Wil M.P. van der Aalst

Process Mining and RPA: How To Pick Your Automation Battles?

Abstract: Robotic Process Automation (RPA) has lowered the threshold for process automation. Repetitive tasks done by people are handed over to software robots. For RPA, there is no need to change or replace the pre-existing information systems. Instead, software robots replace users by interacting directly with the user interfaces normally operated by humans. Actually, RPA can be seen as "the poor man's workflow management solution" because it is cheaper than traditional automation. Therefore, it can be used to automate routine work that would normally not be cost-effective. Process mining plays a key role in deciding what to automate and how. Therefore, RPA is closely related to process mining. Before introducing RPA, one needs to analyze the processes to be automated. Process mining can help to identify promising candidates. Moreover, after RPA has been implemented, process mining can be used to monitor processes and systems even if these use a mixture of RPA, workers, and traditional automation.

Keywords: Robotic Process Automation (RPA), Process Mining, Workflow Management, Business Process Management

1 Introduction

This paper aims to relate *Robotic Process Automation* (RPA) and *process mining* and put both in a historical context. *Workflow Management* (WFM) has been around for several decades [6]. In the mid-nineties, the term *Straight Through Processing* (STP) was used to describe the ultimate goal of WFM: Making operational processes cheaper, faster, and better by avoiding manual intervention. This turned out to be challenging and many WFM projects failed. WFM was subsequently replaced by *Business Process Management* (BPM), which had a broader scope and put more emphasis on management aspects [2, 8, 20]. However, traditional BPM often relied on modeling, leading to a "disconnect" with reality. We have all seen the idealized process models expressed in languages like BPMN that completely failed to capture the real problems. Moreover, the goal should not be to model, but to improve the process at hand. This often did not happen

Wil M.P. van der Aalst, Process and Data Science, RWTH Aachen University, Aachen, Germany & Fraunhofer FIT, Sankt Augustin, Germany. E-mail: wvdaalst@rwth-aachen.de. WWW: www.vdaalst.com

because it would be too expensive to change the information systems or the actual inefficiencies and compliance problems remained invisible.

Some will argue that RPA is not new at all, thereby referring to "screen scraping" (capturing data by reading text from a computer display and transferring it to a new application) and "Taylorism" (i.e., analyzing and improving work processes systematically). However, the combination of process mining and RPA provides new ways of learning and automating routine processes.

The goal of this paper is not to discuss specific RPA or process mining techniques. Instead, we focus on the relations between both worlds and possible interfaces. Therefore, we elaborate on the specifics of event data used in an RPA context. Moreover, we discuss possible use cases for this combination. These show that process mining and RPA complement each other: The former learns about processes and the latter automates them.

In this paper, we first sketch the history of process automation (Section 2). In this context, we position RPA as "The Poor Man's WFM" in Section 3. Then we introduce process mining as a way to exploit event data (Section 4). Section 5 connects process mining and RPA by discussing the specifics of RPA-based event data. This section shows that many design choices are needed to bridge the gap between both. Section 6 elaborates on the interplay between both worlds. Section 7 concludes the paper.

2 A Brief History of WFM and BPM

Since the industrial revolution, productivity has been increasing because of technical innovations, improvements in the organization of work, and the use of information technology [2]. Adam Smith (1723-1790) showed the advantages of the division of labor. Frederick Taylor (1856-1915) introduced the initial principles of scientific management. In the seventies, people like Skip Ellis and Michael Zisman already worked on so-called office information systems, which were driven by explicit process models [2]. Skip Ellis developed the Officetalk system at Xerox PARC in the late 1970s using Information Control Nets (ICN), a variant of Petri nets, to model processes [10]. Also, the office automation system SCOOP (System for Computerizing of Office Processes) developed by Michael Zisman used Petri nets to represent business processes. These systems can be seen as early *Workflow Management (WFM) systems*. However, it took another 15 years until WFM technology was ready to be applied at a large scale. In the mid-nineties, many commercial WFM systems were available and there was the

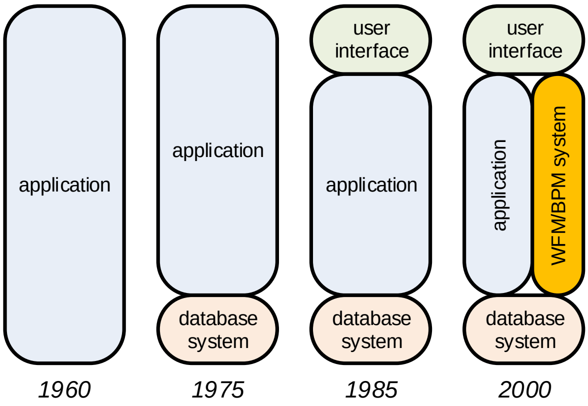

Fig. 1: Positioning of WFM/BPM systems in a historical context (based on [1, 2]).

expectation that WFM systems would be an integral part of any information system [6].

Figure 1 shows the development of information systems over time, explaining the initial great optimism related to WFM technology. Initially, information systems were developed from scratch, i.e., everything had to be programmed, even storing and retrieving data. Soon people realized that many information systems had similar requirements with respect to data management. Therefore, this generic functionality was subcontracted to a database system. Later, generic functionality related to user interaction (forms, buttons, graphs, etc.) was subcontracted to tools that can automatically generate user interfaces. The trend to subcontract recurring functionality to generic tools continued in different areas. Workflow Management (WFM) systems are similar to Database Management (DBM) systems but focus on processes rather than data. In the mid-1990s, many WFM systems became available. These systems focused on automating workflows with little support for process analysis, process flexibility, and process management. Nevertheless, many expected that WFM systems would be as common as DBM systems. However, this did not happen. WFM systems were succeeded by *Business Process Management (BPM) systems* that were broader in scope. The BPM discipline combines knowledge from information technology and knowledge from management sciences and applies this to operational business processes [2, 8, 20]. BPM systems are generic software systems that are driven by explicit process designs to enact and manage operational business processes. Examples of BPM systems include the software products from Pegasystems, Appian, IBM, Bizagi, Oracle, Software AG, TIBCO Software, Bonitasoft, Kofax, and Signavio. However, despite the availability of WFM/BPM systems, process management is not subcontracted to such systems at a scale comparable to DBM systems. The application of "pure" WFM/BPM systems is still limited to specific industries such as banking and insurance. However, WFM/BPM technology is often hidden inside other systems. For example, ERP systems like SAP and Oracle provide workflow engines. Therefore, the landscape is not so clear. Organizations such as Gartner also invent new terms such as "Intelligent Business Process Management Suites" (iBPMS), yet the actual usage of such systems remains limited.

There seem to be three main reasons why the adoption of WFM/BPM technology is low.

- Applying WFM/BPM technology is rather *expensive*. Processes are hard-coded in application software or not supported at all. Many processes also involve software from different vendors, making integration difficult and time-consuming.
- Although the "M" in WFM and BPM refers to "Management", the focus is on modeling and automation rather than management. Traditional WFM/BPM systems fail to learn from the event data they collect.
- Real-life processes are more *complex* than people like to believe. The well-know 80-20 rule applies to processes, i.e., 80% of all cases are rather simple, but explain only 20% of the complexity of the process. The remaining 20% of cases tend to be neglected by software and management, but consume 80% of the resources of an organization.

The above three obstacles for WFM/BPM explain the current interest in Robotic Process Automation (RPA) and process mining.

3 RPA: The Poor Man's WFM

Robotic Process Automation (RPA) is a form of automation using software robots (bots) replacing humans. The three main RPA vendors are UIPath (founded in 2005), Automation Anywhere (founded in 2003), and Blue Prism (founded in 2001). Other vendors include Workfusion, Kryon Systems, Softomotive, Contextor, EdgeVerve, Nice, and Redwood Software. The key difference between RPA and traditional WFM/BPM is that RPA does not aim to replace existing (back-end) information systems. Instead, software robots interact with the existing information systems in the same way as humans do. In traditional WFM/BPM systems, the process is specified precisely and the WFM/BPM system orchestrates the modeled process by implementing simple activities and calling pre-existing applications through Application Programming Interfaces (APIs). In contrast, RPA software interacts with the pre-existing applications through (graphical) user interfaces directly replacing humans, i.e., automation is realized by taking over tasks from workers directly through the user interface. A typical RPA scenario is a sequence of copy-and-paste actions normally performed by a human. *Since there is no need to replace the existing information systems, RPA can be seen as "The Poor Man's WFM".* Figure 2 shows the situation before (left) and after (right) introducing RPA.

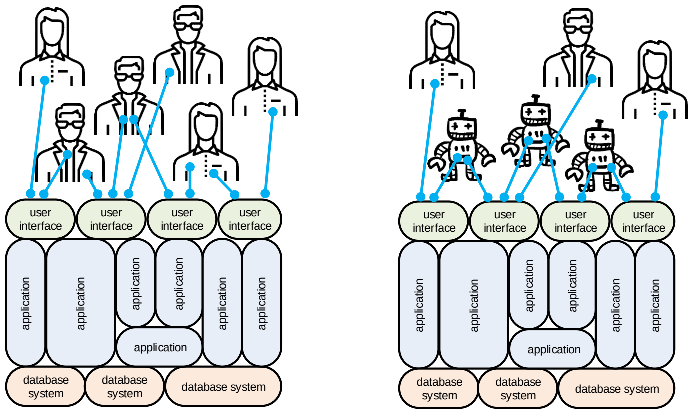

Fig. 2: People tend to be the glue between different applications (left). RPA does not change the "back-end" like in traditional automation (compare with Figure 1). Robots interact with the information systems as if they are people (right).

To understand RPA, it is important to realize that workers and information are "dancing" together. An information system may trigger its users and provide information. Similarly, people start applications and enter information. Consider, for example, the usage of forms. Most forms are partly prefilled with information and users complete the missing information, thereby possibly triggering new actions. Sometimes the user takes the initiative and sometimes the system. When there are multiple information systems, people are often the "glue" between the different parts (cf. Figure 2). See, for example, the scenario where a user copies address information from one information system to another one.

Figure 3 further illustrates the positioning of RPA with respect to the traditional setting and the situation where WFM/BPM software is used. Both RPA and WFM/BPM automate simple tasks and provide the glue between existing information systems. WFM/BPM connects to these systems via the "back-end" using APIs. RPA connects to these systems via the "frontend" using (graphical) user interfaces. In [5], the terms "inside-out" and "outside-in" are

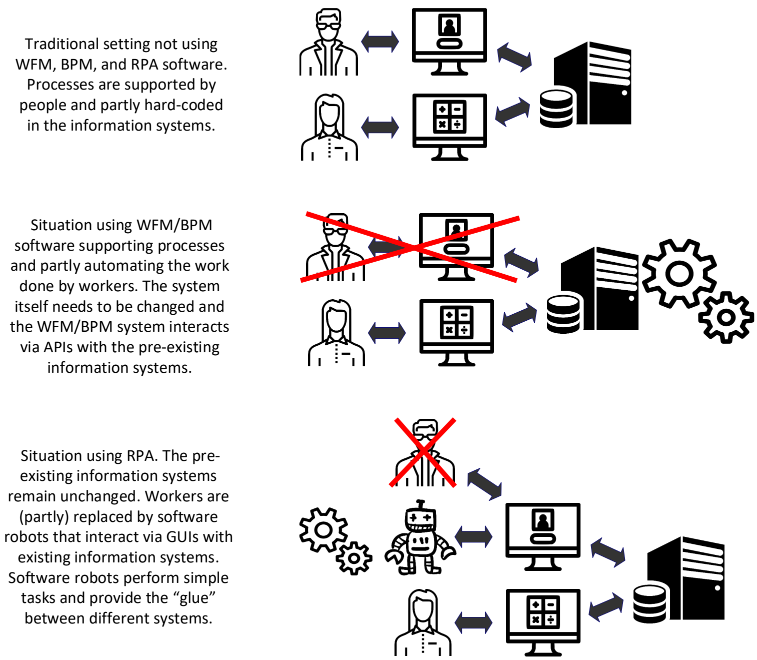

Fig. 3: Three situations: (a) traditional setting, (b) WFM/BPM setting, and (c) RPA setting.

used for respectively the backend WFM/BPM approach and the frontend RPA approach. RPA can be much more cost-effective than traditional automation because the information systems do not need to be changed or replaced. RPA can automate various mundane and routine tasks in the workplace. At the same time, there are some risks. RPA can handle processes and tasks that are repetitive and deterministic. However, these should require little to no judgment and have few exceptions. Technical glitches, exceptions, changing user interfaces, or changing contextual factors provide problems for software robots. There are also obvious security risks, and the lack of communication may conceal important issues (e.g., recurring problems are detected too late). Therefore, sometimes it is better to only use RPA as an "auto-completion tool" where a human still needs to confirm the suggested solution. In [7] the relation between RPA and Enterprise Architecture (EA) is discussed in more detail.

Most of the RPA vendors emphasize the link between RPA and *Artificial Intelligence* (AI) and *Machine Learning* (ML). Classical RPA applications are rule-based and are basically programmed by people. More innovative RPA approaches, sometimes called *cognitive RPA*, aim to learn from humans by observing repetitive tasks [15]. For example, Natural Language Processing (NLP) techniques are used to classify text and routed to the right resource. Image recognition can be used to recognize a button or an edit field, Optical Character Recognition (OCR) can retrieve handwritten text. However, the examples reported are typically focusing on a single well-defined task (like classification). Note that it is relatively easy to recognize buttons, etc. and program actions like clicking such a button and entering a username and password. However, all of this is done without understanding the semantics of the actions. Moreover, AI and ML are rarely used for learning dynamic behavior.

In [17], the authors propose an NLP-based approach that automatically identifies and classifies tasks from textual process descriptions as manual, user, or automated. The goal of the approach is to reduce the effort that is required to identify suitable candidates for robotic process automation. However, the work highly depends on the presence of such descriptions. Often such information is missing, over-simplified, or outdated. Therefore, we focus on the actual behavior observed.

4 Using Process Mining To Pick Your Automation Battles

Process mining techniques use event data to show what people, machines, and organizations are really doing. Process mining provides novel insights that can be used to identify and address performance and compliance problems [3]. Just like spreadsheets can do anything with numbers, process mining can do anything with event data, i.e., it is a generic, domain-independent, technology to improve processes. The application of process mining is much broader than RPA. However, let us first relate both using Figure 4. The diagram sketches the typical Pareto distribution found in event logs. Often, a small percentage of activities account for most of the events and a small percentage of traces variants account for most of the traces [5]. For example, 20% of the activities may account for 80% of the events. Similarly, the 20% most frequent process variants may explain 80% of the cases. Traditional process automation focuses on the most frequent activities and process variants. Only for high-frequent activities and process variants, it may be cost-effective to automate tasks and introduce WFM/BPM. Less frequent activities and process variants need to be handled by workers that exploit human flexibility and creativity. RPA focuses on the middle part, i.e., routine work that is not frequent enough to be automated in the traditional sense. *Process mining is a key technology to identify routine work that can be supported using RPA.*

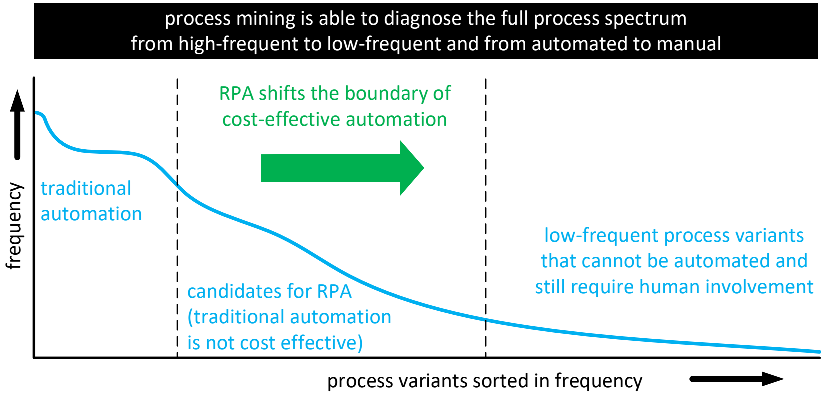

Fig. 4: Relating RPA and process mining (based on [5]).

Therefore, we claim that process mining can be used to pick the "automation battles" that are cost-effective and feasible.

Process mining starts from *event data*, typically stored in an *event log* (see Section 5). An event log views a process from a particular angle. Each event in the log refers to at least (1) a particular process instance (called a *case*), (2) an *activity*, and (3) a *timestamp*. There may be additional event attributes referring to resources, people, costs, etc., but these are optional. With some effort, such data can be extracted from the information systems used by the organization. For example, an SAP system may hold thousands of tables with information about hundreds of processes. In real-life information systems, there may be many possible case identifiers. Therefore, it is often better to use an intermediate logging format where events may refer to any number of objects (cf. Definition 3).

Process mining uses such event data to answer a variety of process-related questions. Process mining techniques such as process discovery, conformance checking, model enhancement, and operational support can be used to improve performance and compliance [3]. Currently, there are over 30 commercial offerings of process mining software (e.g., Celonis, Disco, ProcessGold, myInvenio, PAFnow, Minit, QPR, Mehrwerk, Puzzledata, LanaLabs, StereoLogic, Everflow, TimelinePI, Signavio, and Logpickr). They all can discover so-called *Directly-Follows Graphs* (DFGs) showing frequencies and bottlenecks. DFGs can be seamlessly simplified by removing nodes and edges based on frequency thresholds. DFGs are simple and provide interesting insights, but only provide a starting point. More advanced discovery algorithms like the inductive miner discover better process models, also showing concurrency (e.g., Petri nets, BPMN diagrams, and UML activity diagrams) [3]. Typically, four types of process mining are identified [3].

– *Process discovery*: learning process models from event data. A discovery technique takes an event log and produces a process model without using additional information. An example is the well-known Alpha-algorithm, which takes an event log and produces a Petri net explaining the behavior recorded in the log. Most of the commercial process mining tools first discover DFGs before conducting further analysis.
– *Conformance checking*: detecting and diagnosing both differences and commonalities between an event log and a process model. Conformance checking can be used to check if reality, as recorded in the log, conforms to the model and vice versa. The process model used as input may be descriptive or normative. Moreover, the process model may have been made by hand or learned using process discovery.
– *Process reengineering*: improving or extending the model based on event data. Like for conformance checking, both an event log and a process model are used as input. However, now, the goal is not to diagnose differences. The goal is to change the process model. For example, it is possible to repair the model to better reflect reality. It is also possible to enrich an existing process model with additional perspectives. For example, replay techniques can be used to show bottlenecks or resource usage. Process reengineering yields updated models. These models can be used to improve the actual processes.
– *Operational support*: directly influencing the process by providing warnings, predictions, or recommendations. Conformance checking can be done 'on-the-fly' allowing people to act the moment things deviate. Based on the model and event data related to running process instances, one can predict the remaining flow time, the likelihood of meeting the legal deadline, the associated costs, the probability that a case will be rejected, etc. The process is not improved by changing the model, but by directly providing data-driven support in the form of warnings, predictions, and/or recommendations.

All techniques start from the so-called control-flow perspective, which focuses on the ordering of activities. Then the time perspective (bottlenecks, delays, and frequencies), the data perspective (understanding decisions), and the resource and organization perspective (social networks, roles, and authorizations) are added.

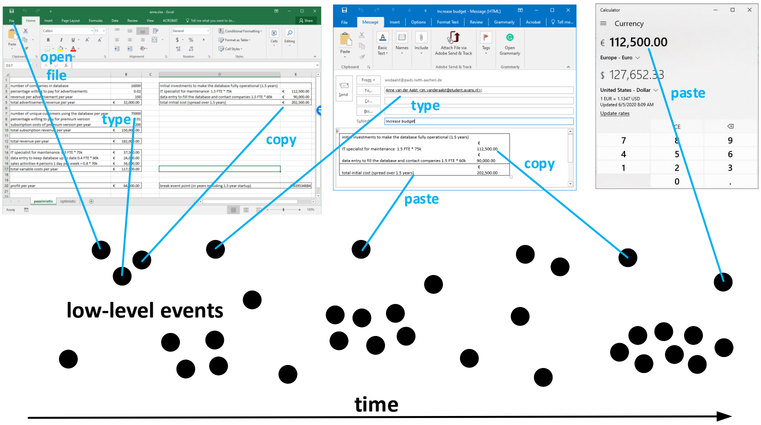

Fig. 5: To learn processes in an RPA context, we need to record all relevant user interactions. Actions performed by users (typing, clicking, etc.) can be seen as low-level events.

5 Formalizing the Input

To be able to learn from people performing activities that should be taken over by software robots, we need to record the interactions between users and the interfaces they use. Figure 5 provides a simplistic illustration where every low-level interaction is represented by a black dot. Such a dot may refer to a mouse click, closing a window, typing an address, selecting a name from a pull-down menu, etc. Existing RPA solutions are able to capture such events. For example, UiPath Studio provides several types of recording (e.g., basic, desktop, web, image). The so-called “Universal Recorder” of Automation Anywhere also supports getting events from various applications (e.g., SAP, Office, and Oracle), web browsers, and operating systems (e.g., windows desktop). Such recordings are mostly used for manual analysis and scenario building. Often screenshots can be recorded to better contextualize events. This helps to understand why users perform certain actions in exceptional situations (e.g., in case of a system failure). Skan CPX is an example of software that is focusing on capturing events using computer vision. Collecting events from the user interface is relatively easy, but it is much more difficult to attach semantics automatically.

The low-level interactions recorded by RPA software can be viewed as events, but cannot be directly used for process mining. The data recorded by RPA software are ad hoc and highly system-dependent. Low-level interactions need to be aggregated and correlated to create event logs that can be used for process mining. To discuss this in a meaningful way, we first need to formalize the elements of an event log.

In its simplest form, an event log can be seen as a multiset of traces where each trace is a sequence of activities, e.g., $L = [\langle a,b,c,e\rangle^{45}, \langle a,c,b,e\rangle^{38}, \langle a,d,e\rangle^{27}]$. This view may be adequate for control-flow discovery, but is too simple for RPA applications that lack a clear case notion. Therefore, we introduce so-called *object-centric event logs* [4]. An event in such a log may refer to any number of objects and attribute values.

Definition 1 (Universes and Events)**.** *To define events, we introduce the following universes:*

- $\mathbb{U}_{ei}$ *is the universe of event identifiers,*
- $\mathbb{U}_{act}$ *is the universe of activity names,*
- $\mathbb{U}_{time}$ *is the universe of timestamps,*
- $\mathbb{U}_{ot}$ *is the universe of object types (also called classes),*
- $\mathbb{U}_{oi}$ *is the universe of object identifiers (also called entities),*
- $type \in \mathbb{U}_{oi} \to \mathbb{U}_{ot}$ *assigns precisely one type to each object identifier,*
- $\mathbb{U}_{omap} = \{omap \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi}) \mid \forall_{ot \in dom(omap)} \forall_{oi \in omap(ot)}\ type(oi) = ot\}$ *is the universe of all object mappings indicating which object identifiers are included per type,*[1]
- $\mathbb{U}_{att}$ *is the universe of attribute names,*
- $\mathbb{U}_{val}$ *is the universe of attribute values,*
- $\mathbb{U}_{vmap} = \mathbb{U}_{att} \not\rightarrow \mathbb{U}_{val}$ *is the universe of value assignments,*[2] *and*
- $\mathbb{U}_{event} = \mathbb{U}_{ei} \times \mathbb{U}_{act} \times \mathbb{U}_{time} \times \mathbb{U}_{omap} \times \mathbb{U}_{vmap}$ *is the universe of events.*

An event $e = (ei, act, time, omap, vmap) \in \mathbb{U}_{event}$ is characterized by a unique event identifier ei, the corresponding activity act, the event's timestamp $time$, and two mappings $omap$ and $vmap$ for respectively object references and attribute values.

Definition 2 (Event Projection)**.** *Given* $e = (ei, act, time, omap, vmap) \in \mathbb{U}_{event}$, $\pi_{ei}(e) = ei$, $\pi_{act}(e) = act$, $\pi_{time}(e) = time$, $\pi_{omap}(e) = omap$, *and* $\pi_{vmap}(e) = vmap$.

1 $\mathcal{P}(\mathbb{U}_{oi})$ is the powerset of the universe of object identifiers, i.e., objects types are mapped onto sets of object identifiers. $omap \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi})$ is a partial function. If $ot \notin dom(omap)$, then we assume that $omap(ot) = \emptyset$.

2 $\mathbb{U}_{att} \not\rightarrow \mathbb{U}_{val}$ is the set of all partial functions mapping a subset of attribute names onto the corresponding values.

Consider a event e with $\pi_{act}(e)$ = "place order" and $\pi_{time}(e)$ = "2020-10-07 08:23:19". $\pi_{omap}(e) \in \mathbb{U}_{ot} \not\rightarrow \mathcal{P}(\mathbb{U}_{oi})$ maps a subset of object types onto sets of object identifiers for an event e. For example, $\pi_{omap}(e)(\mathit{Order}) = \{o_{4567}\}$, $\pi_{omap}(e)(\mathit{Item}) = \{i_{786}, i_{888}, i_{923}\}$, and $\pi_{omap}(e)(\mathit{Payments}) = \emptyset$ (i.e., the place order event e refers to one order, three items, and no payments). $\pi_{vmap}(e) \in \mathbb{U}_{att} \not\rightarrow \mathbb{U}_{val}$ maps a subset of attribute names onto attribute values. For example, $\pi_{vmap}(e)(\mathit{cost}) = 75$ and $\pi_{vmap}(e)(\mathit{location})$ = "Berlin".

An *object-centric event log* is a collection of *partially ordered events* [4]. Event identifiers are unique, i.e., two events cannot have the same event identifier.

Definition 3 (Object-Centric Event Log)**.** *$L = (E, \preceq_E)$ is an event log with $E \subseteq \mathbb{U}_{event}$ and $\preceq_E \subseteq E \times E$ such that:*
- *$\preceq_E$ defines a partial order (reflexive, antisymmetric, and transitive),*
- *$\forall_{e_1,e_2 \in E}\ \pi_{ei}(e_1) = \pi_{ei}(e_2) \;\Rightarrow\; e_1 = e_2$, and*
- *$\forall_{e_1,e_2 \in E}\ e_1 \preceq_E e_2 \;\Rightarrow\; \pi_{time}(e_1) \leq \pi_{time}(e_2)$.*

Object-centric event logs generalize the traditional event log notion where each event has precisely one case identifier. We can mimic such logs using a special object type $\mathit{Case} \in \mathbb{U}_{ot}$ such that $|\pi_{omap}(e)(\mathit{Case})| = 1$ for any event $e \in E$. Since traditional process mining techniques assume this, it is common practice to convert event data with events referring to a variable number of objects to classical event logs by "flattening" the event data. Assume that we take a specific object type as a case identifier. If an event has multiple objects of that type, then we can simply create one event for each object. If an event has no objects of that type, then we simply omit the event. If an event has precisely one object of the selected type, then we keep that event. Hence, by selecting an object type as the case identifier, we can "flatten" the log and apply standard process discovery and conformance checking techniques.

Let us assume that we want an event log $L = (E, \preceq_E)$ in order to apply various process mining techniques in an RPA setting as described before. *How to get such an event log in the context of RPA?* As illustrated in Figure 6 we cannot directly use the low-level events and need to aggregate and correlate user interactions.

Aggregation. First, we need to decide at what level we would like to record user activities. Examples of *low-level activities* include click, double click, select item, type text, copy, paste, close window, etc. It is possible to see each of these as individual events or they can be grouped into *higher-level events* such as filling out a form. It is also possible to think of hierarchical recordings having multiple levels. Only low-level events can be seen as atomic. For example, it may take a few minutes to fill out a form in one system while gathering information from

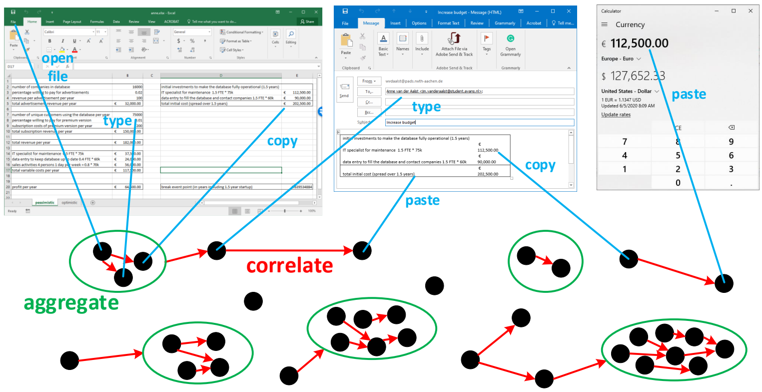

Fig. 6: Low-level user-interactions need to be aggregated and correlated to build event logs.

other systems. How to segment low-level events and create such a hierarchy is situation dependent.

Correlation. Related to aggregation is the topic of correlation. A user may use different systems at the same time and work on multiple cases. Copying an address from SAP and pasting the address in a web-form are clearly related. However, the user may also simply type the address in the web-form manually (while looking at the SAP screen). Correlation is often based on comparing values, e.g., the zip code “D-52074” or URL “pads.rwth-aachen.de” appearing in two different windows. In object-centric event logs, events can have multiple object identifiers without picking a specific case notion. This provides the required flexibility. However, the mapping from values and identifiers in the user interface to event attributes and objects remains something that is application and situation dependent. This is unavoidable given the ad-hoc nature of low-level user-interaction recordings.

The process sketched in Figure 6 is far from trivial. Earlier, we defined events to be of the form: $e = (ei, act, time, omap, vmap) \in \mathbb{U}_{event}$. The correlation between events (aggregated or not) needs to take place via *omap* (i.e., the objects the event is referring to). For example, events e_{356} and e_{412} are related because $\pi_{omap}(e_{356})(Zip) = \{\text{“D-52074”}\}$ and $\pi_{omap}(e_{412})(Zip) = \{\text{“D-52074”}\}$. Events may have standard attributes and object types, e.g., *vmap* and *omap* may contain mandatory information on user name, computer name, window id, session id, etc. When aggregating events, it makes sense to have two times ($time_{start}$ and $time_{end}$) for each event. Similarly, it may make sense to split *omap* and *vmap* into input and output, i.e., $omap_{in}$, $vmap_{in}$, $omap_{out}$, and $vmap_{out}$. This way one can infer create, read, update, and delete actions in forms. For example, if $omap_{in}(Price) = 500$ and $omap_{out}(Price) = 600$, then we know that the price was increased by 100. Hence, high-level events could be of the form $e = (ei, act, time_{start}, time_{end}, omap_{in}, vmap_{in}, omap_{out}, vmap_{out})$ to better capture the duration, input, and output. However, the resulting log can still be viewed as an object-centric event log that can be used to generate different flattened event logs depending on the questions that need to be answered.

The above discussion shows that it is far from trivial to create meaningful event logs from low-level user interactions. However, this step is essential when deciding on what to automate.

6 On The Interplay Between Process Mining and RPA

The connection between process mining and RPA was first discussed in [5]. In [13] it is shown how a commercial process mining system like Celonis can be used to support the whole lifecycle of RPA initiatives. In [16] the term Robotic Process Mining (RPM) is introduced to refer to "a class of techniques and tools to analyze data collected during the execution of user-driven tasks in order to support the identification and assessment of candidate routines for automation and the discovery of routine specifications that can be executed by RPA bots". The authors propose a framework and RPM pipeline combining RPA and process mining, and identify challenges related to recording, filtering, segmentation, simplification, identification, discovery, and compilation. In [12] a RPA-rule deduction approach is presented combining process mining and captured user behavior in the form of Input-Output (IO) logs.

As mentioned earlier, *the scope of process mining extends far beyond RPA* since it also covers process steps fully handled by humans or automated in the traditional way. However, RPA is not just related to process mining and influences the broader Business Process Management (BPM) discipline. The role of RPA in BPM architectures was already elaborated in [15]. The paper focuses on the use of RPA in public administration (e.g., automatically classifying documents). In [19] a review of the state of the art in RPA and 15 challenges are given. Both papers identify a gap between the inflated expectations and the actual tool support provided. RPA vendors tend to present general purpose artificial intelligence and machine learning techniques as breakthroughs in process automation. However, process mining shows that even structured processes like Purchase-To-Pay (P2P) and Order-To-Cash (O2C) tend to be much more complex than anticipated. Such reality checks are essential to make proper RPA decisions.

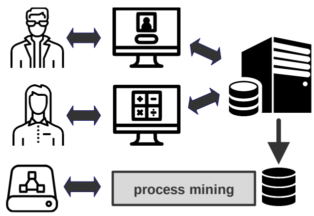

(a) Traditional process mining using event data extracted from the information systems supporting the process.

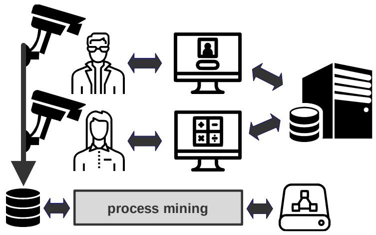

(b) Process mining using event data obtained by observing the workers using the information systems.

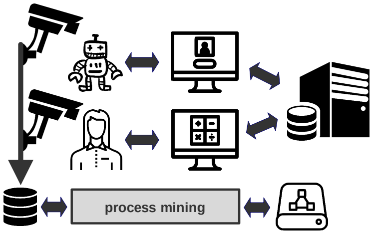

(c) Process mining after introducing RPA. Event data are obtained by observing workers and robots.

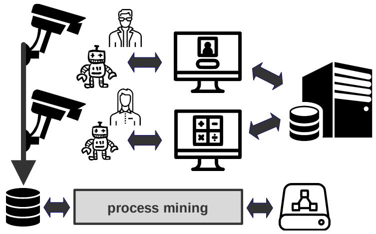

(d) Process mining after introducing RPA using an adaptive distribution of work over workers and robots.

Fig. 7: Process mining can be used before and after the introduction of RPA. Robots and workers use the same (graphical) user interfaces and the role-distribution may be flexible and change over time. Fortunately, process mining provides a holistic view of the processes at hand and interplay between robots and workers.

To conclude the paper, we discuss the *relationship between process mining and RPA* in more detail using Figure 7. In Figure 7(a), the traditional usage of process mining is described. In this scenario, event data are extracted from the information systems supporting the process. Workers are not observed directly. In Figure 7(b), process mining is applied to event data collected directly from the (graphical) user interfaces, i.e., the interactions between workers and information systems are directly recorded. This scenario is particularly useful in the phase before RPA is introduced. Process mining can be used to detect routine work that can be automated by mimicking the behavior of workers. Rather than manually programming robots, process discovery can be used to configure the robots correctly. In Figure 7(c), process mining is used after introducing RPA. Part of the work formerly done by workers is now done by software robots. In this scenario, process mining is used to check whether the processes run as planned. If a software robot malfunctions due to technical glitches, exceptions, changing user interfaces, or changing contextual factors, then this can be detected using conformance checking techniques. Note that a lack of human oversight of the work produced by robots constitutes a real risk of catastrophic outcomes. Figure 7(d) describes the most advanced scenario. In this scenario, the work is flexibly distributed over workers and software robots. For example, tasks are initially performed by robots and are escalated to workers the moment there is a complication of exception. Similarly, workers can hand off work to robots using an “auto-complete” option. Moreover, the RPA solution may adapt due to changes in the underlying process (e.g., concept drift).

7 Conclusion

Process automation has a long history. WFM and BPM systems have been around for decades, but their application is limited to high-volume structured processes. RPA has lowered the threshold for automation. The phrase “RPA is the Poor Man’s WFM” (coined in this paper) illustrates this. Due to RPA, it is possible to automate many mundane repetitive routines in an economically viable manner. Process mining helps to identify process fragments that can be supported using RPA. This is the reason that process mining and RPA vendors have joined forces. For example, in October 2019, process mining vendor ProcessGold was acquired by RPA vendor UiPath. Similarly, vendors like Celonis started to support “task mining” and “action automation” (using the action engine) to boost RPA-related capabilities. Skan is combining computer vision and machine learning capabilities with process mining.

According to Deloitte and EY, up to 30 to 50% of RPA projects fail, and most are more expensive and time-consuming than planned [9, 21]. Process mining can be used to avoid such failures. As Figure 4 shows, the scope of process mining includes everything from routine activities and processes automated using WFM, BPM, and RPA to one-of-a-kind activities and processes that require human interventions and creativity. Moreover, process mining helps to support the different phases of RPA as highlighted in Figure 7.

Hence, there is huge potential. However, many challenges need to be addressed. Actually, the uptake of RPA triggers many interesting research questions.

– *What event data to store and how to structure these?* Computer vision, image recognition, OCR, and NLP can be used to capture events. However, how to add semantics and how to decide that event are relevant for the process.
– *What characteristics make processes suitable to be supported by RPA?* Many RPA projects fail because automation turns out to be infeasible or they try to automate processes that are too infrequent or changing too fast. RPA

needs to be approached more systematically using data-driven cost-benefit analyses.

– *How to control software robots and avoid security, compliance, and economic risks?* The ISO 10218-1 standard defines safety requirements for industrial robots. Such standards are missing for software robots. However, malfunctioning robots (e.g., due to changing circumstances) may have devastating effects for an organization (e.g., leaking sensitive information or making costly decisions).
– *How can software robots and people seamlessly work together?* The border between tasks best done by humans and tasks best done by machines will continue to shift. Intelligence amplification (IA) (also referred to machine augmented intelligence or enhanced intelligence) aims to enhance the human worker using AI. This results in processes where robots and people seamlessly work together.

Process mining plays a key role in answering these questions and can be placed in a larger context where work is distributed among machines and people.

The frontier between the tasks performed by humans and those performed by machines and algorithms is continuously moving and changing global labor markets. In [14] three waves of automation (algorithmic, augmentation, and autonomous) are predicted replacing much of the work previously done by people. In [11], Frey and Osborne predict the degree of computerization for 702 occupations. They estimate that 47 percent of jobs in the US will be replaced by (software) robots. In [18] three types of roles are identified: stable roles (work that remains), new roles (new types of work that did not exist before), and redundant roles (work that is taken over by e.g. robots). These broader trends highlight the economic and social impact of RPA and process mining.

Acknowledgments: We thank the Alexander von Humboldt (AvH) Stiftung for supporting our research.

Bibliography

[1] W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21–66, 1998.

[2] W.M.P. van der Aalst. Business Process Management: A Comprehensive Survey. *ISRN Software Engineering*, pages 1–37, 2013. doi:10.1155/2013/507984.

[3] W.M.P. van der Aalst. *Process Mining: Data Science in Action*. Springer-Verlag, Berlin, 2016.

[4] W.M.P. van der Aalst. Object-Centric Process Mining: Dealing With Divergence and Convergence in Event Data. In P.C. Ölveczky and G. Salaün, editors, *Software Engineering and Formal Methods (SEFM 2019)*, volume 11724 of *Lecture Notes in Computer Science*, pages 3–25. Springer-Verlag, Berlin, 2019.

[5] W.M.P. van der Aalst, M. Bichler, and A. Heinzl. Robotic Process Automation. *Business and Information Systems Engineering*, 60(4):269–272, 2018.

[6] W.M.P. van der Aalst and K.M. van Hee. *Workflow Management: Models, Methods, and Systems*. MIT Press, Cambridge, MA, 2004.

[7] G. Auth, C. Czarnecki, and F. Bensberg. Impact of Robotic Process Automation on Enterprise Architectures. In C. Draude, M. Lange, and B. Sick, editors, *INFORMATIK 2019*, volume 295 of *Lecture Notes in Informatics*, pages 59–65. GI, 2019.

[8] M. Dumas, M. La Rosa, J. Mendling, and H. Reijers. *Fundamentals of Business Process Management*. Springer-Verlag, Berlin, 2013.

[9] D. Dutta, A. Gillard, and G. Kaczmarskyj. *Get Ready for Robots: Why Planning Makes the Difference Between Success and Disappointment*. Ernst and Young, 2016. https://eyfinancialservicesthoughtgallery.ie/wp-content/uploads/2016/11/ey-get-ready-for-robots.pdf.

[10] C.A. Ellis. Information Control Nets: A Mathematical Model of Office Information Flow. In *Proceedings of the Conference on Simulation, Measurement and Modeling of Computer Systems*, pages 225–240, Boulder, Colorado, 1979. ACM Press.

[11] C.B. Frey and M.A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114(C):254–280, 2017.

[12] J. Gao, S.J. van Zelst, X. Lu, and W.M.P. van der Aalst. Automated Robotic Process Automation: A Self-Learning Approach. In H. Panetto, C. Debruyne, M. Hepp, D. Lewis, C.A. Ardagna, and R. Meersman, editors,

On the Move to Meaningful Internet Systems, International Conference on Cooperative Information Systems (CoopIS 2019), volume 11877 of *Lecture Notes in Computer Science*, pages 95–112. Springer-Verlag, Berlin, 2019.

[13] J. Geyer-Klingeberg, J. Nakladal, F. Baldauf, and F. Veit. Process Mining and Robotic Process Automation: A Perfect Match. In *Proceedings of the Industrial Track at the 16th International Conference on Business Process Management (BPM 2018)*, pages 124–131, 2018.

[14] J. Hawksworth, R. Berriman, and S. Goel. Will Robots Really Steal Our Jobs? An International Analysis of the Potential Long Term Impact of Automation. Technical report, PricewaterhouseCoopers, 2018.

[15] C. Houy, M. Hamberg, and P. Fettke. Robotic Process Automation in Public Administrations. In R. Michael, H. Sebastian, D. Ratz, D. Richter, and E. Schweighofer, editors, *Digitalisierung von Staat und Verwaltung*, volume 291 of *Lecture Notes in Informatics*, pages 62–74. GI, 2019.

[16] V. Leno, A. Polyvyanyy, M. Dumas, M. La Rosa, and F.M. Maggi. Robotic Process Mining: Vision and Challenges. *Business and Information Systems Engineering*, 2020.

[17] H. Leopold, H. van der Aa, and H.A. Reijers. Identifying Candidate Tasks for Robotic Process Automation in Textual Process Descriptions. In J. Gulden, I. Reinhartz-Berger, R. Schmidt, S. Guerreiro, W. Guédria, and P. Bera, editors, *Enterprise, Business-Process and Information Systems Modeling*, pages 67–81. Springer-Verlag, Berlin, 2018.

[18] T.A. Leopold, V. Ratcheva, and S. Zahidi. The Future of Jobs Report. Technical report, Centre for the New Economy and Society, World Economic Forum, 2018.

[19] R. Syeda, S. Suriadia, M. Adamsa, W. Bandaraa, S. Leemans, C. Ouyang, A. ter Hofstede, I. van de Weerd, M. Wynn, and H.Reijers. Robotic Process Automation: Contemporary themes and challenges. *Computers in Industry*, 115(1):103162, 2020.

[20] M. Weske. *Business Process Management: Concepts, Languages, Architectures.* Springer-Verlag, Berlin, 2007.

[21] D. Wright, D. Witherick, and M. Gordeeva. *The Robots Are Ready. Are You? Untapped Advantage in Your Digital Workforce.* Deloitte, 2017. https://www2.deloitte.com/uk/en/pages/consulting/articles/the-robots-are-ready-are-you.html.