

PM²: a Process Mining Project Methodology

Maikel L. van Eck, Xixi Lu, Sander J.J. Leemans, and Wil M.P. van der Aalst

Eindhoven University of Technology, The Netherlands

{m.l.v.eck,x.lu,s.j.j.leemans,w.m.p.v.d.aalst}@tue.nl

Abstract. Process mining aims to transform event data recorded in information systems into knowledge of an organisation’s business processes. The results of process mining analysis can be used to improve process performance or compliance to rules and regulations. However, applying process mining in practice is not trivial. In this paper we introduce PM², a methodology to guide the execution of process mining projects. We successfully applied PM² during a case study within IBM, a multinational technology corporation, where we identified potential process improvements for one of their purchasing processes.

Keywords: process mining, methodology, case study, business process management.

1 Introduction

Process mining techniques can be used to automatically discover process models, check the conformance of process models to reality, and extend or improve process models using data of actual process executions [1]. Process mining analysis results can be used to improve the performance of processes or an organisation’s compliance to rules and regulations. Hence, process mining provides the bridge between data mining or machine learning techniques and the business process management discipline.

Within the field of data mining, efforts have been made to establish methodologies to support organisations with their data mining projects [9, 12]. The aim of these methodologies is to guide the planning and execution of such projects in order to save time and costs, e.g. by helping to avoid the presentation of irrelevant insights. This also results in a better understanding and acceptance of data mining projects [9]. Two widely used methodologies are CRISP-DM [16], developed by a consortium led by SPSS, and SEMMA, developed by SAS [12].

Efforts have also been made to create project methodologies that are tailored toward supporting process mining projects, as methodologies like CRISP-DM and SEMMA are very high-level and provide little guidance for process mining specific activities [1]. To the best of our knowledge, there are two well-known process mining methodologies: *Process Diagnostics Method* (PDM) [6], which has also been adapted for healthcare environments [14], and the *L* life-cycle model* [1]. PDM is designed to quickly provide a broad overview of a process, while L* covers many different aspects of process mining and touches on broader topics like process improvement and operational support.

Unfortunately, these methodologies are not suitable for every project. The scope of PDM is limited, covering only a small number of process mining techniques and emphasises on avoiding the use of domain knowledge during the analysis [6], which makes

it less applicable for larger, more complex projects [18]. L* covers more techniques, but was primarily designed for the analysis of structured processes and aims at discovering a single integrated process model. Neither L* nor PDM explicitly encourages iterative analysis, which proved vital for both our own case study as well as the case study performed in [18]. Moreover, both methodologies can benefit from additional practical guidelines to help inexperienced practitioners to overcome common challenges.

To address these issues, we present PM²: a Process Mining Project Methodology. PM² is designed to support projects aiming to improve process performance or compliance to rules and regulations. It covers a wide range of process mining and other analysis techniques, and is suitable for the analysis of both structured and unstructured processes. For each *stage* of PM², we define its *inputs* and *outputs* and discuss the concrete steps to be executed, referred to as *activities*. PM² supports quick analysis iterations and evolving insights, taking existing best practices into account. We provide practical guidance on using the methodology by discussing a case study performed together with IBM. There we applied PM² and used various process mining techniques to answer research questions related to the performance of a purchasing process.

The structure of the paper is as follows. In Sect. 2 we discuss the PM² methodology and explain each of its stages. The case study is discussed in Sect. 3 and the paper is concluded in Sect. 4.

2 The PM² Methodology

In this section we present the PM² methodology. We first give an overview of PM² and then discuss each stage of the methodology in detail.

PM² guides organisations performing process mining projects aimed at improving process *performance* or *compliance* to rules and regulations. The goals of a process mining project can be very concrete, e.g. achieving a cost reduction of 10% for a given process, or more abstract, e.g. obtaining valuable insights regarding the performance of several processes. Through PM², these goals are translated into concrete *research questions* which are iteratively refined and answered, resulting in findings that are the basis of improvement ideas for the selected process.

An overview of the PM² methodology is shown in Fig. 1. The methodology consists of six stages that relate to several different input and output objects of the following types: goal-related objects, data objects, and models. The four goal-related objects are (1) *research questions* derived from project goals, which are answered by (2) *performance findings* and (3) *compliance findings*, leading to (4) *improvement ideas* to achieve the goals. The data objects denote the three different representations of process-related data: (1) *information systems* contain live process data in various forms, which can be extracted and linked to discrete events to form (2) *event data*. Event data can be transformed into (3) *event logs* by defining a case notion and event classes. We consider two types of models: (1) *process models* and (2) *analytic models*. Process models describe the ordering of activities in a process, possibly enhanced with additional information e.g. temporal constraints, resource usage or data usage. Moreover, we also consider business rules as (abstract) process models formally defining constraints with respect to

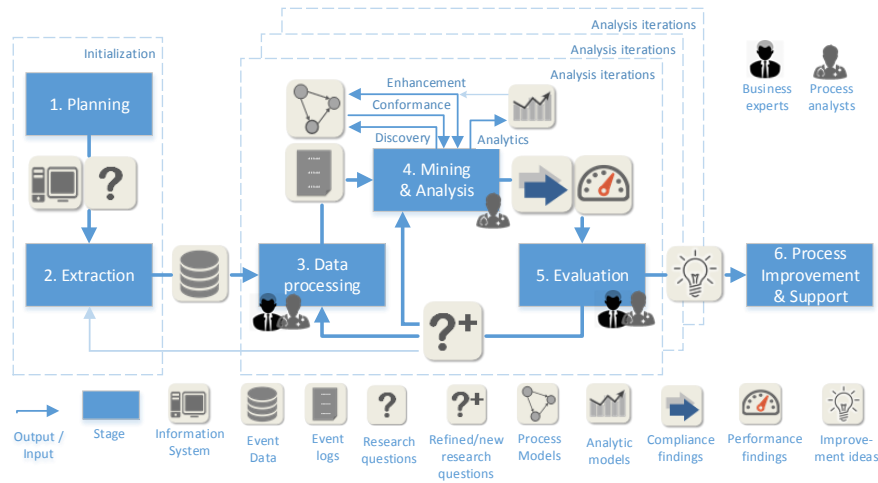


Fig. 1: An overview of the PM² methodology.

the execution of business processes. Analytic models are any other type of models that give insight into the process, e.g. decision trees.

The first two stages of the methodology are (1) *planning* and (2) *extraction*, during which initial research questions are defined and event data are extracted. After the first two stages, one or more *analysis iterations* are performed, possibly in parallel. In general, each analysis iteration executes the following stages one or more times: (3) *data processing*, (4) *mining & analysis*, and (5) *evaluation*. An analysis iteration focusses on answering a specific research question by applying process mining related activities and evaluating the discovered process models and other findings. Such an iteration may take anywhere from minutes to days to complete, mainly depending on the complexity of the mining & analysis. If the findings are satisfactory then they can be used for (6) *process improvement & support*.

In the following we discuss each stage, its input and output, and the activities that are performed in it.

2.1 Stage 1: Planning.

The objective of the *planning* stage is to set up the project and to determine the research questions. We consider two main goals for starting process mining projects: improving performance of a business process, or checking its compliance with respect to certain rules and regulations.

The inputs of this stage are the organisation's *business processes*. The outputs are goal-related *research questions* and a set of *information systems* supporting the execution of the business processes to be analysed.

We identified three activities for this stage: *identifying research questions* (R.Q.), *selecting business processes* (B.P.), and *composing project team*. The order in which these activities are executed may sometimes vary, as there may already be a specific goal or research question before starting the process mining project.

- *Selecting business processes.* A process mining project generally starts with selecting the business processes to be analysed and improved. Both the *process characteristics* as well as the *quality of event data* should be taken into account since they have large influence on the achievable results of the project [4]. Bose et al. [4] identified the four categories of problems related to the quality of event data: missing data, incorrect data, imprecise data and irrelevant data. For example, imprecise timing information for events affects the results of performance measurements, while the absence of unique identifiers to link all related events makes the event log creation harder [13]. In addition to these two factors, we also consider the *changeability* of the business processes, i.e. the organisation needs to be able to influence or adapt process executions based on the findings. This is important if process improvement is the main project goal. After selecting the business processes, the set of information systems that store the relevant process execution data is identified.
- *Identifying research questions.* During this activity, the goals are identified and translated into *research questions*, which we defined as *questions related to the selected process that can be answered using event data*. Research questions can be related to different aspects of business processes, e.g. quality, time, resource, cost. Various case studies [18] showed the importance of defining concrete research questions for a successful process mining project. However, we demonstrate in our case study that abstract research questions from the initialization phase can be refined through explorative analysis, resulting in concrete improvement ideas and valuable insights.
- *Composing project team.* The last activity involves selecting the people that work on the project. Earlier case studies and our own study show that project teams need experts with different backgrounds [18]. We define the following roles: business owners (who are in charge of the business processes), business experts (who know the business aspect and executions of the processes), system experts (who are familiar with the IT aspect of the processes and the systems supporting the processes), and process analysts (who are skilled in analysing processes and applying process mining techniques). The most important roles are the business experts and the process analysts, between which collaboration is essential to evaluate the analysis findings and to ensure that the findings are relevant and usable.

2.2 Stage 2: Extraction.

The *extraction* stage aims to extract *event data* and, optionally, *process models*. Inputs for this stage are the *research questions* and the *information systems* that support the execution of the selected business processes to be analysed. The outputs of this stage are *event data*, i.e. a collection of events without predefined case notion or event classes, and possibly *process models*.

We identified three activities for this stage: *determining scope*, *extracting event data*, and *transferring process knowledge*.

- *Determining scope.* This activity involves determining the scope of the data extraction, based on which the event data is to be created. We give four examples of questions to be considered: (1) with which *granularity* should be event data extracted (e.g. considering events related to purchase orders but neglecting events related to the

items of purchase orders); (2) within which period; (3) which data attributes should be extracted; (4) which correlation between data should be used to collect them.

- *Extracting event data.* Once the extraction scope is determined, event data can be created by collecting the selected process related data from the relevant information systems and joining them into a single collection of events, for example, a table in which each entry represents an event.
- *Transferring process knowledge.* This activity can be executed simultaneously with the creation of event data. Tacit knowledge related to the selected business processes and the data attributes is exchanged between business experts and process analysts, through e.g. interviews or brainstorm sessions, which enables the analysts to be effective in the data processing and mining stages. Such process knowledge may include written process documentation or hand-made process models. Process knowledge is shared throughout the project, but understanding of the process is essential for an effective data processing stage.

In contrast to existing process mining methodologies, we explicitly divided the event data extraction and the log creation and processing into two stages. One reason is that the event data extraction is time-consuming and less frequently repeated than data processing activities like filtering [13]. Another reason is that it is possible to create different views on the same event data that result in different event logs, as discussed in the next section.

2.3 Stage 3: Data Processing.

The main objective of the *data processing* stage is to create event logs as different views of the obtained event data and to process event logs in such a way that it is optimal for the mining and analysis stage. In addition to the *event data* as our main input, one can also use *process models* as an input to filter the event data. The outputs are *event logs* that are used in the *mining and analysis* stage.

We identify four types of activities for this stage: *creating views*, *aggregating events*, *enriching logs* and *filtering logs*. Fig. 2 shows an overview of these activities and how they are applied.

- *Creating views.* Event logs are specific views on event data, created by defining case notions and event classes. Case notion relate events such that together they form a process instance, while event classes distinguish different activities within a process instance. Which view to create depends on the goal of the analysis, e.g. an order is a logical case notion to analyse throughput times, while the resource is a better case notion to analyse resource utilisation. A similar situation holds for defining the event classes: if every order is a process instance and the resource involved in an event is its event class, then process discovery algorithms produce handover of work graphs.
- *Aggregating events.* Aggregating events can help to reduce complexity and improve structure of mining results [5]. We distinguish two types of aggregation: *is-a* and *part-of*. The *is-a* aggregation considers different types of events belonging to an equivalent but more general event class while the number of events remains the same. For example, two events labeled with *Simple Manual Analysis* and *Complex Manual Analysis* are considered instances of event class *Manual Analysis* but remain as two events. In




Fig. 2: An overview of different types of data processing activities.

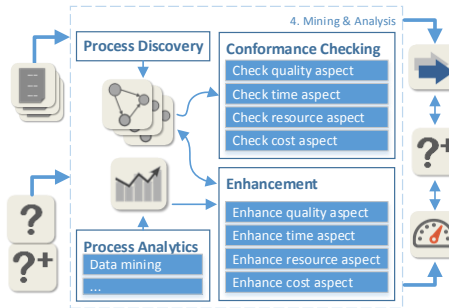


Fig. 3: An overview of the activities in mining and analysis stage.

contrast, the *part-of* aggregation merges multiple events into larger events, as is the case with sub-processes. Both types of aggregation can also be applied in reverse, i.e. defining a specialisation. A more general technique is to define a hierarchy based on event attributes that can be used to aggregate events, as is discussed in [2], e.g. considering location at a city, country or continent level.

- *Enriching logs.* Event logs, as any other data, can be enriched with various additional attributes [11]. We discuss two ways of enriching an event log: (1) deriving or computing additional events and data attributes based on the log itself, or (2) adding external data. The throughput time of a case can be a computed data attribute, while adding information on the weather at the time an event occurred is an example of including external data.
- *Filtering logs.* Finally, filtering is a well-known and frequently used data processing step to reduce complexity or focus the analysis on a specific part of the dataset. This activity is often performed multiple times in an analysis iteration to obtain different perspectives on the event data. We distinguish three types of filtering techniques: *slice and dice* (also known as attribute filtering), *variance-based*, and *compliance-based*.
 - *Slice and dice* can be used to remove events or traces based on the values recorded for a specific attribute, e.g. activity name, resource identifier or timestamps of events, or based on simple statistics, e.g. number of events of a trace or case durations.
 - *Variance based filtering* groups similar traces, e.g. through clustering, which can be used to partition the event log in order to discover simpler process models for each of the partitions of a complex process [14].
 - *Compliance based filtering* can be used to remove traces or events that do not comply with a given rule or fit a given process model, which is a very flexible form of filtering.

2.4 Stage 4: Mining & Analysis.

In the *mining & analysis* stage, we apply process mining techniques on event logs and aim to answer answer research questions and gain insight into processes *performance* and *compliance*. If the research questions are more abstract, explorative techniques

combined with process discovery can be applied on event logs to get an overall view of the *business process*, e.g. its control-flow. Once more specific research questions have been defined, the analysis can focus on answering concrete research questions, e.g. the difference between the throughput times of the cases executed the activity *Manual analysis* and the cases that skipped this activity.

Inputs for this stage are *event logs*. In addition, if *process models* are available, they can also be used for conformance checking and enhancement activities. Output for this stage are findings that answer research questions related to *performance* and *compliance* goals.

We identify four types of activity for this stage: *process discovery*, *conformance checking*, *enhancement* and *process analytics*. The first three activities are well-known process mining techniques [1]. *Process analytics* are other complementary analysis techniques, e.g. data mining and visual analytics, which can be applied in the context of business processes [11]. Fig. 3 shows an overview the four activities.

- *Process Discovery*. Given an event log as input, we generally start with process discovery techniques, which return a fact-based process model as output. For discussions on different process discovery techniques, see e.g. [7].
- *Conformance Checking*. Given a process model, discovered or documented, describing intended behaviour, and an event log recorded real behaviour, conformance checking techniques *aim at the detection of inconsistencies between a process model and its corresponding execution log* [15]. Research questions related to the compliance of business processes to different aspects, such as quality, time, resource and cost, can be checked using conformance checking techniques, the results of which can also be used to *enhance* process models.
- *Enhancement*. The enhancement activity is defined as *extending or improving an existing process model using information about the actual process recorded in an event log* [1], for example, by extending process model with performance information related to time or cost, or repairing the process model according to current executions shown by the corresponding event log. The results of enhancement are process models, e.g. enhanced with different aspects, whereas the results of conformance checking can be considered without any process model.
- *Process Analytics*. In addition to the three process mining activities, other analysis techniques can be applied in the context of event logs and process models, such as data mining techniques [11] or visual analytics (e.g. histograms of events per case), of which the results can be used to enhance process models with additional aspects.

2.5 Stage 5: Evaluation.

The objective of the *evaluation* stage is to relate the analysis findings to improvement ideas that achieve the project's goals. The inputs are the *process models*, *performance* and *compliance* findings from the analysis stage. The outputs are *improvement ideas* or new *research questions*.

The activities for this stage are: *Diagnose*, and *Verify & Validate (V&V)*.

- *Diagnose*. Diagnosing the findings obtained through mining and analysis includes the following: (1) correctly interpreting the results (e.g. understanding the process

model discovered), (2) distinguishing interesting or unusual results from the expected ones (e.g. large set of abnormal executions), and (3) identifying or refining research questions for possible further iterations.

- *Verify & Validate*. The correctness of the (unexpected) findings is investigated. *Verification* compares the findings obtained to the original data and system implementations, while *validation* compares the findings to the claims of process stakeholders, e.g. interviewing the resources involved in processes. Both verification and validation may help identifying the underlying root causes and designing ideas for possible process improvements.

One of the challenges in process mining projects is often that the process analysts are not domain experts for the process they are analysing [6, 18], which means that they may have difficulties determining the causes of unexpected analysis results. Therefore, it is essential that process experts are involved in the verification and validation of the results. Ideally, they would already be involved during the previous mining stage, guiding the analysis to make sure that the results are useful for the organisation.

2.6 Stage 6: Process Improvement & Support.

The objective of the *process improvement & support* stage is to use the gained insights to modify the actual process execution. The inputs of this stage are the *improvement ideas* from the evaluation stage. The outputs of this stage are *process modifications*.

The activities are: *implementing improvements* and *supporting operations*.

- *Implementing improvements*. Achieving process improvements is often the main motivation for a process mining project. However, the actual implementation of process modifications is generally a separate project and a different area of expertise. The results of a process mining project then form the fact-based input of such process improvement efforts. Approaches that focus on this area include business process re-engineering and Six Sigma [8]. After changing the process, the improvements can be measured in another analysis project.
- *Supporting operations*. Process mining can provide operational support by detecting problematic running cases, predicting their future or suggesting recommended actions. To use process mining for operational support it is essential that the results are of high quality, and that there is an IT infrastructure in place that links these results to live event data. It is a challenging form of process mining, suitable only for very structured processes [1].

3 IBM Case study

In this section we describe how PM² was applied in a concrete case study. The case study has been conducted at IBM, a leading multinational technology and consulting corporation. Among many other services, IBM provides hardware service plans. We have analysed a supporting process in this area: the purchasing process for spare parts. This process starts with the creation of a purchase requisition, after which an order is

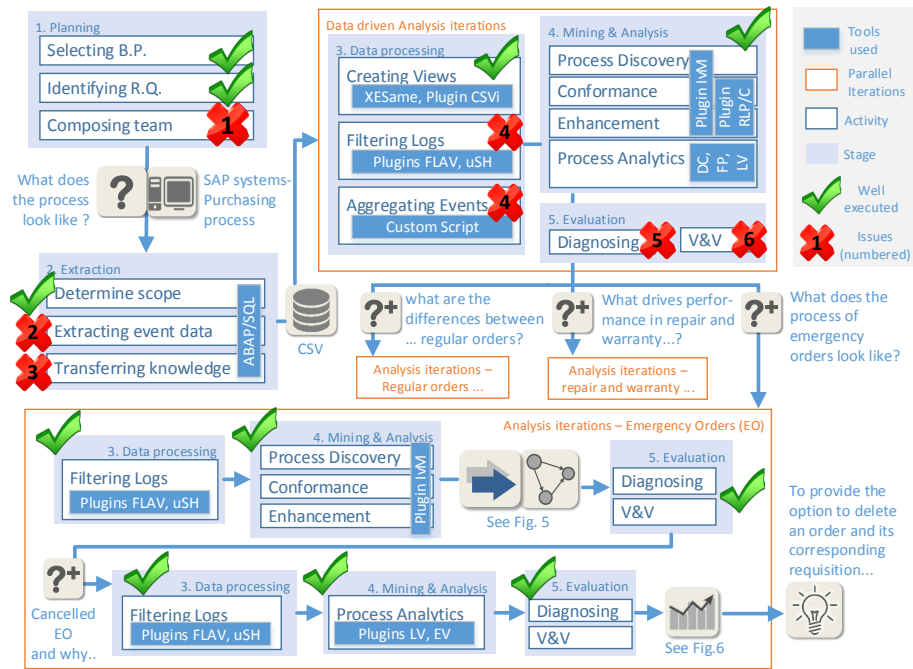


Fig. 4: An overview of the case study execution according to the stages of PM².

sent out to a supplier. The process ends when all ordered items are delivered, which occasionally requires multiple deliveries. There are three different types of orders: regular replenishment orders, emergency orders, and orders for the repair of parts. The process is performed independently at several different IBM facilities around the world.

For our case study, we mainly used the open-source process mining framework ProM toolkit [19] to support applying our methodology. The ProM is freely available¹ and contains some of the latest developments in process mining research, implemented as plug-ins.

In the following, we first discuss the activities executed in each stage of the methodology and the analysis results, listing only the tools and plugins used that were most important for our analysis. Following that we summarize the lessons learned from this case study.

3.1 Execution

An overview of the case study's activities is shown in Fig. 4. It gives the concrete inputs and outputs of each stage, which tools were used, and it shows where we deviated from PM² or encountered issues.

Planning. IBM's primary goal was to get detailed insights in the spare parts purchasing process' performance. This process was selected mainly due to the availability of good

¹ <http://promtools.org/>; November 2014

quality event data. Initially, there was only an abstract research question: *what does the process look like?* A project team was created consisting of a team leader, two business experts, two process analysts, and a system expert. Only the process analysts had previous experience with process mining.

Extraction. The scope of the extraction was limited to all purchase orders created in a single month. For each order, all events related to this order and its requisition were extracted from the SAP system, i.e. including events outside the specified month. Events related to the individual items of an order were not taken into account. This scope ensured that there was sufficient data to get a realistic overview of the workload of the entire process, while still being of a manageable size. The extracted event data contained hundreds of thousands of events related to thousands of orders.

Explorative Analysis Iteration. The first analysis iteration focussed on getting general insight into the process. In the *data processing* stage, multiple event logs were created. A view with the order as a case notion and the event types as the event classes was used for most analyses. Other views were created as well, e.g. the resource or supplier as a case notion for social network analysis. *Part-of* aggregation was used to divide the process into three sub-processes (related to requisitions, orders and deliveries) to obtain more structured models. Filtering was used for various purposes, e.g. simplifying the process by focussing on the events belonging to one sub-process.

The main activities performed during the *mining & analysis* stage were process analytics and discovery, with minor use of enhancement. Visual analytics and statistics provided basic insights related to e.g. the number of events per order or the case duration. The results showed clear weekly patterns and batch processing, such as orders being sent to the suppliers on a specific day. Process discovery on the aggregated log and the filtered logs of the sub-processes returned fact-based process models, which were enhanced with time information to show several bottlenecks.

We discuss some of the tools used during the data processing and mining stages. Event logs for different views are created using the XESame toolkit or by simply importing a CSV file (*CSVi*) in ProM [19]. Event log filtering is available through various plug-ins, e.g. *Filter Log by Attribute Values (FLAV)* or *using Simple Heuristics (uSH)*. The sub-process aggregation was performed using a custom script. For process analytics, the *Log Visualizer (LV)* provides basic statistics and can be used to inspect individual cases and events. Using the *Dotted Chart (DC)* plug-in [17], the events and cases of the log are visualised against time, revealing time patterns, concept drift and batch processing. There are many process discovery algorithms, but here the *Inductive visual Miner (IvM)* [10] was mainly used because it is fast and produces structured models that can be enhanced or analysed by other plug-ins. The *Replay a Log for Performance/Conformance Analysis (RLP/C)* [3] plug-in was used to enhance a model with time or model quality information. The *Feature Prediction (FP)* plug-in enriches the event log with additional attributes, e.g. case duration, and provides a general framework for deriving and correlating process characteristics [11].

The evaluation of the results by business experts, done without involving the process analysts, led to a clear definition of three separate processes (i.e. emergency orders,




Fig. 5: Two variants of cancelling an emergency order, generally mutually exclusive.

regular orders, and repair and warranty orders), and a list of very concrete research questions, which were mostly answered using process analytics techniques. In the following, we discuss how the three processes were investigated.

Analysis iteration - Emergency Orders. We started a new analysis iteration with the refined research question: *what is the process model of emergency orders according to the events recorded?* During the *data processing*, the event log was filtered on the order type and purchasing organisation to obtain the emergency orders. In addition, we only retained the event classes indicated as relevant by the business experts. During the *mining and analysis* stage, we discovered a structured process model and enhanced the model with conformance results using *IvM* to show the frequencies of cases taking a certain path through the process model. During the *evaluation* stage, business experts identified an interesting combination of four event classes involved in two variants of cancelling an order that are usually mutually exclusive, shown in Fig. 5.

To further investigate the two variants, we started a second analysis iteration with the research question *what has exactly happened during these variants of executions and why did these executions happen?* The *data processing* activities were executed to focus on the two variants of cancelled emergency orders. We used analytic techniques in the *mining and analysis* stage to view variants of the remaining cases (using the ProM plug-in *Explore Variants (EV)*), while examining their corresponding complete executions in the original event log, shown by Fig. 6. We observed that the cancellation of some orders happened closely before the cancellation of the corresponding requisitions, which indicates users had to start two separate workflows to cancel the order and its related requisition, i.e. double work. Furthermore, we also observed that users sometimes sent a second order cancellation request even though the initial request was still pending in the workflow. During the *evaluation* stage, the business expert validated the observations by using the SAP system to see the workflows executed. The business expert concluded that an idea for improvement would be to provide the option to cancel an order and its corresponding requisition in one manual action and to inform users if an order is still pending to be cancelled, to prevent double work.

Analysis iteration - Regular Orders. We started another analysis iteration with the refined research question: *what are the differences between the business processes han-*

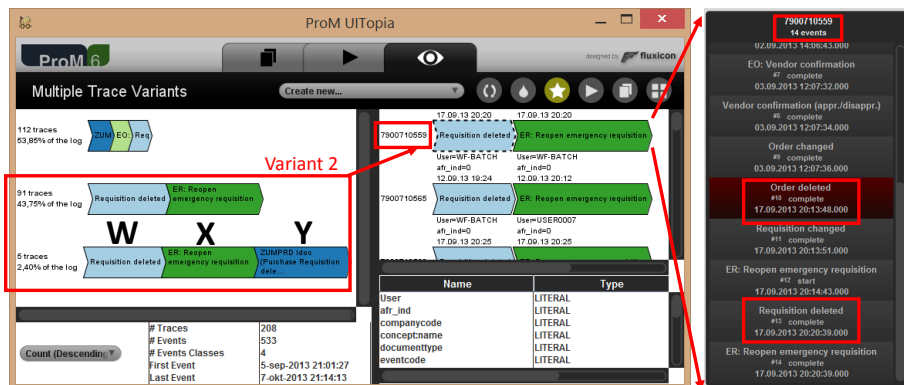


Fig. 6: Inspection of traces containing one variant of cancelling emergency orders.

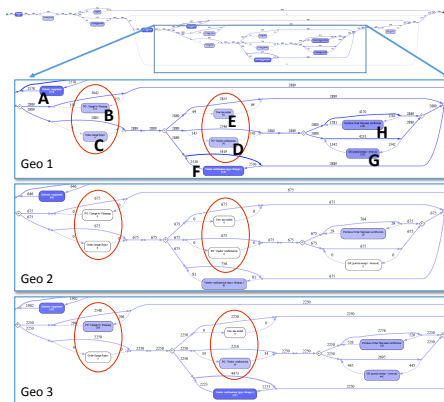


Fig. 7: The logs of three geographies, played-out on the model of the first. Lighter activities are less executed.

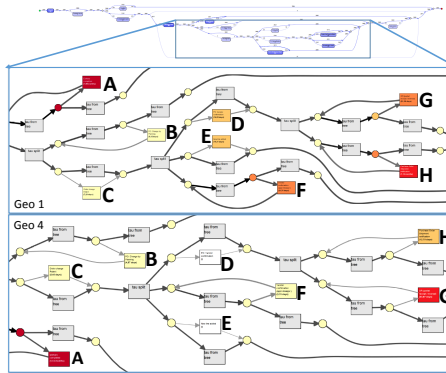


Fig. 8: A performance comparison for two geographies. Darker activities take longer.

...diligent the regular orders of four different geographies? We first used filtering to obtain one event log for each geography. Moreover, we only considered the 80% of cases with mainstream behaviour (known as the “happy flow”) to make the differences more visible and filter out noise. We then used *IvM* to discover process models, and used the plug-in *Show deviations on process tree* to show the deviations between each process model and the other logs. The obvious differences between the mainstream processes, shown in Fig. 7, triggered the refined research question: *what are then the differences between them with respect to the time aspect?* Again, we processed the logs and applied the plug-in *RLP/C*, indicating different bottlenecks as shown in Fig. 8.

Analysis iteration - Repair and Warranty Orders. This iteration’s research question was: *what drives performance differences for the repair and warranty orders for different geographies?* In the data processing stage logs were created for each geography. Structured models were discovered with *IvM* and then manually compared, showing

clear differences between geographies in the likelihood of orders being rejected by the suppliers. Detailed analysis revealed that one geography uses a pricing model where IBM pays a fee to the supplier even if a part is not successfully repaired, while other geographies use a pricing model where they only pay on successful repairs. Finally, the process models showed that in one geography an order confirmation from the supplier is not always recorded. This results in a risk that parts are not available when needed, so a concrete improvement idea is to implement order confirmations with all suppliers.

Process Improvement and Support. At the time of writing, this stage is still ongoing. The process owners are interested in further analysis and there are discussions on implementing the suggested process improvements based on the project findings.

3.2 Discussion

In this section we discuss the lessons learned and the open challenges encountered during the case study, as well as several good practices when using PM².

By applying PM² during the case study, the project team has successfully executed the process mining project. Detailed insights in the performance of IBM's spare part purchasing process have been delivered, achieving the project's goal. In addition, several concrete process improvement ideas have been generated.

A lesson learned is that process mining is most effective when process analysts work closely together with business experts in a highly iterative and interactive manner. This was observed when comparing the data-driven analysis iteration with the later analysis iterations. Initially, there were no business experts in the project team (Issue 1 in Fig. 4) and the transfer of knowledge was limited (Issue 3). The lack of domain knowledge in the data processing stage resulted in incorrect filtering and aggregation (Issue 4), leading to findings that were not representing the processes correctly (Issue 6). Business experts were added to the project team at the end of the data-driven iteration and the analysis that followed was executed in tandem by the business experts and process analysts. This tandem execution was considered to be "a golden combination" according to the stakeholders at IBM, as it led to faster analysis iterations and concrete process improvement ideas.

Another learning point is that a basic understanding of process mining is beneficial for all those involved in the evaluation stage. However, the business experts that joined the team had no previous experience with process mining. Hence, interpretation of the findings was difficult and time-consuming (Issue 5). A full-day process mining workshop was organised to improve process mining understanding for the business experts and business understanding for the process analysts. The evaluation stage became faster and more effective as a result.

We also learned that abstract research questions can be refined during the analysis to obtain valuable insights. It is known that concrete research questions guide a process mining project [18], however sometimes coming up with good research questions at the start of a project is difficult. Data-driven exploration can generate unexpected findings, leading to concrete research questions to explain the findings during further analysis.

A challenge we encountered during the project is that comparing process models is difficult. Some research questions were related to the process execution for different

geographies. To answer these questions, process models were discovered and compared for each geography. However, manual comparison of these models is labour-intensive and existing tool support is limited.

Tool support for interactive analysis is important as well. In ProM, it is currently time-consuming to switch between different views or filter applications on the same data and to compare results. One possible solution would be to use process cubes [2].

A good practice that we identified during the project is to check the event data for errors using statistical techniques and manual inspection. In our project the event timestamps were not always created correctly (Issue 2), which led to incorrect results later in the analysis. Once this was discovered, the event data extraction and the analysis had to be redone. Checking the event data for errors would have prevented this.

Finally, a good practice for the data processing stage is to discuss ideas to simplify future analysis. Identifying sub-processes and process variants helped to reduce the complexity of the models and stimulated specific research questions in our project. Similarly, discussing the importance of events or the expected relations between events helped with the identification of unexpected patterns during the analysis evaluation.

4 Conclusion

In this paper, we have presented the PM² process mining project methodology. PM² is highly iterative and emphasises the need for close collaboration between process analysts and business experts. We have discussed and explained the inputs and outputs of each stage of the methodology, as well as the concrete activities that can be executed.

To illustrate the feasibility of PM² in practise and to provide practical guidance for its application, we performed a process mining project together with IBM. Their spare parts purchasing process was analysed in the case study using PM². We described a range of tools and techniques that were used, the results obtained, the challenges encountered and the lessons learned in the case study. We showed that we ran into issues when we deviated from our methodology, while applying PM² led to valuable insights and concrete process improvement ideas. IBM is currently discussing the implementation of the process improvement ideas generated in the project and considering the use of process mining in a wider and more structural manner.

We plan on continuing to work with IBM to further refine PM², and to provide even more discussion on its use. We also believe that there is still a need for more detailed practical guidance to help process mining practitioners tackle various challenges, e.g. determining how to choose a case notion, or when to use which mining algorithms. Such guidance could be given in the form of some kind of process mining “cookbook”. Additionally, we plan to apply PM² in process mining projects within other organisations. Finally, we encountered challenges for which tool support is missing, e.g. model comparison, and we suggest further research into this.

Acknowledgements. We would like to thank the employees of IBM and in particular: Bart Pastoor, Jorn Kerkhof, Hans Nijssen, and Michiel Kuipers, who have worked closely with us during the case study described in this paper.

References

1. Aalst, W.v.d.: *Process Mining: Discovery, Conformance and Enhancement of Business Processes*. Springer (2011)
2. Aalst, W.v.d.: Process cubes: Slicing, dicing, rolling up and drilling down event data for process mining. In: *Asia Pacific Business Process Management*, pp. 1–22. Springer (2013)
3. Aalst, W.v.d., Adriansyah, A., Dongen, B.v.: Replaying history on process models for conformance checking and performance analysis. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2(2), 182–192 (2012)
4. Bose, R.P.J.C., Mans, R.S., van der Aalst, W.M.P.: Wanna improve process mining results? In: *Computational Intelligence and Data Mining (CIDM), 2013 IEEE Symposium on*. pp. 127–134. IEEE (2013)
5. Bose, R.J.C., van der Aalst, W.M.: Abstractions in process mining: A taxonomy of patterns. In: *Business Process Management*, pp. 159–175. Springer (2009)
6. Bozkaya, M., Gabriels, J., Werf, J.: Process Diagnostics: a Method Based on Process Mining. In: *Information, Process, and Knowledge Management, 2009. eKNOW'09. International Conference on*. pp. 22–27. IEEE (2009)
7. De Weerd, J., De Backer, M., Vanthienen, J., Baesens, B.: A multi-dimensional quality assessment of state-of-the-art process discovery algorithms using real-life event logs. *Information Systems* 37(7), 654–676 (2012)
8. Harmon, P.: *Business process change: A guide for business managers and BPM and Six Sigma professionals*. Morgan Kaufmann (2010)
9. Kurgan, L.A., Musilek, P.: A Survey of Knowledge Discovery and Data Mining Process Models. *The Knowledge Engineering Review* 21(01), 1–24 (2006)
10. Leemans, S., Fahland, D., van der Aalst, W.: Exploring processes and deviations. In: *Business Process Management Workshops*. p. to appear (2014)
11. de Leoni, M., van der Aalst, W.M., Dees, M.: A general framework for correlating business process characteristics. In: *Business Process Management*, pp. 250–266. Springer (2014)
12. Mariscal, G., Marbán, Ó., Fernández, C.: A Survey of Data Mining and Knowledge Discovery Process Models and Methodologies. *The Knowledge Engineering Review* 25(02), 137–166 (2010)
13. Nooijen, E.H., van Dongen, B.F., Fahland, D.: Automatic discovery of data-centric and artifact-centric processes. In: *Business Process Management Workshops*. pp. 316–327. Springer (2013)
14. Rebuge, Á., Ferreira, D.R.: Business Process Analysis in Healthcare Environments: a Methodology based on Process Mining. *Information Systems* 37(2), 99–116 (2012)
15. Rozinat, A., van der Aalst, W.M.P.: Conformance checking of processes based on monitoring real behavior. *Inf. Syst.* 33(1), 64–95 (2008)
16. Shearer, C.: The crisp-dm model: the new blueprint for data mining. *Journal of data warehousing* 5(4), 13–22 (2000)
17. Song, M., van der Aalst, W.: Supporting process mining by showing events at a glance. In: *Workshop on Information Technologies and Systems*. pp. 139–145 (2007)
18. Suriadi, S., Wynn, M.T., Ouyang, C., ter Hofstede, A.H., van Dijk, N.J.: Understanding Process Behaviours in a Large Insurance Company in Australia: A Case Study. In: *Advanced Information Systems Engineering*. pp. 449–464. Springer (2013)
19. Verbeek, H., Buijs, J.C., Van Dongen, B.F., van der Aalst, W.M.: Xes, xesame, and prom 6. In: *Information Systems Evolution*, pp. 60–75. Springer (2011)