A Taste of the Different Flavors of Tiramisù

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

[1]University of Tartu, Estonia
[2]Eindhoven University of Technology, The Netherlands
[3]Utrecht University, The Netherlands
[4]Technical University of Denmark, Kgs. Lyngby, Denmark
[5]Universidad de Sevilla, Spain

Abstract

The amalgamation of process mining and visual analytics holds the promise for a fruitful new frontier of process exploration, where both fields provide crucial ingredients to the resulting delicacy. The field of process mining brings the concepts of events, traces, and process thinking in general, while the field of visual analytics brings the idea of interactive and cyclical analysis, accounting for the expectations and desires of specific domains and the individuals therein. In this paper, we focus on one such delicacy, aptly named the `Tiramisù` framework, presenting two flavors of the corresponding multi-layered recipe. More specifically, we build on our earlier work to provide refined servings of `Tiramisù` for the healthcare and the personal informatics domain, with both being taste-tested by the end users, and the latter being significantly enhanced as a result.

Keywords

Process mining, Visual analytics, Knowledge-intensive processes, Visualization

Metadata description	Value
Tool name	Tiramisù
Current version	1.0
Legal code license	Apache 2.0
Languages, tools and services used	VueJS, Python, Streamlit
Supported operating environment	Microsoft Windows, GNU/Linux, Mac
Download/Demo URL	https://tiramisuframework.github.io/healthcare/, https://tiramisu-calendar.streamlit.app
Documentation/Source code repository	https://github.com/tiramisuframework/healthcare, https://github.com/tiramisuframework/tiramisu-calendar
Screencast video	https://www.youtube.com/watch?v=WoTT85I9-sI

1. Introduction

Process mining aims at extracting information from events recorded by information systems, with the ultimate goal to provide insights into the executed processes based on the analysed

ICPM 2024 Tool Demonstration Track, October 14-18, 2024, Kongens Lyngby, Denmark
*Corresponding author.
†These authors contributed equally.
anti.alman@ut.ee (A. Alman); a.arleo@tue.nl (A. Arleo); i.m.beerepoot@uu.nl (I. Beerepoot); andbur@dtu.dk (A. Burattin); claudio.diciccio@uniroma1.i (C. D. Ciccio); resinas@us.es (M. Resinas)
0000-0002-5647-6249 (A. Alman); 0000-0003-2008-365 (A. Arleo); 0000-0002-6301-9329 (I. Beerepoot); 0000-0002-0837-0183 (A. Burattin); 0000-0001-5570-0475 (C. D. Ciccio); 0000-0002-0837-0183 (M. Resinas)
© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Pre-print copy of the manuscript published by CEUR-WS.org (available at ceur-ws.org)

data [1]. Event data may vary based on the process domain, and typically covers different dimensions (control flow, time, resources, object lifecycles). A critical, yet largely unaddressed, issue of process mining is the inability to navigate through distinct, and possibly domain-specific, dimensions [2]. This limitation reportedly hinders process improvement investigation [3], as it narrows the analysis scope to a partial view over an inherently complex search space, which is typical of knowledge-intensive processes [4, 5]. In the discipline of visual analytics, the problem of representing complex, multi-faceted phenomena has been extensively studied [6]. This is also the type of phenomena often encountered within real-life business processes.

In previous work, we thus introduced `Tiramisù` [7, 8], a novel conceptual framework leveraging visual analytics principles [6] to enhance process mining with an interactive investigation of data integrating process-specific and context-related information. `Tiramisù` is designed in a pluri-layered fashion, to equip end users with a context-aware visualization mixing classical process mining representation elements —such as workflow nets or directly-follows graphs [1]— 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).

In this demo paper, we illustrate `Tiramisù`'s core ingredients and two variations of the recipe, i.e., prototypical implementations grounding its components into usable toolkits.

2. Recipe (Architecture)

Figure 1 illustrates the layered architecture of `Tiramisù`, a framework consisting of a *backdrop* providing context, and one or more dimension layers. From a visualization standpoint, the backdrop establishes a common context for the other layers in our framework. This backdrop might serve as a spatial reference, such as a geographic location or the layout of a building (see Sect. 3.1), or as a temporal reference, like a calendar (see Sect. 3.2). In the tiramisù metaphor, the backdrop is akin to the cream that infuses all the layers of the dessert. Superimposed layers can represent both process-related and non-process-related data. We distinguish between the following two types of dimension layers: **Overlay layers** enhance the visualization by mapping data onto new layers that are placed over the backdrop; **On-demand layers** provide additional contextual information for individual elements or small groups of elements. These layers can be activated by interactions such as hovering over or clicking on elements of the backdrop or those representing the process. Layers can encapsulate the behavior of the process being studied, similar to how the coffee-soaked ladyfingers provide structure and texture to the dessert. Overlay layers explicitly anchor to the backdrop, thus contextualizing the information.

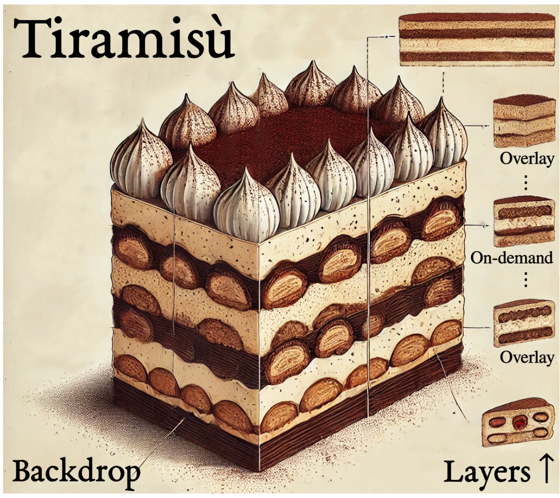

Figure 1: The `Tiramisù` architecture at large (the dessert drawings were generated with the assistance of AI)

For example, in case of a floor plan, the data on overlay layers would directly link to that floor, facilitating the user's understanding of how events are related within the same location.

3. Serving Suggestions (Functionality)

Here we describe two implementations of `Tiramisù`, focused on different classes of knowledge-intensive processes. Their names are inspired by variations of the tiramisù dessert: **Berries**, pertaining to healthcare [9]: and **Banana**, focusing on personal information management [10].

3.1. Berries (Behavioral Deviation Analysis in Healthcare)

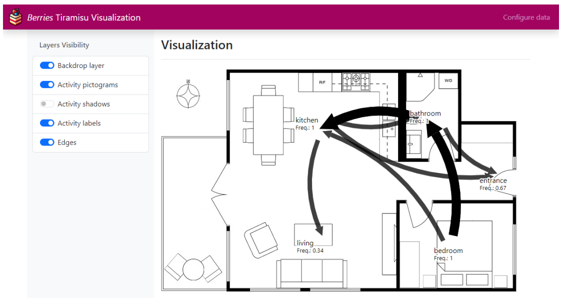

Figure 2: Screenshot of `Tiramisù` application depicting a morning routine.

The first instance of `Tiramisù` has been designed to help understanding deviations from personal behavior in the context of healthcare. In this case, in particular, we can analyze whether the daily routines of a person in a nursing home are maintained over time or not (as possible symptoms of worsening of neurodegenerative diseases). Simply applying existing process mining techniques to this task might result in basic process maps where nodes represent various patient activities and arcs indicate dependencies or temporal constraints among these activities [11].

The system , which is designed to emphasize the spatial dimension of processes and activities happening in different locations, is implemented as a JavaScript application using the Vue.js framework[1]. A screenshot of it is reported in Fig. 2. It takes a URL of a JSON configuration file and a DFG file describing the process as input. The configuration file specifies the backdrop as the floor map and the placement of different activity representations on top of it.

3.2. Banana (Personal Work Process Analysis)

The second tool we have developed following the `Tiramisù` recipe has been built to provide insights about personal work processes. For instance, for an academic, her personal work processes performed during her daily work might involve writing a paper, preparing a course, reviewing research papers, or supervising a PhD student, among many other processes. As can be observed by these examples, depending on the characteristics of the work, personal work processes can be knowledge-intensive and significantly unstructured, which means they are well-suited to benefit from the characteristics of the `Tiramisù` architecture [8].

To analyze personal work processes, the tool must be provided with collected personal information in the form of an event log that can be analyzed using process mining techniques. There are several techniques that can be used to collect personal information such as timesheet

[1]https://vuejs.org/

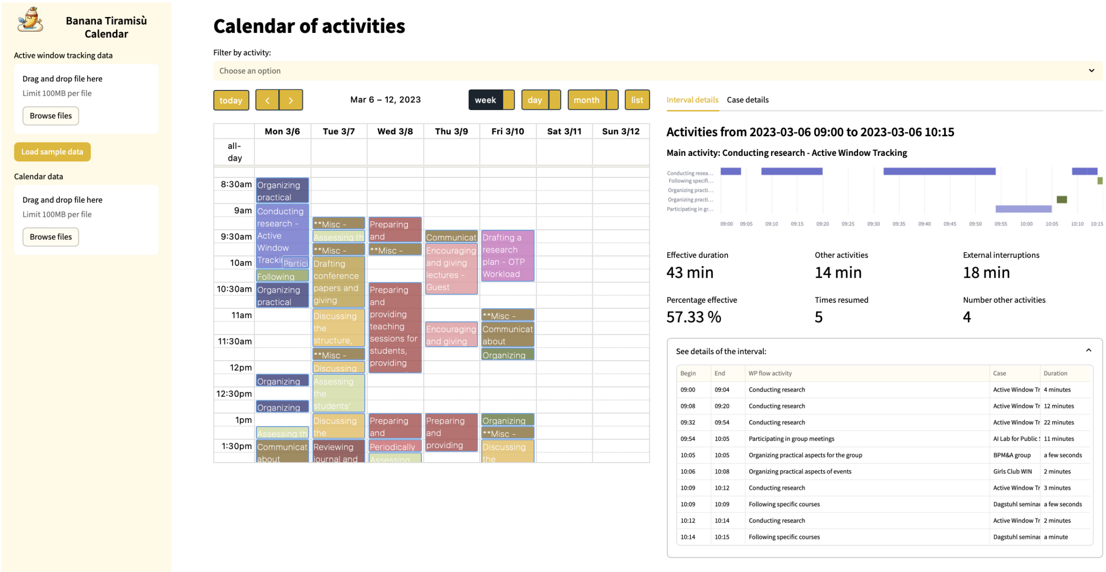

Figure 3: Screenshot of `Tiramisù` calendar depicting the Active Window Tracking data against the calendar backdrop and showing the details of an interval using the details-on-demand principle.

techniques or screen recordings, each with their own advantages and disadvantages [12]. Our tool supports two types of data, namely: Active Window Tracking data and calendar data.

Active Window Tracking data [13] can be seen as an event log that records the active window in a computer at any moment in time. Each event includes the title of the active window, the name of the corresponding app, and the timestamps when the window became active and stopped being active. There are several different tools that record this information. One of them is Tockler,[2] which is open source and stores all data locally, avoiding privacy concerns. We also assume that the user has labelled each event in the log with information about (i) the activity that was being done at that moment, e.g., conducting research, preparing lectures, etc., and (ii) the case the activity belongs to, e.g., the ICPM `Tiramisù` demo paper. The Worktagger tool[3] can be used to support the user in the labelling task. We have chosen this type of data because it has proven to be very useful to obtain insights [13, 14] from personal work processes. Calendar data can be obtained from any calendar management application and it involves the events scheduled in one's calendar including the start and end time, and the title of the event.

Since the intention of the tool is to give the user insights into a working day and how it develops, we have chosen the calendar as the backdrop. A calendar is a very intuitive visualization for every type of user, which is another reason why it is a good choice as a backdrop. Against this backdrop, both the Active Window Tracking data and the Calendar data are visualized as different layers. At this moment, the tool shows Active Window Tracking data all the time, while allows the user to show or hide Calendar data.

Figure 3 depicts a screenshot of the `Tiramisù` calendar. Active Window Tracking data is

[2] https://maygo.github.io/tockler/
[3] https://github.com/project-pivot/worktagger

visualized on the left-hand side on top of the calendar backdrop. The activities are represented by boxes, titled as: [activity name] - [case name]. The position and size of each box represents when the activity started and its duration, respectively. The red buttons above the calendar support common calendar functionalities such as navigating to different time periods and choosing the time period length (or showing all of the activities in a list view). The checkbox labelled *Include calendar data* allows the user to show or hide the calendar data layer on top of the backdrop. The Active Window Tracking data cannot be hidden, but it is possible to show only the activities selected in the *Filter by activity* drop box that appears in the top of the figure.

In many cases, for instance, when a person multitasks or when there are frequent interruptions in the environment, the duration of activities can be very short, e.g., around two or three minutes. Instead, calendars work best with a granularity of at most 15 minutes. Therefore, representing those fine-grained activities would make the calendar too cluttered, and it would negatively impact the quality of the visualization. For this reason, we abstract the activities obtained from the Active Window Tracking data to 15-minute time slots. Specifically, we assign each 15-minute time slot to the activity that has been performed for the longest time in that period as long as it is above a certain threshold. Otherwise, the 15-minute time slot is assigned to a "misc" activity or to no activity if the user has been inactive for the majority of the time slot.

Besides the overview provided by the calendar backdrop and the layers of information on top of it, the tool is designed to work following the details-on-demand principle. Specifically, when the user clicks in any of the intervals in the calendar, a full set of details and metrics regarding that interval appears, as shown in the right-hand side of Fig. 3. The information provided includes (i) the details of the activities performed in the interval that were abstracted away in the calendar; (ii) some metrics about interruptions, effectiveness, and other activities performed in the same interval, and (iii) the details of the events and windows active in the interval. Similarly, by clicking in *Case details* the user can find details of the case performed in the interval selected. For both, the interval and the case details the metrics computed are based on those proposed in [14] to analyze the interruptions during scientific research collaborations.

4. Storage Tips and Shelf Life (Maturity and Availability)

In our earlier work [8], the original implementations of the healthcare and personal information management have undergone an extensive evaluation in the form of semi-structured interviews with end users. The healthcare tool was considered intuitive and fulfilled the defined needs. As such, only cosmetic changes were made to the tool.

The Banana implementation, on the other hand, is a much evolved version of the initial proof-of-concept that was detailed in [8]. In that proof-of-concept, the Active Window Tracking data was simply pre-processed, before a calendar was generated that was then integrated into Google Calendar. The biggest limitation of that approach was that this generic tool did not allow for providing further details on the data. Instead, the tool we are presenting here is a dedicated tool that provides many details on demand, which was one of the most requested features we gathered from the evaluation of the original proof-of-concept.

Both implementations follow the same core concepts of the `Tiramisù` framework. Differences in the use cases, visual design, and the technology stack are deliberate, intended to showcase the versatility of the framework for developing process mining tools for practical use cases.

Acknowledgments

This work started at Schloss Dagstuhl (Leibniz-Zentrum für Informatik), seminar 23271 "Human in the (Process) Mines", and was partially supported by grants PID2021-126227NB-C21 funded by MCIN/AEI /10.13039/501100011033/FEDER, EU, and TED2021-131023B-C22 funded by MCIN/AEI/10.13039/501100011033 and by the European Union "NextGenerationEU"/PRTR.

References

[1] W. M. P. van der Aalst, Process mining: A 360 degree overview, in: Process Mining Handbook, Springer, 2022, pp. 3–34.
[2] I. Beerepoot, et al., The biggest business process management problems to solve before we die, Comput. Ind. 146 (2023) 103837.
[3] K. Kubrak, F. Milani, A. Nolte, A visual approach to support process analysts in working with process improvement opportunities, Bus. Process Manag. J. 29 (2023) 101–132.
[4] J. De Weerdt, A. Schupp, A. Vanderloock, B. Baesens, Process mining for the multi-faceted analysis of business processes - A case study in a financial services organization, Comput. Ind. 64 (2013) 57–67.
[5] C. Di Ciccio, A. Marrella, A. Russo, Knowledge-intensive processes: Characteristics, requirements and analysis of contemporary approaches, J. Data Semant. 4 (2015) 29–57.
[6] D. Keim, G. Andrienko, J.-D. Fekete, et al., Visual analytics: Definition, process, and challenges, in: Information Visualization, Springer, 2008, pp. 154–175.
[7] A. Alman, A. Arleo, I. Beerepoot, A. Burattin, C. D. Ciccio, M. Resinas, Tiramisù: A recipe for visual sensemaking of multi-faceted process information, in: Process Mining Workshops - ICPM 2023 International Workshops, Springer, 2023, pp. 19–31.
[8] A. Alman, A. Arleo, I. Beerepoot, A. Burattin, C. Di Ciccio, M. Resinas, Tiramisù: making sense of multi-faceted process information through time and space, Journal of Intelligent Information Systems (2024) 1–21.
[9] J. Munoz-Gama, N. Martin, et al., Process mining for healthcare: Characteristics and challenges, J. Biomed. Informatics 127 (2022) 103994.
[10] T. Catarci, A. Dix, A. Katifori, G. Lepouras, A. Poggi, Task-centred information management, in: DELOS Conference, volume 4877, Springer, 2007, pp. 197–206.
[11] G. Di Federico, A. Burattin, Do you behave always the same?, in: ICPM Workshops Proceedings, Springer, 2023, pp. 5–17.
[12] T. Sinik, I. Beerepoot, H. A. Reijers, A peek into the working day: Comparing techniques for recording employee behaviour, in: RCIS, Springer, 2023, pp. 343–359.
[13] I. Beerepoot, et al., A window of opportunity: Active window tracking for mining work practices, in: ICPM, IEEE, 2023, pp. 57–64.
[14] I. Beerepoot, A. del Río-Ortega, M. Resinas, H. A. Reijers, Interruptibility during scientific

research collaborations: The effect of pressure, proximity and quiet time, in: CSCW Companion '24, 2024, p. To appear. doi:10.1145/3678884.3681848.

This document is a pre-print copy of the manuscript
(Alman et al. 2024)
published by CEUR-WS.org (available at ceur-ws.org).

References

Alman, Anti, Alessio Arleo, Iris Beerepoot, Andrea Burattin, Claudio Di Ciccio, and Manuel Resinas (2024). “A Taste of the Different Flavors of Tiramisù”. In: *ICPM Doctoral Consortium / Demo*. Ed. by Jochen De Weerdt, Giovanni Meroni, Han van der Aa, and Karolin Winter. Vol. 3783. CEUR Workshop Proceedings. CEUR-WS.org, pp. 154–159. URL: https://ceur-ws.org/Vol-3783/paper_357.pdf.

BibTeX

```
@InProceedings{   Alman.etal/ICPM-D2024:TasteDifferentFlavorsTiramisu,
  author        = {Alman, Anti and Arleo, Alessio and Beerepoot, Iris and
                  Burattin, Andrea and Di Ciccio, Claudio and Resinas,
                  Manuel},
  booktitle     = {ICPM Doctoral Consortium / Demo},
  title         = {A Taste of the Different Flavors of Tiramis\`{u}},
  year          = {2024},
  pages         = {154--159},
  crossref      = {ICPMDCDemos2024},
  keywords      = {Process mining; Visual analytics; Knowledge-intensive
                  processes; Visualization},
  url           = {https://ceur-ws.org/Vol-3783/paper_357.pdf}
}
@Proceedings{     ICPMDCDemos2024,
  title         = {Doctoral Consortium and Demo Track at the International
                  Conference on Process Mining 2024 co-located with the 6th
                  International Conference on Process Mining ({ICPM} 2024),
                  Copenhagen, Denmark, October 15, 2024},
  year          = {2024},
  editor        = {De Weerdt, Jochen and Meroni, Giovanni and van der Aa, Han
                  and Winter, Karolin},
  publisher     = {CEUR-WS.org},
  series        = {{CEUR} Workshop Proceedings},
  volume        = {3783},
  urn           = {urn:nbn:de:0074-3783-7}
}
```