CrossMark

Robotic Process Automation

Wil M. P. van der Aalst · Martin Bichler · Armin Heinzl

1 Introduction

A foundational question for many BISE (Business and Information Systems Engineering) authors and readers is "What should be automated and what should be done by humans?" This question is not new. However, developments in data science, machine learning, and artificial intelligence force us to revisit this question continuously. Robotic Process Automation (RPA) is one of these developments. RPA is an umbrella term for tools that operate on the user interface of other computer systems in the way a human would do. RPA aims to replace people by automation done in an "outside-in'' manner. This differs from the classical "inside-out" approach to improve information systems. Unlike traditional workflow technology, the information system remains unchanged. Gartner defines Robotic Process Automation (RPA) as follows: "RPA tools perform [if, then, else] statements on structured data, typically using a combination of user interface interactions, or by connecting to APIs to drive client servers, mainframes or HTML code. An RPA tool operates by mapping a process in the RPA tool language for the software robot to follow, with runtime allocated to execute the script by a control dashboard." (Tornbohm 2017). Hence, RPA tools aim to reduce the burden of repetitive, simple tasks on employees (Aguirre and Rodriguez 2017).

Commercial vendors of RPA tools have witnessed a surge in demand. Moreover, many new vendors entered the market in the last 2 years. This is no surprise as most organizations are still looking for ways to cut costs and quickly link legacy applications together. RPA is currently seen as a way to quickly achieve a high Return on Investment (RoI). There are dedicated RPA vendors like AutomationEdge, Automation Anywhere, Blue Prism, Kryon Systems, Softomotive, and UiPath that only offer RPA software (Le Clair 2017; Tornbohm 2017). There are also many other vendors that have embedded RPA functionality in their software or that are offering several tools (not just RPA). For example, Pegasystems and Cognizant provide RPA next to traditional BPM, CRM, and BI functionality. The goal of this editorial is to reflect on these developments and to discuss RPA research challenges for the BISE community.

Prof. Dr. Ir. W. M. P. van der Aalst (✉)
RWTH Aachen, Lehrstuhl für Informatik 9 - Process and Data Science (PADS), 52056 Aachen, Germany
e-mail: wvdaalst@pads.rwth-aachen.de

Prof. Dr. M. Bichler
Department of Informatics, Decision Sciences and Systems, Technical University of Munich (TUM), Boltzmannstr 3, 85748 Munich, Germany
e-mail: bichler@in.tum.de

Prof. Dr. A. Heinzl
Chair of General Management and Information Systems, University of Mannheim, 68161 Mannheim, Germany
e-mail: heinzl@uni-mannheim.de

2 Motivating Example

The first author recently changed employer (from Eindhoven University of Technology to RWTH Aachen University) and relocated from the Netherlands to Germany, thereby changing health insurance, pension system, tax system, banks, etc. Anyone that has made such a move will be aware of the many databases one's name is in and that changing status and address turns out to be a manual task. It often involves making multiple phone calls to get things right. Even within the same organization, there are

often multiple information systems containing information about a single person.

Consider for example Eindhoven University of Technology (TU/e) and RWTH Aachen University. Both universities have information systems for salary administration, pension, personnel, teaching, research, projects, finance, etc. When a person leaves TU/e or is hired by RWTH, information needs to be added, removed, or changed in all of these systems. At a larger university, hundreds of new staff members are hired each year, and it is impossible to integrate all information into a single system. For example, SAP is used for finance and Moodle is used as a learning management system, but these two systems are completely disconnected. This requires people entering information into multiple systems and trying to maintain consistency. However, the work is fairly simple and tedious. RPA provides agents that interact with different information systems thus partly replacing humans. Using Artificial Intelligence (AI) and Machine Learning (ML), this can be done in a fairly robust manner. For example, the RPA agent should continue to work properly when the web interface of some application changes. The use of AI/ML help to interpret changing interfaces, thus making it very different from traditional "screen scraping" (also called web scraping or web harvesting). AI/ML can also help to mimic human behavior in combining different applications (solution recipes).

To understand the relevance of RPA, we consider Fig. 1. This figure shows the "long tail of work". The x-axis shows the different types of cases. Two cases are of the same type if they are similar and can be handled in the same way. The y-axis shows the frequencies of these case types. Typically, one sees a Pareto distribution. This means that 80% of the cases can be explained by 20% of the case types. This means that there are many case types that are rather rare. Automation aims to address the most frequent cases types (say 20% of all case types). Less frequent cases are not considered because automation is too expensive. Costs further increase when different proprietary systems need to be integrated. Therefore, the remaining 20% of the cases is often handled manually by humans entering information repeatedly and making decisions. In such settings, humans serve as the "glue" between different IT systems. However, these remaining 20% of the cases, cover 80% of the case types and are much more time-consuming than the frequent ones. Using RPA it is possible to support the middle part by having agents that interact with the different information systems as if they were human. This is not always possible or economically viable. Therefore, the "end of the long tail" (right-hand-side of Fig. 1) still needs to be handled by human workers.
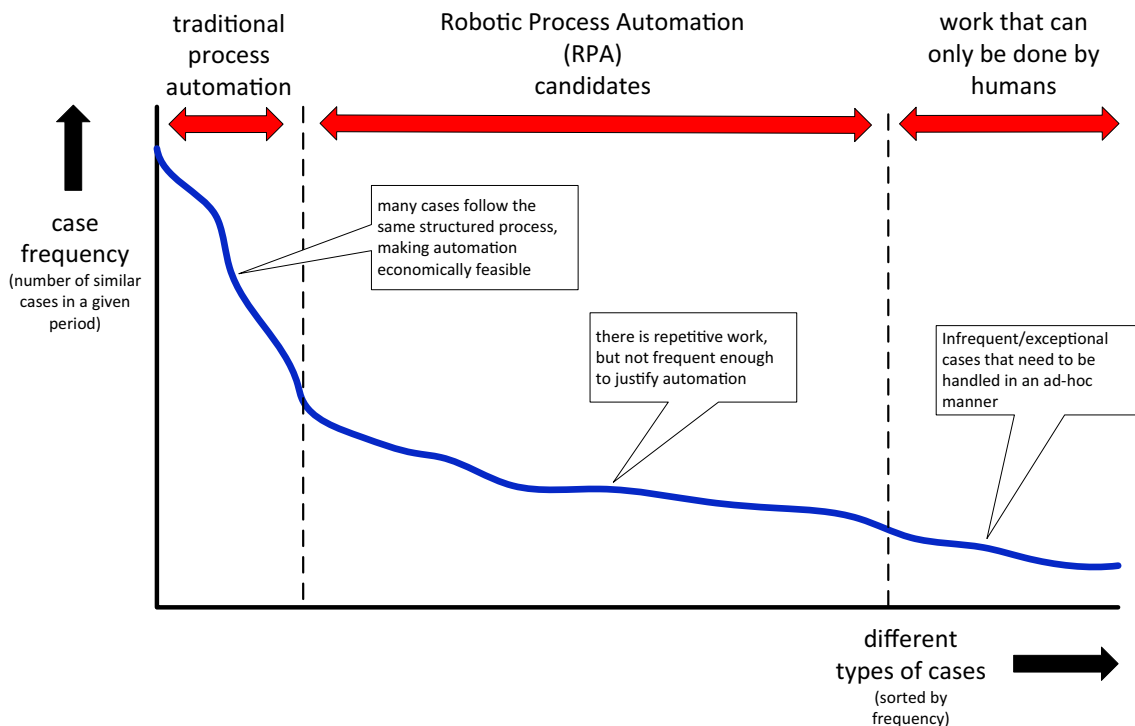

Fig. 1 Positioning RPA

3 Straight Through Processing (STP) Reinvented?

According to Gartner, RPA tools are at the "peak of inflated expectations" in the so-called Hype Cycle (Kerremans 2018). There are many vendors offering RPA tools, including AutomationEdge, Automation Anywhere, Blue Prism, Cognizant, Conduent, Kofax, Kryon Systems, Pegasystems, Softomotive, and UiPath. The sudden uptake suggests that RPA is new. However, "process management veterans" will remember the hype around Straight Through Processing (STP) in the mid-nineties (van der Aalst and van Hee 2002; ter Hofstede et al. 2010). STP was first used in the financial industry. The term refers to processes that can be performed without any human involvement. For example, allowing information that has been electronically entered to be transferred from one party to another in the settlement process without manually re-entering the same information repeatedly. This was one of the key sales features of Workflow Management (WfM) systems in the nineties (e.g., Staffware). Unfortunately, STP turned out to be applicable to only a few processes (only the left-hand-side of Fig. 1). Therefore, WfM systems evolved into Business Process Management (BPM) systems focusing more on the management aspects. BPM projects are often considered too expensive because of "inside-out" approach that is used (the system has to be developed from scratch, and system integration is expensive).

So what is new? RPA differs from STP in two ways. First of all, RPA uses an "outside-in'' approach where the existing information systems remain unchanged. Instead of redesigning the system, humans are replaced by agents. Second, RPA aims to be robust with respect to changes of the underlying information systems. When the layout of an electronic form changes, but the key content remains unchanged, the RPA software should adapt just as humans do. According to Forrester, BPM has a legacy of long implementations and fuzzy business cases (Le Clair 2017) and RPA aims to do the opposite (quick wins that require little investment).

4 Data-Driven Intelligence

Today, call centers and large "administrative factories" use RPA. To achieve more widespread adoption, RPA needs to become "smarter". The promise is that with the use of Artificial Intelligence (AI) and Machine Learning (ML) techniques, more complex and less defined tasks can be supported. Humans learn by doing and learn from a coach. The goal is that RPA tools learn in the same way. For example, by observing human problem resolving capabilities (e.g., in case of system errors, unexpected system behavior, changing forms) RPA tools can adapt and handle non-standard cases.

In addition, the interplay between RPA agents and humans is interesting. When a case turns out to be exceptional, the RPA agent may handover the case to a human. By observing the human handling complex cases, the RPA system can learn. There is also an obvious link with process mining (van der Aalst 2016; Kerremans 2018). For example, RPA vendor UiPath and process mining vendor Celonis collaborate to automatically visualize and select processes with the highest automation potential, and subsequently, build, test, and deploy RPA agents driven by the discovered process models. Other vendors report similar use cases. Process discovery can be used to learn processes "by example" and process fragments that are suitable for RPA can be detected subsequently. Conformance checking can be used to check for deviations, predict problems, and signal handovers from agents to humans.

Of course, one should be very careful. RPA agents mimicking people can start making incorrect decisions because of contextual changes. This may remain unnoticed for some time, leading to disastrous situations. There are also ethical and security risks when RPA agents impersonate people.

The uptake of RPA provides many interesting research questions. Some of them are not new, but addressing them has become more urgent. Some example questions include (van der Aalst and van Hee 2002; Chandler et al. 2017; Le Clair 2017; Kirchmer 2017; Tornbohm 2017):

1. What characteristics make processes suitable to be supported by RPA?
2. How to let RPA agents learn? How to coach RPA agents?
3. How to control RPA agents and avoid security, compliance, and economic risks?
4. Who is responsible when an RPA agent "misbehaves"?
5. How can RPA agents and people seamlessly work together?

The above questions are key topics for the BISE community. Therefore, the uptake of RPA provides interesting research opportunities. The BISE community could, and also should, play an active role in driving RPA research.

References

Aguirre S, Rodriguez A (2017) Automation of a business process using robotic process automation (RPA): a case study. Appl Comput Sci Eng Commun Comput Inf Sci. https://doi.org/10.1007/978-3-319-66963-2_7

Chandler S, Power C, Fulton M, van Nueten N (2017) Who minds the bots? Why organisations need to consider risks related to Robotic Process Automation. PricewaterhouseCoopers, London

Kerremans M (2018) Gartner market guide for process mining. Report G00353970. Gartner

Kirchmer M (2017) Robotic process automation – pragmatic solution or dangerous illusion? bpm-d.com/bpm-d-exhibiting-at-btoes-2/. Accessed 22 April 2018

Le Clair C (2017) The Forrester wave: robotic process automation: the 12 providers that matter most and how they stack up. Forrester, Cambridge

ter Hofstede AHM, van der Aalst WMP, Adams M, Russell N (2010) Modern business process automation: YAWL and its support environment. Springer, Heidelberg

Tornbohm C (2017) Gartner market guide for robotic process automation software. Report G00319864. Gartner

van der Aalst WMP (2016) Process mining: data science in action. Springer, Heidelberg

van der Aalst WMP, van Hee KM (2002) Workflow management: models, methods, and systems. MIT Press, Cambridge