

Process Miner, Are You Sure? Conducting Valid and Reliable Experiments in Process Mining

Jana-Rebecca Rehse¹[0000-0001-5707-6944], Sander J.J. Leemans²[0000-0002-5201-7125], Peter Fettke^{3,4}[0000-0002-0624-4431], and Jan Martijn E.M. van der Werf⁵[0000-0002-7264-381X]

¹ University of Mannheim, Germany

`rehse@uni-mannheim.de`

² RWTH Aachen University, Germany

`s.leemans@bpm.rwth-aachen.de`

³ German Research Center for Artificial Intelligence (DFKI), Germany

⁴ Saarland University, Germany

`peter.fettke@dfki.de`

⁵ Utrecht University, The Netherlands

`J.M.E.M.vanderWerf@uu.nl`

Abstract. With the increasing availability of large-scale event logs, process mining has become a data-oriented research discipline, in which experiments play an essential role. Experiments are used to demonstrate the feasibility of algorithms, establish properties of algorithms, and compare new algorithms to existing ones. However, unlike other disciplines, process mining has not established conventions for how experiments should be conducted. Hence, researchers typically design their experiments ad-hoc, with little guidance on how they can ensure the validity and reliability of the conclusions. When comparing papers, even if the designed experiments are reliable and valid, the differences in set-up often prohibit a comparison of results across papers, which in turn makes it difficult to draw conclusions for the field. To address these problems, this tutorial advocates for a systematic approach for conducting experiments in process mining. Based on a discussion of reliability and validity as two main principles of scientific inquiry, we present a checklist that researchers can follow to design and execute process mining experiments in a reliable and valid way. Furthermore, we invite attendees to engage in a discussion to establish experimental conventions for process mining.

Keywords: Process mining · Experiments · Reliability · Validity.

1 Motivation

Process mining has its roots in (theoretical) computer science and automata theory, where the main objective is to design algorithms that exhibit certain formal properties. With the increasing availability of large-scale real-life event logs, process mining has shifted towards a data-oriented research discipline, where the main goal is to design algorithms that are applicable and useful in practice. This

shift from purely formal to real-world properties also entails a change in research methods: whereas the former can be mathematically proven, the latter require different evaluation methods, grounded in empirical science [2]. However, unlike other disciplines, process mining has not established conventions for how experiments should be conducted. This lack of conventions may cause experimental results to be unreliable or only partially valid. Also, there are often significant differences between individual papers, which makes it hard or even impossible to compare their results and draw conclusions for the field as a whole.

The goal of this tutorial is to raise awareness about these potential methodical problems in process mining research, to present a novel methodology that can help researchers in conducting reliable and valid experiments, and to engage the process mining community in a reflective discourse on the importance of research methods and the role that reliability and validity play in our experiments.

2 Tutorial Outline

Our tutorial will consist of two major parts:

Reliability & Validity After a short motivation on the need for methodical guidance in the process mining field, we introduce reliability and validity as fundamental principles of scientific inquiry. We discuss the different types of reliability and validity, as outlined in Figure 1, and explain how they might impact experimental results. Adapting these generic principles to the process mining domain, we provide examples of how these types of reliability and validity can be ensured or threatened in process mining experiments.




Fig. 1. Different types of reliability and validity

To conclude the first part, attendees are welcomed to pick an (experimental) process mining paper of their choice and analyze whether and how the different types to reliability and validity are addressed or threatened in this paper.

Process Discovery Engineering In the second part of the tutorial, we introduce “Process Discovery Engineering” [3] as a new research methodology for process

mining.⁶ This methodology is based on algorithm engineering [4,2] and is supposed to help researchers design and conduct experiments in process mining. It consists of two parts [1]: (i) a shared terminology, which defines terms to characterize the underlying ontological and epistemological position and generic notions of research, and (ii) a checklist, which guides researchers in considering potential threats to reliability and validity when designing and executing experiments. After being introduced to both parts, attendees will get the chance to apply the process mining checklist to their selected paper, to reflect on how it might have changes the design or communication of the experiment.

3 Discussion Questions

The tutorial will end with a plenary group discussion on the contents of the tutorial and a potential way forward. To this end, we will discuss the following questions:

- Why is it important from a theoretical point of view to have high-quality experiments in process mining?
- What role does theory play in process mining?
- How do generic notions of research, like “research method” or “hypothesis” occur in process mining?
- In what extent are reliability and validity well-addressed in the process mining community? How are they threatened?
- What are best practices for conducting experiments?
- How could experimental conventions for process mining look like?
- What additional research is required in this area?

Ideally, this discussion will reveal new insights and reflections on how to perform experiments systematically and rigorously, in order to advance the field as a whole. Therefore, this tutorial is intended for all researchers, who have conducted some experimental process mining research and are willing to reflect on that process. Both students and experienced researchers might benefit from this reflection process, so we welcome all levels of experience.

References

1. Frank, U.: Towards a pluralistic conception of research methods in information systems research. Tech. Rep. 7, ICB-Research, Universität Duisburg-Essen (2006)
2. Mendling, J., Leopold, H., Meyerhenke, H., Depaire, B.: Methodology of algorithm engineering. arXiv preprint arXiv:2310.18979 (2023)
3. Rehse, J.R., Leemans, S., Fettke, P., Van der Werf, J.M.E.: On process discovery experimentation. *ACM Transactions on Software Engineering and Methodology* (2024)
4. Sanders, P.: Algorithm engineering—an attempt at a definition. In: *Efficient algorithms*, pp. 321–340 (2009)

⁶ Note that although this research methodology is mainly focused on process discovery as an application area, most of our insights can be applied to other areas of process mining as well.