers in process mining. Inf. Syst. **46**, 123–139 (2014). https://doi.org/10.1016/j.is.2014.04.004
6. Bayomie, D., Di Ciccio, C., Mendling, J.: Event-case correlation for process mining using probabilistic optimization. Inf. Syst. **114**, 102167 (2023). https://doi.org/10.1016/j.is.2023.102167
7. Beerepoot, I., Barenholz, D., Beekhuis, S., Gulden, J., Lee, S., Lu, X., Overbeek, S., van de Weerd, I., van der Werf, J.M., Reijers, H.: A window of opportunity: Active window tracking for mining work practices. In: ICPM. pp. 57–64. IEEE (2023)
8. Beerepoot, I., et al.: The biggest business process management problems to solve before we die. Comput. Ind. **146**, 103837 (2023). https://doi.org/10.1016/j.compind.2022.103837
9. Bertrand, Y., den Abbeele, B.V., Veneruso, S., Leotta, F., Mecella, M., Serral, E.: A survey on the application of process mining to smart spaces data. In: ICPM Workshops. pp. 57–70. Springer (2022). https://doi.org/10.1007/978-3-031-27815-0_5
10. Caban, J.J., Gotz, D.: Visual analytics in healthcare - opportunities and research challenges. J. Am. Medical Informatics Assoc. **22**(2), 260–262 (2015). https://doi.org/10.1093/jamia/ocv006
11. Catarci, T., Dix, A., Katifori, A., Lepouras, G., Poggi, A.: Task-centred information management. In: DELOS Conference. vol. 4877, pp. 197–206. Springer (2007)
12. de Leoni, M., Suriadi, S., ter Hofstede, A.H.M., van der Aalst, W.M.P.: Turning event logs into process movies: animating what has really happened. Softw. Syst. Model. **15**(3), 707–732 (2016). https://doi.org/10.1007/s10270-014-0432-2
13. De Weerdt, J., Schupp, A., Vanderloock, A., Baesens, B.: Process mining for the multi-faceted analysis of business processes - A case study in a financial services organization. Comput. Ind. **64**(1), 57–67 (2013). https://doi.org/10.1016/j.compind.2012.09.010
14. Di Ciccio, C., Marrella, A., Russo, A.: Knowledge-intensive processes: Characteristics, requirements and analysis of contemporary approaches. J. Data Semant. **4**(1), 29–57 (2015). https://doi.org/10.1007/s13740-014-0038-4
15. Di Federico, G., Burattin, A.: Do you behave always the same? In: Process Mining Workshops. pp. 5–17 (2023)
16. Di Federico, G., Burattin, A., Montali, M.: Human behavior as a process model: Which language to use? In: ITBPM@BPM. pp. 18–25. CEUR-WS.org (2021), https://ceur-ws.org/Vol-2952/paper_293a.pdf
17. Dixit, P.M., Caballero, H.S.G., Corvò, A., Hompes, B.F.A., Buijs, J.C.A.M., van der Aalst, W.M.P.: Enabling interactive process analysis with process mining and visual analytics. In: HEALTHINF. pp. 573–584. SciTePress (2017). https://doi.org/10.5220/0006272605730584

18. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: Fundamentals of Business Process Management, Second Edition. Springer (2018). https://doi.org/10.1007/978-3-662-56509-4
19. Gschwandtner, T.: Visual analytics meets process mining: Challenges and opportunities. In: SIMPDA. pp. 142–154. Springer (2015). https://doi.org/10.1007/978-3-319-53435-0_7
20. Hadlak, S., Schumann, H., Schulz, H.J.: A survey of multi-faceted graph visualization. In: EuroVis (STARs). pp. 1–20 (2015)
21. Hildebrandt, T.T., Mukkamala, R.R.: Declarative event-based workflow as distributed dynamic condition response graphs. arXiv preprint arXiv:1110.4161 (2011)
22. Keim, D., Andrienko, G., Fekete, J.D., Görg, C., Kohlhammer, J., Melançon, G.: Visual analytics: Definition, process, and challenges. Springer (2008)
23. Kubrak, K., Milani, F., Nolte, A.: A visual approach to support process analysts in working with process improvement opportunities. Bus. Process Manag. J. **29**(8), 101–132 (2023)
24. Li, I., Dey, A.K., Forlizzi, J.: A stage-based model of personal informatics systems. In: CHI. pp. 557–566. ACM (2010). https://doi.org/10.1145/1753326.1753409
25. Miksch, S.: Visual analytics meets process mining: Challenges and opportunities. In: ICPM. pp. xiv–xiv (2021). https://doi.org/10.1109/ICPM53251.2021.9576854
26. Munoz-Gama, J., Martin, N., et al.: Process mining for healthcare: Characteristics and challenges. J. Biomed. Informatics **127**, 103994 (2022). https://doi.org/10.1016/j.jbi.2022.103994
27. Munzner, T.: A nested process model for visualization design and validation. IEEE Trans. Vis. Comput. Graph. **15**(6), 921–928 (2009). https://doi.org/10.1109/TVCG.2009.111
28. Raidou, R.G.: Visual analytics for the representation, exploration, and analysis of high-dimensional, multi-faceted medical data. Biomedical Visualisation: Volume 2 pp. 137–162 (2019)
29. Shneiderman, B.: The eyes have it: A task by data type taxonomy for information visualizations. In: VL. pp. 336–343. IEEE (1996). https://doi.org/10.1109/VL.1996.545307
30. Sinik, T., Beerepoot, I., Reijers, H.A.: A peek into the working day: Comparing techniques for recording employee behaviour. In: RCIS. pp. 343–359. Springer (2023). https://doi.org/10.1007/978-3-031-33080-3_21
31. Sirgmets, M., Milani, F., Nolte, A., Pungas, T.: Designing process diagrams - A framework for making design choices when visualizing process mining outputs. In: OTM. pp. 463–480. Springer (2018). https://doi.org/10.1007/978-3-030-02610-3_26
32. van der Aalst, W.M.P.: The application of petri nets to workflow management. Journal of Circuits, Systems, and Computers **8**(1), 21–66 (1998). https://doi.org/10.1142/S0218126698000043
33. Wynn, M.T., Lebherz, J., et al.: Rethinking the input for process mining: Insights from the XES survey and workshop. In: ICPM workshops. pp. 3–16. Springer (2021). https://doi.org/10.1007/978-3-030-98581-3_1