

6-8-2022

Artificial Intelligence in Information Systems: State of the Art and Research Roadmap

Pär J. Ågerfalk

Uppsala University, Sweden, par.agerfalk@im.uu.se

Kieran Conboy

National University of Ireland, kieran.conboy@nuigalway.ie

Kevin Crowston

Syracuse University, crowston@syr.edu

Jenny S. Z. Eriksson Lundström

Uppsala University Informatics and Media, jenny.eriksson@im.uu.se

Sirkka Jarvenpaa

University of Texas at Austin, sirkka.jarvenpaa@mcombs.utexas.edu

See next page for additional authors

Follow this and additional works at: <https://aisel.aisnet.org/cais>

Recommended Citation

Ågerfalk, P. J., Conboy, K., Crowston, K., Eriksson Lundström, J. S., Jarvenpaa, S., Ram, S., & Mikalef, P. (2022). Artificial Intelligence in Information Systems: State of the Art and Research Roadmap. *Communications of the Association for Information Systems*, 50, pp-pp. <https://doi.org/10.17705/1CAIS.05017>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in *Communications of the Association for Information Systems* by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Artificial Intelligence in Information Systems: State of the Art and Research Roadmap

Authors

Pär J. Ågerfalk, Kieran Conboy, Kevin Crowston, Jenny S. Z. Eriksson Lundström, Sirkka Jarvenpaa, Sudha Ram, and Patrick Mikalef



Artificial Intelligence in Information Systems: State of the Art and Research Roadmap

Pär J. Ågerfalk
Uppsala University
par.agerfalk@im.uu.se

Kieran Conboy
National University of Ireland, Galway
kieran.conboy@nuigalway.ie

Jenny Eriksson Lundström
Uppsala University
jenny.eriksson@im.uu.se

Sudha Ram
University of Arizona
sram@email.arizona.edu

Kevin Crowston
Syracuse University
crowston@syr.edu

Sirkka L. Jarvenpaa
University of Texas at Austin
sirkka.jarvenpaa@mcombs.utexas.edu

Patrick Mikalef
Norwegian University of Science and Technology
patrick.mikalef@ntnu.no

Abstract:

Many would argue that artificial intelligence (AI) is not only technology but also a paradigmatic shift in the relationship between humans and machines. Much literature assumes that AI-powered practices substantially differ from and profoundly change organizational structures, communication, affordances, and ecosystems. However, AI research remains fragmented and often lacks clarity. While the information systems (IS) discipline can play a pivotal role in AI's emergence and use, the discipline needs a clear direction that specifies how it can contribute and its key research themes and questions. This paper draws on a professional development workshop that we organized at the 2020 International Conference on Information Systems and the discussions that followed. We summarize and synthesize how AI has impacted organizational practices over five decades and provide views from various perspectives. We identify weaknesses in the current AI literature as measured against conceptual clarity, theoretical glue, cumulative tradition, parsimony, and applicability. We also identify direct actions that the IS research community can undertake to address these issues. Finally, we propose a next-step research agenda to guide AI research in the coming years.

Keywords: Artificial Intelligence, Value, Concepts, Coherence, Research Agenda.

This manuscript underwent editorial review. It was received 5/14/2021 and was with the authors for five months for one revision. The Associate Editor chose to remain anonymous.

1 Introduction

This paper reports on a practice development workshop on “Artificial Intelligence: Beyond the Hype”, which we organized at the 2020 International Conference on Information Systems. Specifically, we focus on summarizing the presentations and discussion at the workshop. We also discuss the main takeaways from the event in terms of imminent and emerging areas for research and position them in the current AI discourse in the information systems (IS) discipline.

Given that AI applications constitute algorithmic and computational tools for digitally enabled practices and communication, it makes sense for the IS discipline to focus on AI as a core subject (Ågerfalk, 2020). However, as often occurs with new and emerging IS phenomena, practitioners (and consultants) have almost entirely led the research to create, promote, and disseminate AI technologies. Yet, the application areas and impacts concern virtually every aspect of contemporary society. Many have the sense that, as AI increases in pervasiveness, scholars remain searching for conceptual clarity.

At the workshop, Fred Niederman said:

I think that...most of us most of the time speak in clichés. As the technology becomes more used, will it push us toward more cliché or away from it? If I had to guess, I'd say it would push different people in different ways and the same people in different ways in different circumstances. We already, at times, have difficulty distinguishing truth from falsehood and end up believing those we trust rather than those we do not (assuming that the statements aren't inherently self-contradictory). Will more extensive AI programs make it easier or harder to distinguish “true” from “false” statements?

In more recent times, AI application research has begun to gain momentum, which we can see in the increasing number of dedicated journal special issues, conferences, conference tracks, and workshops on AI. Researchers have made continuous, sustained claims about AI's value and its transformative capability (Mikalef & Gupta, 2021). Many people now claim AI to constitute a paradigmatic shift rather than simply a disruptive technology (Rahwan et al., 2019; Iansiti & Lakhani, 2020). Furthermore, many people have said AI practices to substantially differ from other (non-AI) practice, and predicted that AI will change organizational structures, communication, affordances, and ecosystems in a way that differs from current technologies (Davenport, 2018; Ågerfalk, 2020). In the editorial to the Journal of the Association for Information special issue on AI in organizations, Benbya et al. (2021) reviewed the recent literature on AI in information systems in terms of automation, engagement, insight/decision making, and innovation. Similarly, the “managing AI” MIS Quarterly special issue editors (Berente et al., 2021) discussed AI as “the frontier of computing” that involves autonomy, learning, and inscrutability. Despite these current efforts and apparent progress, AI research involves many challenges, such as:

- 1) How does AI resemble and differ from other material and computational practices?
- 2) Do such similarities and differences matter, and, if so, why?
- 3) What unique methodological and theoretical challenges does one face in carrying out IS research on AI?

IS research typically focuses on phenomena at the intersection of social practices and digital technologies. Therefore, it can offer in-depth and nuanced knowledge about what differentiates various technologies—the familiar and the many novel techniques that one finds under the “AI” umbrella—and their impact on practices and organization structures (Gomes, 2019). Such discussions could include research on understanding and overcoming barriers to the AI adoption in organizations, the future of work in conjunction with AI automation, algorithmic and data bias in AI, how it affects organizations and decision making, and many other topics. In addition, the interface between humans and AI-enabled robots, automation in an age of data deluge, and explainable and accountable AI practices illuminate how different research perspectives can inform the discourse beyond the current hype surrounding AI in IS and throughout the management sciences.

In this report, we examine actions that can strengthen AI research in IS to enable researchers in the discipline to understand the current state of AI, which includes unsolved AI problems, its potential risks and benefits, and the scientific and philosophical questions that it raises for understanding digital agency and human intelligence (Larson, 2010; Aleksander, 2017; Schwartz et al., 2019). As we discuss in Section 2, researchers and practitioners have defined and applied the AI term in different ways. Therefore, we

encourage readers to recognize that the different panel members and authors may use the term somewhat differently.

2 Criteria for Judging “Good” AI Research

To provide a structure for analyzing current AI research, we drew on well-cited tenets about what constitutes a strong concept or theory. These tenets provided the structural basis for the expert opinion pieces, the workshop discussion, and subsequent analysis after the event. Specifically, we considered:

- 1) **Clarity:** one of the most fundamental attributes of concepts and constructs is that they clearly and understandably communicate meaning (Metcalf, 2004; Dubin, 1976, 1978; Weick, 1989).
- 2) **Theoretical glue:** solid underlying logic and rationale should support any good concept or theory. Whetten (1989) refers to such logic as “theoretical glue” that should bind all the factors together.
- 3) **Cumulative tradition:** a good concept or theory should cumulatively build on existing research (Dubin, 1978), yet IS researchers have not done so particularly well (Benbasat & Zmud 2003; Keen, 1991, 1980). Keen (1991) notes that most concepts and concern areas in IS research are not as “new” as authors often claim and “turn out to have long roots”.
- 4) **Parsimony:** authoritative works on concept development usually advocate a parsimonious approach; that is, removing factors that provide little additional value to our understanding (Whetten, 1989).
- 5) **Applicability:** the range of applications of a concept is a key criterion for judging its quality (Metcalf, 2004; Dubin, 1976, 1978; Weick, 1989), so it should be applicable in a wide variety of contexts.

3 Summary of Workshop Activity

In this section, we present the workshop results while referring to the criteria above. Specifically, in Section 3.1, we summarize the presentations that the five speakers—experts with various backgrounds—made to help readers understand the impact that AI has had on organizational practices over five decades. In Section 3.2, we briefly analyze the key concepts presented in and discussed across the seven round table discussion sections.

3.1 The Current State of AI Research

3.1.1 A Synthesis of Five Decades of AI Research (Sudha Ram)

Table 1. Sudha Ram Summary

To set the scene, Sudha Ram provided summarized the history of AI and AI research. Starting with the 1956 Dartmouth workshop that coined the term AI, she explained the development from the symbolic AI paradigm through the weak versus strong AI dispute to today’s deep learning methods and applications.

Sudha Ram began her presentation by summarizing the history of AI and AI research. She pointed to a workshop in Dartmouth in 1956 that John McCarthy, Marvin Minsky, Claude Shannon, and Nate Rochester (whom many would later refer to as the AI founding fathers) organized. At this workshop they, alongside Allan Newell and Herbert Simon, coined the term “artificial intelligence” and predicted great optimism for the domain. They predicted that, within 10 to 20 years, humanity would successfully build a fully intelligent machine. The workshop discussed topics related to natural language processing, machine learning, neural networks, reasoning, and creativity and provided an agenda for AI research. These topics remain major topics that researchers discuss in AI research conferences and journals. However, we have not yet attained or surpassed the prediction they made about creating a machine that possesses the intelligence that a human does. Questions about what “full” intelligence really means and how can we achieve it pertain highly to this discussion. Intelligence has several dimensions to it (e.g., emotional, verbal, logical, social, etc. dimensions). In efforts to develop AI, researchers and practitioners have focused on two distinct approaches: 1) a mathematical approach that relies on deductive reasoning or statistics using inductive reasoning and 2) a biological or psychological approach to create reasoning akin to the human

brain. These two distinct approaches produced myriad methods with little collaboration or conversation between them.

These approaches led to two separate paradigms. The first one focused on using symbolic logic and representing knowledge in the form of symbols and operators. This paradigm fostered the idea that one does not need to mimic the human brain's biology but that one can program rules for reasoning. As a result, expert systems such as MYCIN were developed and rose in application and popularity. The second paradigm, the subsymbolic approach, which psychologists such as Rosenblatt spearheaded, proposed an idea called "perceptrons" based on inspiration from how the human brain functions. The perceptrons were to mimic how neurons in the brain fire depending on assigned weights and thresholds for the inputs. The former approach was transparent and interpretable, the latter was not. The latter approach also had no algorithm to learn the appropriate weights in situations with multiple neuron layers. Thus began the symbolic approach's dominance, which remained in vogue until the early 90s when researchers realized that symbolic AI could not generate general-purpose problem-solving strategies. During this time, they also more generally realized the difficulty in developing AI and in emulating the human brain in a machine. As a result, the AI winter began.

In the 2000s, two major developments spurred contemporary AI development and, thus, began the AI spring: large-scale computational power and vast amounts of available data due to the World Wide Web (WWW) as well as the development of algorithms for training "deep" neural networks. Today, we have multilayer neural networks (also known as deep learning methods) such as convolutional neural nets and recurrent neural nets. While these neural nets started as black boxes, we now have "attention mechanisms" that can open them up to some extent. However, these supervised techniques need data and examples to learn. Thus, emerging areas now include unsupervised methods such as reinforcement learning, which start with a goal and learn to progress toward it. AI has continued to gain strength as we can see with Google's real-time translations, Siri, Alexa, self-driving cars, YouTube's automated video subtitles, Skype's ability to simultaneously translate between languages in video calls, Amazon's cashier-less stores, COBOTS in manufacturing and warehouses, and digital twins.

Researchers have also fervently discussed narrow (or weak) versus general (or strong) AI and how the AI domain needed to focus on the latter. The former refers to AI that can solve specific narrowly defined tasks (e.g., play Go). As an example, one application, AlphaGO, can beat the best human players in the world. However, it could only perform at a mediocre level on other seemingly simpler tasks (e.g., play chess or checkers). In contrast, the latter (also called artificial general intelligence) goes way beyond having the ability to perform only one or even two narrowly defined tasks. More discussion ensued on what human-level AI would comprise, such as sensing, seeing, understanding, thinking, and creating. These need to be integrated and intertwined with each other. How to achieve this with a combination of approaches has now become the mantra for creating AGI.

3.1.2 The Dark Side of AI (Patrick Mikalef)

Table 2. Patrick Mikalef Summary

Patrick Mikalef discussed the need to adopt a dark-side perspective in AI research. He noted that, despite an existing body of research on AI's dark side, we systematic studies that can offer more relevant practical guidance.

In his presentation, Patrick Mikalef discussed the need to adopt a dark-side perspective in AI research. The motivation for this view stems from various noted harmful and unintended consequences that emerged early on when organizations began implementing and using AI (Neubert & Montañez, 2020). According to Mikalef, while AI has resulted in unanticipated effects in several prominent and widely publicized cases, these reports only represent the tip of the iceberg as most problems do not receive any public mention. IS research has predominantly focused on the positive impacts associated with information technology (IT) deployments. However, past attempts to examine unexpected consequences that arise from using IT and their causes and effects have yielded some exciting results and helped advance theorizing (Tarafdar et al., 2013).

Mikalef noted that, despite an existing body of research that has adopted a dark side perspective to examine AI—particularly concerning ethics and biases ingrained in AI algorithms—we lack systematic studies that adopt this perspective as the starting point for their investigation. Focusing primarily on AI use's positive outcomes means that practitioners and researchers incompletely understand reality. More specifically, over the years, other related domains, such as strategic management, organizational research, and marketing research, have developed a more robust and theory-driven approach to examine

dark-side phenomena in their related empirical studies (Bollaert & Petit, 2010; Linstead et al., 2014). According to Mikalef, such an approach helps researchers identify the suboptimal and subliminal consequences and effects that technology can have. By adopting a critical dark-side lens, researchers can problematize phenomena or examine aspects they might otherwise overlook (Alvesson & Sandberg, 2011). Doing so can also help them explore the difficult ethical, political, and ideological issues surrounding AI implementation and use.

Moving forward in this area of inquiry, Mikalef argued that researchers need to build on the cumulative knowledge and theories used in other domains to identify approaches to examine dark-side effects, open up new ways to theorize such phenomena, uncover complementary perspectives, and provide a valuable toolbox of new methods. For instance, Alvesson and Ashcraft (2009) suggest how researchers can employ quantitative approaches in critical methodologies. On the other hand, Linstead et al. (2014) offer a multi-level framework for studying dark-side aspects from different levels of analysis and propose corresponding cross-disciplinary theories to explore them. In closing, Mikalef concluded his presentation by suggesting that, in studying AI deployment in organizations, researchers should consider different dark-side perspectives for the domain to advance and offer more relevant practical guidance.

3.1.3 AI as Digital Agency (Pär Ågerfalk)

Table 3. Pär Ågerfalk Summary

Pär Ågerfalk drew on datafication, machine learning, and digital infrastructuring when calling for IS scholars to reconceptualize agency that sees information systems as computational digital agents. He offered four critical suggestions for IS research: 1) distinguish between different AI types, 2) do not simply replace IT with AI in existing models, 3) do not confuse AI with big data or analytics, and 4) avoid routine-like referring to “recent developments in AI”.

Pär Ågerfalk drew on a recent *European Journal of Information Systems* editorial (Ågerfalk, 2020) to characterize AI as a constituent of contemporary digital practices. In keeping with the previous two presenters, he emphasized that machines remain far from convincingly performing the “cognitive functions typically associated with humans” (Rai et al., 2019). However, applications based on machine learning and pattern matching have become increasingly relevant due to the recent increase in processing power and storage capacity.

According to Ågerfalk, three current developments characterize information systems in contemporary digital practices: 1) datafication, 2) machine learning, and 3) digital infrastructuring. By drawing on large data sets, information systems can “learn” and reconfigure their behavior when acting as digital social agents in complex digital infrastructures that comprise software, people, data, and sensors. Viewing information systems and implementation of digital technology as the design of digital practices opens up several questions that need attention:

- 1) Who is responsible for the automated actions that information systems perform?
- 2) What social relationships do these systems’ activities establish and why?
- 3) What role do human actors have in formulating the regulative actions that govern machine learning actions?
- 4) How do automated actions governed by machine learning enact and shape institutions and institutional logics?

Ågerfalk argued that we might find the answers to such questions might be found if we reconceptualize agency in a way that moves away from the received notion of agency as necessarily human. By seeing our information systems as computational digital agents, not as passive media, we can begin to see how digital technologies become boundary agents in the ontological reversal that Baskerville et al. (2020) and others have called for (Kallinikos 2010; Aakhus et al., 2014; Beynon-Davies, 2018; Eriksson et al., 2018). Furthermore, by viewing these systems as responsible agents, we can begin to comprehend accountable systems beyond notions of explainable AI (Ågerfalk, 2004; Diakopoulos, 2015; Abdul et al., 2018).

Ågerfalk concluded his talk with four suggestions for IS research to achieve the necessary theoretical sensitivity: 1) distinguish between different AI types, 2) do not merely replace IT with AI in existing behavioral research models, 3) do not confuse AI with big data or analytics, and 4) avoid alluding to “recent developments in AI” unless you can show it.

3.1.4 Automation or Augmentation, not AI (Kevin Crowston)

Table 4. Kevin Crowston Summary

Kevin Crowston offered three critical points: 1) the term “AI” does not help frame a research study as dozens of technologies fall under this umbrella term and applications can comprise a mix of them, 2) we should draw on a cumulative tradition of research on automation on rather than treating AI technologies as sui generis, and (3) the exciting space for IS research is not automation for replacing workers but rather the technology for supporting them

In his presentation, Kevin Crowston made three points. The first concerned how IS researchers use the term AI. Crowston argued that the term does not help frame a research study as dozens of technologies fall under this umbrella term—from machine learning to computer vision to natural language processing—and applications can comprise a mix of them (i.e., the term lacks clarity). Further, the prospect of strong or general AI (i.e., a system that can fully replicate a human’s mental abilities) constitutes a distraction for IS researchers at least in the near term. IS research primarily considers cases where a system automates or supports particular tasks, which means that a system can still be impactful even if it has only the narrow intelligence to handle those tasks. He further noted that researchers must stay up to date rather than repeat stylized facts about systems since the technology continues to evolve quickly. For instance, researchers commonly describe machine-learning systems as black boxes even though that only applies to some technologies, and, even for those ones, we have seen recent advances in explainable AI that address this problem.

Second, the term AI focuses too much attention on technology. IS researchers should be the first to recognize that systems include both technology and people. Therefore, we should orient around the particular functionality of interest; that is, systems that display increased autonomy and how that relates to humans. From this perspective, systems’ increased ability to automate a broader range of tasks represents an interesting topic, but IS has studied automation since its inception. What we know about automation (e.g., the paradox of deskilling and upskilling) should also apply in the future. In other words, we should draw on the cumulative tradition of research on automation rather than treating AI technologies as sui generis.

Finally, the exciting space for IS research is not automation for replacing workers but rather the technology supporting them. For instance, researchers have said that, rather than a computer replacing a radiologist, a radiologist using a computer might instead (Davenport & Dreyer, 2018). Similarly, neither a person nor a computer won the the 2005 Freestyle chess championship; rather, a so-called centaur won it: two strong (but not world-class) chess players who worked out how best to use chess programs to analyze plays (ChessBase, 2005). They did not simply play chess or just use a computer: they did something novel that blended the two. How do we conceptualize that interaction?

3.1.5 AI and Time, Trust and Theory (Sirikka Jarvenpaa)

Table 5. Sirikka Jarvenpaa Summary

Sirikka Jarvenpaa talked about how time, trust, and theory provide opportunities for future research. First, shifting one’s temporal agency to advanced technologies may render social and inner time absent. Second, researchers may consider trust and distrust asymmetries to gain benefits and reduce drawbacks from algorithmic computing. Third, we need contextually rich theories that examine work and practices’ broader institutional and regulatory structures.

Sirikka Jarvenpaa talked about time, trust, and theory and how they provide opportunities for future research. In her talk, she considered AI as, very generally, computer algorithms that generate models to organize, categorize, and predict using human developed instructions (Jordan & Mitchell, 2015). AI’s power often relates to the ability to learn from data to derive predictions and facilitate human decision-making.

First, she discussed time. In a recent paper on end time in organizations, Jarvenpaa and Välikangas (2020) argue that researchers have largely not scrutinized how advanced technologies such as algorithmic computing can harm both inner and social time in organizations and collectives. One needs to consider both these times to solve complex organizational and societal problems collaboratively. Inner time refers to a temporal capacity to reflect on actions, meaning, and consequences over time. Social time refers to the time spent with others; for example, when giving and taking multivocal ideas and perspectives. With the pervasive use of advanced technologies, we can standardized the past in digital archives and algorithmically compute the future. The present is also a prediction rather than a representation. The authors argue for the urgency to make “the temporal assumptions of technology

visible so that technology's potential effects on social time and inner time can be better assessed and managed by users of technology seeking to collaborate" (Jarvenpaa & Välikangas, 2020, p. 579). Individuals and groups need to keep a watchful eye when shifting their temporal agency to advanced technologies, which creates the possibility to render social and inner time absent.

Second, she discussed trust. One cannot easily find a paper on algorithmic computing and social and human issues that does not mention trust or confident, positive expectations in one form or another. The IS (IS) literature has a long tradition of considering trust in technology (Lankton et al., 2015). Distrust, or negative expectations, has received less attention in IS research. Researchers accept distrust as a related but distinct construct from trust (Komiak & Benbasat, 2008). When discussing distrust in algorithmic computing, researchers have mainly focused on value incongruence and its detrimental effects, such as producing procedural and distributive injustices and a lack of transparency. They have claimed the algorithms to lack substantive rationality, such as values and morality (Lindebaum et al., 2019). With other technologies such as social media and e-marketplaces, IS researchers have underscored the significance of the simultaneous presence of trust and distrust for greater vigilance (Jarvenpaa & Majchzak, 2010). Jarvenpaa encouraged researchers to consider trust and distrust asymmetries to gain benefits and reduce drawbacks from algorithmic computing.

Third, she discussed theories. Unfortunately, many empirical studies that have leveraged algorithmic computing have made weak theoretical contributions to information systems. As a result, they have failed to inform both the technical and social components simultaneously or what Sarker et al. (2019) call the "sociotechnical axis of cohesion". We can observe that algorithmic computing overlays a third element on this framework: the context's or environment's pervasiveness. The context is not merely a static trigger or impetus for social and technological components; the context itself, including institutional rules and norms, is also fundamentally changed. Hence, there is a need for contextually rich theories that do not just focus only on immediate work and social practices but also examine the broader institutional and regulatory structures.

3.2 Themes Emerging from the Workshop

In this section, we briefly analyze the key concepts presented in and discussed across the seven round table discussion sections (see Figure 1).

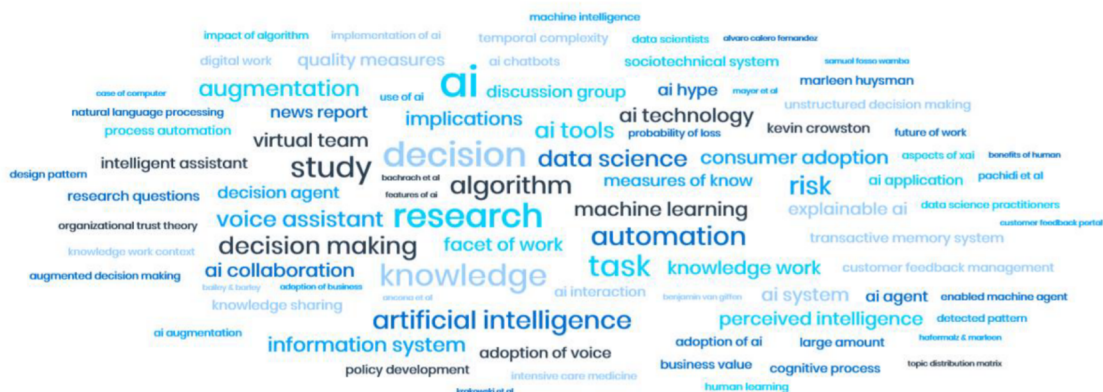


Figure 1. Key Concepts that Emerged from the Round Table Discussions

Taking our departure point as the five presentations and the round table discussions on AI research, the workshop recognized the two AI paradigms (strong and weak): 1) building a "fully intelligent" machine to understand human intelligence or achieve artificial consciousness and 2) augmenting human labor or replacing humans entirely with machines in the workplace with the latter being more immediately relevant to IS research. IS theorizing about machines that embody or exhibit intelligence explores practical questions and brings tenets that concern deeper sensemaking, acceptance, and trust. These issues that, in turn, quickly raise questions related to social aspects of AI pervasiveness in its context or ecosystem regarding agency and accountability, such as:

- 1) Should AI augment or replace humans in the work place?
- 2) How do institutional logics change with AI?
- 3) How important are transparency and interpretability issues (especially when they pertain to personal data)?
- 4) What implications does AI have on the human economy or privacy conditions such as welfare or surveillance?
- 5) Is implementing accountable AI sufficient to handle such issues or do some issue need more than just establishing accountability?

Table 6. Emerging Themes

Themes	Short description
Augmenting or complementing Work	Physical and virtual teams with human and intelligent agents/tools. Questions concern the labor division, whether a system can complement humans at a similar cognition level, and fit between a system and a human worker's cognitive style.
Trust, sensemaking, and AI acceptance	Learning from and by AI entails that one study what, how, and how much we can and should learn from AI agents. It includes studying models, algorithms, outcomes, and knowledge orientation. AI agency and digital practices draw on a broader group of actors.
Explainable AI	Transparency regarding about how AI evaluated inputs to reach a conclusion. Explainable AI constitutes a prerequisite for building trust and facilitating adoption.
Living with AI	Humans' (other than particular end users and developers who interact with AI and automation) lived experience, which builds on AI that collects data on human expressions such as the voice, face and by which human emotions are quantified, analyzed, and used in various settings.
AI agency	AI agency draws on phenomena such as datafication, machine learning, and digital infrastructuring. Featuring the ability to "learn" and reconfigure their behavior when acting as digital social agents in complex digital infrastructures that comprise software, people, data, and sensors. Computational systems processing large data sets entail viewing information systems and implementation of digital technology as the design of digital practices.
Accountable AI	Accountable AI deals with responsibility for actions that AI performs, the social relationships that these activities establish when regulating actions, forming and reconstituting institutions.
Temporal aspects of AI	Advanced technologies such as algorithmic computing can harm both inner and social time in organizations and collectives. Inner time refers to a temporal capacity to reflect on actions, meaning, and consequences over time. Social time refers to the time spent with others, such as by giving and taking multivocal ideas and perspectives.
Dark Side of AI	A dark side AI-perspective focuses on harmful and unintended consequences during AI implementation and use in organizations.

Workshop participants also advocated the need to include a dark-side perspective in AI research. A dark-side view would help to highlight complex ethical, political, and ideological issues surrounding the themes that emerged in the workshop. Similarly, time emerged as an emerging theme that can make organizations and society more or less adept at collaboratively solving problems.

4 Discussion: A Roadmap for AI Researchers

In this section, we capture the contemporary AI issues in IS research and provide recommendations to advance this research domain in the IS discipline.

4.1 The Potential for IS to be a Leading Discipline in the AI Movement

As IS researchers, we need to understand technologies as we study them. In particular, as technology evolves quickly, we need to not repeat stylized facts about systems that may have become obsolete. As an example, Crowston in his presentation and participants in subsequent discussions noted that technologies do not uniformly lack transparency as researchers sometimes claim they do, and advances in explainable AI have yielded more transparent (hybrid) approaches to machine learning. At the workshop, Natalia Levina said: "We find it hard to find the right conceptual language to describe these practices without bringing in a lot of 'philosophical baggage' from prior writings in IS".

Other examples pertained to the particular ethical concerns that come with unsupervised learning approaches built on only a few examples. For instance, Generative Pre-trained Transformer 3 (GPT-3), an autoregressive language model that uses deep learning to produce human-like text, can apply the huge data its creators used to train it to generate text from a few examples. This training activates new ethical challenges concerning biases learned from the training data (Abid et al., 2021). The variations in meaning, ethical implications, and impact of a particular technology or configuration seem all the more evident for organizations and society (Noble, 2018).

For these reasons, the IS community seems exceptionally well equipped to inform the broader contemporary AI discussion given its tradition of contextualizing technologies (Baskerville et al., 2020; Ågerfalk, 2020). In particular, IS researchers can contribute to the AI domain by showing how we can understand AI agency compared to other IT artefacts at the individual, organizational, and societal levels. In this way, they can revisit, augment, and strengthen existing IS theories, methods, and contributions to benefit other disciplines looking for new and better bowls for old petunias. Take, for example, Raisch and Krakowski (2020) who, based on their findings about the interrelations between AI augmentation and AI automation, advised management scholars to fundamentally change AI research as dual AI applications depend on each other across time and space. In return, they may provide some new and exciting research strands that further IS research and the IS community.

4.2 Strengthening AI Research in IS

In pursuing contributions that the IS community can make to the AI discussion, we identified needs for clear research strategies and actionable advice. Rethinking existing theory, methods, and assumptions also requires venturing into adjacent disciplines to establish boundaries between novel and already addressed topics, to identify high-level issues and challenges, and to build on the disciplines' cumulative tradition (Bailey & Barley, 2020). From the discussion, it seems clear that IS researchers need to distinguish the new and unique particulars of AI technology in the approaches they take, to analytically distinguishing AI from analytics, automation, or other technologies even if we often use these technologies and techniques together in practical contexts. Still, does delineating AI from related technologies and their implementation have practical relevance? And, for any definition brought forth to be useful, what should be the guiding foci, level of analysis, outcome, context, and value? What levels or categories should data possess to reach or saturate rigor and relevance criteria? How much of the algorithm or models should AI research in IS present? How can IS research assess AI outcomes? What theory and methods are sound and relevant outside of a particular use context? What should AI studies have to contribute to surpass the current AI hype?

Table 7. Actionable Advice for Strengthening AI Research and Contributions

Criteria	Action
Clarity	<ul style="list-style-type: none"> • Consider context-dependent AI definitions—what does the term AI mean in specific particular contexts, if anything? • Avoid alluding to “recent developments in AI” unless you can show it. Most current AI applications rely on age-old theory made practical by recent computing power and storage capacity increases. • Distinguish between different AI types and situate them in the particular context. • Do not confuse AI with big data or analytics. • Refrain from stylized descriptions.
Theoretical glue	<ul style="list-style-type: none"> • Develop agency theories that connect machines' computational logic, which includes explainable AI, with the institutional logics of organizations. • Regard machine learning as a special case of organizational learning.
Cumulative tradition	<ul style="list-style-type: none"> • How does AI differ in implementation, adoption, and diffusion to other technologies—explore new boundary conditions of existing methods and frameworks. • ISD traditionally emphasized user-developer communication and learning—acknowledge that the system has become a third actor in sensemaking and learning. • Be careful not to adapt an existing research model by simply replacing IT with AI.

Table 7. Actionable Advice for Strengthening AI Research and Contributions

Parsimony	<ul style="list-style-type: none"> • When questioning “artificial” and “intelligence”, break them down into their core elements: what are they, what do they lack, what minimal set defines them? Consider whether they matter and are appropriate? • What is and, more importantly, what is not AI? • How much of an algorithm and its use should IS researchers scrutinize? When does the task no longer become relevant to the IS discipline?
Applicability	<ul style="list-style-type: none"> • Engage with AI practices. Practitioners remain ahead in applying AI yet still struggle in doing so effectively. • Consider the nature of the task being automated, including the workings of hybrid approaches.

The contemporary debate suggests that AI constitutes an umbrella term for various technologies and approaches. As actors in organizations and industry often deploy AI applications in narrow domains and with goals for bespoke situations, they often come with separate logics, limitations, application areas, and implications. Therefore, we need to pay closer attention to whether contributions to and from the IS discipline offer clear definitions, contextual details, and implications. While IS researchers do not necessarily create new algorithms, they use them to solve business or social problems. Hence, we need to understand these technologies and methods before applying them in use and studying their impacts. In theorizing these phenomena with studies that present strong underlying logic and rationale (i.e., theoretical glue), we need empirical studies on AI to simultaneously inform both the technical component and the social component (Sarker et al., 2019). We need to engage with AI practices while recognizing that user behavior also impacts system performance (e.g., overreliance, misuse, or non-use). See, for example, Teodorescu et al. (2021), who argue that fundamental differences between data-driven approaches and traditional code-based information systems give reason to revisit some basic assumptions on which classic IS theories and concepts of human-machine augmentation rest. Still, to some degree, we would advise researchers to take a parsimonious approach to pronouncing their empirical studies' implications. We need to consider whether our studies really concern AI and not, for instance, automation via other means. Classifying various technologies according to technology or application area is not enough. When drawing on AI as digital agency or even computational digital agents, researchers need to address social relationships and the re-shaping of institutions and institutional logics through automated actions. We do, however, caution against pre-mature parsimony. Constructs may represent necessary components but in few niche instances. Dismissing them simply due to low frequency may leave unexplained an important if relatively small domain subset. Such exclusion poses a particular concern in a relatively new, rapidly evolving, multidisciplinary domain such as AI, which features much confusion and concepts that vary in their clarity and maturity. One cannot easily make informed decisions about parsimony in such a confused and evolving context.

Specificity, context, and implications possibly represent key factors that IS researchers need to consider when focusing on the human-autonomy interaction and how users perceive AI regardless of actual capacity or limitations. To do so, they could consider frameworks and theories about ambivalence and trust and distrust asymmetries in gaining benefits and reducing drawbacks from algorithmic computing. Hence, we need contextually rich theories that also examine the broader institutional and regulatory structures that leverage technology. Researchers may need to adopt more expansive approaches and consider auto-ethnography, quantitative methods in critical methodologies (Alvesson & Ashcraft, 2009), various analysis levels, and cross-disciplinary methods (Linstead et al., 2014).

Adopting a dark-side perspective in AI research may aid in building clarity and theorizing as it removes bias in expectations and power displays (Tarafdar et al., 2013). We may need to build on the cumulative knowledge and theories that other domains use. By doing so, we could identify approaches to examine dark-side effects, open up new ways to theorize such phenomena, uncover complementary perspectives, and provide a useful toolbox of new methods. Of course, as a valid counter-argument such an approach, older theories may not be appropriate in the AI context; one may logically call into question earlier theory and the assumptions that underpin it. We agree with this view and do not suggest that AI researchers must always build on existing theories. Instead, when conducting research on AI, we need to identify areas in which we can new systematically apply earlier theories versus areas where such theories fall short.

4.3 Continuing the Discussion: Moving Forward

We need to propose studies in IS that can push AI's frontiers. Whether we have already passed, now near, or never will reach the initial goals for AI (i.e., building a fully intelligent machine) the issue about strong or general AI should not distract from relevant and emerging themes and topics. Due to the societal interest, large-scale computational power, and large amounts of available data, emerging themes for IS deal with accountable AI, AI agency, acceptance and trust of such agency, the division of labor in work and daily life, and temporal AI aspects. In addition, AI agency deals with meta-perspectives such as temporal aspects, performance, underlying assumptions of time, and dark-side perspectives of negative and unintended consequences. As AI has arguably become a consistent part of contemporary digital practices—a kind of computational digital agency—it has enabled an ontological reversal (Baskerville et al., 2020; Kallinikos, 2010; Aakhus et al., 2014; Beynon-Davies, 2018; Eriksson et al., 2018) that opens up emergent themes on accountable systems that stretch beyond explainable AI (Ågerfalk, 2020; Diakopoulos, 2015; Abdul et al., 2018). Questions include: where do chatbots fail, how we make these systems learn, how should humans work with AI systems to get the most benefits, and how do we conceptualize that interaction? What performance metrics do we need to evaluate specific AI methods and improve them? We can also help move the AI domain forward by proposing studies that clarify how we can get machines to learn concepts (rather than perceptual categories) and make analogies and abstractions much like humans.

Table 8. Research Questions related to Emerging Themes

Themes	Research questions
Augmenting or complementing work and the division of labor in hybrid settings	<ul style="list-style-type: none"> • How is labor to be divided in augmenting/complementing or hybrid setups? • How is AI to be activated? • How can AI be adapted to diverse cognitive styles of different human workers? • How are digital and human agents to co-exist (Jarrahi, 2018)? • How are tasks be handed to humans from AI and vice versa? • What tool designs allow researchers to systematically explore and understand AI in work life? • How do we conceptualize AI system and human interaction in situations where both can take initiative?
AI agency acceptance/XAI sensemaking and trust	<ul style="list-style-type: none"> • What, how, how much, and in what way can we understand, accept, and learn from AI agents that differ from settings with only human agents? • To what degree should AI/XAI balance transparency and performance? • As AI becomes a more widespread concern and continues to evolve rapidly, what should IS researchers consider when examining AI: models, algorithms, outcomes, knowledge orientations, XAI, and so on? • What can we learn from social and behavioral theories on acceptance, learning, and their underpinnings?
Explainable AI	<ul style="list-style-type: none"> • How does the extent to which one can explain AI outcomes affect trust in AI? • What aspects promote transparency and trust in AI outcomes? • How can one embed explainability throughout the entire process to develop AI and how can one communicate it to relevant stakeholders?
Living with AI: understanding actors, their needs, and their involvement	<ul style="list-style-type: none"> • What does it mean for individuals, organizations or society to co-exist with AI? • What role do ethics play for AI? • What unique limitations and possibilities does AI have for humans, organizations, and society?
AI agency	<ul style="list-style-type: none"> • Does AI differ from automation, other technologies, and their implementations (e.g., in the phenomena they cover, assumptions about them in the literature, their contextualization)? • In what way do they differ (e.g., in a conceptual, contextual, or value-based manner)? • What role, if any, does AI play in economic, organizational, or societal development?
Accountable AI	<ul style="list-style-type: none"> • Who is responsible for the automated actions that information systems perform? • What social relationships do the activities that these systems conduct establish and why? • What role do human actors have in formulating the regulative actions that govern machine learning actions? • How do automated actions governed by machine learning enact and shape institutions and institutional logics?

Table 8. Research Questions related to Emerging Themes

Temporal aspects of AI	<ul style="list-style-type: none"> • What effects does temporality have on AI's value? For example, should AI decisions consider its input or output rhythm? And how does AI account for time when drawing on variables and data that occur over various time horizons? • How can AI consider event time (e.g., variances in data around specific events) rather than just traditional "clock" calendar time? • Do AI technologies such as algorithmic computing destroy both inner and social time in organizations and collectives and, if so, how?
Dark side of AI	<ul style="list-style-type: none"> • How can AI governance practices mitigate AI use's negative consequences? • Where does accountability lie, and who should be responsible for AI use's unintended consequences? • What effect does AI have on individuals when used to replace or enhance their work (e.g., psychological and physiological impacts)? • How can AI lead to undesirable organizational outcomes, and in which ways do such outcomes manifest? • What effect does AI have on power structures in industries? Does it lead to more centralized control?

5 Conclusion

AI's founding fathers originally envisioned creating a fully intelligent¹ machine and believed natural language processing, machine learning, neural networks, reasoning, and creative as integral to doing so. Today, the AI domain grapples with challenging philosophical concerns about what intelligence means, how we can achieve it, and which really matters since it has several dimensions.

However, we do not need to solve these conundrums and mimic the human brain to automate narrow goals. The research agenda for intelligently automating human tasks has instead become infused with issues concerning the need for transparent and interpretable behavior. Application areas for AI are everywhere in contemporary society. How to live and work with machines become key for acceptance, sensemaking, and trust in AI. From this perspective, systems' increasing ability to automate a broader range of tasks represents an interesting topic. Still, importantly, as a cumulative tradition of research on automation exists, we should draw on it rather than treating AI as *sui generis*. Fruitful approaches relate the technical component and the social component (Sarker et al., 2019). Yet, issues abound concerning the lack of representativeness, low quality or bias in input data, and how technology changes its domain of application, including institutions and their norms and rules.

A dark-side perspective in AI research adds additional research questions to the themes we present in this paper. Due to AI's reliance on vast training datasets, computing power, and real-time execution, issues and questions about accountability, ethics, and responsibility with respect to AI have also emerged. Harmful and unintended consequences during AI implementation and use, their causes, and their effects can help one more comprehensively understand the reality of systems in use. In this vein, Teodorescu et al. (2021) urge researchers to rethink organizational learning in the presence of ML where organizational processes become susceptible to systematic unfairness. Similarly, Gregory et al. (2020) argue for a distinction between positive and negative data network effects concerning perceived platform value for users. They caution against embedding AI in user networks that reside on multi-sided platforms due to intended and unintended consequences that stem from factors that we do not yet understand. Sturm et al. (2021) conclude that, even though human-made ML revisions can be beneficial, one cannot rely on the effect, which can even turn harmful in certain situations.

Studying the unintended consequences of AI implementation and use may also facilitate the sometimes heated discussion about AI as new or old technology and as a disruptive or a paradigmatic shift in society. New issues or resolutions of ethics, politics, ideology, and philosophy are brought to the table, helping to provide understanding or lay these to rest. Fügner et al. (2021) demonstrate a negative impact on the "wisdom of crowds" for human-AI decision environments as human individuality was lost. With studies on AI deployment in organizations that venture into a dark-side perspective, IS researchers will be able to advance and offer more relevant practical guidance about augmenting or complementing labor, living with AI, and perhaps, one day, even issues concerning artificial consciousness. As with other technologies,

¹ Our AI-powered grammar-checking software keeps suggesting "a brilliant"—maybe that is true.

such as social media and e-marketplaces, IS researchers have underscored the value of the simultaneous presence of trust and distrust for AI (Jarvenpaa & Majchrzak, 2010; Komiak & Benbasat, 2008; McKnight & Chervany, 2001; Moody et al., 2017), which has also emerged as a theme to explore or build on for AI. Highlighting the issues that abound in developing, training, and evaluating ML models, Lebowitz et al. (2021) point to a disconnect between AI's know-what and experts' know-how for knowledge work in uncertain domains. Moreover, Kellogg et al. (2019) show how the contemporary manner in which organizations technically and bureaucratically control employees distinctly differs from before and how new emerging practices have contributed to a collective resistance to such algorithmic rule.

Similarly, from the dark-side perspective, Jarvenpaa and Välikangas (2020) argue for a focus on making "the temporal assumptions of technology visible so that technology's potential effects on social time and inner time can be better assessed and managed by users of technology seeking to collaborate" (Jarvenpaa & Välikangas, 2020, p. 579); that is, how AI may harm both inner and social time in organizations and collectives due to myopia or shortsightedness. In the context of pervasive use of advanced technologies, the past can be standardized in a digital archive and the future can be algorithmically computed. Research should investigate emerging themes and issues in concert as to AI's temporal aspects and its impacts on decision making and collaboration.

Acknowledgments

This work was supported, in part, by Science Foundation Ireland grant 13/RC/2094 and co-funded under the European Regional Development Fund to Lero - the Science Foundation Ireland Research Centre for Software (www.lero.ie), and a grant from the US National Science Foundation (Grant 17-45463), and in part, by the Wallenberg Foundations, WASP-HS research in humanities and social science in AI and autonomous systems—grant BioMe: Existential challenges and ethical imperatives of biometric AI in everyday lifeworlds, PI A. Lagerkvist.

References

- Aakhus, M., Ågerfalk, P. J., Lyytinen, K., & Te'eni, D. (2014). Symbolic action research in information systems: introduction to the special issue. *MIS Quarterly*, 38(4), 1187-1200.
- Abdul, A., Vermeulen, J., Wang, D., Lim, B. Y., & Kankanhalli, M. (2018). Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the Conference on Human Factors in Computing Systems*.
- Abid, A., Farooqi, M., & Zou, J. (2021). Large language models associate Muslims with violence. *Nature Machine Intelligence*, 3(6), 461-463.
- Ågerfalk, P. J. (2004). Investigating actability dimensions: A language/action perspective on criteria for information systems evaluation. *Interacting with Computers*, 16(5), 957-988.
- Ågerfalk, P. J. (2020). Artificial intelligence as digital agency. *European Journal of Information Systems*, 29(1), 1-8.
- Aleksander, I. (2017). Partners of humans: A realistic assessment of the role of robots in the foreseeable future. *Journal of Information Technology*, 32(1), 1-9.
- Alvesson, M., & Ashcraft, K. L. (2009). Critical methodology in management and organization research. In A. Bryman & D. Buchanan (Eds.), *The Sage handbook of organizational state of the art and research challenges research methods* (pp. 61-77). Sage.
- Alvesson, M., & Sandberg, J. (2011). Generating research questions through problematization. *Academy of Management Review*, 36(2), 247-271.
- Baskerville, R. L., Myers, M. D., & Yoo, Y. (2020). Digital first: The ontological reversal and new challenges for information systems research. *MIS Quarterly*, 44(2), 509-523.
- Bailey, D. E., & Barley, S. R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization*, 30(2).
- Benbasat, I., & Zmud, R. W. (2003). The identity crisis within the IS discipline: Defining and communicating the discipline's core properties. *MIS Quarterly*, 27(2), 183-194.
- Benbya, H., Pachidi, S., & Jarvenpaa, S. L. (2021). Artificial intelligence in organizations: Implications for information systems research. *Journal of Association for Information Systems*, 22(2), 281-303.
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly*, 45(3), 1433-1450.
- Beynon-Davies, P. (2018). Declarations of significance: Exploring the pragmatic nature of information models. *Information Systems Journal*, 28(4), 612-633.
- Bollaert, H., & Petit, V. (2010). Beyond the dark side of executive psychology: Current research and new directions. *European Management Journal*, 28(5), 362-376.
- ChessBase. (2005). *PAL / CSS report from the dark horse's mouth*. Available from <https://en.chessbase.com/post/pal-c-report-from-the-dark-horse-s-mouth>
- Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.
- Davenport, T. H., & Dreyer, K. J. (2018). AI will change radiology, but it won't replace radiologists. *Harvard Business Review*. Retrieved from <https://hbr.org/2018/03/ai-will-change-radiology-but-it-wont-replace-radiologists>
- Diakopoulos, N. (2015). Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 3(3), 398-415.
- Dubin, R. (1976). Theory building in applied areas. In M. D. Dunnette (Ed.), *Handbook of industrial and organizational psychology* (pp. 17-39). Rand McNally.
- Dubin, R. (1978). *Theory building*. Free Press.
- Eriksson, O., Johannesson, P., & Bergholtz, M. (2018). Institutional ontology for conceptual modeling. *Journal of Information Technology*, 33(2), 105-123.

- Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will humans-in-the-loop become Borgs? Merits and pitfalls of working with AI. *MIS Quarterly*, 45(3), 1527-1556.
- Gomes, C., Dietterich, T., Barrett, C., Conrad, J., Dilkina, B., Ermon S., Fang, F., Farnsworth, A., Fern, A., Fern, X., Fink, D., Fisher, D., Flecker, A., Freund, D., Fuller, A., Gregoire, J., Hopcroft, J., Kelling, S., Kolter, Z., Powell, W., Sintov, N., Selker, J., Selman, B., Sheldon, D., Shmoys, D., Tambe, M., Wong, W.-K., Wood, C., Wu, X., Xue, Y., Yadav, A., Yakubu, A.-A., & Zeeman, M. L. (2019). Computational sustainability: Computing for a better world and a sustainable Future. *Communications of the ACM*, 62(9), 56-65.
- Gregory, R. W., Henfridsson, O., Kaganer, E., & Kyriakou, H. (2020). The role of artificial intelligence and data network effects for creating user value. *Academy of Management Review*, 46(3).
- Iansiti, M., & Lakhani, K. R. (2020). Competing in the age of AI. *Harvard Business Review*. Retrieved from <https://hbr.org/2020/01/competing-in-the-age-of-ai>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577-586.
- Jarvenpaa, S. L., & Majchrzak, A. (2010). Vigilant interaction in knowledge collaboration: Challenges of online user participation under ambivalence. *Information Systems Research*, 21(4), 773-784.
- Jarvenpaa, S. L., & Valikangas, L. (2020) Advanced technology and endtime in organizations: A doomsday for collaborative creativity? *Academy of Management Perspectives*, 34(4), 566-584.
- Jordan, M. I., & Mitchell, T. M. (2015) Machine learning: trends, perspectives, and prospects. *Science*, 349(6245), 255-260.
- Kallinikos, J. (2010). *Governing through technology: Information artefacts and social practice*. Palgrave Macmillan
- Keen, P. G. W., (1980). MIS research: Reference disciplines and a cumulative tradition. In *Proceedings of the International Conference on Information Systems*.
- Keen, P. G. W. (1991). *Shaping the future: Business design through information technology*. Harvard University Press.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2019). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410.
- Komiak, S. Y. X., & Benbasat, I. (2008). A two-process view of trust and distrust building in recommendation agents: A process-tracing study. *Journal of the Association for Information Systems*, 9(12), 727-747.
- Lankton, N. K., McKnight, D. H., & Tripp, J. (2015). Technology, humanness, and trust: Rethinking trust in technology. *Journal of the Association for Information Systems*, 16(10), 880-918.
- Larson, A. (2010). Artificial intelligence: Robots, avatars and the demise of the human mediator. *Ohio Journal of Dispute Resolution*, 25(1), 105-164.
- Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021). Is AI ground truth really “true”? The dangers of training and evaluation AI tools based on experts’ know-what. *MIS Quarterly*, 45(3), 1501-1525.
- Lindebaum, D., Vesa, M., & den Hond, F. (2019). Insights from The Machine Stops to better understand rational assumptions in algorithmic decision-making and its implications for organizations. *Academy of Management Review*, 45(1), 247-263.
- Linstead, S., Maréchal, G., & Griffin, R. W. (2014). Theorizing and researching the dark side of organization. *Organization Studies*, 35(2), 165-188.
- Metcalfe, M. (2004). Theory: Seeking a plain English explanation. *Journal of Information Technology Theory and Application*, 6(2), 13-21.
- McKnight, D. H., & Chervany N. L. (2001). Trust and distrust definitions: One bite at a time. In R. Falcone M. Singh, & Y. H. Tan (Eds.), *Trust in cyber-societies* (LNCS vol. 2246). Springer.

- Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, 58(3).
- Moody, G., Lowry, P., & Galletta, D. (2017). It's complicated: Explaining the relationship between trust, distrust, and ambivalence in online transaction relationships using polynomial regression analysis and response surface analysis. *European Journal of Information Systems*, 26(4), 379-413.
- Neubert, M. J., & Montañez, G. D. (2020). Virtue as a framework for the design and use of artificial intelligence. *Business Horizons*, 63(2), 195-204.
- Noble, S. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- Rai, A., Constantinides, P., & Sarker, S. (2019). Editor's comments: Next-generation digital platforms: Toward human-AI hybrids. *MIS Quarterly* 43(1), iii-ix.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J. F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A., Roberts, M. E., Shariff, A., Tenenbaum, J. B., & Wellman, M. (2019). Machine behavior. *Nature*, 568(7753), 477-486.
- Raisch, S., & Krakowski, S. (2020). Artificial intelligence and management: The automation augmentation paradox. *Academy of Management Review*, 46(1).
- Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (2019). The sociotechnical axis of cohesion for the IS discipline: Its historical legacy and its continued relevance. *MIS Quarterly*, 43(3), 695-719.
- Schwartz, J. Hagel, J. Wool, M., & Monahan, K. (2019). Reframing the future of work. *MIT Sloan Management Review*. Retrieved from <https://sloanreview.mit.edu/article/reframing-the-future-of-work/>
- Suddaby, R. (2010). Editor's comments: Construct clarity in theories of management and organization. *Academy of Management Review*, 35(3), 346-357.
- Sturm, T., Gerlach, J., Pumplun, L., Mesbah, N., Peters, F., Tauchert, C., Nan, N., & Buxmann, P. (2021). Coordinating Human and machine learning for effective organizational learning. *MIS Quarterly*, 45(3), 1581-1602.
- Tarafdar, M., Gupta, A., & Turel, O. (2013). The dark side of information technology use. *Information Systems Journal*, 23(3), 269-275.
- Teodorescu, M. H. M., Morse, L., Awwad, Y., & Kane, G. C. (2021). Failures of fairness in automation require a deeper understanding of human-ML augmentation. *MIS Quarterly*, 45(3), 1483-1500.
- Weick, K. E. (1989). Theory construction as disciplined imagination. *Academy of Management Review*, 14(4), 516-531.
- Weick, K. E. (1995). *Sensemaking in organizations*. Sage.
- Whetten D. (1989). What constitutes a theoretical contribution. *Academy of Management Review*, 14(4), 490-495.

Appendix A: Picture from the Workshop



Figure 1. Group Photo of Some of the Workshop Participants in Zoom

About the Authors

Pär J. Ågerfalk is a Professor in the Department of Informatics and Media at Uppsala University, where he holds the Chair in Information Systems. He received his PhD in Information Systems Development from Linköping University and has held full-time positions at Örebro University, the University of Limerick, Lero—The Irish Software Engineering Research Center, and Jönköping International Business School. He is a former Editor-in-Chief of the *European Journal of Information Systems*, a member of the *Journal of Information Technology* editorial board, and a member of the Distinguished Editorial Advisory Board of the *International Journal of Information Management*. His current research focuses on the orchestration of digital practices, including open source software, digital innovation, information systems development, and artificial intelligence. His work has appeared in journals such as *MIS Quarterly*, *Information Systems Research*, *Journal of Information Technology*, *Journal of the Association for Information Systems*, and *Communications of the ACM*. He is a Distinguished Member Cum Laude and Fellow Award recipient of the Association for Information Systems.

Kieran Conboy is a Professor in Business Information Systems in the School of Business & Economics at the National University of Ireland, Galway, and a co-Principal Investigator in Lero – the Irish Software research center. He previously worked for Accenture Consulting and the University of New South Wales in Australia. He is also on the board of the Irish Research Council and has published over 200 papers in leading international journals and conferences, including *Information Systems Research*, the *European Journal of Information Systems*, *Information Systems Journal*, the *Journal of the AIS*, and *IEEE Software*. He is a Co-Editor-in-Chief of the *European Journal of Information Systems*.

Kevin Crowston is a Distinguished Professor of Information Science in the School of Information Studies at Syracuse University. He received his PhD (1991) in Information Technologies from the Sloan School of Management, Massachusetts Institute of Technology (MIT). His research examines new ways of organizing made possible by the extensive use of information and communications technology. Specific research topics include the development practices of Free/Libre Open Source Software teams and work practices and technology support for citizen science research projects, both with NSF support. His most recent project is a research coordination network on Work in the Age of Intelligent Machines. He is currently Editor-in-Chief for the journals *ACM Transaction on Social Computing and Information, Technology & People*.

Sirkka L. Jarvenpaa is Professor of Information Systems and Bayless/Rauscher Pierce Refsnes Chair in Business Administration at the McCombs School of Business, The University of Texas at Austin. She is a recipient of the Association for Information Systems (AIS) Fellow and LEO Awards (LEO stands for Life Time Achievement of Exceptional Global Contributions in the field of information systems). Her research focuses on inter-organizational and inter-personal collaboration and innovation in fast paced and technologically advanced data and knowledge environments with regulatory and policy implications.

Jenny Eriksson Lundström is a Senior Lecturer and former Head of Department at Uppsala University, Department of Informatics and Media, and Chair of the Swedish Foundation of Legal Information. After her PhD on AI and formal modelling in the legal domain, she has broadened her research scope and agenda. She is currently co-PI for the Wallenberg AI, Autonomous Systems and Software Program in Humanities and Society funded project BioMe: Existential Challenges and Ethical Imperatives of Biometric AI in Everyday Lifeworlds. She is a Guest Co-Editor of the *European Journal of Information Systems* special issue on the Dark Side of Analytics and AI. She is a member of the Academy of Management and the Association for Information Systems.

Patrick Mikalef is an Associate Professor of Data Science and Information Systems at the Department of Computer Science, Norwegian University of Science and Technology, and a research scientist at SINTEF Digital. He is a recipient of the Marie Skłodowska-Curie individual fellowship for his work on data-driven strategies. He received his MSc in Business Informatics from Utrecht University and his PhD in IT Strategy from the Ionian University. His research interests focus on the strategic use of information systems and IT business value in turbulent environments, and the potential of business analytics and artificial intelligence in driving digitalization in organizations. His work has appeared at international conferences and journals, including the *British Journal of Management*, *European Journal of Information Systems*, *Journal of Business Research*, *European Journal of Operational Research*, and *Information and Management*. He serves as a track chair on Business Analytics and Big Data for ECIS and AMCIS. He is a guest co-editor of the *European Journal of Information Systems* special issue on the Dark Side of Analytics and AI.

Sudha Ram is Anheuser-Busch Endowed Professor of MIS, Entrepreneurship & Innovation in the Eller College of Management at the University of Arizona. She has joint faculty appointments as Professor of Computer Science, member of BIO5 Institute, and Institute for Environment. She is the director of the INSITE: Center for Business Intelligence and Analytics at the University of Arizona. She received a Ph.D. from the University of Illinois at Urbana-Champaign in 1985. Her research is in the areas of Enterprise Data Management, Business Intelligence, Network Science, and Machine Learning. She is a Fellow of AIS and a recipient of more than US\$70 million in research grants. She is currently a co-editor for the *Journal of Business Analytics*.

Copyright © 2022 by the Association for Information Systems. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints are via e-mail from publications@aisnet.org.