Tiramisù: Making Sense of Multi-Faceted Process Information Through Time and Space

Anti Alman [1], Alessio Arleo [2,3], Iris Beerepoot [4*], Andrea Burattin [5], Claudio Di Ciccio [4], Manuel Resinas [6]

[1]Institute of Computer Science, University of Tartu, Narva mnt, Tartu, 51009, Estonia.

[2]Department of Mathematics and Computer Science, Eindhoven University of Technology, De Zaale, Eindhoven, 5612 AZ, The Netherlands.

[3]Centre for Visual Analytics Science and Technology, Technische Universität Wien, Favoritenstrasse 9-11, Vienna, 1040-A, Austria.

[4]Process Science, Utrecht University, Heidelberglaan, Utrecht, 3584 CS, Utrecht, The Netherlands.

[5]Software Systems Engineering, DTU Technical University of Denmark, Richard Petersens Plads, Kgs. Lyngby, 2800,Denmark.

[6]SCORE Lab, I3US, Universidad de Sevilla, Spain.

*Corresponding author(s). E-mail(s): i.m.beerepoot@uu.nl;
Contributing authors: anti.alman@ut.ee; a.arleo@tue.nl; andbur@dtu.dk; c.diciccio@uu.nl; 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 they are 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 paper, we propose, describe, and evaluate a framework, named `Tiramisù`, that leverages visual analytics for the interactive visualization of multi-faceted process information, aimed at supporting the investigation and insight generation of users in their process analysis tasks.

`Tiramisù` is based on a multi-layer visualization methodology that includes a visual backdrop that provides context and an arbitrary number of superimposed and on-demand dimension layers. This arrangement allows our framework to display process information from different perspectives and to project this information onto a domain-friendly representation of the context in which the process unfolds. We provide an in-depth description of the approach's founding principles, deeply rooted in visualization research, that justify our design choices for the whole framework. We demonstrate the feasibility of the framework through its application in two use-case scenarios in the context of healthcare and personal information management. Plus, we conducted qualitative evaluations with potential end users of both scenarios, gathering precious insights about the efficacy and applicability of our framework to various application domains.

Keywords: Process mining, visual analytics, knowledge-intensive processes, visualization

1 Introduction

Process mining (PM) is the discipline aimed at extracting information from events recorded by information systems executing processes (van der Aalst and Carmona, 2022). 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 in the pivotal work of Beerepoot et al (2023b), a critical, yet largely unaddressed, issue of process mining is 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 (2023), as it entails a partial view over an inherently complex search space, which is typical of knowledge-intensive processes (De Weerdt et al, 2013; Di Ciccio et al, 2015). Visual Analytics (VA) is the discipline of supporting users' analytical reasoning through the use of interactive visual interfaces (Keim et al, 2008). 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 (Keim et al, 2008; Raidou, 2019), 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 (Miksch, 2021) despite the remarkable achievements that their interplay could bring (Gschwandtner, 2015). 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. We ground our assumption on two facts. First, VA gained popularity due to its user-oriented design meant to leverage the human capability of decoding visual information (Miksch, 2021). Second, there is a significant corpus of visualization literature concerning the representation of event sequence data —see, e.g., the following survey papers by van der Linden et al (2023), and Yeshchenko and Mendling (2024).

We reflected on the following research question: "How to present PM information to different audiences (including non-PM or VA experts), each with different tasks to perform on event sequence data?". This problem is faced in several application domains. In hospitals and nursing homes, for example, concurrent execution of processes (e.g., treatments, patient transfers, operations, etc.) is commonplace, with each process requiring different personnel (e.g., nurses, doctors, caretakers, etc.) all of whom may have widely varying technical literacy and expertise. How can we represent a treatment process to a doctor rather than a nurse or the habits of a patient in a nursing home for the assessment of the progress of neurological diseases (e.g., Alzheimer's, dementia) while, at the same time, accounting for these differences? On a similar note, let us consider the area of personal information management for knowledge workers. In that domain, processes represent concurrent and often intertwined work-related routines, which may be further effected by the workers personal obligations. How would we represent process information in such a setting in general? And more specifically, how can we support knowledge workers in assessing their workflow and prioritizing tasks, when, for example, deadlines approach?

Within this context and motivation, in this paper we theorize (and ground on existing literature) a novel conceptual framework, which we call `Tiramisù`. Akin to an architectural pattern (Rozanski and Woods, 2011), `Tiramisù` pursues the objective of providing a set of conceptual elements and patterns, alongside the interfacing connections and information flow for a class of information systems. It is aimed to apply in the context of process and event sequence data, and is steered particularly by the competing, driving forces of understandability, adaptability, tunability, and aesthetics exerted by the diversified target users. To this end, `Tiramisù` introduces a paradigm shift for the development viewpoint of software architectures employed by future analytic tools for process and event sequence data.

Specifically, `Tiramisù` 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. It also allows the integration of non-process data sources that can provide additional context to the process-based information. Our conceptual framework follows a multi-layered approach, providing end users with a context-aware visualization that allows the integration of classical PM representation elements — such as workflow nets (van der Aalst, 1998), directly-follows graphs (van der Aalst, 2022), or declarative process maps (Alman et al, 2021)— with additional diagrams and visual cues tailored for context-variable and metadata representations (e.g., timelines and calendars for time, geographical or floor 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 domain-specific 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 (Dixit et al, 2017; Sirgmets et al, 2018), which do not consider the representation of domain-specific dimensions, thus limiting their utility for knowledge-intensive processes. Furthermore, providing a domain-specific context to anchor all process-oriented visualizations gives the user a familiar context that facilitates the interpretation of process-oriented visualizations, making `Tiramisù` particularly well suited for users who are not process mining experts.

This paper extends and improves on our preliminary work (Alman et al, 2023) with the following novel contributions. We refine and provide more details on the components of the framework. We introduce and make available two prototypical implementations of PM visualizations developed using our framework for two different use cases. We evaluate our approach through six expert interviews, through which we investigate how potential end users consider our approach in terms of expressiveness and usability. Finally, we discuss the feedback obtained from the study participants to draw conclusions, outline future perspectives, and reflect on `Tiramisù` limitations.

In the remainder of this paper, Section 2 discusses the related work, Section 3 describes our framework including the founding principles, and its components and their interplay. Section 4 showcases the usability of our framework in two different scenarios in the area of knowledge-intensive processes. In Section 5, we describe our evaluation of our framework through semi-structured interviews with potential end users. In Section 6, we reflect on the limitations of our work. Finally, Section 7 discusses the conclusion and draw paths future endeavours.

2 Related Work

In this section, we recapitulate core contributions in the scientific literature concerning the interplay of VA and PM. Thereafter, we summarize the limitations of existing techniques and discuss how our approach aims to overcome them.

Systematic analyses concerning the applicability of VA techniques to PM are relatively recent. In her position paper, Gschwandtner (2015) elucidates the benefits and potentials of visualization approaches for PM tasks, alongside a selection of examples applied to clinical and shopping data. Challenges and opportunities are enumerated, with a clarion call to research on such topics. Other papers extend the discussion on specific aspects of PM, such as conformance checking (Rehse et al, 2022), and advocate the adoption of VA to support business process improvement (Kaouni et al, 2021). Kriglstein et al (2016) investigate how process mining techniques can be *categorized* as VA aspects. The authors apply their classification scheme to a selection of software plug-ins of the open-source process mining toolkit ProM (van der Aalst et al, 2009). These manuscripts informed and motivated our endeavor, which we targeted at a full-fledged approach with a prototypical implementation.

Researchers operating in both the VA and PM areas have proposed tools and platforms conjoining methodologies stemming from the two disciplines. In the following, we describe some notable ones that have a conceptual overlap with our approach. Schuster et al (2024) propose a visualization approach to display process variants based on event data, implemented within the process mining tool Cortado (Schuster et al, 2023). It encompasses low-level and high-level abstractions, and considers different granularity levels of the temporal information provided as input. We share with this solution the objective of differentiating the visualizations depending on the kind of information being under inspection while keeping the different perspectives linked. However, we broaden our investigation beyond the task of variant analysis while considering multiple and multi-faceted information sources.

van der Aalst et al (2011) introduce an approach to enable the replay of event logs through visual representations of models from different perspectives (or "angles," as they are referred to in the paper). This methodology pivots around the idea that process diagrams can be considered as "maps", and cartography is typically well understood by end users. To highlight the elements described in the process model, the authors thus promote visual encodings that are reminiscent of geographic data visualization. Building on this metaphor, the paper introduces the replay feature: as the event log is played out, the system status (and hence its representation) changes, supporting users in identifying bottlenecks and exploring the process in general. A similar approach, broadening the analysis to multi-perspective process mining, is proposed by Mannhardt et al (2015).

Sirgmets et al (2018) propose a framework to support the development of visualizations of process diagrams. At its highest level of the abstraction hierarchy, it builds upon an "Encode" and an "Interact" pillar. These two then present nested classes that drill down to support the selection of the proper encodings and interaction techniques on the basic elements of the visualization. This framework has the advantage of being designed using the principles of the nested visualization model by Munzner (2009), thus providing a solid theoretical background. Compared to `Tiramisù`, this framework focuses on the *process*. While it exposes the "faceting" (i.e., the inclusion of multiple dimensions), it is mainly oriented to the creation of a *diagram* for the process, while `Tiramisù` is a framework for the visualization of multi-faceted complex processes.

Dixit et al (2017) present *InterPretA* (Interactive Process Analytics), a VA tool to perform process-oriented analysis using existing conformance analysis techniques. The user can explore the results, identify issues, and focus only on sub-parts of the whole process if necessary. Auxiliary visualizations quantify and report additional information used to answer some questions about the process or for making reports. Differently from our proposed approach, this tool only focuses on conformance checking, without the possibility to include domain-specific dimensions. Moreover, only basic visualizations (e.g., histograms) are used, while the `Tiramisù` flexible methodology allows for more varied and domain- or user-oriented approaches.

In summary, papers in the literature generally agree on the potential contribution that VA and its methods could bring to PM, in terms both of expressiveness and insight generation. Existing work on VA frameworks for PM formalized and provided reproducible and reliable methods to apply VA principles to the exploration of process models. However, existing solutions and proposals do not account for a comprehensive, higher-level methodology that allows for the seamless integration of orthogonal process information. We remark that this missing aspect should be tailored to different combinations of users and domains. Therefore, we aim to include not only

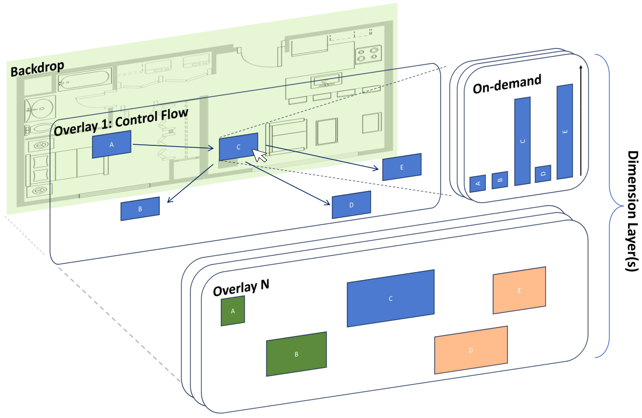

Fig. 1 An exemplification of the `Tiramisù` framework depicting the two types of dimension layers. Several dimension layers can be visualized concurrently, either overlaid on top of other layers or as an on-demand information. In this example, the backdrop resembles the one shown in Scenario 1 (see Section 4.1), wherein it provides a spatial context for the process analysis. In addition, process control flow is highlighted as a typical, but not necessarily required (see Section 4.2), type of overlay to be connected to the backdrop.

multi-faceted data but also the user in the design of VA techniques for PM, in order to make these compound techniques more comprehensible and accessible. In the next section, we present our solution.

3 The Tiramisù Framework

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

3.1 Founding Principles

`Tiramisù` is structured as a framework consisting of four main components: A *backdrop* providing context, one or more dimension layers, several visual mappings, and one visualization configuration. In the following, we illustrate the founding principles (FPs for short) of our approach.

FP1: Visualization Mantra. The `Tiramisù` multi-layer arrangement is rooted in the visualization mantra of "overview first, zoom and filter, details on demand" (Shneiderman, 2003). The combination of a backdrop and process-oriented representation provides a high-level view of the process and its context. The dimension layers, appearing on-demand, provide further information and variables related to the element of the process highlighted by the user. This interaction strategy follows a *focus+context* approach: the backdrop and the process-oriented representation remain visible (context), while the additional dimension layers provide the greater focus. All information is meant to be visible simultaneously with little to no occlusion.

FP2: Data, Tasks, Users. Another founding principle of `Tiramisù` is the *Design Triangle* by Miksch and Aigner (2014). When defining the backdrop and the additional layers, three pillars must be considered: the (type of) *data* to visualize, the *tasks* that visualization is designed to achieve, and the intended *users*. The data often drives the visualization approach selection, as often methods are better suited for some specific data types: for example, geo-referenced data finds its natural correspondence in cartographic maps. The *task* determines what a visualization is designed to support, thereby steering the comparison among methods supporting the same usage goals. In the visualization literature, taxonomies have been introduced to formalize and categorize tasks. The excellent survey by Filipov et al (2023) provides an extensive discussion on the existing task taxonomies in the literature. Last but not least, we devised `Tiramisù` to support a wider variety of potential users. The design of the backdrop and the dimension layers are meant to be understandable by an audience that goes beyond the community of PM experts. To achieve this objective, the conduction of user interviews and the guidance of potential end users are pivotal.

FP3: VA and process-oriented representation. The process-oriented representations used in `Tiramisù` are, in general, rooted in the existing PM methodology. These include node-link diagrams, flow charts, etc. However, depending on the specific domain and user stories to be tackled, other representations may be used as well (cf. Section 4.2). In `Tiramisù`, the presence of VA has a twofold purpose. First, it is meant to augment the process-oriented representation to deepen the user's understanding of the process by providing further context, additional information, and finer-grained detail. The objective is to support a holistic interpretation of the events, not only in the process itself but also in the causes for deviation from the usual or expected behavior. Second, it has the potential to make the representation more generally accessible. This characteristic is especially sought for as the intended users of the system are not limited to PM experts.

3.2 Input and Framework Components

Figure 2 depicts an overview of the inputs and components of `Tiramisù`. The upper part of the picture shows the data that is fed to `Tiramisù`. This data comes from two different sources, namely process data, colored in green, and non-process data, in orange. The lower part of the picture displays the components of the `Tiramisù` framework (in purple). Next, we explain each of these elements in detail.

Input

We classify the input into two separate classes: process data and non-process data. We describe these categories in the following.

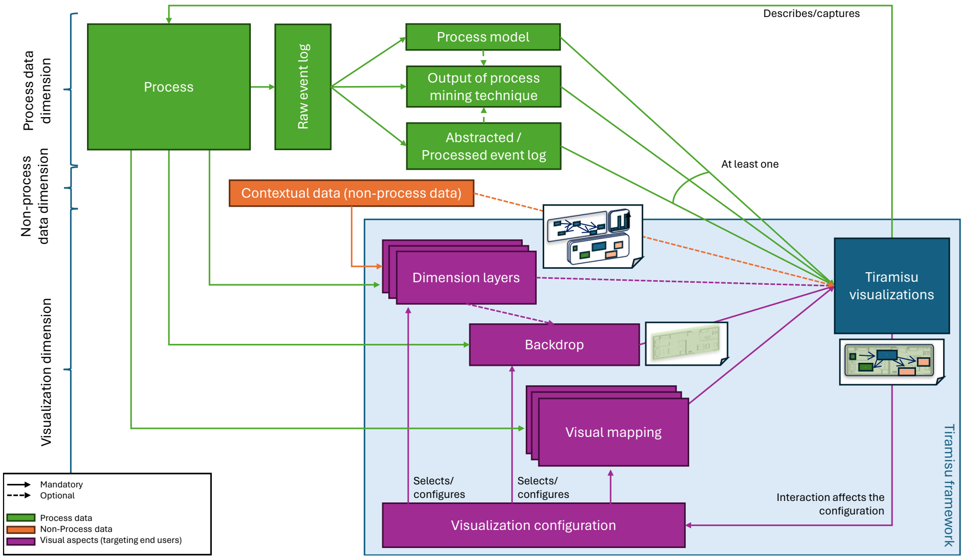

Fig. 2 A detailed overview of the `Tiramisù` framework, accurately representing the connections and interactions between the individual elements of the framework. In the "sticky notes" next to Dimension layers, Backdrop, and the whole `Tiramisù` visualization, examples from Fig. 1 are reported.

Process data.

Process data refers to the data obtained from the execution of the process we want to describe with `Tiramisù`. As it is typically done in process mining, we assume that process execution information is collected in a raw process event log. We use the term *raw* to emphasize that the collected data may not be at the right level of abstraction for visualization using `Tiramisù`. For instance, a raw event log may record the execution of the process by tracking the position of a process actor captured by movement sensors. To obtain insight into the execution of the process from this data, one would typically want to abstract the raw event log at the level of process activities and discover a process model from it (Di Federico and Burattin, 2023a; Soffer et al, 2019). This process model will then be one of the inputs that `Tiramisù` will use to create the visualization. In general, we envision that to build visualizations that describe the process, `Tiramisù` will use data from at least one of three possible source types: *(i)* the aforementioned process models; *(ii)* the abstracted or preprocessed event logs; *(iii)* the output of process mining analysis techniques computed from either the raw event log or the abstracted or preprocessed event log. This output could range from simple statistics or process performance indicators and diagnostics (e.g., execution frequencies, process cycle time, activity durations) to the output of more complex techniques like conformance checking metrics or alignments (e.g., the deviation depicted in Fig. 4).The choice between the type of input depends on the characteristics of the use case addressed, in accordance with FP2.

Non-process data.

With non-process data we refer to contextual information that does not directly pertain to the process execution itself, yet provides additional details of the environment in which the process unfolds. An example of non-process data is the location from which a person was working during the moment of process execution (see Scenario 2, Section 4.2). Unlike process data, which is mandatory for building `Tiramisù` visualizations, non-process data is optional (hence the dashed line in Fig. 2).

`Tiramisù` framework components.

The `Tiramisù` framework consists of four main components (purple boxes in Fig. 2), which are combined together in the "Tiramisu visualizations" (the dark blue box).

Backdrop.

The backdrop (also known as *base*) is a layer serving as a common context for the process and for all the other dimensions that we include in the final interactive visualization. In the metaphor with the tiramisù, the backdrop is the cream that permeates all the dessert's layers. The backdrop is specific for each domain and must be 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-oriented representations and additional information views. In Fig. 2, this relationship is depicted with the arrow connecting the *Process* and the *Backdrop*. From a visualization perspective, it provides a common 2D context to all the other layers of our framework. For example, a backdrop can be a spatial reference (a geographic location or the layout of a building, see Section 4.1), or a temporal reference (a calendar, see Section 4.2).

Dimension layers.

Our framework supports further layers to be superimposed on the backdrop, in a details-on-demand fashion (Shneiderman, 1996) as per FP1. These dimensions can represent either process data or non-process data (FP3), as displayed in Fig. 2 with the incoming arrows to *Dimension layers*. As depicted in Fig. 1, we identify two major categories of dimension layers:

- Overlay layers, which augment the visualization by mapping data to new layers on top of the backdrop;
- On-demand layers, which provide contextual information on individual or arbitrarily small groups of nodes and edges in the process model (as in the scenario presented in Section 4), implementing FP1. They can be triggered using mechanisms such as hovering or clicking over elements of the backdrop or representing the process.

Overlay layers bring the advantage of providing an explicit anchor to the backdrop. This feature is aimed to put the information in context. For instance, if the backdrop is a floor map, the information shown by the overlay layers is linked to that floor, thereby helping the user understand the relationships between reported facts that occur in the same location.

Among the overlay layers, it is possible that one of them directly projects a representation of the process, for example its control flow. This layer would encapsulate the gist of the process behavior under study. In the tiramisù metaphor, this would consist of the coffee-imbued *Savoiardi* biscuits, which give structure and texture to the dessert. The choice of which specific dimension layers to use is strongly correlated with the choice of the backdrop and the use cases being addressed by the visualization as per FP2. The current iteration of the `Tiramisù` framework envisions substantial flexibility in this regard. For example, the overlay layers in Section 4.1 are mainly inspired by process maps, while Section 4.2 completely foregoes any graph-based process representation.

Visual mappings.

Each overlay dimension layer is determined by a mapping from the process data, which can be a process model, statistics, an abstracted event log, or a combination of them, or non-process data to the backdrop. This mapping is defined in the visual mapping component. A common example is mapping graph-based process models into the visual context provided by a backdrop that encodes spatial information (e.g., a floor map). 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. This differs from the majority of PM models, where no information is encoded in the positioning of the nodes in the plane. 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 4.1. 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 other dimension layers of the framework.

Visualization configuration.

The visualization configuration selects and configures the other three components of the framework. This serves to two main purposes. First, it configures the visibility of the specific dimension layers, so that the user can toggle them to avoid excessive visual clutter, thus following the details-on-demand principle. Second, it configures how visual mapping works. For instance, setting the threshold of the minimum frequency for a node or edge to be displayed. The visualization configuration changes through the interaction of the user with `Tiramisù` visualizations.

4 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 (Munoz-Gama et al, 2022) and personal information management (Catarci et al, 2007). 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 of Di Federico and Burattin (2023b), is described in Section 4.1 and illustrates the application of our solution with a particular focus on the spatial dimension. The latter, inspired by the investigation of Beerepoot et al (2023a), is discussed in Section 4.2 and involves the temporal dimension.

4.1 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 (Fernandez-Llatas et al, 2015) and in the field of VA (Caban and Gotz, 2015). Indeed, the interplay of the two research fields in this context has the potential to yield remarkable outcomes, as shown by Dixit et al (2017). 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 neurodegenerative diseases. Naively applying 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 (Di Federico and Burattin, 2023b). While the above models may be sufficient for a process mining expert to extrapolate meaningful process insights, it can quickly become challenging for target users, who may lack training in interpreting these formalisms: 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 (Di Federico and Burattin, 2023b; Di Federico et al, 2021).

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, similarly to the visualization approach adopted by Dogan et al (2019) for sensor mapping. On top of it, we can project an overlay dimension layer that represents the sleeping routine followed by the patient in a way that more closely matches what nurses are familiar with, i.e., the environment that the patient interacts with. This can be seen in Fig. 3, which is a visualization of the process (in this case, the sleeping routine) followed by one patient over several nights. In this scenario, each case is composed of the activities done by the patient in one single night.

The system is currently implemented as a Javascript application, and it relies on the Vue.js[1] framework. The system takes as input the URL of a JSON configuration file and a DFG file describing the process. The configuration file allows the specification of the backdrop and where the representations of the different spaces should be on top of the backdrop. The system expects a correspondence between the location names and the activities of the process. The configuration must also provide information regarding

[1] See https://vuejs.org/.

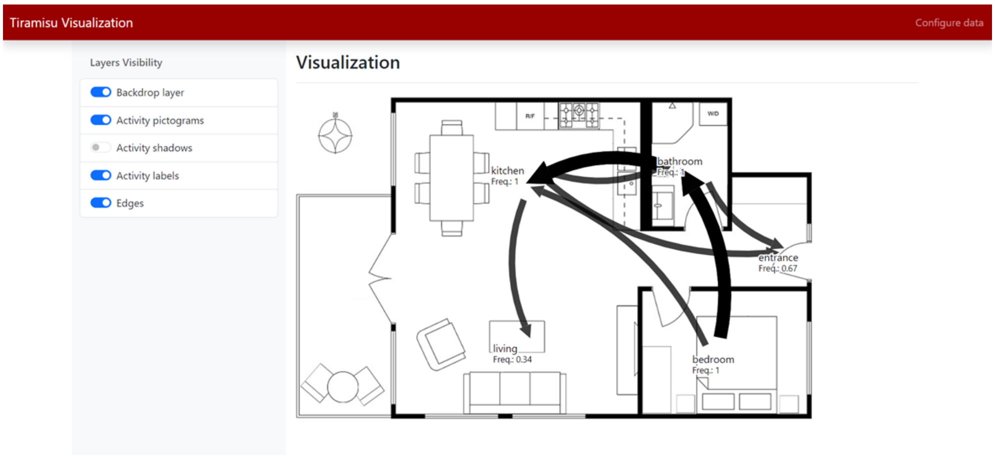

Fig. 3 The backdrop used for Scenario 1, with the floor map where patients live, overlaid with the sleeping routine of a patient.

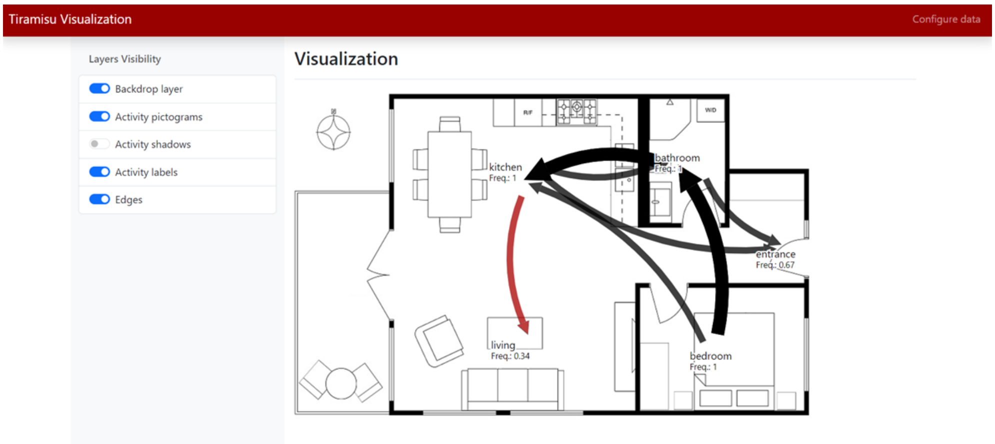

Fig. 4 The backdrop overlaid with a deviation from the typical routine, highlighted in red.

the activities' details, that have to be loaded when an activity is clicked on. The source code, as well as deployment and some example data, are publicly available[2].

As discussed in Section 3, the backdrop serves as a common context onto which it is possible to project process-related information. Figure 3 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 activities. This approach brings visualization

[2]The source code is available at https://github.com/tiramisuframework/healthcare, and a deployment is available at https://tiramisuframework.github.io/healthcare.

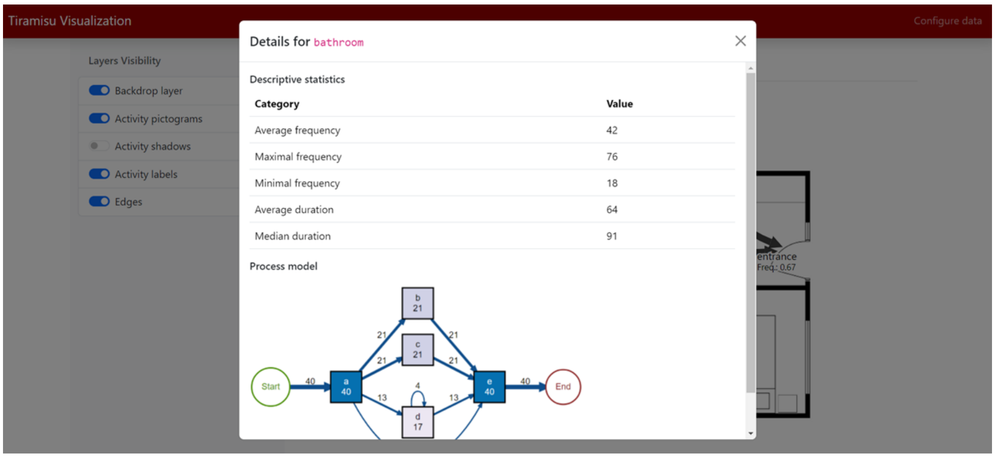

Fig. 5 An illustration of an interaction to dive deeper into the details of a deviation.

closer to what the analysts (i.e., nurses) 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 in Fig. 4 would be specifically designed in a manner that allows nurses 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. 5, by interacting with the visualization, nurses will have the opportunity to select specific activities 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 (Shneiderman, 1996), such pop-up will report details regarding the underlying process execution, such as the actual Workflow net model (van der Aalst, 1998) possibly with colors suggesting where deviations happened, a DCR model (Hildebrandt and Mukkamala, 2011), and statistical aspects such as the average duration. This, in turn, provides a possible solution for **US3**.

4.2 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 information to

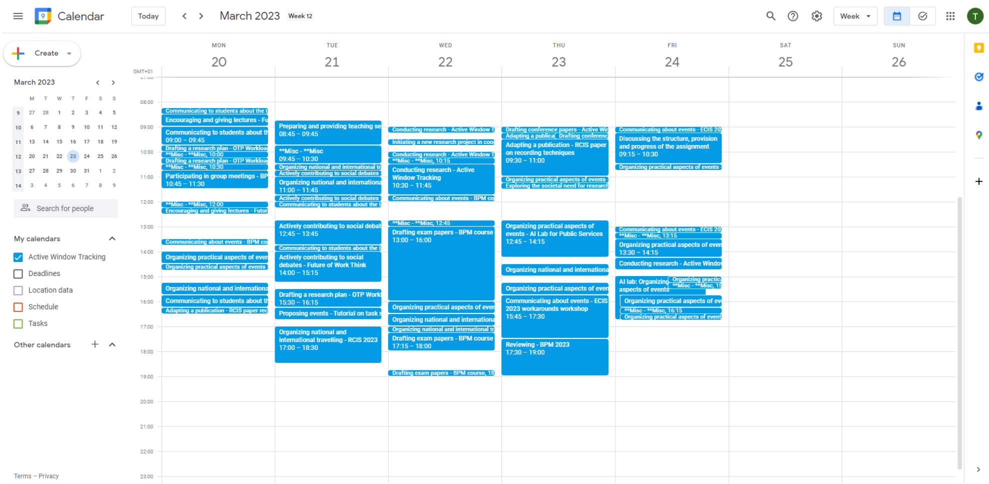

Fig. 6 The backdrop of Scenario 2 overlaid with the weekly summary of the activities executed by the researcher.

pursue the objective of aiding people to collect and reflect on their personal information (Li et al, 2010). 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 (Sinik et al, 2023). 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 (Bertrand et al, 2022).

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 lectures, grading students or reviewing research papers, among many other activities. Specifically, our focus is on the retrospective analysis of the personal work processes a person has followed during a certain period of time. 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 research papers;
- **US5:** As a researcher, I want to shed light on deadlines that are relevant in the context of my paper writing;
- **US6:** As a researcher, I want to dive deeper into my work on the days preceding the deadlines.

To this end, we use a calendar as the backdrop of the visual representation. Using Google Calendar, we project information about the researcher's work practices onto the calendar backdrop (see Fig. 6). In this example, four calendars were imported to 'my calendars', which represent four different overlay dimension layers. One of the

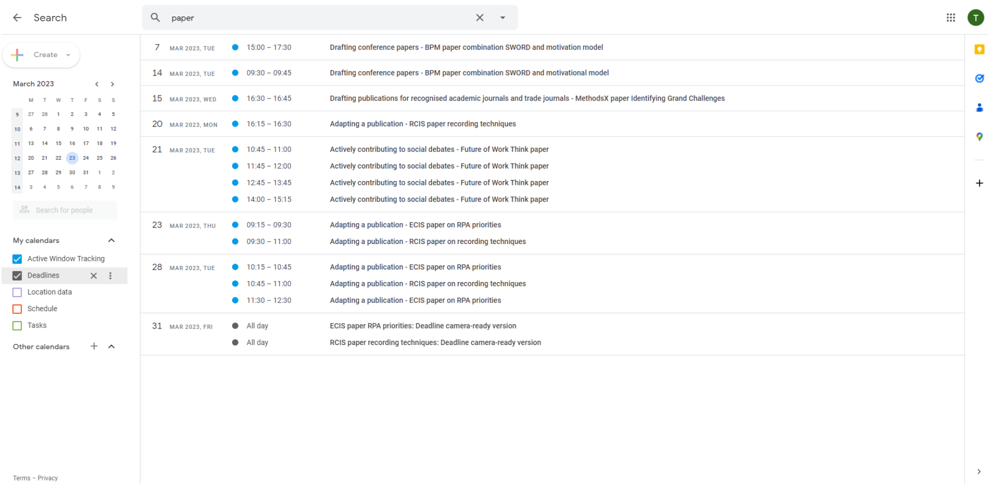

Fig. 7 Usage of the calendar's search functionality to provide an overview of behavior related to a particular (type of) activity or case, in this case the work on research papers. In addition to the illustration of Active Window Tracking data, the 'deadlines' calendar is turned on to study relevant deadlines.

dimension layers uses process data collected as Active Window Tracking data (Beerepoot et al, 2023a). This is actual data extracted from a researcher's computer behavior and is abstracted to 15-minute time slots, i.e., we assign each 15-minute time slot to the activity performed a longer time in that period as long as it is above a certain threshold. Otherwise, the 15-minute time slot is assigned to a generic "misc" activity or to no activity if inactive time of the user in the time slot is above another threshold. We do this because without abstraction the duration of many activities could be around two or three minutes, which would make the calendar too cluttered. The source code of the abstraction mechanism is publicly available[3].

The activities are represented as boxes, titled according to the following format: [activity name] - [case name]. The position and size of each box represent when the activity started and its duration, respectively. Additional calendars serve as additional dimension layers that can be turned on and off according to specific needs. In Fig. 6, only the Active Window Tracking data of the researcher, i.e., the actual computer behavior, is turned on.

The calendar view serves as a starting point for the researcher to analyze their behavior. In order to address **US4**, the search functionality in the calendar can be utilized. Figure 7 illustrates how searching for 'paper', enables the analyst to be presented with an overview of all relevant behavior related to the writing of research papers. It shows the days and time slots in which the researcher worked on five unique research papers. To cater for **US5**, the 'deadlines' calendar is turned on to show the relevant deadlines in the same view. Two deadlines were relevant in the context of the writing of the papers, both set on March 31st.

[3]The source code is available at https://github.com/tiramisuframework/tiramisu-calendar.

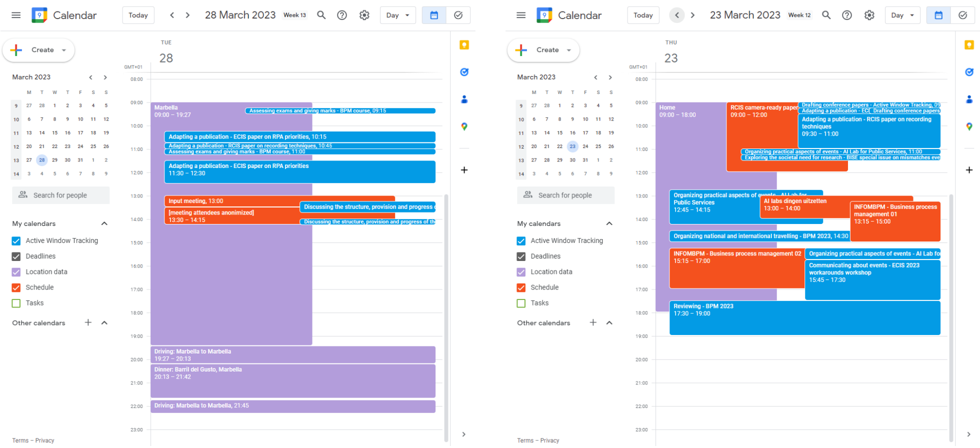

Fig. 8 Examining the context of work in the days before the deadline.

In order to dive deeper into the behavior on the days preceding the deadlines (i.e., **US6**), the calendar's daily view provides the needed scope of information. In Fig. 8, location data from Google Maps is included as well as the researcher's calendar data from Microsoft Outlook. This shows how the researcher was working on the two research papers from her holiday location in Marbella, Spain (left part of the figure). The time blocked for the RCIS paper on Thursday morning (right, in orange) reveals that the plan was to be finished on that day, but to no avail. Following again the details-on-demand principle (Shneiderman, 1996). The user accesses more detailed information by clicking on that activity. As it turns out, the submission system was still closed on that day. A similar situation was found on Friday, where a reminder in the morning served to check whether the paper was ready to be submitted, but it was not, ultimately resulting in some holiday work.

Similarly to Section 4.1, 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.

5 Expert Evaluation

We evaluated the `Tiramisù` framework through semi-structured interviews with potential end users (in accordance with FP2), adopting a walkthrough technique. The panel of experts (Table 1) included a variety of roles and age range.

For our evaluation, we followed the same structure and context as our presented use cases (see Section 4). Concerning the behavioral deviation analysis in healthcare (see Section 4.1), we interviewed three nurses active in different areas: a nursing home, a hospital, and a home care setting. For the second scenario, personal productivity analysis in work processes (see Section 4.2), we interviewed three researchers at various career levels and scientific disciplines. For the latter, we specifically selected

Table 1 Panel of experts

Reference	Description	Age range	Interview date
R1	Associate professor in psychology	50-60	04/03/2024
N1	Recently-retired caregiver nursing home	60-70	06/03/2024
R2	Senior PhD candidate veterinary science	20-30	06/03/2024
N2	Centre head cardiology and neurology	40-50	08/03/2024
N3	District nurse	30-40	08/03/2024
R3	Postdoctoral researcher sociology	30-40	12/03/2024

researchers without a process (mining) background (in line with FP3). The aim of the interviews was to determine (1) whether the presentation of information is believed useful and intuitive by the experts, and (2) whether the user stories are relevant according to them. Therefore, we systematically walked each expert through the scenario, encouraging them to reflect and comment on the user stories and the presentation of information. The interviews were recorded and transcribed, before reflecting on the results, common themes, and additional user stories mentioned by the experts.

5.1 Behavioral deviation analysis in healthcare

Across the three nursing experts, the user stories were well-understood and the presentation of information for this particular setting was found intuitive. Notably, while explaining Fig. 3, N1 was already finishing our sentences, e.g., about the thickness of edges representing the number of times a patient moved from one room to another. N1 also recognized the use of the described data sources, as the care home in which she worked made use of similar sensors that could be attached to strollers and notified the nurses when a patient left the building. Similarly, they made use of smart mats that were placed next to beds, that notified nurses when patients left their beds. N1 saw particular value in visualizing such multi-faceted process information for patients with dementia: “Otherwise, you often don’t know what someone is doing in the meantime. You close the door and what does someone do? How do they do it in that apartment? Of course, you don’t see that. So that’s obviously a good idea to map that out there for people who are demented. (...) Especially at night, for example. Then I think you have much less insight into what is happening.”

Interestingly, N2 shared that she recognized the value of the presented information for home care and care home settings, but less so for a hospital setting. Although there is currently a major campaign ongoing in the hospital about the importance of moving around and exercising (which includes the designation of walking routes and other initiatives), the initiatives are mostly aimed at detecting and acting on single events such as patients falling or walking away. For example, they make use of detection poles that trigger an alarm as soon as they are passed by the patient. Patients are closely monitored in the hospital already - there is a specific medical reason for them to be there. They are typically in for a short period of time, and some of them are bed-ridden. As such, there is less interest in visualizing physical routines and deviations thereof. If a patient is not progressing or even deteriorating, their therapists will quickly notice while they are doing their assessments.

Just like N1 and N2 recognized the use of movement sensors to detect patients falling or walking away, N3 was also familiar with this type of control software. In her district nursing team, they had had discussions about potential use of smart monitoring devices, particularly for people with cognitive problems, but they had not (yet) implemented it. One obstacle that both N3 and N1 talked about concerns privacy issues. Although it is interesting to gather insight into the routines of people with cognitive problems such as dementia, they are also often the ones who are suspicious, especially those in the early stages: "They realize that they forget things and also disguise things, so to speak. They are not eager for everyone to keep an eye on them." If those issues can be solved, however, she believes it would be of added value. Similar to N1, N3 mentioned that as a (district) nurse, you only see the patient daily or a few times a week, meaning one is provided with a snapshot of the patient's behavior. One is not aware of the rhythm of a patient throughout the day, while this gives a lot of information.

When discussing the user stories, N3 mentioned that this mirrors how they are taught to think. Clinical reasoning is about determining problems and getting to the core of the problem: "For example, what is the reason that person no longer goes outside? So, what's behind that? Is it fear or is it mobility? Or is it not a need? That's what you're trying to figure out.". Whether it is a patient having a wound or that is incontinent, they reason from the whole towards the cause of the problem.

In terms of additional user stories, N1 gave an interesting example of wanting to know about a patient's change in behavior over time. It was about an old lady who could not voice her thoughts very well and was bullied by other residents in the care home. N1 only found out when she was sitting next to the lady and observed strange interactions between her and another lady. She could imagine that it would be valuable to monitor movements to the different social areas over time. For example, if one would go to the restaurant less and less, this could trigger further investigation of why that is the case. As such, we define the following additional uses stories:

- **US7:** As a nurse, I want to spot the presence of a patient's change in behavior over time;
- **US8:** As a nurse, I want to dive deeper into the context of a behavioral change;

5.2 Personal productivity analysis in work processes

As all three researchers are familiar with email clients such as Microsoft Outlook and Google Calendar, the presentation of information on the calendar backdrop is found intuitive by each of them. Turning calendars on and off, and clicking on an activity to pull up additional details, is also quickly grasped by the participants. Interesting to note is that the participants would use the information for different ends. For example, R1 was interested in the presence of breaks and distractions, and how this affects work processes. A related aspect that she mentioned was around the perception of having a very busy schedule: "What I hear very often around me is that people have the feeling that they have been very busy all day and have been busy with all kinds of things and think at the end of the day: but I haven't made any progress". A system such as the one presented would provide insights into this lack of progress and possible causes for this such as distractions and a resulting lack of concentration. She also mentioned an interest in insights into the distribution of time between education-related and research-related tasks.

Where R1 was interested in the distribution of time between teaching and other tasks, R2 described a particular interest in insights into specific teaching processes. Not so much at the level of individual lectures or other teaching activities, but looking at the overall coordination efforts involved in a course. Generally, she would benefit from more awareness about the amount of time she spends on tasks and the ability to anticipate on work, in order to make more rational choices: "Because with students especially, I'm way too nice. I spend way too much time on it actually and I want to, but I think you can also do that smarter, so that you don't necessarily detract from your student relationship, but that you can organize it a bit more conveniently for yourself". In addition, she believed these insights could potentially lead to better personal resignation. She described the feeling of guilt she sometimes experiences when exercising during work hours. Getting insight into the work she has done and the progress she has made would serve as a reminder that she is in fact doing enough.

R3 confirmed the interest in awareness around teaching duties, distractions, and the perception of feeling busy and unproductive. In addition, she talked about the wish to include more contextual information in analysis of tasks. For example, she is currently working on a task that she typically does toward the end of the day: sometimes starting at 2, sometimes much later, or anywhere in between. Sometimes working from home, other times working at the office. She mentioned an interest in knowing what influence these time and location aspects have on productivity. The way the information is presented in the calendar allows her to investigate this.

One interest that all researchers mentioned is the ability to compare planned work with actual behavior and determining the reasons for discrepancies between the two. R1: "I don't know if the system is so smart that it can already indicate from which day you that you really took the wrong turn, say that you should have worked on project A because the deadline for project A is imminent. But then you start working on project B, so that's been kind of a critical day". Naturally, a requirement is that the researcher would need to register their plans, either in the calendar or elsewhere. Similar to the researcher's data used in the scenario, R2 and R3 recognize using a calendar to block time for tasks. R2 talked about scheduling tasks up front and shifting them forward to another day if the task could not be finished, effectively acting as a work organization device and providing an opportunity to include the following user stories to our existing set:

- **US9:** As a researcher, I want to investigate discrepancies between scheduled work and actual work;
- **US10:** As a researcher, I want to investigate interruptions in behavior;

5.3 `Tiramisù` for Spatial and Temporal Analysis

Based on our evaluation of the two detailed scenarios, we are able to derive general insights about the use of `Tiramisù` for the analysis of multi-faceted process information. In Table 2, we provide an overview of situations in which we believe a spatial or temporal backdrop may be especially valuable, and which challenges still need to be addressed.

Table 2 Expected value of and challenges for analysis using spatial and temporal backdrops

	Spatial backdrop	Temporal backdrop
Value	In settings with heavy emphasis on movements, repeated across time or cases	For knowledge workers engaged in different types of tasks and with autonomy over their work
Challenge	Dealing with privacy issues around location data	Combining different data sources (in a secure way)

Based on the evaluation of the behavioral deviation scenario with nurses we conclude that the spatial backdrop will be most valuable in settings that are movement-heavy, and where these movements take place repetitively across days or cases. Home care or nursing home settings are great examples of this, whereas the movements in a hospital may be too restricted and too short-lived to make a floormap the best fitting backdrop available. Other settings with repetitive movements where a floor map may be of relevance are bakeries or restaurant kitchens, where people move to different stations during their shifts. One step larger on the scale could be factories or farms, where workers move across larger areas (fields, pastures, storage buildings, vehicle depots, etc.). On a regional level, a map could be relevant to study the processes of package deliverers or taxi drivers. One challenge that we identify here is dealing with the privacy of workers when using location data. Workers may not be very fond of their location being tracked and analyzed.

Based on the evaluation of the personal productivity scenario we conclude that the temporal backdrop is especially relevant when studying the processes of knowledge workers involved in a variety of tasks and who have autonomy over their work. A temporal backdrop such as a calendar allows for in-depth analyses of their distribution of time across tasks and the presence of breaks and distractions. Knowledge workers for whom such analyses may be of value include, for example, consultants, accountants, developers, designers, and writers. They spend a significant amount of time on a computer, engage in meetings, and schedule their work. Projecting recorded data of their work practices on a calendar backdrop allows for analyses of work towards deadlines, discrepancies between scheduled work and actual work, and the influence of different locations and interruptions on their productivity. To do so, however, one would need to combine data from different sources, such as computer interaction logs, location history, and calendar data. This would need to be recorded and stored in a secure manner.

6 Discussion

`Tiramisù` allows the representation of multi-faceted information while providing at the same time a common domain-specific context, thus enabling users to navigate through distinct, and possibly domain-specific, dimensions. In addition, `Tiramisù` can be used to depict both process-level visualizations like most regular process mining tools, but also case- and event-level visualizations, which are especially relevant in knowledge-intensive processes where understanding the particularities of single cases play a more prominent role (Di Ciccio et al, 2015). An example of process-level visualization is depicted in Fig. 3 where the process is the sleeping routine of one patient during several nights, and each case is composed of the activities done by the patient in one single night. Case-level visualizations can be found in Fig. 4, where the deviation of a single case is visualized. Finally, event-level visualizations are used in the second use case, which depicts event information from multiple cases (e.g., reviews for different conferences) without aggregating them. All these features make `Tiramisù` particularly suitable for making sense of the inherently complex search space of knowledge-intensive processes, as we have shown in our evaluation. Concerning more classical process mining scenarios (i.e., other than those related to knowledge-intensive processes), we believe that `Tiramisù` can be useful to complement other process mining operations especially in settings where domain-specific information plays a significant role. Exploring the role of `Tiramisù` in these settings is an interesting avenue for future work.

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 (Dumas et al, 2018). Our framework is, in this sense, no different. The injection of additional facts stemming from target users is complementary to this investigation and paves the path for future work.

The design of the visualizations on the backdrop, and on the other layers is a challenge of its own, and should be tackled by visualization experts in close cooperation with 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 (2018), which presents a framework aimed to guide design choices for the effective visualization of analytical data, and to the paper by Munzner (2009), 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 (van der Aalst, 2016). 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 (Di Ciccio et al, 2015). 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 (Baier et al, 2014; Bayomie et al, 2023). 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` (Wynn et al, 2021).

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 (Di Ciccio et al, 2023).

7 Conclusion and Future Work

In this paper, we have presented `Tiramisù`, a conceptual 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 achieve this by organizing the visualizations of the different dimensions around a backdrop that provides the user with a domain-specific representation of the context of the process that facilitates the interpretation of process-oriented visualizations. At the same time, it provides the user with a familiar context, which makes `Tiramisù` particularly useful for users who are not process mining experts.

We have demonstrated our solution through its prototypical implementation in two scenarios illustrating its suitability in the context of knowledge-intensive processes. Also, we have evaluated `Tiramisù` through semi-structured interviews with potential end users adopting a walkthrough technique. For both scenarios, the way the process information was presented was found intuitive and the user stories were recognizable to the participants. In some settings, though, this information is believed to be more relevant (e.g., in care homes) than in others (e.g., hospitals). The evaluation also showed how target users may use the information for new ends, and based on those discussions, we defined four additional user stories. Regardless of the specific user story of interest, a major obstacle is the availability of the needed multi-faceted data.

We foresee several research endeavors originating from our current idea. First and foremost, we have presented a conceptual framework. We tested its applicability in two challenging scenarios, but the prototypical implementations specifically tailored for the use cases under analysis. Building a `Tiramisù` system in the form of a general-purpose software service for VA-powered process mining is an exciting piece of work we foresee as a clear step ahead in our plans. To reach this goal, establishing a set of development guidelines that are cross-domain is a key intermediate milestone. The guidelines could generalize and build upon the insights we gathered with the present research effort. The customization and refactoring of layers, their aggregation rules, and backdrops is also a key challenge that needs to be addressed in the future. It would also be interesting to analyze the integration of filtering mechanisms with the elements of `Tiramisù` to facilitate data exploration. Furthermore, 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.

8 Statements and Declarations

8.1 Ethical approval

We declare that this submission follows the policies as outlined in the Guide for Authors. The interviewed participants provided consent for recording the interview and sharing their insights. As for data used in the two scenarios, we hereby declare the following. The healthcare scenario described in Section 4.1 is based on solely synthetic data. The personal informatics scenario in Section 4.2 contains real-world data extrapolated from user interaction logs and registered locations from one of the authors involved. This author provided written consent to make the insights and figures available for the purpose of this study, after masking sensitive information.

8.2 Availability of supporting data

The source code of our software prototypes can be accessed here: https://github.com/tiramisuframework. A running instance of the tool showcasing the healthcare scenario (Section 4.1) is available here: https://tiramisuframework.github.io/healthcare/. The data used for the illustration of the personal informatics scenario (Section 4.2) are subject to restrictions as they were used under licence for the current study and are thus not publicly available.

8.3 Competing interests

Claudio Di Ciccio is a co-author of this manuscript and a member of the Editorial Board of the Journal of Intelligent Information Systems. Therefore, he did neither assume responsibility for nor oversaw the peer review process. The authors have no other competing interests as defined by Springer, or further interests that might be perceived to influence the results and/or discussion reported in this paper.

8.4 Funding

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 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.

8.5 Authors' contributions

The authors contributed equally to this work.

8.6 Acknowledgements

This work started at Schloss Dagstuhl (Leibniz-Zentrum für Informatik), seminar 23271 "Human in the (Process) Mines" (Di Ciccio et al, 2023). We thank the institute and the organizers for this precious opportunity and inspiring experience. We also thank the nurses and researchers who participated in the evaluation for their valuable insights.

References

Alman A, Di Ciccio C, Maggi FM, et al (2021) RuM: Declarative process mining, distilled. In: BPM. Springer, pp 23–29, https://doi.org/10.1007/978-3-030-85469-0_3

Alman A, Arleo A, Beerepoot I, et al (2023) Tiramisù: A recipe for visual sense-making of multi-faceted process information. In: Process Mining Workshops - ICPM 2023 International Workshops. Springer, pp 19–31, https://doi.org/10.1007/978-3-031-56107-8_2

Baier T, Mendling J, Weske M (2014) Bridging abstraction layers in process mining. Inf Syst 46:123–139. https://doi.org/10.1016/j.is.2014.04.004

Bayomie D, Di 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

Beerepoot I, Barenholz D, Beekhuis S, et al (2023a) A window of opportunity: Active window tracking for mining work practices. In: ICPM. IEEE, pp 57–64

Beerepoot I, et al (2023b) The biggest business process management problems to solve before we die. Comput Ind 146:103837. https://doi.org/10.1016/j.compind.2022.103837

Bertrand Y, Van den Abbeele B, Veneruso S, et al (2022) A survey on the application of process mining to smart spaces data. In: ICPM Workshops. Springer, pp 57–70, https://doi.org/10.1007/978-3-031-27815-0_5

Caban JJ, Gotz D (2015) Visual analytics in healthcare - opportunities and research challenges. J Am Medical Informatics Assoc 22(2):260–262. https://doi.org/10.1093/jamia/ocv006

Catarci T, Dix A, Katifori A, et al (2007) Task-centred information management. In: DELOS Conference, vol 4877. Springer, pp 197–206, https://doi.org/10.1007/978-3-540-77088-6_19

De Weerdt J, Schupp A, Vanderloock A, et al (2013) Process mining for the multi-faceted analysis of business processes - A case study in a financial services organization. Comput Ind 64(1):57–67. https://doi.org/10.1016/j.compind.2012.09.010

Di Ciccio C, Marrella A, Russo A (2015) Knowledge-intensive processes: Characteristics, requirements and analysis of contemporary approaches. J Data Semant 4(1):29–57. https://doi.org/10.1007/s13740-014-0038-4

Di Ciccio C, Miksch S, Soffer P, et al (eds) (2023) Human in the (Process) Mines (Dagstuhl Seminar 23271), Dagstuhl Reports, vol 13, https://doi.org/10.4230/DAGREP.13.7.1

Di Federico G, Burattin A (2023a) CvAMoS - Event Abstraction Using Contextual Information. Future Internet 15(3):113. https://doi.org/10.3390/FI15030113, URL https://doi.org/10.3390/fi15030113

Di Federico G, Burattin A (2023b) Do you behave always the same? In: Process Mining Workshops, pp 5–17, https://doi.org/10.1007/978-3-031-27815-0_1

Di Federico G, Burattin A, Montali M (2021) Human behavior as a process model: Which language to use? In: ITBPM@BPM. CEUR-WS.org, pp 18–25, URL https://ceur-ws.org/Vol-2952/paper_293a.pdf

Dixit PM, Caballero HSG, Corvò A, et al (2017) Enabling interactive process analysis with process mining and visual analytics. In: HEALTHINF. SciTePress, pp 573–584, https://doi.org/10.5220/0006272605730584

Dogan O, Martinez-Millana A, Rojas E, et al (2019) Individual behavior modeling with sensors using process mining. Electronics 8(7). https://doi.org/10.3390/electronics8070766, URL https://www.mdpi.com/2079-9292/8/7/766

Dumas M, La Rosa M, Mendling J, et al (2018) Fundamentals of Business Process Management, Second Edition. Springer, https://doi.org/10.1007/978-3-662-56509-4

Fernandez-Llatas C, Lizondo A, Monton E, et al (2015) Process mining methodology for health process tracking using real-time indoor location systems. Sensors 15(12):29821–29840. https://doi.org/10.3390/s151229769

Filipov V, Arleo A, Miksch S (2023) Are we there yet? a roadmap of network visualization from surveys to task taxonomies. Computer Graphics Forum 42(6):e14794. https://doi.org/10.1111/cgf.14794

Gschwandtner T (2015) Visual analytics meets process mining: Challenges and opportunities. In: Ceravolo P, Rinderle-Ma S (eds) Data-Driven Process Discovery and Analysis - 5th IFIP WG 2.6 International Symposium, SIMPDA 2015, Revised Selected Papers, Lecture Notes in Business Information Processing, vol 244. Springer, pp 142–154, https://doi.org/10.1007/978-3-319-53435-0_7

Hildebrandt TT, Mukkamala RR (2011) Declarative event-based workflow as distributed dynamic condition response graphs. arXiv preprint arXiv:11104161 https://doi.org/10.48550/arXiv.1110.4161

Kaouni A, Theodoropoulou G, Bousdekis A, et al (2021) Visual analytics in process mining for supporting business process improvement. In: NiDS, pp 166–175, https://doi.org/10.3233/FAIA210089

Keim D, Andrienko G, Fekete JD, et al (2008) Visual analytics: Definition, process, and challenges. In: Information Visualization. Springer, pp 154–175, https://doi.org/10.1007/978-3-540-70956-5_7

Kriglstein S, Pohl M, Rinderle-Ma S, et al (2016) Visual Analytics in Process Mining: Classification of Process Mining Techniques. In: Andrienko N, Sedlmair M (eds) EuroVis Workshop on Visual Analytics (EuroVA). The Eurographics Association, https://doi.org/10.2312/eurova.20161123

Kubrak K, Milani F, Nolte A (2023) A visual approach to support process analysts in working with process improvement opportunities. Bus Process Manag J 29(8):101–132. https://doi.org/10.1108/BPMJ-10-2021-0631

Li I, Dey AK, Forlizzi J (2010) A stage-based model of personal informatics systems. In: CHI. ACM, pp 557–566, https://doi.org/10.1145/1753326.1753409

Mannhardt F, de Leoni M, Reijers HA (2015) The multi-perspective process explorer. In: BPM (Demos), pp 130–134, URL https://ceur-ws.org/Vol-1418/paper27.pdf

Miksch S (2021) Visual analytics meets process mining: Challenges and opportunities. In: ICPM, pp xiv–xiv, https://doi.org/10.1109/ICPM53251.2021.9576854

Miksch S, Aigner W (2014) A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data. Computers & Graphics 38:286–290. https://doi.org/10.1016/j.cag.2013.11.002

Munoz-Gama J, Martin N, et al (2022) Process mining for healthcare: Characteristics and challenges. J Biomed Informatics 127:103994. https://doi.org/10.1016/j.jbi.2022.103994

Munzner T (2009) A nested process model for visualization design and validation. IEEE Trans Vis Comput Graph 15(6):921–928. https://doi.org/10.1109/TVCG.2009.111

Raidou RG (2019) Visual analytics for the representation, exploration, and analysis of high-dimensional, multi-faceted medical data. Biomedical Visualisation: Volume 2 pp 137–162. https://doi.org/10.1007/978-3-030-14227-8_10

Rehse JR, Pufahl L, Grohs M, et al (2022) Process mining meets visual analytics: the case of conformance checking. arXiv preprint URL arXiv:2209.09712

Rozanski N, Woods E (2011) Software Systems Architecture. Addison Wesley

Schuster D, van Zelst SJ, van der Aalst WMP (2023) Cortado: A dedicated process mining tool for interactive process discovery. SoftwareX 22:101373. https://doi.org/10.1016/j.softx.2023.101373

Schuster D, Zerbato F, van Zelst SJ, et al (2024) Defining and visualizing process execution variants from partially ordered event data. Information Sciences 657:119958. https://doi.org/10.1016/j.ins.2023.119958

Shneiderman B (1996) The eyes have it: A task by data type taxonomy for information visualizations. In: VL, IEEE, pp 336–343, https://doi.org/10.1109/VL.1996.545307

Shneiderman B (2003) The eyes have it: A task by data type taxonomy for information visualizations. In: The craft of information visualization. Elsevier, p 364–371, https://doi.org/10.1016/B978-155860915-0/50046-9

Sinik T, Beerepoot I, Reijers HA (2023) A peek into the working day: Comparing techniques for recording employee behaviour. In: RCIS. Springer, pp 343–359, https://doi.org/10.1007/978-3-031-33080-3_21

Sirgmets M, Milani F, Nolte A, et al (2018) Designing process diagrams - A framework for making design choices when visualizing process mining outputs. In: OTM. Springer, pp 463–480, https://doi.org/10.1007/978-3-030-02610-3_26

Soffer P, Hinze A, Koschmider A, et al (2019) From event streams to process models and back: Challenges and opportunities. Inf Syst 81:181–200. https://doi.org/10.1016/J.IS.2017.11.002, URL https://doi.org/10.1016/j.is.2017.11.002

van der Aalst WMP (1998) The Application of Petri Nets to Workflow Management. Journal of Circuits, Systems, and Computers 8(1):21–66. https://doi.org/10.1142/S0218126698000043

van der Aalst WMP (2016) Process Mining - Data Science in Action, Second Edition. Springer, https://doi.org/10.1007/978-3-662-49851-4

van der Aalst WMP (2022) Foundations of process discovery. In: van der Aalst and Carmona (2022), p 37–75, https://doi.org/10.1007/978-3-031-08848-3_2

van der Aalst WMP, Carmona J (eds) (2022) Process Mining Handbook, vol 448. Springer, https://doi.org/10.1007/978-3-031-08848-3

van der Aalst WMP, van Dongen BF, Günther CW, et al (2009) ProM: The process mining toolkit. In: BPM (Demos), URL http://ceur-ws.org/Vol-489/paper3.pdf

van der Aalst WMP, de Leoni M, ter Hofstede AH (2011) Process mining and visual analytics: Breathing life into business process models. BPM Center Report BPM-11-15, BPMcenter org 17:699–730. URL http://bpmcenter.org/wp-content/uploads/reports/2011/BPM-11-16.pdf

van der Linden S, de Fouw E, van den Elzen S, et al (2023) A survey of visualization techniques for comparing event sequences. Computers & Graphics 115:522–542. https://doi.org/https://doi.org/10.1016/j.cag.2023.05.016

Wynn MT, Lebherz J, et al (2021) Rethinking the input for process mining: Insights from the XES survey and workshop. In: ICPM workshops. Springer, pp 3–16, https://doi.org/10.1007/978-3-030-98581-3_1

Yeshchenko A, Mendling J (2024) A survey of approaches for event sequence analysis and visualization. Information Systems 120:102283. https://doi.org/10.1016/j.is.2023.102283