

# Implementation of Automation as Distributed Cognition in Knowledge Work Organizations: Six Recommendations for Managers

Completed Research Paper

**Aleksandre Asatiani**

University of Gothenburg  
SE-412 96 Gothenburg, Sweden  
aleksandre.asatiani@ait.gu.se

**Esko Penttinen**

Aalto University School of Business  
FI-00076 AALTO Espoo, Finland  
esko.penttinen@aalto.fi

**Tapani Rinta-Kahila**

The University of Queensland  
Business School & AIBE  
QLD 4072 Brisbane, Australia  
t.rintakahila@uq.edu.au

**Antti Salovaara**

Aalto University School of Arts, Design  
and Architecture & School of Science  
FI-00076 AALTO Espoo, Finland  
antti.salovaara@aalto.fi

## Abstract

*Knowledge work organizations are increasingly leveraging automation to enhance and transform their business processes. Many types of automation tools are being deployed in a large variety of information processing tasks, requiring effective management of human–automation co-operation. Yet, conceptual understanding of human–automation hybrid work remains thin and current literature lacks practical recommendations for managers. To address this gap, we synthesize findings from our three earlier case studies with organizations pursuing a wide array of automation tools and examine them through the lens of distributed cognition. We demonstrate how distributed cognition informs about the organizing for human–automation interaction when deploying automation. Our contribution lies in the presentation of six recommendations on three issues: human–automation task allocation, mitigation of the risk of deskilling, and management of collective knowledge across human and automation.*

**Keywords:** Automation, RPA, Distributed Cognition, Deskilling, Task allocation, Knowledge

## Introduction

The rapid emergence of new kind of automation, enabled by information technology (IT), has challenged today's organizations to reconsider the division of work between human agents and automation tools. This issue has become especially pertinent in the context of knowledge work, i.e., the work of producing and reproducing information and knowledge (Schultze 2000). This is because IT-based systems seem increasingly capable of taking over a large share of cerebral activities previously conducted solely by humans (Frey and Osborne 2017). While the advances in machine learning have made automation increasingly sophisticated and thereby seemingly resembling intelligent behavior, automation's intelligence still

depends on human involvement both during automation's training and configuration as well as its on-line operation.

Automation has been traditionally used in organizations to make repeatable tasks more efficient, to regulate manufacturing quality, or to perform computations that are too difficult or error-prone for humans. While automation can and will continue to be used for these tasks, the remit of the technology is rapidly expanding to new areas of work. Although configuring automation might require large datasets coupled with human touch, the final trained system can be embedded into small-scale business tasks, even as individual workers' cognitive aids. In this sense, automation is increasingly seen as a companion to human workers that enhances their capabilities.

One of the most rapidly adopted forms of automation in knowledge work organizations is Robotic Process Automation (RPA). It is a rule-based approach to mimicking human actions in knowledge work processes (Lacity and Willcocks 2016; Penttinen et al. 2018). Even though RPA was introduced relatively recently, organizations have already gained implementation experience and accrued knowhow on how to operate their army of RPAs.

The conceptual understanding of the best ways to integrate automation in knowledge work has not kept pace with this progress. Most of the discussion on implications of automation has been circling around the issue of human professions being replaced by computers altogether (Carr 2015; Ford 2015). There is much less discussion on more subtle questions such as the use of automation in these professions' subtasks, and on the nature of those subtasks that are most amenable for takeover by automation. The following question then arises: *how should these opportunities of automation be put to best use in organizations?*

Whereas automation can be leveraged to "superpower" human workers (Daugherty and Wilson 2018), current literature provides few practical guidelines for organizing work in a manner that makes this possible. As the focus has mostly been on inspecting either human-only activities or machine-only activities, the lacking understanding of hybrid work in human-automation teaming has been referred to as "the missing middle" (ibid.). With this gap in understanding, organizations may miss out on the potential these technologies have to offer. In more extreme cases, the very same characteristics of automation that equip humans with unprecedented capabilities can also backfire as negative consequences on workers and their organization, if no proper attention is given to how the automation is being used. Probably the most often mentioned negative consequences are the fear that human workforce is demoted to low-skill service work which still remains difficult for machines to master (Ford 2015) and that technology invites overreliance on its performance (Butler and Gray 2006).

Therefore, in this paper we provide managers with new tools to think about implementation and management of automation within a knowledge work organization in order to be able to make the most of its potential while avoiding the common pitfalls. To further our understanding of how such automation systems are utilized in organizations, we take a nuanced and close-range approach: we suggest that a view that sees organizations as networks of cooperating agents can provide a fruitful approach for management of automation. To this end, we apply the lens of *distributed cognition* (Hutchins 1995). This framework considers information processing agents, whether humans or automation, as parts of a larger distributed network that together constitutes a shared, distributed cognition. In line with these perspectives, our treatise on the topic highlights the complementary nature of human-automation collaboration. Our approach is applicable both to autonomously learning systems – the quintessential examples of artificial intelligence (Kaplan and Haenlein 2019) – as well as systems built on human-devised rule-based automation – which constitute a more traditional class of machine-augmented knowledge processes such as intelligent decision aids.

We will introduce distributed cognition in the following section. After this we analyze three previously completed studies (Asatiani et al. 2019; Rinta-Kahila et al. 2018; Salovaara et al. 2019) through the lens of distributed cognition. We identify and discuss three key issues with implementing automation in knowledge work organizations: 1) task allocation between humans and automation, 2) mitigation of the risk of deskilling, and 3) management of collective knowledge across human and automation. We introduce pertinent issues faced by the three case organizations, discuss the key learnings gained from them, and propose managerial recommendations to help overcome these challenges.

## Distributed Cognition

Distributed cognition (Hollan et al. 2000; Hutchins 1995) is a conceptual framework developed in cognitive science and anthropology for an analysis of information processing systems where the processing is distributed across several agents. Distributed cognition attends to the computational work that these agents carry out in a shared fashion when they transform information towards increasingly usable and actionable representations, and in this way accomplish complex tasks. The agents in a distributed cognitive system can be both humans and information processing artifacts, all of which take in information, process it, and produce output that other agents may then use to further refine the information.

A classic example of distributed cognition is ship navigation close to a shore (Hutchins 1995). Before GPS was widely adopted, plotting a ship's location and velocity required a coordinated effort of several people and use of several information-processing tools. A simplified description of this information processing task is following. At the sides of a ship, "watchstanders" use gyrocompasses to take bearings to landmarks, and deliver the readings to the ship deck, where an officer adjusts a device called *hoey* to obtain the ship's orientation, expressed in map coordinates. Using several bearings together, the officer finds a "fix" on the map chart and thereby can visualize the ship's likely location on the map. Finally, based on the previous fixes at known time intervals, the plotter is able to compute, using a calculator or a computational heuristic called a "three-minute rule", how fast the ship is moving. Several agents, both humans and information processing artifacts, thereby jointly accomplish the task of ship navigation. Since its introduction, distributed cognition has been successfully applied to several settings of distributed work, including air traffic control (Walker et al. 2010), organizational memory (Ackerman and Halverson 2004; Perry et al. 1999) and hospitals (Rogers and Ellis 1994).

This lens is highly resonant with the organizational implementation of automation, where automation systems of various types and sizes operate as information processing units that interact with humans. In the following sections we demonstrate how conceptualization of organizational implementation of automation through distributed cognition can help managers to achieve a clearer picture of automation and its capabilities.

## Three Viewpoints to Human–Automation Collaboration

In this section, we introduce three viewpoints to implementation of automation in knowledge work organizations: human–automation task allocation, mitigation of the risks of deskilling, and management of collective knowledge across human and automation. For each issue, we make recommendations for managers, on the one hand, to harness the benefits and, on the other hand, to avoid the potential problems associated with organizational implementation of automation tools. As we will see, an analysis based on distributed cognition on an organization's operational structure allows for critical appreciation on the merits of different human–automation collaboration networks, depending on the environment in which the organization operates and the objectives it seeks to maximize.

From the three aforementioned qualitative case studies (Asatiani et al. 2019; Rinta-Kahila et al. 2018; Salovaara et al. 2019), we have already used distributed cognition as our theoretical lens in one (Salovaara et al. 2019) and found it informative. In other two studies (Asatiani et al. 2019; Rinta-Kahila et al. 2018) we have done so only now, to synthesize overarching recommendations. Although the three case studies have pursued different research questions, they address the same phenomenon in comparable contexts: implementation of automation in knowledge work organizations. We reviewed each study's findings via the lens of distributed cognition by focusing on the configurations in which human and automation agents feed, process, share, and refine information in the case organizations. Short case summaries are provided in Table 1 below. The details of the case studies, including research methods, can be found in Appendices A–C. We have not presented a synthesis of these studies before, and to our knowledge, analyzing automation in organizations using the framework of distributed cognition has not been carried out before by others either.

Case	Case company description	Study description
Case company A	Malware-protection software company with 20 global offices and approximately 950 employees.	The aim was to study how the company deals with the frame problem (i.e., that AIs cannot recognize their own errors) in organizing its operations as a collectively mindful digital high-reliability organization. Findings were based on a qualitative inductive analysis of in-depth interviews.
Case company B	IT service provider specializing in digital business solutions and financial processes, with almost 1 000 employees. The company operates internationally, focusing on the Nordic countries.	This single-case study set out to explore how automating knowledge work tasks may lead to unintentional deskilling of workers. The study involved qualitative in-depth interviews conducted in three data-gathering stages.
Case company C	Large telecommunications service provider operating in Northern Europe with over 20 000 employees.	This study explored the use of a federated governance model for RPA projects, and uncovered challenges and opportunities associated with the governance model. Also, this study applied qualitative in-depth interviews with key informants involved in RPA implementation at the case company.

**Table 1. Case summaries**

### ***Human–Automation Task Allocation***

One of the fundamental issues that managers need to ponder when implementing automation tools is the task allocation between human agents and automation. Contemporary organizations tend to prefer modular information system structures by employing best of breed strategies and loose coupling of systems (Serrano et al. 2014). This approach favors the adoption of a distributed cognition approach to task allocation between human agents and automation where agents assume complementary information processing responsibilities.

To probe the meaningful division of tasks between human agents and automation, we have conducted case studies in two knowledge work organizations that have successfully implemented automation in their operations, albeit optimizing for different objectives: a software security company emphasizing reliability in its operations (Appendix A) and an accounting firm whose objectives are efficiency-related (Appendix B). We analyzed the task-level work allocations in both organizations, arriving at three characteristics regarding their human–automation task allocations.

First, automation should be considered by attending to the level of *mindfulness* (Weick et al. 1999) that is required from the respective operations. Mindless tasks (i.e., routine tasks that do not require cognitively demanding processes) are more amenable for automation than mindful tasks (e.g., tasks that require creativity or out-of-the-box thinking) which, in turn, are better to retain for humans. In line with this distinction, the software security company had offloaded to automation such mindless tasks as event log monitoring and malware sample collection. Mindful tasks, which included tasks such as careful replication of malware’s behavior and post-mortem analyses of difficult cases, remained human-executed, thereby ensuring sensitivity in the identification of new security threats. In the accounting firm, similarly, mindless tasks offloaded to software robots consisted of manual re-keying of data into various systems while mindful operations were related to analyses of accounting data and the resulting decision support to the client.

Second, human–automation collaborations should be complementary instead of substitutive in nature. The automation tools we researched in the two companies were primarily rule-based; for instance, they involved pattern matching in malware detection and task specification in software robot programming. We observed several instances where implementing automation in isolation from human work would have been prone

to failures. Even though the recent advances in machine learning might give the automation tools capabilities to autonomously author new rules, even the most sophisticated rule-based engines can face exceptional situations that need to be resolved through escalation procedures where humans rectify the emergent problems. Complementarity here means that humans and automation may be assigned to operate on the same problems where one helps the other by validating one another's decisions, verdicts, and actions.

Third, the distinction between epistemic vs pragmatic tasks in distributed cognition (Kirsh and Maglio 1994) is useful for the analysis of human–automation collaborations. This distinction attends to *epistemic* tasks that analyze and produce information and *pragmatic* tasks that make decisions to act and accomplish them. Epistemic information-processing tasks are highly amenable for implementations of automation while pragmatic tasks may pose threats to organizations because an autonomous, unsupervised automation making decisions may cause considerable harm in a very little time (consider, e.g., errors in high frequency trading). However, we identified two conditions under which also pragmatic tasks can be assigned to automation. First, this is possible if the pragmatic tasks are continuously mindfully controlled by humans. Second, automation can be used to interrupt potential human slips, lapses, and mistakes.

Following from the third characteristic, we identified three forms of using automation: epistemic analyses, pragmatic decision-making, and pragmatic human error prevention (Salovaara et al. 2019). These can be orchestrated to work together if an organization adopts a layered constellation of operations where higher-level mindful human-based tasks control and improve lower-level mindless automation-based tasks. For example, in the malware protection company, the core operations were run by a rule-based system which was updated and informed by human-based improvement layers. Within these outer layers, automation either epistemically informed humans or pragmatically sought to prevent their elementary errors. This socio-technical structure enabled the company to address scalability issues emerging from the rapid growth of malware threats.

Based on the learnings from these case studies, we present two recommendations to managers to meaningfully allocate tasks between human workers and automation.

**Recommendation 1:** Be careful when assigning pragmatic tasks to automation but seek for opportunities to find uses for automation in epistemic tasks.

This recommendation draws from the second characteristic (see above) that made a distinction between epistemic tasks (where information is processed) and pragmatic tasks (where actions are executed) as well as from the third characteristic that identified a principle by which this distinction can guide complementary task division between humans and automation. Managers should take a granular look into their operations on the level of activities and tasks to analyze the inherent properties of these tasks and assign them to either humans or automation. Managers should remember that in some cases it might even be necessary (due to, for example, scalability issues) to assign, in addition to epistemic tasks, also pragmatic tasks to automation, provided that they are closely controlled by humans.

**Recommendation 2:** Divide tasks into their mindful and mindless components, and offload the mindless part to automation while keeping the mindfulness-requiring part to humans.

This recommendation returns the attention to the first characteristic above. Instead of division between epistemic and pragmatic tasks, here the attention is on their repetitiveness and mindlessness. The level of mindlessness provides a complementary dimension to think about human–automation task division. Thus, independent of whether tasks are epistemic or pragmatic, its routine-like (i.e., mindless) parts are potentially amenable for automation. Mindfulness–mindlessness dimension can be applied as a principle for task division where mindful sub-parts (e.g., decision-making) and mindless parts (e.g., repetitive information processing), may be separated conceptually and then allocated to humans and automation that together can accomplish them using their complementary strengths. This complementary approach may also provide relief from the anxieties and fears associated with job losses and instead instill positive initial worker reactions towards the potential implementation of automation. By choosing the complementary approach, companies can also tackle some of the risks associated with deskilling of workers. This will be discussed in the following section.

### **Mitigating the Risk of Deskilling**

Closely connected to human–automation task allocation, deskilling represents an outcome of governance gone wrong. Paradoxically, while automation is often expected to liberate its operators from repetitive tasks, humans tend to offload also complex, multidimensional reasoning to algorithms (Bainbridge 1983) – thus, giving up what has been always considered as the key competitive advantage of the human species. For instance, consider the eagerness of consumers to adopt GPS navigation applications, which often comes at the expense of deteriorating innate orientation skills. Ideally, automation of repetitive routine tasks, such as back-office tasks, should not lead to such deskilling that would damage the core value-generating activities. However, this may be one of automation’s effects: mastery of repetitive routine tasks can turn out being surprisingly important: it has been argued that diligent conduction of detailed tasks plays a key role in accumulating skills and forming hard-earned expertise (Arnold and Sutton 1998). Taking such repetition away would then prevent the acquisition of skills and hamper their maintenance.

Using automation is enchanting and, for the most part, workers tend to welcome actions that decrease their routine-like workload or make wicked tasks less cumbersome (Carr 2015; Millman and Hartwick 1987; Silverman 1966). Yet, neither organizations nor their workers may be fully aware of the pitfalls of automation. For organizations, this represents a danger to the longevity of their collective intellectual capital, which is a key asset of many organizations today. For individual knowledge workers, the threat is similar: deskilling threatens to strip them from their most prominent merchandise in the job market. Ultimately, the danger of deskilling represents a specific type of automation governance problem, posing us questions such as: Will automation implementation affect humans’ skills? What kind of effect and on which skills? Are those skills important to retain?

The lens of distributed cognition suggests that a seamless, collaborative interaction between human agents and automated tools, as opposed to using automation as an isolated “black box”, could help to prevent the ill effects of deskilling. Indeed, the pertinent questions regarding skill maintenance problems can be discussed using the concepts that we introduced in the previous section. We stated that pragmatic actions can be delegated to automation provided that they are continuously and mindfully controlled by humans. Deskilling, seen in this view, is a threat to such a supervisory control. We base this view on our second case study (Rinta-Kahila et al. 2018, Appendix B), where an accounting firm had had to recover from gradually increased deskilling. For several years, sustained human control had not been needed in one of their core operations. The findings from that study provide lessons for reflection on the challenges related to automated pragmatic actions and possible mitigating strategies.

The company had implemented an intelligent system to automate a notoriously cumbersome part of accounting, fixed assets management (FAM). The software automated the depreciation allocations and calculations as well as the tax report creations, which are generally considered being prone to manual mistakes. Over the years, the accountants became used to the automated FAM process. However, its use was suddenly discontinued as a part of a managerial decision to streamline the overall information systems (IS) architecture. This resulted in surprising problems: the accountants were unable to conduct FAM anymore as the automated functions they had gotten used to were gone. It became apparent that their skills had eroded over the years. But why did deskilling happen and could it have been prevented?

In our analysis based on accountants’ retrospective stories about deskilling and recovery from it, we found that the key distinction in FAM lied in this pragmatic task’s *control* and its residual part – *execution*. We believe that the distinction between control and execution helps further understand vulnerabilities involved in distribution of pragmatic operations between humans and automation. Control depends on the understanding of the task: knowing the activities needed to conduct the task, maintaining the required skills and abilities, and being able to verify the correctness of the task’s outputs. With this competence, the agent (who should be a human, following the discussion above) will be capable of high-quality decision-making. Execution, on the other hand, reflects the actual conduction of the task; the mechanistic work. The difference of this dichotomy to the mindfulness–mindful dichotomy above is subtle but important. Mindfulness and control appear to pair naturally with each other, and so do mindlessness and execution. However, mindfulness and mindlessness complement each other while control and execution have a subordinate relationship: one needs to have some extent of control over a task to be able to execute it.

Our case company had implemented FAM in a manner that separated the task’s control and execution. It handed off the execution to the software and assigned the accountants as ‘controllers’ of the process.

However, the software had certain capabilities to exercise task control too: even procedures that tend to require some amount of human interpretation were effectively automated, and the software's third-party developers kept it up to date so as to respond to any legislative changes in a timely manner. In addition, the software included output-verification features that ensured the correctness of FAM reports. With such advanced features, mindful human control was no longer necessary for error-free execution. This made the accountants gradually complacent and willing to relinquish also the control of the task to the software. As no organizational initiatives were taken to ensure the accountants' knowledge maintenance and development, the software slowly turned into a black box. The accountants found little motivation to engage in learning or even in maintaining their existing stock of knowledge – they were happily sitting in the backseat while letting the automation take the steering wheel. Eventually both the control and execution of FAM processes was in the automation's responsibility, which led to a temporary trouble when the software was replaced with the new streamlined IS architecture.

To summarize, in this case study we found that pragmatic operations can be further divided into control and execution components, of which humans are necessary in the former part should the latter be automated. A sophisticated automation tool, however, may start to overtake also the control responsibilities. Via this mechanism, automation can become a causal link to the threat of deskilling. Based on the learnings from this case, we provide two recommendations aimed at harnessing the benefits of automation while preventing its negative effects on human skills:

**Recommendation 3:** Implement automation to execute pragmatic tasks but make sure human workers retain the control of those tasks in order to sustain their skill level.

While automating complex, knowledge-intensive procedures can yield significant cost savings to an organization, managers should consider whether this comes at the expense of losing the organization's human-embedded knowledge capital. Algorithms operate mindlessly with no understanding of why they perform the appointed tasks and what their larger implications are. Relying on an automation with task-control capabilities, humans may lose their conceptual grip on operations, become complacent, and subsequently also lose their mindful control of tasks. Although automation may produce the desired outputs when operating as a black box, it can be dangerous for an organization to forget how exactly those outputs came to be and how to verify their correctness. When the development and maintenance of one's knowledge regarding the automated work task is offloaded solely to automation and its provider, control of the work task gradually disappears. Rather, automation should be implemented as an electronic colleague (Arnold and Sutton 1998), that operates in an interactive manner with human workers. There, automation would assume the role of a subordinate instead of a peremptory authority, supporting humans who remain in charge of controlling the automation, essentially by developing their own expertise and mindfully verifying automation's outputs.

**Recommendation 4:** To retain task control, organize automated operations' inner workings so that they use the same information representations that also human workers will use in their supervisory control roles.

An organization's distributed cognition, composed of humans and automation, is vulnerable to deskilling if the modular division of pragmatic responsibilities into control and execution components masks the inner workings of those responsibilities, i.e., black-boxes the process knowledge. Therefore, while humans may no longer manually execute tasks, they should still retain a conceptual understanding of execution. Since human workers typically act as automation's supervisors, considering different types of supervisory control (previously discussed in terms of supervising human workers' performance, see Challagalla and Shervani 1996) can inform human-automation teaming, too. Whereas capability control (maintaining and developing one's skills and abilities) and output control (monitoring the output quality) represent higher levels of supervisory control, *activity control* is closely connected to the hands-on execution of the task as it refers to the specification and monitoring of work-task activities. After all, even though a human supervisor would not do hands-on execution of a work task, he or she must retain good amount of conceptual process knowledge to be able to effectively manage human workers who do the manual work. The same should hold for human-automation relationships.

Retaining activity control would allow humans to maintain some of the competence of the executing tasks, to counter-balance the possibility that the automation is competent in some of the control tasks of the human (as was the case in our case study). But how to achieve this? In the terminology of the distributed

cognition, it is beneficial if the computational transformations of information within the network of agents have “representational overlap,” meaning that both the humans and automation are able to operate on the same representations of information. Such representations may manifest as conceptual illustrations on the automation’s interface or in organization’s work procedures that connect to automation’s workings. This naturally requires automation to provide humans with an access to the process knowledge through meaningful representations. Similar idea has been applied in studying effective use of information systems (Burton-Jones and Grange 2012), where transparent interaction between user and system enable representational fidelity, which in turn allows user to take informed actions.

To navigate through the disruption caused by deskilling, the case company made determined attempts to recover its knowledge capital. Yet, the company also wanted to increase the extent of automation in the newly implemented system. Whereas the management recognized the need for more automation, they did not want to fall into the trap of complacency and deskilling again. Thus, to achieve the desired representational overlap, the company engaged into a process of informing<sup>1</sup>. In essence, the company deciphered the knowledge that had been black-boxed by the previously implemented automation into explicit process descriptions and instructions.

In practice this meant that each accountant had to start learning the basics of FAM from scratch by working both individually and in teams. The relearned processes were then documented into explicit operating instructions intended to ensure that incumbent workers would retain the relearned knowledge and that new workers would gain a fundamental understanding of the FAM process right from the start. The relearning process was followed by developing more automation into the new system. However, the company made sure to automate only repetitive routine tasks, ensuring that accountants would maintain task control while leveraging the automation primarily for execution. The necessity of engaging in an informing process highlights the importance of strategically managing organizations’ collective knowledge capital in the era of increasing implementation of automation. This makes the question of how exactly one should manage knowledge across humans and automation agents a prevalent dilemma for managers. In another case study, which will be discussed next, we encountered crucial issues to be taken into account when managing knowledge embedded in and produced by automation.

### ***Managing Collective Knowledge across Humans and Automation Agents***

We have found that automation also has an impact on organizations’ stock of knowledge beyond the above-presented risk of deskilling and complacency. As more and more cognitive processes may be offloaded to automation, the computational agents become a greater part of organization’s collective knowledge. Thus knowledge needs to be managed across both humans and automation (Stone et al. 2016).

In the perspective of distributed cognition, organizational knowledge is embedded in information artifacts, cognitive agents, and their interactions. Pieces of information and knowledge are also often replicated across several entities (Ackerman and Halverson 2004). In order to generate value, organizational knowledge needs to be externalized beyond individual agents and coordinated to serve organizational goals (Alavi and Leidner 2001; Davenport et al. 1998). Externalization is however challenging, because knowledge should not be abstracted too far beyond its application; otherwise it loses its actionability (Hecker 2012). Therefore, while employing centralized knowledge management models may sound appealing because they may seem simpler to manage and maintain, local, re-contextualized viewpoints on the applicability of knowledge “in action” should also be respected.

We studied (Asatiani et al. 2019) a telecommunications company (Appendix C) operating across Nordic region facing a knowledge management dilemma in its implementation of RPA across local country units. The company wanted to capitalize on the collective knowledge that the RPA implementation generated and use it to increase new RPA projects’ efficiency and effectiveness. To achieve this, the company empowered local country units to develop solutions that worked best for their local context, but also maintained the macro-level control over the implementation processes. In this federated, also known as hybrid, governance structure, a central hub took care of procurement and maintenance while six local centers had the

---

<sup>1</sup> A process of production of data about the automated process through textual symbols, essentially converting tacit knowledge into explicit knowledge (Zuboff 1988).

responsibility for identifying areas suitable for automation, gathering requirements, installing, and adapting the automation for the tasks at hand.

We found that the federated approach encountered several difficulties (Asatiani et al. 2019). The central hub attempted to gather various projects' requirements, automation rulebooks, algorithms and process designs into a central repository, to be reused in future projects. Building up this collective body of knowledge would have allowed new projects to avoid common pitfalls, and leverage already existing solutions, then feeding these improvements back to the central hub for others to adopt. However, while the central hub succeeded in gathering some information artifacts, it failed in making the knowledge contained in the repository a shared organizational memory. Two primary problems emerged. First, the central repository contained many requirements, automation rulebooks, algorithms, and process designs that did not seem to address the immediate requirements of local projects. Second, there was a lack of communication and collaboration between the central hub and local units. In short, the link between action and knowledge was broken in this centralization process. We will now reconsider these findings from the perspective of distributed cognition – a viewpoint that we did not apply in our original publication because of its different focus, but which relates closely to distributed organizational knowledge.

Interpreted with the framework of distributed cognition, the experiences in the telecommunications company show how knowledge, embedded in agents' processes and individual memories, is not abstract and easily transferable across organizational units – in our case between the central hub and local units (Alavi and Leidner 2001; Davenport et al. 1998; Nonaka and Takeuchi 1995). In this case, the units had distinct organizational memories: each was largely a separate distributed cognition of its own. Due to a lack of cooperation and coordination across units, the organizational memories did not merge sufficiently. Joint projects, worker exchanges, and “worker-in-residence” arrangements were not created. Boundary objects (Star and Griesemer 1989) that could have served as information carriers between units had not yet evolved to facilitate communication between units, because their evolution takes time.

For knowledge sharing and intellectual capital management to succeed, it is important that members of the organization created networks of distributed cognitions by *externalizing* their information, and make sure to do this in an actionable manner. Humans distribute their cognition naturally: when performing complex cognitive tasks, such as in mathematical problems, most people use pen and paper to “offload” information from their memory and to visualize thought processes. Thus, pen and paper become cognitive artifacts that amplify agents' cognitive capacity (Hutchins 2006; Nemeth et al. 2004; Risko and Gilbert 2016) and mediate their collective work (Nemeth et al. 2004). This exemplifies externalization on a local level, in the immediate surroundings of the actor. Now, in the era of automation, computational implementations of automation can become elaborate forms of such externalization: instead of pen and paper, the process (or a part of it) is externalized to a computational agent.

Our case study shows that an externalization of knowledge has its limits: externalization to a remote hub proved too difficult. While the centralized repository collected externalized explicit information from human agents (e.g., process documentation) and knowledge contained within automation (e.g., algorithms and rulebooks), the information was detached from action. Management of knowledge capital needs to tread a middle ground and respect the stickiness of knowledge in actions and interactions between agents, humans and automation alike. Based on these experiences at the case company, we offer the following recommendations.

**Recommendation 5:** Make sure that the distributed knowledge can be re-contextualized into a local, actionable form in remote units.

This recommendation directs the focus on a rarely-attended issue in the mainstream knowledge management research. Previous literature (Alavi and Leidner 2001; Nonaka and Takeuchi 1995) has presented knowledge transfer as a process where local knowledge needs to be first externalized into generalized form, thereby making it transferable across organizational units. In the second step, the knowledge can be taken into use in a different part of the organization by recontextualizing it for the local needs of the setting. Therefore, managers are faced with two different challenges – generalization and re-contextualization – both of which have been found particularly thorny to solve. The challenge of the first step has stemmed from the difficulty of incentivizing workers into documenting their knowledge. In addition, the generalization of knowledge may have been suboptimal: written from a point of view which does not address those details that are relevant for the users of the knowledge, for example. RPA

implementation may change this picture for better, however, by reducing the need for workers to manually document their knowledge. When work tasks are coded into RPA-executable form, the tedious generalization effort is bypassed and the outcome is not only a document of knowledge, but also a directly executable rulebook or script module. There is no need for generalizing the knowledge in an abstract manner with a human user primarily in mind. Instead, externalized knowledge, represented as code and scripts, can be immediately capitalized by other automated agents. However, following Recommendation 4, this externalization into executable code and data should nevertheless be done in a manner that provides humans a conceptual understanding of automation's inner workings.

The second challenge – that of knowledge localization; i.e., how the RPA rulebook can be taken into use elsewhere in the organization – however does remain, and is the reason for the recommendation above. As it currently stands, RPA agents residing within different local contexts are unable to directly *learn* from each other. Therefore, knowledge externalized from one RPA implementation to a central repository cannot be directly applied to an RPA instance elsewhere. Traditional hazards of knowledge transfer, such as fragmentation, de-contextualization, and information overload (Massimo and Mariano 2007) are not mitigated by introducing automation. Even in this context the problem of knowledge transfer remains largely sociological rather than technical one (Jasimuddin et al. 2012).

However, the managers can now exert more attention to solving the second challenge. This is because automation provides the necessary incentives for the first challenge – externalization – to succeed. The second challenge's requirement to recontextualize and localize the knowledge should be solved through active interaction between both humans and automation across different units. Managers should put in place a hybrid mechanism to transfer explicit knowledge (e.g., rulebooks and algorithms) through codification and tacit knowledge (e.g., best practices of implementation of automation) through personalized interaction.

**Recommendation 6:** Recruit knowledge facilitators to facilitate knowledge transfer across local contexts.

The immediate upshot from the previous recommendation is the need for experts in organizations' changes. There is an increasing need for people who are able to, interested in, and mandated to tailor existing RPA scripts and solutions to local contexts in other parts of the organization. Earlier knowledge management scholars (Davenport et al. 1998; Jasimuddin et al. 2012) have suggested roles of knowledge facilitators, knowledge managers, or knowledge administrators for this task. Responsibilities of such a knowledge facilitator are to maintain knowledge repository, support knowledge externalization and re-contextualization, and assist local units in finding the right sources of knowledge.

A knowledge facilitator in the context of automation has a dual role. In addition to serving as a conduit for creating a shared organizational memory among humans, the knowledge facilitator should possess technical expertise (e.g., programming skills) to interact with automation and customize it. Even if RPA scripts are intended for reuse, there will remain details in their implementation in new contexts where localization is needed. The knowledge facilitator should be able to understand knowledge contained within automation, and be able to adapt it to the local needs.

## **Discussion**

While automation increasingly permeates contemporary organizations, as of today we do not seem to be moving towards a future where automated machines replace humans at large scale. Instead, in our case studies we have observed increasing levels of symbiosis (Licklider 1960) between human workers and a variety of smaller, specialized automation tools, as organizations adopt human-machine hybrid activities (Daugherty and Wilson 2018) – with varying success. As such, we believe that an automation revolution in knowledge work organizations will not be accomplished through a simplistic handover of work tasks from human workers to machines. Rather, managers need to think work tasks in terms of their constituent activities and functions, and determine which parts benefit from humans' inherent strengths, and where the advantages of automation can be most effectively leveraged. Further, while the key to successful human-automation partnerships lies in an effective division of labor between these agents, it is also important to ensure that the agents do not operate in isolation but inform each other's activities. We have taken a managerial perspective on the phenomenon, paying particular attention to the distribution of tasks and transplantation of skills and intellectual capital between humans and automation.

Viewing human–automation teaming as a distributed network of agents instigates certain complexity. In an ideal world, the human–machine symbiosis would take place seamlessly, with automation taking away voluminous routine tasks, helping humans focus on ones requiring creativity and decision-making. However, the symbiosis does not come without inherent challenges. It requires carefully crafted governance structures and managerial vigilance to materialize the gains promised by the automation advocates. In the example of ship navigation (Hutchins 1995), the task was successfully completed through extremely precise coordination, taking place in three-minute cycles, involving human and non-human agents, acting as one distributed unit with a shared goal. Contemporary organizations implementing automation may need to develop analogous precision for the human–machine collaboration. Managers need to develop appropriate strategies for addressing the distribution and ownership of the tasks and associated intellectual capital, they must aim to accrue automation expertise to serve long-term success of the organization, and they need to ensure continuous development of human and automation agents. Whereas automation tools may enable an organization to leverage human resources with a lower level of intellectual capital than before, this is probably not a fruitful long-term strategy, especially with knowledge work organizations. Instead, we stress the importance of assessing their long-term implications of automation implementation decisions by taking the human element into account. As such, managers should think beyond investing in mere technological capabilities and consider how the abilities of human workers could be maintained and improved on par with the technological advances. Indeed, while implementing smart technologies may yield benefits on the short term, in the long run their users need to be smart, too, if an organization wishes to remain successful (Arnold 2018).

An overarching goal for this paper has been to present starting points for a long-lasting conceptual foundation for automation implementation in knowledge work organizations. Such a foundation should retain its relevance even in the fast-paced automation development context. In the environment where “*what used to be considered as intelligent behavior exhibited by machines five years ago is now considered barely noteworthy*” (Kaplan and Haenlein 2019), it may seem hard to project the long-term outlook of automation. Yet we believe that the distributed cognition research has a potential for offering this foundation, and that the case studies and recommendations presented in this paper are convincing evidence of this potential. Automation implementation also offers a new test for distributed cognition’s concepts and has a potential to accelerate the development of new concepts and models within its conceptual framework. By being nuanced in terms of automation’s cooperative role, guided by conceptual understanding about their transformative potential in organizational practices, it is possible to sustainably manage the ongoing transformation of knowledge work at workplaces and business. In particular, the possibility to consider both humans and automation as cognitive agents capable of shared information processing provides a level of abstraction that is both independent of individual breakthroughs in automation development and concrete enough for being useful.

## Conclusion

In this paper, we have viewed automation implementation in knowledge work organizations through the lens of distributed cognition. This analytical lens has guided us to consider organizations’ operations as networked processes where agents – both humans and automation – jointly process information, represent it in more informative ways, and act in order to achieve joint goals. In such a view, humans and automation necessarily complement each other’s strengths and weaknesses when jointly performing cognitive tasks. Our past experience and ongoing research activities in organizations highlight the complexly interwoven relationship between humans and intelligent technologies, where automation would ideally manifest itself as a mosaic of transparently-operating computational agents instead of a black-boxed monolith. While our research demonstrates the benefits of thinking automation implementation from the distributed cognition perspective, it also points out potential pitfalls of failing to do so. The six recommendations put forward in this article are an attempt to draw attention to some of the key issues in organizational automation implementation and to provide concrete managerial guidance on them. Recommendations 1 and 2 concern the task division between human and automation. First, we recommend managers to seek for opportunities to apply automation in epistemic tasks but be careful when automating pragmatic tasks. Second, we suggest managers to divide work into their mindful and mindless components, and to offload the mindless part to automation while keeping the mindfulness-requiring part to humans. Recommendations 3 and 4 are put forth to mitigate automation’s potentially detrimental effects on human workers’ skills. Building on our insights on task division, we contend that automation should be implemented to execute pragmatic tasks

while making sure human workers retain the control of those tasks. Moreover, we suggest that leveraging representational overlap in the control and execution of the task may help humans to maintain their activity control which thereby mitigates deskilling. Finally, recommendations 5 and 6 help in managing collective knowledge across human and automation agents. We suggest that managers should strive for reusability of distributed knowledge (whether produced by humans or automation) by re-contextualizing it for the use of local units. Recruiting knowledge facilitators can help to ensure the localization of knowledge is successful.

As a result, our insights may inform managers looking to materialize the full potential of automation tools in a productive and sustainable manner that also attends to existing social capital that workers possess. While the precise strategy and technological tools will largely depend on the context of each organization, these recommendations can help managers to strike the right balance between pushing for innovative automation technology implementations and empowering their workers.

## References

- Ackerman, M. S., and Halverson, C. 2004. "Organizational Memory as Objects, Processes, and Trajectories: An Examination of Organizational Memory in Use," *Computer Supported Cooperative Work (CSCW)* (13:2), pp. 155–189.
- Alavi, M., and Leidner, D. 2001. "Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues," *MIS Quarterly* (25:1), pp. 107–136.
- Arnold, V. 2018. "The Changing Technological Environment and the Future of Behavioural Research in Accounting," *Accounting & Finance* (58:2), pp. 315–339.
- Arnold, V., and Sutton, S. G. 1998. "The Theory of Technology Dominance: Understanding the Impact of Intelligent Decision Aids on Decision Maker's Judgments," *Advances in Accounting Behavioral Research* (1:3), pp. 175–194.
- Asatiani, A., Kämäräinen, T., and Penttinen, E. 2019. "Unexpected Problems Associated with the Federated IT Governance Structure in Robotic Process Automation (RPA) Deployment," No. 2/2019, Business+Economy Working Papers, Aalto University.
- Bainbridge, L. 1983. "Ironies of Automation," *Automatica* (19:6), pp. 775–779.
- Burton-Jones, A., and Grange, C. 2012. "From Use to Effective Use: A Representation Theory Perspective. Information Systems Research," *Information Systems Research* (24:3), pp. 632–658.
- Butler, B. S., and Gray, P. H. 2006. "Reliability, Mindfulness, and Information Systems," *MIS Quarterly* (30:2), pp. 211–224.
- Carr, N. 2015. *The Glass Cage: How Our Computers Are Changing Us*, W. W. Norton & Company; 1 edition (September 8, 2015).
- Challagalla, G. N., and Shervani, T. A. 1996. "Dimensions and Types of Control: Supervisory Effects on Salesperson Performance and Satisfaction," *Journal of Marketing* (60:1), pp. 89–105.
- Daugherty, P., and Wilson, J. 2018. *Human+Machine - Reimagining Work in the Age of AI*, Boston, Massachusetts: Harvard Business Review Press.
- Davenport, T. H., Prusak, L., and others. 1998. *Working Knowledge: How Organizations Manage What They Know*, Harvard Business Press.
- Ford, M. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*, Basic Books.
- Frey, C. B., and Osborne, M. A. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change* (114), Elsevier B.V., pp. 254–280.
- Hecker, A. 2012. "Knowledge Beyond the Individual? Making Sense of a Notion of Collective Knowledge in Organization Theory," *Organization Studies* (33:3), pp. 423–445.
- Hollan, J., Hutchins, E., and Kirsh, D. 2000. "Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research," *ACM Transactions on Computer-Human Interaction* (7:2), pp. 174–196.
- Hutchins, E. 1995. *Cognition in the Wild*, Cambridge, MA: The MIT Press.
- Hutchins, E. 2006. "The Distributed Cognition Perspective on Human Interaction," *Roots of Human Sociality: Culture, Cognition and Interaction*, pp. 375–398.
- Jasimuddin, S. M., Connell, N., and Klein, J. H. 2012. "Knowledge Transfer Frameworks: An Extension Incorporating Knowledge Repositories and Knowledge Administration," *Information Systems Journal* (22:3), pp. 195–209.
- Kaplan, A., and Haenlein, M. 2019. "Siri, Siri, in My Hand: Who's the Fairest in the Land? On the Interpretations, Illustrations, and Implications of Artificial Intelligence," *Business Horizons* (62:1),

- Elsevier, pp. 15–25.
- Kirsh, D., and Maglio, P. 1994. “On Distinguishing Epistemic from Pragmatic Action,” *Cognitive Science* (18:4), pp. 513–549.
- Lacity, M., and Willcocks, L. 2016. “Robotic Process Automation at Telefónica O2,” *MIS Quarterly Executive* (15:1), pp. 21–35.
- Licklider, J. C. R. 1960. “Man-Computer Symbiosis,” *IRE Transactions on Human Factors in Electronics* (HFE-1:1), pp. 4–11.
- Massimo, F., and Mariano, S. 2007. “Information Technology Repositories and Knowledge Management Processes: A Qualitative Analysis,” *VINE - the Journal of Information and Knowledge Management Systems* (37:4), pp. 440–451.
- Millman, Z., and Hartwick, J. 1987. “The Impact of Automated Office Systems on Middle Managers and Their Work,” *MIS Quarterly* (11:4), p. 479.
- Nemeth, C. P., Cook, R. I., O’Connor, M., and Klock, P. A. 2004. “Using Cognitive Artifacts to Understand Distributed Cognition,” *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*. (34:6), pp. 726–735.
- Nonaka, I., and Takeuchi, H. 1995. *The Knowledge Creation Company: How Japanese Companies Create the Dynamics of Innovation*, New York: Oxford University Press.
- Penttinen, E., Kasslin, H., and Asatiani, A. 2018. “How to Choose Between Robotic Process Automation and Back-End System Automation ?,” in *26th European Conference on Information Systems, ECIS 2018*, pp. 1–14.
- Perry, M. J., Fruchter, R., and Rosenberg, D. 1999. “Co-Ordinating Distributed Knowledge: A Study into the Use of an Organisational Memory,” *Cognition, Technology & Work* (1:3), pp. 142–152.
- Rinta-Kahila, T., Penttinen, E., Salovaara, A., and Soliman, W. 2018. “Consequences of Discontinuing Knowledge Work Automation: Surfacing of Deskilling Effects and Methods of Recovery,” in *Proceedings of the 51st Hawaii International Conference on System Sciences*, pp. 5244–5253.
- Rogers, Y., and Ellis, J. 1994. “Distributed Cognition -an Alternative Framework for Analyzing and Explaining Collaborative Working,” *Journal of Information Technology* (9:2), pp. 119–128.
- Salovaara, A., Lyytinen, K., and Penttinen, E. 2019. “High Reliability in Digital Organizing: Mindlessness, the Frame Problem, and Digital Operations,” *MIS Quarterly* (43:2), pp. 555–578.
- Schultze, U. 2000. “A Confessional Account of an Ethnography about Knowledge Work,” *MIS Quarterly* (24:1), pp. 3–41.
- Serrano, N., Hernantes, J., and Gallardo, G. 2014. “Service-Oriented Architecture and Legacy Systems,” *IEEE Software* (31:5), pp. 15–19.
- Silverman, W. 1966. “The Economic and Social Effects of Automation in an Organization,” *American Behavioral Scientist* (9:10), pp. 3–8.
- Star, S. L., and Griesemer, J. R. 1989. “Institutional Ecology, Translations’ and Boundary Objects: Amateurs and Professionals in Berkeley’s Museum of Vertebrate Zoology, 1907-39,” *Social Studies of Science* (19:3), Sage Publications London, pp. 387–420.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Julia, H., Kalayanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., and Teller, A. 2016. “Artificial Intelligence and Life in 2030,” *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*.
- Walker, G. H., Stanton, N. A., Baber, C., Wells, L., Gibson, H., Salmon, P., and Jenkins, D. 2010. “From Ethnography to the East Method: A Tractable Approach for Representing Distributed Cognition in Air Traffic Control,” *Ergonomics* (53:2), pp. 184–197.
- Weick, K. E., Sutcliffe, K. M., and Obstfeld, D. 1999. “Organizing for High Reliability: Processes of Collective Mindfulness,” in *Research in Organizational Behavior* (Vol. 1), R. S. Sutton and B. M. Staw (eds.), pp. 81–123.
- Zuboff, S. 1988. *In the Age of the Smart Machine: The Future of Work and Power*, Basic books.

## Appendix A: Fact Sheet on Case Study A

<b>Case A: Malware protection company</b>	
Publication	Salovaara, A., Lyytinen, K., and Penttinen, E. 2019. "High Reliability in Digital Organizing: Mindlessness, the Frame Problem, and Digital Operations," <i>MIS Quarterly</i> (43:2), pp. 555–578
Research question	How does a digital high-reliability organization combat the frame problem and organize its operations as a collectively mindful sociotechnical system?
Automation artefact	A rule-based engine that handles the incoming influx of over one million malware samples per day.
Theory	High-reliability organizations (HROs) are exceptional sociotechnical systems that can operate in a nearly error-free manner in environments wherein most other organizations constantly run the risk of major accidents. The HRO theory posits that the unique cognitive mindset by which HROs operate is manifested in five system-level characteristics: preoccupation with failure, reluctance to simplify interpretations, under-specification of structures, commitment to resilience, and sensitivity to operations. We complemented the theorizing on high-reliability operations with an analysis of the threats that properties of digital material pose for high reliability, and built on the framework of distributed cognition.
Method and empirical data	Qualitative exploratory revelatory case study on a malware protection company based on four rounds of interviews and one round of participant observation. Open thematic qualitative analysis was applied to analyze how malware threats were detected and dealt with. Three analytical distinctions were employed to analyze malware protection operations: actors that performed them (human-based or digital), the nature of cognition (mindful or mindless), and the purpose (epistemic or pragmatic).
Findings	The findings of the present paper build directly on the findings on the latter two of the distinctions (see above). The study produced a three-layer structure of operations: core layer, improvement layer, and anticipation layer. The malware company's core layer is essentially an algorithmic protection block consisting of pragmatic operations and is entirely digital. The improvement layer continuously repairs the core layer's operations to mitigate its unavoidable errors and related failures. The anticipation layer counters threats of entrenchment. It consists of epistemic efforts to turn "unknown unknowns" (black swans) into "known unknowns." This layer's operations are outward-facing, to anticipate threats proactively and reveal needs to update rules or restructure operations. The study also shed light on the nature of digital operations in achieving high reliability. First, epistemic digital operations essentially inform operations. Second, pragmatic digital control operations prevent physical operations in case of danger of grave human error. Third, pragmatic executing operations influence digital inputs directly without humans in the action loop.

**Table A1. Case A: Malware protection company**

## Appendix B: Fact Sheet on Case Study B

<b>Case B: Accounting company</b>	
Publications	<p>Rinta-Kahila, T., Penttinen, E., Salovaara, A., and Soliman, W. 2018. "Consequences of Discontinuing Knowledge Work Automation – Surfacing of Deskilling Effects and Methods of Recovery," in <i>Proceedings of the 51st Hawaii International Conference on System Sciences</i>, pp. 5244–5253.</p> <p>Rinta-Kahila, T. 2018. "Pulling the Plug: The Concept, Process, and Outcomes of Organizational Information System Discontinuance," Dissertation, Aalto University.</p>
Research question	How does automation lead to the deskilling of knowledge workers in an organizational context?
Automation artefact	A highly automated information system to handle classifications, calculations, depreciations, and report production associated with fixed assets management.
Theory	Theoretical review on the effect of automation on workers' skill repertoire. Prior literature posits that hands-on routines often constitute the basis of humans' expertise, a phenomenon known as the generation effect; conversely, ceasing to conduct such hands-on routine work has been argued to lead to a degeneration effect where the accumulated expertise disappears. Prior research also informs on automation complacency (handing task execution to automation and leaving a human to supervise it, and refers to a decline in human's supervisory performance) and automation bias (operator gives undue weight to the information provided by the system, ignoring other sources of information, such as environmental cues that the automation cannot detect).
Method and empirical data	Case study on an accounting company where workers had become deskilled due to relying on automated fixed assets management system over the course of seven years. Altogether 19 interviews were conducted in three phases. Interview data were analyzed to identify temporal stages in the process of deskilling with the aim of producing a process understanding of how environment, organization and individual level events affected workers' skill repertoire.
Findings	The study identified the notions of control and execution as useful analytical devices to understand what happens in a gradual deskilling process. In the case company, the transfer of execution to the automation was followed by a detrimental, unintended transfer of control of different kinds (activity control, output control, and capability control) to the automation. When the automation was discontinued and workers' manual skills were needed again, this caused a shock as it revealed the deskilling and increased the workers' work burden disrupting the overall accounting process. The company finally employed several coping strategies and navigated through the difficult period and re-established workers' skills. Coping strategies included emotion-focused coping by ventilating frustration but still submitting to passive acceptance of the situation, and problem-focused coping by relearning the fixed assets management processes, both individually and in groups. Problem-focused coping occurred also at the organizational-level through buying consultancy services from the provider of the discontinued system and sending the accountants to formal training events offered by the local Chamber of Commerce.

**Table B1. Case B: Accounting company**

## Appendix C: Fact Sheet on Case Study C

<b>Case C: Telecommunications company</b>	
Publication	Asatiani, A., Kämäräinen, T., and Penttinen, E. 2019. "Unexpected Problems Associated with the Federated IT Governance Structure in Robotic Process Automation (RPA) Deployment," No. 2/2019, Business+Economy Working Papers, Aalto University.
Research questions	How to set up a federated governance structure for Robotic Process Automation (RPA)? What are the opportunities and challenges associated with the federated governance of RPA?
Automation artefact	Robotic Process Automation (RPA)
Theory	Prior literature on IT governance of RPA suggests that a federated IT governance could help organizations to harness benefits of both centralized governance (e.g. centralized procurement, common IT policy) and decentralized governance (e.g. IT solutions addressing specific local business needs).
Method and empirical data	Case study conducted in a telecommunications company (Telco) which is one of the early adopters of RPA in Northern Europe. Telco operates its RPA development and deployment through a federated governance model. Eight interviews examined how the federated governance model was implemented at Telco and what the benefits and drawbacks associated with the governance model were.
Findings	The study revealed two unexpected challenges in benefit realization. Although Telco aspired to enable efficient reuse of RPA components through its centralized center of excellence, this type of knowledge transfer across RPA projects was found to be unrealistic due to the variance in IT legacy systems and work practices across the local units. Also, striking a balance between local needs and what the centralized center of excellence could offer was seen problematic. Evidence of this was the fact that the training and technical implementation support by the center of excellence was perceived subpar by the local units.

**Table C1. Case C: Telecommunications company**