Conceptualizing Change Activities in Process Mining

Steven Knoblich[1[0009−0002−8509−7042]],
Lukas Pfahlsberger[1[0000−0002−1367−9441]], and
Jan Mendling[1[0000−0002−7260−524X]]

Humboldt-Universität zu Berlin,
Rudower Chaussee 25, 12489 Berlin-Adlershof, Germany
{steven.knoblich,lukas.pfahlsberger,jan.mendling}@hu-berlin.de

Abstract. In an increasingly dynamic world, business processes must be able to respond to frequently occurring and random changes during their execution. Consequently, this means that the process models must be able to handle this complexity and enable process analysts to derive the right conclusions quickly. However, current approaches in the field of process mining do not distinguish between process activities associated with *change* and those with *routine*. This condition leads to more complicated, overloaded, and sometimes misguided process visualizations that make it difficult for analysts to evaluate them. In this paper, we address the research problem by conceptualizing a new type of process activity that we call *change activity* which we base on causal knowledge. We thereby extend the *causal process mining* approach with another important aspect for handling random occurrences of events. We evaluated our findings through a survey of process mining experts from research and practice. Our results indicate that a dedicated visualization of *change activities* reduces the complexity of process visualizations. In addition, unimportant information is hidden and important information is highlighted so that analysts can make better assessments.

Keywords: Change Activities · Causal Process Mining · Visual Analytics.

1 Introduction

In today's rapidly evolving business environment, the ability of business processes to adapt to unexpected and frequent changes is crucial for organizational success. *Classic* process mining approaches neglect to explicitly designate such *out-of-the-ordinary* changes and treat them as such. In most cases, *changes* are treated the same as *routine* tasks which are on the happy path of the process. This lack of differentiation and explicit visualization poses challenges for process analysts since they are confronted with complex, cluttered, and sometimes misleading visualizations (e.g., so-called spaghetti models) [1, 2].

Business process professionals can benefit greatly from analyses and visualizations that go beyond simple directly-follows representations and enrich models with insights on *changes* [3, 4]. On the one hand, it can help to reduce the time and cost for the process model analysis, but also for the training and onboarding of new analysts. On the other hand, it can increase the quality of the analysis due to a more precise detection of *real* problem root causes, decrease wrong and misleading statements but also increase the flexibility during the analysis by switching the focus on different undesired changes and their impact. Despite the stated merits, the literature has only dealt with the subject in a rudimentary way. For example, they investigated robust process discovery [5] or probabilistic approaches to event-case correlation trying to connect events to the same case with the challenge of including *change* behaviour [6]. Moreover, Lu et al. [3, p. 1] detect contextual activity and claim that it "can affect the performance of any process discovery algorithm". We thus pose the following research question:

How can process mining approaches be conceptually enhanced to effectively differentiate change activities from routine activities in business processes?

In this paper, we address this research question by conceptualizing and visualizing a new type of process activity, called the *change activity*. This activity differentiates from what we call a *routine activity* which has, per definition, one or more causal relationships with other *routine activities*. We build our concept on the basis of the *causal process mining* approach by Waibel et al. [7]. We further evaluate our concept conducting an online survey with business process management professionals.

The remainder of this paper is structured as follows. Section 2 presents the research background against which we position our work. Section 3 presents our concept for change activities by drafting general assumptions and proposing a visual representation. Section 4 presents our evaluation design, data collection procedure, and results. Chapter 5 gives the conclusion of our paper.

2 Background

In this section, we present the background of our research. First, we illustrate the problem. Then we summarize research on change and random occurrences by contextual factors in business process management. Finally, we describe causal process mining as a foundation for our solution.

2.1 Problem Statement

Business processes are typically not executed in isolation. Rather, they are embedded in specific contexts that can trigger *changes* at unexpected times that affect the outcome of the process. *Changes* in processes are one of the most significant factors for the increased complexity in process visualizations. As a theoretical thought experiment, imagine a process with 20 successive process activities which are executed in a row as one process variant. By adding a single change activity, which can theoretically take place after each of these 20 activities, the number of variants in the process model grows rapidly. In fact, this results in a binary decision (add or not add) after each of the 20 activities independently, potentially leading to $2^{20} = 1,048,576$ process variants:

This complexity triggers numerous problems that affect the time, quality, and costs of process analysis in the context of business process management.

The consequences of a lack of differentiation between *routine activities* and *change activities* leads to the following problems. First, the exponential increase in the number of process variants leads to significantly longer times required for process analysis and the training of new process analysts. Each additional variant introduces a new branch of potential actions and outcomes, which complicates the understanding and documentation of the process. Analysts must consider and evaluate each possible variant to determine whether the behavior is desired or undesired, which can be exceedingly time-consuming.

Second, quality can suffer when process complexity becomes unmanageable. Analysts might overlook certain process variants or fail to identify critical issues or causal relationships, leading to suboptimal process improvements. Moreover, the intricate nature of complex processes can make it challenging to maintain consistent quality across all variants. The risk of errors and inconsistencies increases, which can degrade the overall quality of process outcomes.

Third, with increased complexity comes higher costs. The time required to analyze numerous process variants directly translates to increased labor costs. Additionally, complex processes may require a higher skill level and cognitive capabilities from analysts, adding to the overall expense for senior experts. Businesses must also consider the costs associated with potential errors or inefficiencies that arise from inadequate process understanding and management.

2.2 Change and Context in Business Processes

Prior research on context and change of business processes has proposed techniques for modeling and mining. Van der Aalst and Dustdar [8] mention four types of *contexts*: Case context, process context, social context, and external context. The case context includes properties directly related to individual process instances, such as customer type or order size. The process context involves the interactions and competition among multiple instances of the same process, such as resource availability and workload. The social context refers to human and organizational factors, including social networks and individual performance variations. The external context encompasses broader environmental factors, such as weather, economic conditions, and regulatory changes that influence process handling [8].

Lu et al. [3, p. 108] use the term *context activities* indicating that these process activities do not follow a causal order. They attribute this to the fact that it is not the control-flow that influences its execution, but rather random external factors. It is often unclear whether such context activities should be regarded as noise or as part of the control flow. In practice, this results in the generation of the so-called spaghetti or flower models, which are often too complex for analysts to comprehend [3].

Guo et al. [9] develop an algorithm to detect what they call *invisible tasks* that are difficult to determine and sort relate to a *routine* process. Viewed from a similar angle, Goedertier et al. [5] stress that process analysis must deal with challenges such as expressiveness, noise, incomplete data, and the inclusion of prior knowledge. They propose the inclusion of so-called *artificial negative events* to better contrast the ordinary from the extraordinary.

Di Ciccio and Montali [10] introduce the concept of *declarative process mining*, which focuses on behavioral rules and uses the DECLARE language and graphical notations. Building on this work, van Dongen et al. [11] developed a mixed paradigm approach to conformance checking that combines the strengths of procedural and declarative representations. The authors applied their solution to real-world event logs, using a common software to visualize their result, but focusing on other aspects such as execution time or fitness.

None of this previous work has provided a specific classification of different activities such as *change* or *routine activities*.

2.3 Causal Process Knowledge in Process Mining

Process mining is a data-driven technique that involves extracting insights from event data to analyze and improve business processes. It combines principles from computational intelligence, data mining, and process management to visualize, monitor, and optimize business processes within an organization [12]. By revealing discrepancies between intended and actual processes, process mining facilitates better decision making and process optimization [13].

Causal process mining is an approach described by Waibel et al. [7]. It is different from classic process mining in that it takes causal knowledge into account. More specifically, it transforms relational data structures based on the causal template into a causal event graph that internalizes the complex interrelationships between data objects that trigger other objects based on *causation*. Figure 1 illustrates the differences between both approaches.

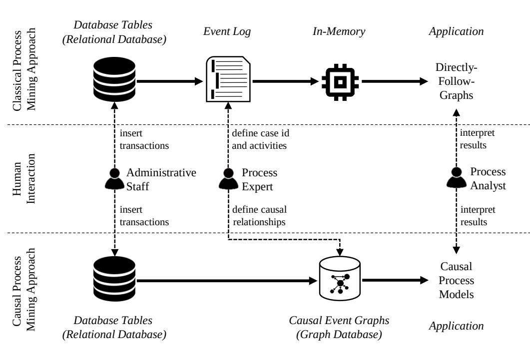

Fig. 1. Difference between *classic* process mining approaches and causal process mining [7].

As a foundational figure in the philosophy of causation, Hume [14] asserted that all knowledge is derived from experience and hinges on the associations between perceived events. Building on this concept, Waldmann [15] explored knowledge-based causal induction, emphasizing causal directionality as the crucial element in interpreting statistical correlations.

Regarding causal knowledge in business processes, experts with extensive domain-specific experience are invaluable for process improvement. Their experience equips them with an accurate understanding of the causal relationships between individual activities within business processes. For instance, a process owner of an order-to-cash process intuitively understands that a customer order will eventually result in an invoice being created. In contrast, it is evident to the process owner that an invoice followed by a customer order would contradict the causal logic of the process [7].

Translating these concepts to process mining algorithms is mostly missing in research [7, 16]. An experiment by Rembert et al. [16], which incorporated prior knowledge, shows that this approach enhances robustness against noise, thereby reducing the likelihood of measurement and ordering errors, especially in processes with a high degree of infrequent behavior. In a related study, Diamantini et al. [17] show that domain knowledge repairs event logs and generates more accurate models for complex and highly variable processes. A template developed by Waibel et al. [7] supports the integration of causal sequences to discover the structure of processes, especially of control flows. This creates simpler process models, with less self-loops and spurious arcs, compared to the classic approach.

Based on this, Pfahlsberger et al. [18] present multi-perspective path semantics based on causal knowledge differentiating between *desired* and *undesired* behavior. We introduce three of these path semantics here, as we adapt them for conceptualizing the representation of change activities and their causal relationships.

– Conformance path:
   The conformance path merges desired and observed behaviors in a process. It represents the expected flow based on the analyst's hypothesis and actual behavior from data. Any deviation is seen as unexpected. Visually, it is shown as a gray angular arrow to subtly indicate the desired behavior.
– Hypothetical path:
  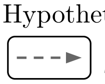 The hypothetical path represents unobserved yet desired behavior, based on causal process knowledge. It implies alternative paths exist, allowing parallel activities or arbitrary follow-up choices. Visually, it is depicted as a gray dashed arrow with a filled arrowhead, indicating indefinite nonconforming behavior.
– Prohibited path:
  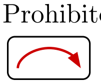 The prohibited path represents undesired but observed behavior, where the process violates against causal process knowledge. Visually, it is

shown as a solid red arrow with a curvilinear course, contrasting with the allowed shortcut path, indicating undesirable behavior.

This previous work by Pfahlsberger et al. [18] focused on path semantics. A conceptual approach for different types of activities, especially change activities, is missing.

3 Conceptualizing Change Activities in Process Mining

In this section, we conceptualize the term *change activity* by formulating general assumptions and designing a visual representation. In this regard, we build on the fundamental concept of *causal process mining* developed by Waibel et al. [7]. Our proposed approach further extends the visual components for multi-perspective path semantics by Pfahlsberger et al. [18].

We define the term *change* in the context of a business process as a principally undesired event that cannot be clearly sequenced within a chain of process activities, as its execution time can occur randomly during execution. For instance, a *change* in an order-to-cash process can be an adaptation of a price for a specific item. Alternatively, during the execution of a purchase-to-pay process, a *change* could manifest itself in the form of adapted supplier terms. What is defined as *change activity* as part of the process mining analysis, always depends on the context and must be specified by a process domain expert. With reference to the previous example, a price change may be part of the standard process in one company, hence not be considered as a *change activity* because it is desired. On the other hand, in another company, such an event clearly qualifies as a *change activity* since its occurrence is undesired.

3.1 Assumptions about Change Activities

In this section, we formulate four central assumptions to delineate the concept of *change activities*. We are placing our focus on the analysis and visual representation of such *change activities* in regard to its impact on the structural effects of the process execution. By structural effects, we mean the triggering of unwanted process patterns such as *rework*, *correction*, *disarray*, or *negligence*, as conceptualized in the approaches of Pfahlsberger et al. [18]. On the one hand, we define a *neutral impact* on the process as a nonmeasurable influence on unwanted process patterns. On the other hand, we define a *negative impact* as a subsequent triggering of such an undesired process pattern.

- Change activities are only relevant for the analysis if they have a negative impact on the process[1]. It is therefore necessary to identify the earliest possible point in the process at which a change has a negative impact. As a consequences, the change also has a negative impact on any subsequent activity.

[1] In this case, negative impact primarily refers to structural effects with regard to rework, correction, disarray, or negligence [18]

- Change activities that have a neutral impact on the process[2], do not have to be analyzed and visualized in relation to the exact point in time at which they were performed. The quantity of their occurrences should be displayed aggregated at the last point in the process, from when a negative impact could be expected.
- Change activities can not be directly triggered by a specific preceding process activity, meaning that the analysis and visualization of incoming paths serves is not necessary. If there is a causal relationship with another preceding process activity, it can not be considered a change activity.
- Change activities can not be directly triggered by another change activity of an identical or different type. This means that it is assumed that there is no causal relationship between the temporal sequence of multiple change activities.

3.2 Visualizing Change Activities

In order to make the assumptions made previously visually accessible to process analysts, we conceptualize four different visual constellations all illustrated in Table 1. The constellations are divided into two categories. First, with or without negative change activities, and with or without neutral change activities. Each pattern is visually depicted for easy recognition of an underlying behavior. We also exemplify the patterns from the perspective of a simple process.

The changes are represented by different visual elements, such as solid lines for observed and dashed lines for unobserved/ hypothetical change. Rectangular black lines for accepted behavior and curved red lines for unaccepted observed behavior, indicating the nature of the change. The diagram represents all possible combinations in four quadrants:

- The top-left quadrant shows the scenario without neutral and without negative change activities. The hypothetical path from the change activity to last activity where the change has a neutral impact is represented.
- The top-right quadrant shows the scenario without neutral change but with negative change activity. The hypothetical path represented and the prohibited path, linking the change activity to the following routine activity.
- The bottom-left quadrant shows the scenario with neutral change activity and without negative change activity. The conformance path from the change activity to last activity where the change has a neutral impact is represented.
- The bottom-right quadrant shows the scenario with both neutral and negative change activity. The conformance path and the prohibited path is represented.

These patterns and the different visual representations of *routine* and *change activities* help to understand the different impacts and implications of changes

[2] In this case, neutral impact primarily refers to no structural effects with regard to rework, correction, disarray, or negligence [18]

within a process instance, providing a clear visual representation of the combinations and their potential consequences. In addition, the lack of a link to the *change activity* reduces the complexity of the visualization. Finally, it helps analysts improve root cause analysis for process inefficiencies.

Overview change constellations	**Without negative change activities**	**With negative change activities**
Without neutral change activities	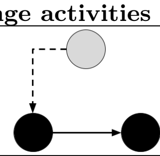	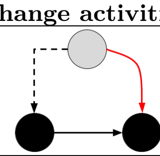
With neutral change activities	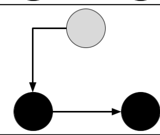	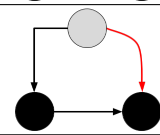

Legend: Routine Activity, Change Activity,
Hypothetical path, Conformance path, Prohibited path

Table 1. Overview of visual components for representing different change constellations

4 Research Method

In this section we describe the research method used for the evaluation, including the survey tasks. This is followed by a description of the data collection and the results. This section concludes with an acknowledgement of the limitations.

4.1 Evaluation Case Setting

To evaluate our concept, we drafted a fictitious case of a food ordering process with 87 orders executed. We chose this case because it is easy to understand and both experienced and unexperienced participants are familiar with such a setting. The overall process contains six *routine activities*, namely, *Answer Call*, *Receive Order*, *Request Delivery Address*, *Prepare Order*, *Deliver Order*, and *Receive Payment* as well as one *change activity*, namely, *Update Ordered Quantity*. The process always starts with the activity *Answer Call* and is sequentially followed by the remaining five routine activities. Every process instance is terminated with *Receive Payment*. During the execution of the process sequence, the *change activity Update Ordered Quantity* was randomly triggered 13 times. We visualized the process in two variants. Figure 2 shows a visual representation of a classic process mining approach. Figure 3 depicts the same process visualized with the extended *causal* process mining approach.

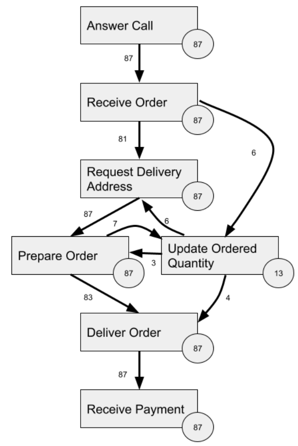

Fig. 2. Visualization of a directly-follows graph from *classic* process mining

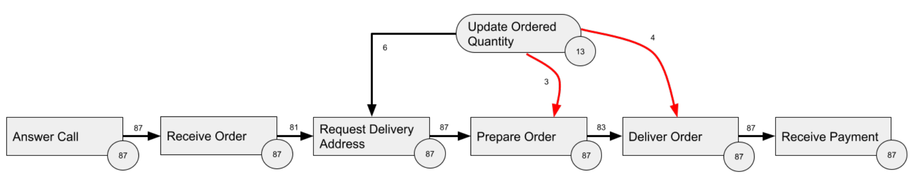

Fig. 3. Visualization of a causal event graph from *causal* process mining

4.2 Survey Tasks

To evaluate the performance effect of our concepts, we designed an online survey. The objective of the survey was to test the understanding of visualizations of the *classic* process mining approach (based on directly-follows graphs) in comparison with the *causal* process mining approach (based on causal event graphs with our extended elements for change activities). Both approaches differ visually through different path semantics, coloring, and activity symbols. We presented all participants with a visual representation of both approaches that both cover the same fictional case:

– The *classic* process mining approach (based on directly-follows graphs): Each node in the graph represents an activity, while the directed edges between nodes indicate the immediate succession of activities based on event logs [12].
– The *causal* process mining approach (based on causal event graphs): Each node in the graph represents an activity, while the gray directed edges between nodes indicate causal relations and the rounded red relations indicate temporal violations [7].

In the survey, we asked participants to solve the following two tasks:

Task 1: What is the change activity in this processes?
Task 2: How many undesired executions of the change activity happened in the process?

Answer options were randomized to mitigate order effects. Two primary metrics were recorded. First, the number of correct responses for identifying change activities and counting undesired executions. Second, the time taken by participants to complete each question. The data collected was analyzed to compare the effectiveness of the *classic* and *causal* approaches. Specifically, we assessed the accuracy of participants' responses as well as the efficiency required to answer each question. These metrics are critical for validating our new visualization technique and addressing the research question.

The visualizations of the *classic* approach served as the control condition, maintaining the current visualization logic without specific highlights or causal assessments. The experimental condition employed our proposed visualization, which incorporated specific highlights for change activities and visual assessments of causal relationships. The purpose of this comparison was to determine whether the new visualization improved the participants' ability to accurately and quickly identify change activities and/or their undesired executions.

4.3 Data Collection

Participants were recruited from a diverse pool of individuals to ensure a comprehensive analysis. Recruitment efforts were carried out through various channels, including direct email contact with two groups from business and research sectors at Humboldt University of Berlin, as well as two LinkedIn posts by the paper's authors. These posts were widely shared on the platform. This multifaceted approach resulted in 27 completed questionnaires. Prior to administering the main survey, participants' experience with business process management and their professional background were collected to contextualize the findings. The online survey was conducted anonymously for a one-week period from Monday, July 24, 2024, to Sunday, July 30, 2024. The majority of responses were received within the first three days, and the final response recorded on July 28, 2024.

The survey was send out via mass emailing to contacts from the authors' network. Additionally, it was shared on social channels such as *LinkedIn* and *X*. A total of 57 people visited the survey page. 37 of them started filling out the survey, of which 27 completed the survey to the end. One participant did not answer the second question regarding the directly-follows graph visualization. Among the participants, 19 had a background in business and eight assigned themselves to the academic field. One of the participants did not identify with either an academic or a business background. In terms of experience in business process management domain, ten participants reported having one to three years of experience, 11 had three to seven years, and five participants had more than seven years of experience. The descriptive statistics are depicted in Figure 4.

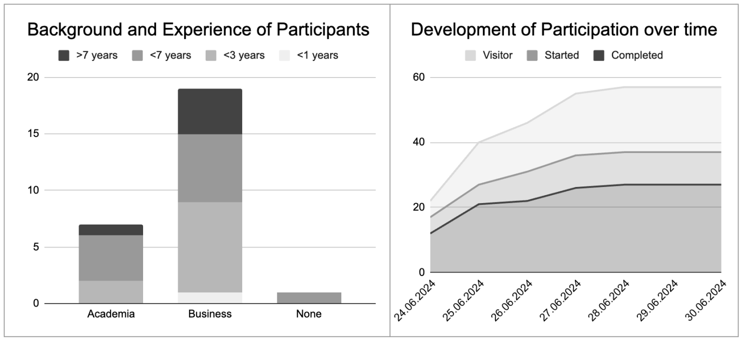

Fig. 4. Development of participants and their background.

4.4 Results

The study yielded several noteworthy findings. First, participants using the causal event graph generally required more time to interpret the visualizations compared to those using directly-follows graphs. This indicates that causal event graphs might be more complex or less intuitive than directly-follows graphs for the participants. Academics were more familiar with the visualizations than business professionals, responding significantly faster than the latter. This finding underscores the need for clearer process mining visualizations in practical applications, as business professionals may require more time and effort to interpret the same data. For the directly-follows graph, the median response time was 50 seconds, whereas for the causal event graph, it was 61 seconds. When broken down by background, the median response time for participants with a business background was consistently 85 seconds across both types of graphs.

Second, the variability in response times was greater among business professionals compared to academics. This suggests differing levels of knowledge between the two groups and highlights a clearer understanding of process visualizations in the academic domain, particularly among those using directly-follows graphs.

Third, the level of experience did not have a strong significant impact on the participants' response times. However, there was a slight trend indicating that individuals with more than three years of experience exhibited less fluctuation in their response times, with most responding within 100 seconds. The data suggests that less experienced participants had more difficulty interpreting the causal event graphs. Of the three participants who took more than 200 seconds to respond, two had less than three years of business process management experience.

Fourth, when answering Question 1, participants were able to identify the sought-after *change activity* in the fictional example at almost the same speed, likely due to the straightforward nature of the activity label. Notably, there was only one incorrect response from a participant with a business background and one to three years of business process management experience for the directly-follows graph. This uniformity suggests that some tasks in process mining may be universally intuitive regardless of the visualization used.

Lastly, a clear pattern emerged in the responses to Question 2. The directly-follows graph seemed to prompt participants to respond very quickly; however, their interpretations were significantly less accurate than those using the causal event graph. Only 3 out of 27 participants answered correctly with the directly-follows graph, whereas 16 participants provided correct answers with the causal event graph. This discrepancy indicates the importance of selecting the appropriate visualization method to ensure both efficiency and accuracy in process mining tasks.

In summary, while directly-follows graphs may facilitate quicker responses, the accuracy of interpretations is higher with causal event graphs. This finding highlights the trade-off between speed and accuracy in the interpretation of process mining visualizations and suggests that different contexts may require different approaches to visualization. Consequently, our research suggests that enhancing process mining approaches with more intuitive and clear visualizations of *changes* can effectively improve their differentiation from *routine activities*. This differentiation increases the quality of the process analysis and improves decision-making.

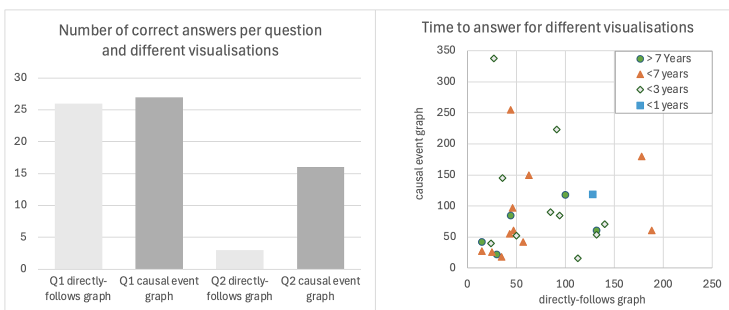

Fig. 5. Results from survey.

4.5 Limitations

We acknowledge the following three limitations. First, our survey compared the classic with the causal process approach, including the conceptualization for change activities. It is important to note that previous research has already demonstrated a positive impact of the causal approach on process discovery and interpretation [7, 18], so the effect of our extension needs to be seen in this context. Moreover, the causal representation in our study was supported by color-coding, which may have influenced the results. Future research should also consider the effect of the orientation of process discovery visualizations, as this factor could significantly affect user comprehension and interaction. Second, the demographics of the study participants were not fully balanced. Variables such as user experience levels and their professional or educational background were not uniformly distributed. Future studies should aim to include a more diverse and representative sample to better generalize the findings. Third, the fictitious case used in our study was relatively simple, consisting of only seven activities and up to eight connections. Although this simplicity helped control for confounding variables, it may not accurately reflect the complexity of real-world process discovery scenarios. We anticipate that more complex data sets and visualizations could amplify the observed effects and provide a more comprehensive understanding of the relative benefits of our concept.

5 Conclusion

In this paper, we conceptualize a new *type* of process activity – the *change activity* – that helps analysts to better differentiate between *change* and *routine* activities. We thus address a still existing research problem that classic process mining approaches often lead to complex and misleading visualizations, such as Spaghetti models. Our aim is to encourage future research to confront persistent representational bias in process mining, which often skews the true nature of the underlying process [2]. Our paper contributes in two ways. First, we propose general assumptions for change activities and then outline a visual representation based on the *causal process mining* approach. Second, we evaluate our concepts based on a fictitious case of a food delivery service with an online survey of process experts from practice and science. Our results indicate that process analysts can benefit significantly from a differentiation between *change* and *routine* activities by pointing out root causes of the problem related to changes during the process execution more accurately compared with *classic* process mining representations. Nonetheless, at least according to our survey, they need slightly more time for this. In fact, by incorporating a more context-aware visualization in regard to changes, analysts can reduce cost associated with process model analysis and enhance the overall quality of their analysis.

Acknowledgments. This work was supported by the Einstein Foundation Berlin [grant number EPP-2019-524, 2022].

References

1. Van der Aalst, W.M.P.: Process mining: Discovering and improving Spaghetti and Lasagna processes. In: 2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM), pp. 1–7. IEEE (2011)
2. Van der Aalst, W.M.P., Buijs, J., Van Dongen, B.: Towards improving the representational bias of process mining. In: Data-Driven Process Discovery and Analysis: First International Symposium, SIMPDA 2011, Campione d'Italia, Italy, June 29–July 1, 2011, Revised Selected Papers, vol. 1, pp. 39–54. Springer Berlin Heidelberg (2012)

3. Lu, Y., Chen, Q., Poon, S.K.: Detecting Context Activities in Event Logs. In: International Conference on Business Process Modeling, Development and Support, pp. 108–122. Springer International Publishing, Cham (2022)
4. Knoblich, S., Mendling, J., Jambor, H.: Review of Visual Encodings in Common Process Mining Tools. In: 1st Visual Process Analytics Workshop, vol. 1, (2024)
5. Goedertier, S., Martens, D., Vanthienen, J., Baesens, B.: Robust process discovery with artificial negative events. In: Journal of Machine Learning Research, vol. 10, pp. 1305–1340 (2009)
6. Bayomie, D., Di Ciccio, C., La Rosa, M., Mendling, J.: A probabilistic approach to event-case correlation for process mining. In: International Conference on Conceptual Modeling, pp. 136–152. Springer International Publishing, Cham (2019)
7. Waibel, P., Pfahlsberger, L., Revoredo, K., Mendling, J.: Causal Process Mining from Relational Databases with Domain Knowledge. CoRR, abs/2202.08314 (2022), https://arxiv.org/abs/2202.08314
8. Van der Aalst, W.M.P., Dustdar, S.: Process mining put into context. In: IEEE Internet Computing, vol. 16, no. 1, pp. 82–86 (2012)
9. Guo, Q., Wen, L., Wang, J., Yan, Z., Yu, P.S.: Mining invisible tasks in non-free-choice constructs. In: International Conference on Business Process Management, pp. 109–125. Springer International Publishing, Cham (2015)
10. Di Ciccio, C., Montali, M.: Declarative process specifications: reasoning, discovery, monitoring. In: van der Aalst, W.M.P., Carmona, J. (eds) Process Mining Handbook, pp. 108–154. Springer Nature Switzerland AG, Cham (2022)
11. van Dongen, B.F., De Smedt, J., Di Ciccio, C., Mendling, J.: Conformance checking of mixed-paradigm process models. Preprint submitted to Information Systems, pp. 1–65. arXiv:2011.11551v1 [cs.FL] (2020)
12. Van der Aalst, W.M.P., Weijters, A.J.M.M.: Process Mining. In: Dumas, M., van der Aalst, W., ter Hofstede, A.H.M. (eds.) Process-Aware Information Systems: Bridging People and Software Through Process Technology, pp. 235–254. John Wiley & Sons, Hoboken, New Jersey (2005)
13. Van der Aalst, W.M.P.: Process Mining: Overview and Opportunities. ACM Transactions on Management Information Systems, vol. 99, no. 99, Article 99, 16 pages (February 2012). https://doi.org/10.1145/0000000.0000000
14. Hume, D.: An Enquiry Concerning Human Understanding. 2nd edn. Hackett Publishing Company, Indianapolis, IN (1977). Original work published 1748
15. Waldmann, M.R.: Knowledge-based causal induction. In: Psychology of Learning and Motivation, vol. 34, pp. 47–88. Elsevier (1996)
16. Rembert, A.J., Omokpo, A., Mazzoleni, P., Goodwin, R.: Process Discovery Using Prior Knowledge. In: Basu, S., Pautasso, C., Zhang, L., Fu, X. (eds.) Service-Oriented Computing - 11th International Conference, ICSOC 2013, Berlin, Germany, December 2-5, 2013, Proceedings. LNCS, vol. 8274, pp. 328–342. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-45005-1_23
17. Diamantini, C., Genga, L., Potena, D., van der Aalst, W.M.P.: Building instance graphs for highly variable processes. Expert Syst. Appl., vol. 59, pp. 101–118 (2016)
18. Pfahlsberger, L., Rubensson, C., Knoblich, S., Vidgof, M., Mendling, J.: Multi-perspective path semantics in process mining based on causal process knowledge. In: Companion Proceedings of the 16th IFIP WG 8.1 Working Conference on the Practice of Enterprise Modeling (2023)