prof.dr.ir Wil van der Aalst

RWTH Aachen University

Hybrid Intelligence

There used to be a clear separation between tasks done by machines and tasks done by people. Applications of machine learning in speech recognition (e.g., Alexa and Siri), image recognition, automated translation, autonomous driving, and medical diagnosis, have blurred the classical divide between human tasks and machine tasks. Although current AI technology outperforms humans in many areas, tasks requiring common sense, contextual knowledge, creativity, adaptivity, and empathy are still best performed by humans. Hybrid Intelligence (HI) blends human intelligence and machine intelligence to combine the best of both worlds.

Hybrid Intelligence

To Automate or Not to Automate, That is the question

Machine Learning (ML) and Robotic Process Automation (RPA) have lowered the threshold to automate tasks previously done by humans. Yet organizations are struggling to apply ML and RPA, effectively causing many digital transformation initiatives to fail. Process Mining (PM) techniques help to decide what should be automated and what not. Interestingly, most processes work best using a combination of human and machine intelligence. Therefore, we relate Hybrid Intelligence (HI) to process management and process automation using RPA and PM.

As Niels Bohr once said "It is difficult to make predictions, especially about the future" and, of course, this also applies to process management and automation. In 1964, the RAND Corporation published a report with predictions about technological development based on the expectations of 82 experts across various fields. For 1980, the report predicted that there would be a manned landing on Mars and families would have robots as household servants. We are still not any way close to visiting Mars and, 40 years later, we only have robot vacuum cleaners. For 2020, the expectation was that we would breed apes and other animals to carry out our daily chores. None of this happened.

When it comes to predictions about Artificial Intelligence (AI) we can witness periods with great optimism and periods with great skepticism (called "AI winters"). In 1950, Alan Turing introduced the well-known Turing test centering around the following question: Can a human evaluator distinguish between a human and a machine using only natural language conversations? This question is still controversial and triggered questions like: Can a machine have a mind, mental states, and consciousness in the same sense that a human being can? Independent

of this philosophical debate, we can see that more and more tasks are taken over by software trained based on examples. Alan Perlis wrote in 1982 "A year spent in Artificial Intelligence is enough to make one believe in God" and, indeed, it is amazing how AI technology can recognize images and sound, translate texts, and play games like go and chess without using a predefined strategy. However, there are still many tasks that are too difficult for AI. In 2015, Elon Musk stated that: "The Tesla that is currently in production has the ability to do automatic steering autopilot on the highway. That is currently being beta tested and will go into a wide release early next month. So, we are probably only a month away from having autonomous driving at least for highways and for relatively simple roads. My guess for when we will have full autonomy is approximately three years." In 2016, Turing award winner Geoffrey Hinton stated that "it is quite obvious that we should stop training radiologists" expecting that image recognition algorithms would outperform humans very soon. However, we are still driving our cars, and there is still a shortage of human radiologists. In short, we still need human intelligence.

Hybrid Intelligence (HI), sometimes also called Augmented Intelligence, emphasizes the assistive role of Machine Learning (ML), i.e., deep neural nets and other data-driven techniques are there to enhance human intelligence rather than to replace it (just like telescopes are there to enhance human vision).

RPA & PM: What is New?

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. RPA can be seen as "the poor man's workflow management solution" because it is cheaper than traditional automation.

In the 1970s, people like Skip Ellis and Michael Zisman already worked on so-called office information systems, which were driven by explicit process models. Systems such as Officetalk and SCOOP can be seen as early Workflow Management (WFM) systems. However, it took another 15 years until WFM technology was ready to be applied on a large scale. In the mid-nineties, many commercial WFM systems were available and there was the expectation that WFM systems would be an integral part of any information system. Many people, including the author, expected

that WFM systems would be as common as database management systems. However, this did not happen. WFM systems were succeeded by Business Process Management (BPM) systems that were broader in scope, but were also never widely adopted. 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 was never subcontracted to such systems at a scale comparable to database management systems. Actually, a few years ago, many considered the area of Business Process Management (BPM) to be dead. Organizations associated BPM with making process models rather than diagnosing and improving processes. There were three main reasons for this skepticism:

- Applying WFM/BPM technology was rather expensive. Processes are hardcoded 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-known 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.

These obstacles explain why organizations embraced Robotic Process Automation (RPA) and Process Mining (PM). RPA and PM revived the interest in Business Process Management (BPM). RPA 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.

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".

Using Process Mining Before and After RPA

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.

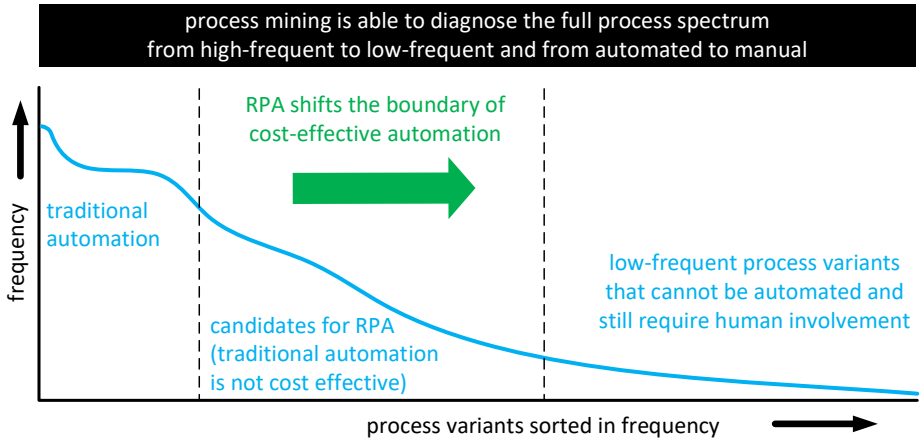

Process mining techniques use event data to show what people, machines, applications, and organizations are really doing. Process mining provides novel insights that can be used to identify and address performance and compliance problems. Just like spreadsheets can do anything with numbers, process mining can do anything with event data, i.e., it is a generic, domain-independent technology to improve processes. There are over 35 commercial offerings of process mining software (e.g., Celonis, Disco, ProcessGold, myInvenio, PAFnow, Minit, QPR, Mehrwerk, Puzzledata, LanaLabs, Process Diamond, Everflow, TimelinePI, Signavio, and Logpickr).

Often, a small percentage of activities account for most of the events, and a small percentage of trace variants account for most of the traces. 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 classic WFM/BPM software. 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 crucial technology to identify routine work that can be supported using RPA. Therefore, we claim that process mining can be used to pick the "automation battles" that are cost-effective and feasible.

This vision matches well with the notion of Hybrid Intelligence (HI). We should not aim for a strict divide between work done by software robots

and work done by humans. 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. Part of the work formerly done by workers is now done by software robots. Process mining can be 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.

Using combinations of process mining and machine learning, it is possible to flexibly distribute work over workers and software robots. For example, tasks are initially performed by robots and are escalated to workers the moment there is a complication or 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).

Action-Oriented Process Mining (AOPM)

The goal of RPA is to partially automate tasks in the process, and process mining can help identify where this makes the most sense. However, RPA builds on top of existing systems ranging from SAP and Salesforce to homegrown applications. It is unrealistic to assume that RPA, ML, and AI will replace these systems. Hybrid Intelligence (HI) should not only combine human intelligence and machine intelligence; it should also do this in a complex landscape of existing systems. Hence, it is naïve to assume that process-mining results will replace existing systems handling the operational tasks. However, there are opportunities to use process-mining results to automatically manage the process better.

Action-Oriented Process Mining (AOPM) turns observed events into management actions when needed. The goal is not to support the operational process itself (that already exists in some form), but to support the management of the process. Process mining diagnostics related to compliance and performance combined with process knowledge and reinforcement learning provide the ingredients for a reactive system that automatically triggers management workflows improving the process. The goal of AOPM is not to automate the tasks, but the management of the process.

Contact

Wil van der Aalst

Process and Data Science
RWTH Aachen University
Aachen, Germany
www.vdaalst.com

Prof.dr.ir. Wil van der Aalst is a full professor at RWTH Aachen University leading the Process and Data Science (PADS) group. He is also part-time affiliated with the Fraunhofer-Institut für Angewandte Informationstechnik (FIT) where he leads FIT's Process Mining group. His research interests include process mining, Petri nets, business process management, workflow management, process modeling, and process analysis. Wil van der Aalst has published over 800 articles and books and is typically considered to be in the top-15 of most cited computer scientists with an H-index of over 155 and more than 110.000 citations. Next to serving on the editorial boards of over ten scientific journals, he is also playing an advisory role for several companies, including Celonis, Fluxicon, and UiPath. Van der Aalst is an IFIP Fellow, IEEE Fellow, ACM Fellow, and received honorary degrees from the Moscow Higher School of Economics (Prof. h.c.), Tsinghua University, and Hasselt University (Dr. h.c.). He is also an elected member of the Royal Netherlands Academy of Arts and Sciences, the Royal Holland Society of Sciences and Humanities, the Academy of Europe, and the North Rhine-Westphalian Academy of Sciences, Humanities and the Arts. In 2018, he was awarded an Alexander-von-Humboldt Professorship.

Epilogue

RPA and process mining can help to overcome the limitations of traditional WFM/BPM approaches. They also complement each other. RPA is often used in a bottom-up manner realizing quick wins. Process mining can be used for identifying RPA opportunities. However, process mining also views processes in a more holistic top-down manner. The question of what to automate is not new. However, with the uptake of Machine Learning (ML) and Artificial Intelligence (AI), the tradeoffs are changing rapidly. Hybrid Intelligence shows that despite advances in ML and AI, many tasks require a combination of human intelligence and ML/AI. Process management is for sure one of such tasks. Currently, process management is done by humans while the tasks are supported by a complex landscape of systems. Using RPA and Action-Oriented Process Mining (AOPM), we do not try to replace these systems. Instead, we selectively augment tasks currently done by people at both the operational level and the management level.

Learn more?

- W.M.P. van der Aalst. Process Mining: Data Science in Action. Springer-Verlag, Berlin, 2016. http://www.springer.com/978-3-662-49850-7
- W.M.P. van der Aalst, M. Bichler, and A. Heinzl. Robotic Process Automation. Business and Information Systems Engineering, 60(4):269-272, 2018.
- Coursera Course Process Mining: Data science in Action https://www.coursera.org/learn/process-mining
- W.M.P. van der Aalst. Development of the Process Mining Discipline. In L. Reinkemeyer, editor, Process Mining in Action: Principles, Use Cases and Outlook, pages 181-196. Springer-Verlag, Berlin, 2020.
- W.M.P. van der Aalst and A. Berti. Discovering Object-Centric Petri Nets. Fundamenta Informaticae, 175(1-4):1-40, 2020.