

Hybrid Intelligence: To Automate or Not to Automate, That is the Question

Wil M.P. van der Aalst

RWTH Aachen University

Fraunhofer-Institut für Angewandte Informationstechnik

Ahornstr. 55, D-52074, Aachen, Germany

wvdaalst@pads.rwth-aachen.de

Abstract:

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 Artificial Intelligence (AI) and Machine Learning (ML) technologies outperform 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. Hence, current and future Business Process Management (BPM) initiatives need to consider HI and the changing boundaries between work done by people and work done by software robots. Consider, for example, the success of Robotic Process Automation (RPA), which demonstrates that gradually taking away repetitive tasks from workers is possible. In this viewpoint paper, we argue that process mining is a key technology to decide what to automate and what not. Moreover, using process mining, it is possible to systematically monitor and manage processes where work is distributed over human workers and software robots.

Keywords:

Hybrid intelligence; data science; process science; machine learning; business process management.

1. Introduction

Machine Learning (ML) and Robotic Process Automation (RPA) have lowered the threshold to automate tasks previously done by humans [5,7,10,11,17,23]. Yet organizations are struggling to apply ML and RPA, effectively causing many digital transformation initiatives to fail. Process mining 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 process mining.

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 [5]. For 1980, the report predicted that there would be a human-crewed 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 way 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 humans to do many tasks despite the amazing progress in AI and ML.

In this viewpoint paper, we focus on the question “To Automate or Not to Automate?” thereby linking Hybrid Intelligence (HI) [1,2,12,15], Artificial Intelligence (AI), and Machine Learning (ML) [9,13,16] to Business Process Management (BPM) [3,18] and Robotic Process Automation (RPA) [4,17,23]. This question is highly relevant because there is consensus that AI/ML will dramatically change the workplace [5,7,11]. Figure 1 shows the results of a PwC study based on OECD data collected in the context of the Program for the International Assessment of Adult Competencies (PIAAC) of the Organization for Economic Cooperation and Development (OECD) [7]. Recalling Niels Bohr’s quote and the RAND Corporation report mentioned before, one should take such analyses with a grain of salt. Nevertheless, it is worthwhile to try and identify jobs that might be of high risk of automation. The PwC study anticipates three waves of automation until mid-2030: (1) *algorithm wave* (early 2020s), (2) *augmentation wave* (late 2020s), and (3) *autonomy wave* (mid 2030s) [7]. The first wave focuses on the automation of simple computational tasks and analysis of structured data in areas like finance, insurance, information, and communications. This wave is already a reality considering, for example, the closing of local banks in most countries. The second wave focuses on the automation of repeatable tasks such as filling in forms, communicating, and exchanging information using technologies such as RPA. The third wave will automate of physical labor and problem-solving in manufacturing and transport. Figure 1 shows the expected impact of the three waves. For sure, the three automation waves will disrupt labor markets. Initially, mostly administrative work (e.g., in banking and insurance) is impacted, but over time, also a substantial

fraction of physical labor will disappear. For example, autonomous vehicles will soon become a reality in transportation, storage, manufacturing, and construction.




Figure 1: PwC analysis of the PIAAC OECD data predicting the proportion of existing jobs that may disappear due to automation by the mid-2030s in three overlapping waves [7].

Figure 2 shows the two types of automation considered in this paper. *Task automation* is limited to a single task, e.g., automatically performing a credit check or making a payment. *Process automation* considers end-to-end processes. The Purchase-to-Pay (P2P) process shown in Figure 2 includes multiple activities. Some of these activities may be automated, but independent of this, processes need to be coordinated, controlled, and continuously improved and adapted. BPM and process mining focus on end-to-end processes. For example, process mining can be used to detect performance and compliance problems. Such problems can automatically trigger corrective workflows. Both types of automation may benefit from a human-machine symbiosis where human intellect is complemented by machine intelligence. How work is divided exactly remains a challenging question in years to come (see the three waves in Figure 1).

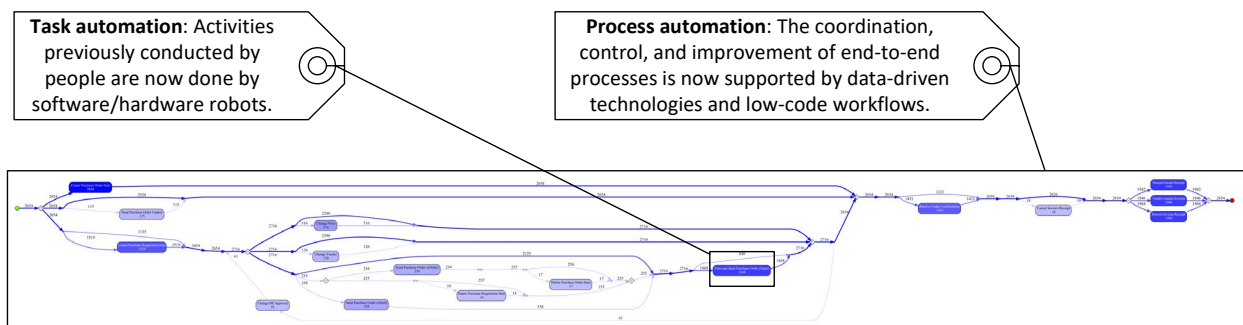


Figure 2: Two types of automation: (1) task automation and (2) process automation. Data-driven technologies such as AI, ML, RPA, and process mining can support both types of automation. However, often a combination of human intelligence and machine intelligence leads to the best results.

Hybrid Intelligence (HI) is one of the key elements in *digital transformation* initiatives, i.e., the adoption of digital technology to transform services or businesses by replacing non-digital or manual processes with digital processes or replacing older digital technology with newer digital technology. This extends beyond traditional automation and may include new types of innovation and creativity, e.g., new business models, new sales channels, new products, and new services. Such changes typically require, but also accelerate, task and process automation.

The remainder of this viewpoint paper is organized as follows. Section 2 introduces Hybrid Intelligence (HI). Section 3 provides a critical analysis of traditional BPM initiatives highlighting three main problems. Section 4 introduces RPA as a technology for task automation. Section 5 introduces process mining as a technology to make BPM more data-driven and enable new forms of process automation. Section 6 relates the different topics, advocating a convergence of HI, BPM, RPA, ML, and process mining. Section 7 concludes the paper.

2. Hybrid Intelligence

Deep Blue, a chess-playing computer developed by IBM, won its first game against world champion Garry Kasparov in 1996. AlphaGo, a Go-playing computer developed by DeepMind Technologies, defeated the best-ranked Go player Ke Jie in 2017. The more powerful AlphaGo Zero learned by just playing games against itself, but was able to defeat any human player by the end of 2017. Speech recognition software has been around since the 1950s when Bell Laboratories presented the “Audrey” system that was able to recognize the numbers 1 to 9. IBM’s “Shoebbox” system presented in 1962 was able to recognize 16 words. Until a decade ago, speech recognition software would not function very well. However, today we are surrounded by Amazon’s Alexa, Apple’s Siri, Microsoft’s Cortana, Google Assistant, etc. A similar development can be seen in image recognition and many tasks that before could only be done by humans. These successes can be attributed to progress in *deep learning*, where *Artificial Neural Networks* (ANNs) having multiple layers progressively extract higher-level features from the raw input [9,16]. Although neural networks had been around for decades, these techniques started to outperform classical approaches around 2012. Today, there is a lot of excitement about the amazing possibilities of deep learning. However, also the limitations become increasingly visible, especially in organizational settings and situations with limited data or many changes.

It is not easy to clearly define terms related to “intelligence” and “learning”. The “AI Effect”, commonly known as Tesler’s Theorem, says that “Artificial Intelligence is whatever hasn’t been done yet”. (Actually, Larry Tesler said “Intelligence is whatever machines haven’t done yet”.) Tesler’s Theorem shows that things that were previously seen as Artificial Intelligence (AI) are removed from the definition of AI when they become standard. When people use the term AI today, they often refer to Machine Learning (ML) based on ANNs. However, for most of its history, AI was dominated by *symbolic AI*, also known as “classical AI”, “rule-based AI”, and “good old-fashioned AI”, and associated with expert systems and logical reasoning. In recent years, AI got increasingly associated with ML.

ML techniques are data-driven and learn from data without explicitly being programmed. We typically distinguish between training data and test data. For example, we train an ANN to distinguish dog and cat pictures that are labeled. While training, the ANN updates the weights in the internal representation until the number of incorrectly classified pictures is minimized. Then the trained ANN is used to classify test data, i.e., unseen dog and cat pictures that need to be classified correctly. Given enough training data, such an ANN may perform amazingly well in practice, although it was never programmed to do so and has no explicit knowledge of cats and dogs. DeepMind’s AlphaGo Zero learned to play Go in a superior manner by just knowing the rules and playing against itself. There are many machine learning techniques ranging from classical approaches such as regression, decision trees, logistic regression, k-means clustering, and principal component analysis to support vector machines, convolutional neural networks, autoencoders, long short-term memory networks, and generative adversarial networks. Approaches can be classified into supervised learning (using labeled data, e.g., for classification), unsupervised learning (using unlabeled data, e.g., to discover unknown patterns), and reinforcement learning (finding the balance between the exploration of uncharted territory and the exploitation of current knowledge).

ML can be seen as part of *data science*, i.e., the broader interdisciplinary field aiming to turn data into real value. Data may be structured or unstructured, big or small, static or streaming. Value may be provided in the form of predictions, automated decisions, models learned from data, or any type of data visualization delivering insights. Data science includes data extraction, data preparation, data exploration, data transformation, storage and retrieval, computing infrastructures, various types of mining and learning, presentation of explanations and predictions, and the exploitation of results taking into account ethical, social, legal, and business aspects [19].

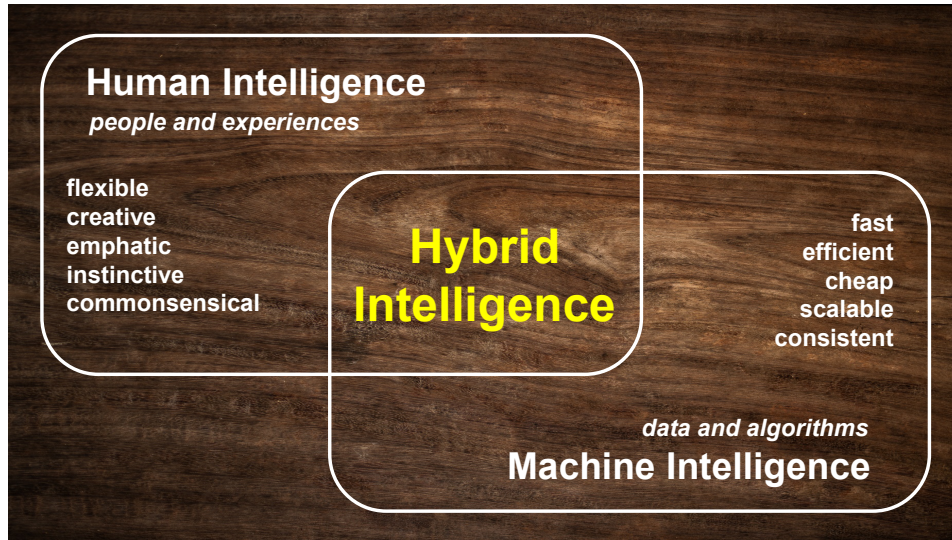


Figure 3: Hybrid Intelligence (HI) aims to combine the best of human intelligence and machine intelligence.

Hybrid Intelligence (HI), sometimes also called *Augmented Intelligence*, emphasizes the assistive role 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). In [2], Dellermann et al. define Hybrid Intelligence (HI) as “the ability to achieve complex goals by combining human and artificial intelligence, thereby reaching superior results to those each of them could have accomplished separately, and continuously improve by learning from each other”.

Figure 3 illustrates that *Hybrid Intelligence* (HI) combines both two forms of intelligence.

- *Human intelligence* is about people and experiences and can be characterized by terms such as flexible, creative, emphatic, instinctive, and commonsensical.
- *Machine intelligence* is about data and algorithms and can be characterized by terms such as fast, efficient, cheap, scalable, and consistent.

HI aims to combine the best of both worlds. The spectacular developments in ML have extended the reach of software and hardware robots. Once a robot is able to perform a repetitive task at a similar level of quality, it is often also more cost-effective. The rise of the “platform economy” has accelerated this. Transaction platforms that match supply and demand (e.g., Amazon, Alibaba, Airbnb, Uber, and Baidu) and technology platforms (e.g., Microsoft’s software platform and the App stores of Google and Apple) have the characteristic that they grow very fast and that, in the end, often one winner remains (due to the traditional economy of scale, low marginal costs, and network effects). Due to these platforms new technologies can be adopted fast at a global scale. However, humans still have unique capabilities. Consider, for example, disruptive events like the COVID-19 pandemic where one is confronted with completely new challenges that require flexibility, creativity, and intuition. People have the ability to transfer experiences from one problem domain to another. Moreover, empathy (i.e., the capacity to understand or feel what another person is

experiencing) and ethics (i.e., reasoning about moral concepts such as good and evil, right and wrong, virtue and vice, justice and crime) require human intelligence [21]. In HI, human intelligence and machine intelligence complement each other.

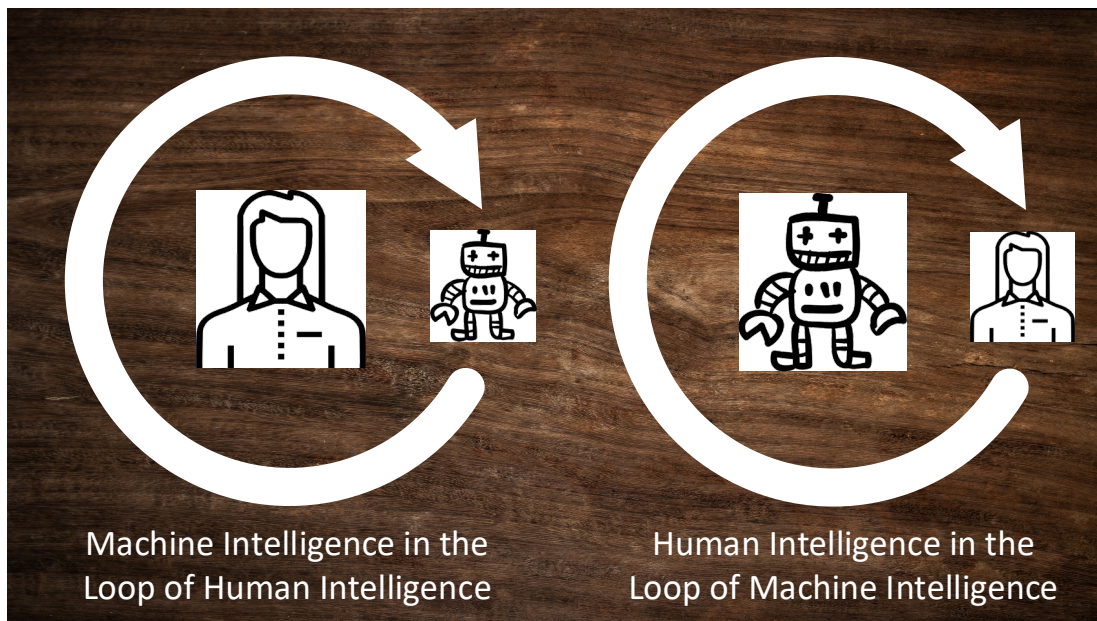


Figure 4: Human in the loop or machine in the loop?

Figure 4 shows how human intelligence and machine intelligence can be combined. The left-hand side shows the traditional use of AI/ML in organizations. AI/ML is used to provide decision support or assist in performing repetitive tasks. For example, a data-driven sales forecast supports decision-making in logistics and production, or ML algorithms help to speed up standard routines of image acquisition in radiography. The human is in control, and AI/ML is used as a tool. The right-hand side of Figure 4 shows the opposite situation. Machine intelligence is used to automatically process cases without human intervention. However, the machine can call the human for help in exceptional cases. For example, credit scoring or X-ray diagnostics are performed automatically, but boundary cases are evaluated by human experts.

The interplay between human intelligence and machine intelligence may lead to new insights. AlphaGo showed human players new strategies for playing Go, as has been acknowledged by the world's leading Go players. Shi Yue said "AlphaGo's game last year transformed the industry of Go and its players. The way AlphaGo showed its level was far above our expectations and brought many new elements to the game." Zhou Ruiyang said "I believe players more or less have all been affected by Professor Alpha. AlphaGo's play makes us feel more free and no move is impossible to play anymore. Now everyone is trying to play in a style that hasn't been tried before." This example shows that humans can learn from machines. This also applies to operational processes e.g., in healthcare or sales. Therefore, organizations need to embrace HI and actively manage the constantly shifting distribution of work between workers and robots.

3. Business Process Management: A Critical Analysis

In this viewpoint paper, we focus on the relation between HI and *Business Process Management* (BPM) [3,18], considering new technologies such as ML, RPA, and process mining.

Already in the 1970s, people like Skip Ellis and Michael Zisman 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 these systems would become 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 use software from different vendors, making a seamless 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.

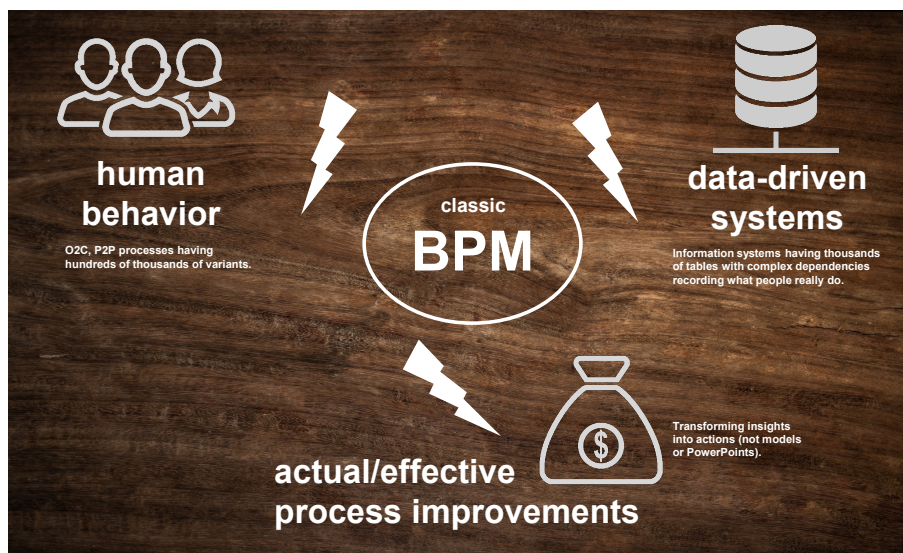


Figure 5: Old-school BPM tends to be (1) unable to capture human behavior, (2) unable to deal with the complexity of real-life systems, and (3) unable to realize actual improvements.

Figure 5 hypothesizes about possible reasons for the limited success of traditional WFM/BPM systems and approaches. Human behavior and information systems tend to be oversimplified, leading to a disconnect with reality. Moreover, it is often impossible to show that the process actually improved. In the remainder of this section, we elaborate on each of the three problems highlighted in Figure 5.

3.1 Inability to capture human behavior

Simple processes such as Order-to-Cash (O2C) and Purchase-to-Pay (P2P) tend to be much more complicated than expected. It is not uncommon to find thousands of process variants, i.e., unique ways of executing O2C or P2P processes (just considering the ordering for activities). Some of these variants may be undesirable. However, deviations often have good reasons. People are adapting to contextual factors, not present in the process model. It is also very difficult to create simulation models that exhibit the real behavior of an organization. It is possible to create so-called “digital twins” of highly-structured production processes. However, in processes where human actors are in the lead and where people need to distribute attention over multiple processes, it is still impossible to create meaningful “digital twins” that have an acceptable predictive value. These difficulties show why traditional WFM/BPM systems and approaches failed. Assuming that reality can be captured in the form of a BPMN (Business Process Model and Notation) and implemented using a WFM/BPM system is a recipe for disaster.

3.2 Inability to deal with the complexity of real-life systems

Moreover, real-life information systems are more complicated than stakeholders like to think. When making process models and talking about new systems, people tend to underestimate the complexity of the underlying data. Standard systems like SAP’s S/4HANA contain hundreds of thousands of tables. Of course, a typical organization uses only a subset of these tables. However, it shows the complexity of real-life information systems. Data related to one process may be scattered over dozens or even hundreds of tables connected through primary key and foreign key relationships. Although relational databases are well-understood and more structured than NoSQL-based non-relational databases, these cannot be described using a UML class diagram and BPMN model. Nevertheless, many WFM/BPM vendors suggested that it would be easy to replace existing systems using a properly configured WFM/BPM system. This is, of course, not the case. It is very naïve to think that existing systems can be replaced easily. There are numerous examples of failed ERP systems implementations that drove companies into bankruptcy (e.g., Shane Co., American LaFrance, FoxMeyer Corp., etc.). These bankruptcy cases had in common that people underestimated the complexity. The author has witnessed numerous organizations that selected a WFM/BPM system that never went into production. Therefore, it is important to try and realize process improvements while keeping the existing information systems. RPA (see Section 4) builds on top of existing information systems while automating repetitive work.

3.3 Inability to realize actual improvements

The third problem highlighted in Figure 5 is the limited ability to provide actionable results. Making process models, organizing workshops/meetings, and implementing new information systems do not necessarily lead to process improvements. Some of the larger organizations have invested in creating repositories of process models. However, such repositories become outdated quickly and do not necessarily impact the operational processes. Actually, most workers are not aware of their existence. Wallpaper-sized BPMN models that aim to be close to reality are too abstract because they are not connected to the actual data, and stakeholders can always question their validity. Process mining (see Section 5) addresses this by showing continuously updated process maps that show the current situation.

The problems highlighted in Figure 5 explain why organizations embraced RPA and process mining during the last decade. Both helped to revive the interest in 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. Moreover, process mining

helps to capture the actual end-to-end processes while acknowledging their complexity and focusing on the real problems.

4. Robotic Process Automation: Focusing on Individual Tasks

Robotic Process Automation (RPA) has lowered the threshold for process automation [23]. 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 (e.g., SAP). 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 often much cheaper than traditional automation [23]. Figure 6 show the main idea of RPA. In most organizations, one can easily find people whose main job is to connect information systems using copy-and-paste actions and simple repetitive tasks. These provide the required “glue” between applications and the outside world. Despite the repetitive nature of the work, it is not cost-effective to replace the information systems used. Systems may be provided by different vendors and may be too old to change (legacy software). Therefore, it is cheaper to copy-and-paste address information or send e-mails manually. RPA does not aim to change the existing systems but take over the repetitive work of people. It is a form of automation using software robots (bots) replacing humans.

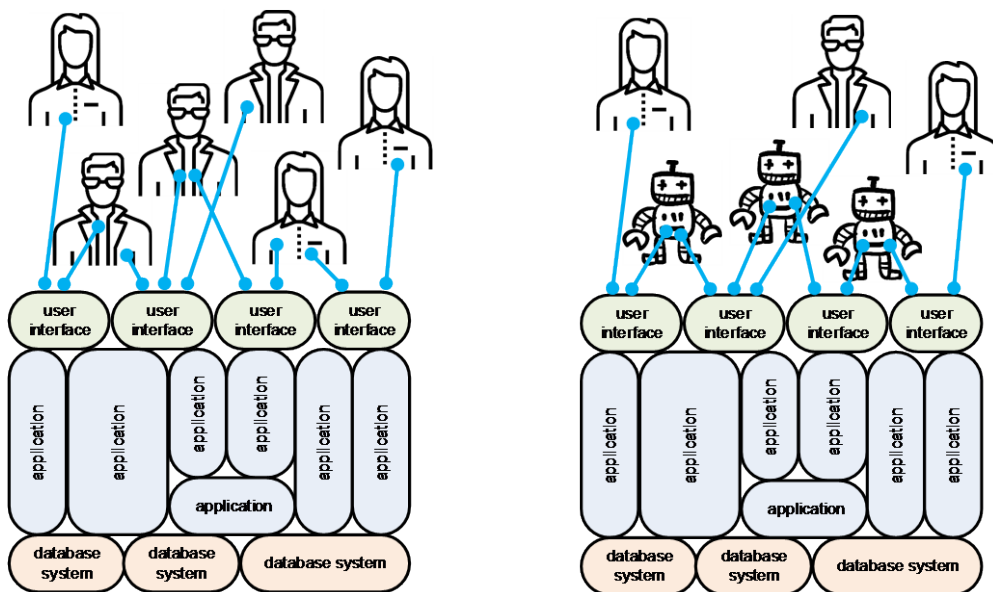


Figure 6: In many organizations, humans are the glue between applications and the outside world (left). This leads to repetitive tasks that can also be done by software robots without changing the underlying information systems.

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.

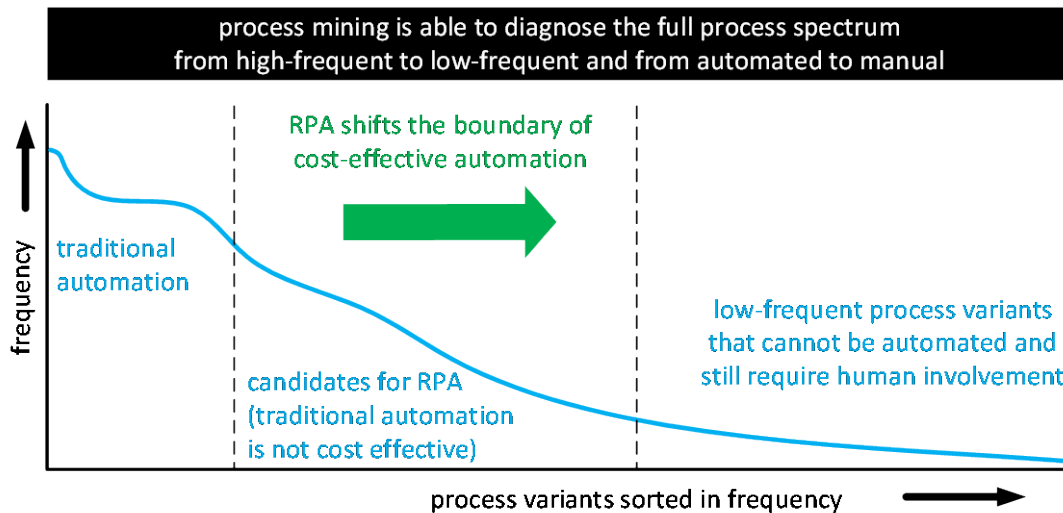


Figure 7: RPA shifts the boundary of cost-effective automation and can therefore be seen as “the poor man’s workflow management solution”. Process mining complements RPA by identifying routine work and monitoring processes before and after the introduction of RPA.

Before introducing RPA, one needs to analyze the processes to be automated. Process mining can help to identify promising candidates [3,21,24]. 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. This is illustrated by Figure 7. The figure shows that processes and process variants can be sorted by frequency. Obviously, one would like to automate the most frequent processes and process variants first. Because traditional automation (e.g., using WFM/BPM systems or writing ABAP code to change SAP) is rather expensive, there often exist many repetitive tasks which are not automated. This corresponds to the middle of the spectrum depicted in Figure 7. After introducing RPA there are three types of tasks: (1) tasks handled by the information system using traditional automation, (2) tasks handled by software robots, and (3) low-frequency tasks still done manually. The whole can be monitored and analyzed using process mining as is discussed next.

5. Process Mining: Focusing on End-to-End Processes

RPA can be seen as a bottom-up activity, i.e., removing repetitive tasks. Process mining can help to identify and automatically learn such tasks [3,21,24]. However, the primary use case of process mining is the top-down analysis of end-to-end processes [19,22,23]. 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, Apromore, Minit, QPR, Mehrwerk, Puzzledata, LanaLabs, Process Diamond, Everflow, TimelinePI, Signavio, and Logpickr), next to open-source tools like ProM, PM4Py, bupaR, and RapidProM.

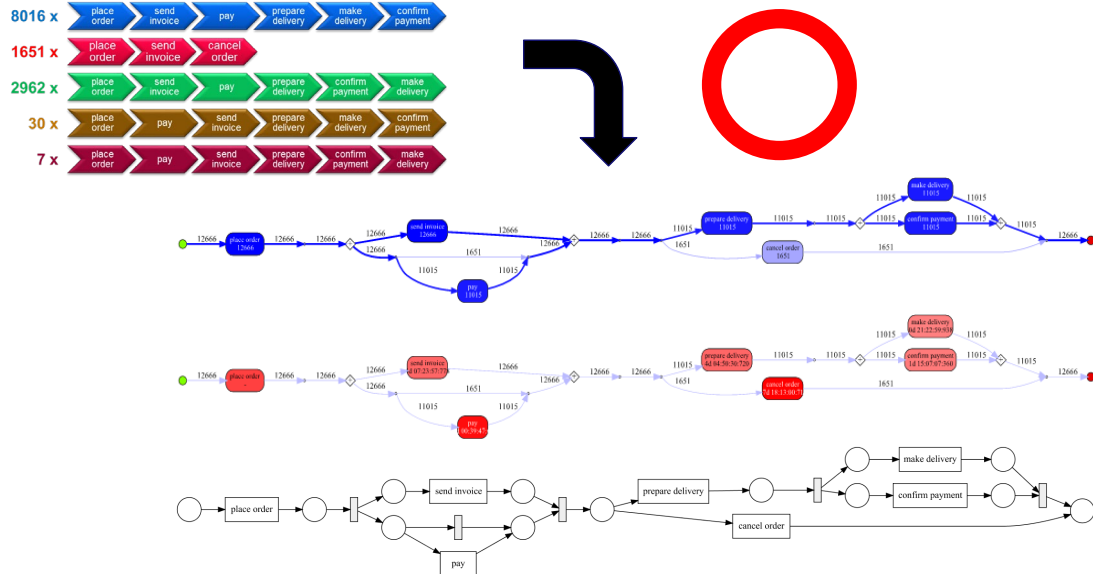


Figure 8: Event data serve as the starting point for process mining. Such data can be used to discover what the real process is, where the bottlenecks are, where the process deviates, and what the root-causes are. If there is enough data and the process is stable, then it is even possible to predict performance and conformance problems.

All process-mining tools start from *event data*. An event log is a collection of events stored using a format like XES (xes-standard.org). An event may have many different attributes, but at least a *case identifier*, an *activity name*, and a *timestamp*. Additional attributes may refer to locations, resources, costs, transactional information, and energy consumed. Events are grouped using the case identifier and sorted using the timestamps. Hence, each case corresponds to a trace, i.e., a sequence of events. Focusing on the activity names only, these traces can be grouped into variants, i.e., sequences of activities. Figure 8 illustrates the typical use of process mining using a small event log with 71,043 events, 12,666 cases, and 7 unique activities. A possible trace is the sequence $\langle \text{place order, send invoice, pay, prepare delivery, make delivery, confirm payment} \rangle$. There are over 8000 cases corresponding to this activity sequence. Using process mining, one can uncover *compliance and performance problems*. Initially, process mining efforts focused on process discovery. However, over time it has become clear that process discovery is just the starting point to process improvement. One can witness an uptake in conformance checking and performance analysis techniques. Moreover, process mining is often combined with ML techniques to find root causes for inefficiencies and deviations. As was illustrated by Figure 7, event logs often follow a Pareto distribution, i.e., a few variants explain a large proportion of the event log.

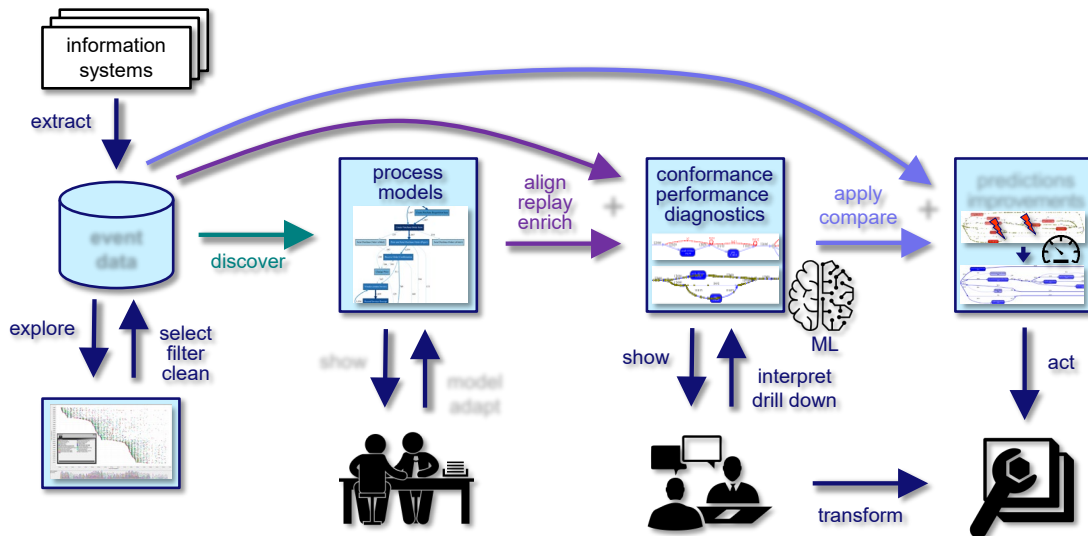


Figure 9: A high-level overview of process mining, showing also the different activities.

Figure 9 provides another, more high-level, view on process mining. Process mining starts by extracting event data from information systems. This may be quite involved since traditional process mining techniques assume an event log with a single case notion. Using object-centric process mining, this requirement can be relaxed, i.e., each event may refer to any number of objects and it is possible to discover more holistic process models [24]. However, most approaches still assume a single case notation and logs in the form of an XES file or a similar database table. Such event data are used to discover process models showing the real process. Such models can be enriched with frequency and timing information. Given a discovered or normative process model it is also possible to do conformance checking and highlight deviations. Next to visualizing conformance and performance problems, it is possible to explain and predict these. Note that process discovery and conformance checking are unrelated to mainstream AI/ML techniques. However, process mining can be used to generate standard classification problems, e.g., what are the characteristics of the cases that deviate, fail, or get delayed. This can be used to predict such problems and recommend actions.

One can witness a shift in focus from backward-looking to forward-looking process mining [22]. Organizations are more interested in what is happening now or what is going to happen next. Backward-looking process mining can be used to fundamentally improve processes, but provides little support for the day-to-day management of processes. Therefore, event data need to be updated continuously and process-mining techniques need to be able to analyze cases that are still running. This is needed to control and influence the running process instances. Techniques for operational support (i.e., detecting compliance and performance problems at runtime, predicting such problems, and recommending actions) have been around for more than a decade. However, the challenge is to make these techniques more reliable and also trigger the actions needed.

Action-Oriented Process Mining (AOPM) focuses on automated corrections actions based on process mining diagnostics [14]. 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. Note that, like RPA, AOPM does not aim at replacing the original information system. The acquisition of Integromat, a low-code online

automation platform, by Celonis illustrates this development. Integromat provides over 500 application connectors to interact with the most widely used information systems (Salesforce, Office, Teams, Twitter, etc.). When the Celonis process mining system detects a known problem, Integromat can trigger the required corrective actions.

6. Towards the Convergence of HI, BPM, RPA, ML, and Process Mining.

In Section 3, we argued that old-school BPM tends to be (1) unable to capture human behavior, (2) unable to deal with the complexity of real-life systems, and (3) unable to realize actual improvements. Process mining helps to address (1) and (3) by looking at the real processes in an objective manner before and after interventions. Moreover, just like RPA process mining does not try to replace existing systems and face the complexity of real-life systems (2).

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. As shown using Figure 7, 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).

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 and ML 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 many opportunities to use process-mining results to automatically manage the process better.

7. Conclusion

This viewpoint paper discussed Hybrid Intelligence (HI) from the viewpoint of task and process automation. We started with the question “To automate or not to automate?”. 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. Due to advances in AI/ML, the answer to the question will change continuously. HI suggests that for many of the more challenging

tasks, we will need to mixture of human and machine intelligence to get the best results. Although deep learning has had an amazing success in areas such as speech recognition, automated translation, image recognition, smart maintenance, and sentiment analysis, there are also obvious limitations. Machine intelligence tends to fast, efficient, cheap, scalable, and consistent, but also inflexible, non-creative, non-emphatic, non-instinctive, and lacking common sense. In HI, human intelligence (i.e., people having experience and domain expertise) complements machine intelligence. We introduced HI and indicated the relevance for task and process automation.

We also provided a critical analysis of traditional WFM/BPM approaches. We identified weaknesses of traditional approaches that were not data-driven while trying to replace existing systems based on process models. In hindsight, these approaches can be considered naïve for two reasons. First of all, real processes have a lot of variability due to human behavior. Simple P2P or O2C processes may have thousands of variants, and this is in stark contrast with the oversimplified models produced by humans. Second, information systems like SAP's ERP system are extremely complex with thousands of database tables. Therefore, attempts to simply replace such systems are destined to fail. Robotic Process Automation (RPA) and process mining address these limitations by better using the available data and systems. RPA builds upon existing systems by taking over repetitive tasks from humans. 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. A recent development in the field of process mining is that performance and conformance problems automatically trigger corrective workflows leveraging both the data and systems present. However, data-driven techniques should also be able to say "I do not know" or "I'm not sure" and leave decisions to people. This is the true spirit of HI where people, data, and software augment each other.

Acknowledgments

The author thanks the Alexander von Humboldt (AvH) Stiftung for supporting his research.

References

- [1] D. Dellermann, A. Calma, N. Lipusch, T. Weber, S. Weigel, and P. Ebel. The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems. Hawaii International Conference on System Sciences (HICSS 2019), pages 1-10, 2016.
- [2] D. Dellermann, P. Ebel, M. Söllner, J.M. Leimeister. Hybrid Intelligence. *Business and Information Systems Engineering*, 61(5):637-643, 2019
- [3] M. Dumas, M. La Rosa, J. Mendling, H. Reijers. *Fundamentals of Business Process Management*. Springer, 2018.
- [4] J.G. Enríquez, A. Jiménez-Ramírez, F.J. Domínguez-Mayo, and J.A. García-García. Robotic Process Automation: A Scientific and Industrial Systematic Mapping Study. *IEEE Access*, Volume 8, pages 39113-39129, 2020.
- [5] C.B. Frey and M.A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* 114:254-280, 2017.
- [6] T.J. Gordon and O. Helmer-Hirschberg. Report on a Long-Range Forecasting Study. RAND Corporation, 1964.
- [7] 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.

- [8] E. Kamar. Directions in Hybrid Intelligence: Complementing AI Systems with Human Intelligence. Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI 2016), pages 4070–4073, 2016.
- [9] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436-444, 2015.
- [10] 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.
- [11] 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.
- [12] M. Mirbabaie, S. Stieglitz, and N. Frick. Hybrid intelligence in hospitals: Towards a research agenda for collaboration. *Electron Markets*, 2021.
- [13] T. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [14] G. Park and W. M. P. van der Aalst. A General Framework for Action-Oriented Process Mining. *Business Process Management Workshops 2020, Lecture Notes in Business Information Processing 397*, Springer, 206-218, 2020.
- [15] M. Peeters et al. Hybrid Collective Intelligence in a Human–AI Society. *AI & Society*, 36(1):217–238, 2021.
- [16] D.E. Rumelhart, G. Hinton, and R.J. Williams. Learning Representations by Back-Propagating Errors. *Nature*, 323:533-536, 1986.
- [17] R. Syed et al. Robotic Process Automation: Contemporary Themes and Challenges. *Computers in Industry*, Volume 115, 2020.
- [18] W.M.P. van der Aalst. *Business Process Management: A Comprehensive Survey*. ISRN Software Engineering, pages 1-37, 2013.
- [19] W.M.P. van der Aalst. *Process Mining: Data Science in Action*. Springer-Verlag, Berlin, 2016.
- [20] W.M.P. van der Aalst. Responsible Data Science: Using Event Data in a “People Friendly” Manner. In *Enterprise Information Systems. ICEIS 2016. Lecture Notes in Business Information Processing, Volume 291*. Springer, 2016.
- [21] W.M.P. van der Aalst. On the Pareto Principle in Process Mining, Task Mining, and Robotic Process Automation. In *Proceedings of the 9th International Conference on Data Science, Technology and Applications (DATA 2020)*, pages 5-12. SciTePress, 2020.
- [22] 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.
- [23] W.M.P. van der Aalst, M. Bichler, and A. Heinzl. Robotic Process Automation. *Business and Information Systems Engineering*, 60(4):269–272, 2018.
- [24] W.M.P. van der Aalst and A. Berti. Discovering Object-Centric Petri Nets. *Fundamenta Informaticae*, 175(1-4):1-40, 2020.

Biographical notes



Wil M.P. van der Aalst

Wil van der Aalst is a full professor at RWTH Aachen University leading the Process and Data Science (PADS) group. His research interests include process mining, Petri nets, business process management, workflow management, process modeling, and process analysis. He is also part-time affiliated with the Fraunhofer FIT and a member of the Board of Governors of Tilburg University. 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.