Artificial Intelligence and Robots in Individuals' lives: How to Align Technological Possibilities and Ethical Issues

Ting-Peng (TP) Liang
trenz@uni-goettingen.de

Lionel Robert
lprobert@umich.edu

Suprateek Sarker
ss6wf@comm.virginia.edu

Christy Cheung
ccheung@hkbu.edu.hk

Christian Matt
christian.matt@iwi.unibe.ch

Manuel Trenz
trenz@uni-goettingen.de

Ofir Turel
oturel@fullerton.edu

8 . 6 . 2021

Artificial intelligence and robots in individuals' lives:

How to align technological possibilities and ethical issues

Abstract

Purpose:

This paper reports the panel discussion on the topic of artificial intelligence (AI) and robots in our lives. This discussion was held at the *Digitization of the Individual (DOTI)* workshop at the International Conference on Information Systems in 2019. Three scholars (in alphabetical order: Ting-Peng Liang, Lionel Robert, and Suprateek Sarker) who have done AI- and robot-related research (to varying degrees) were invited to participate in the panel discussion. The panel was moderated by Manuel Trenz.

Design/methodology/approach:

This paper introduces the topic, chronicles the responses of the three panelists to the questions the workshop chairs posed, and summarizes their responses, such that readers can have an overview of research on AI and robots in individuals' lives, and insights about future research directions.

Findings:

The panelists discussed four questions with regards to their research experiences on AI- and robot-related topics. They expressed their viewpoints on the underlying nature, potential, and effects of AI in work and personal life domains. They also commented on the ethical dilemmas for research and practice, and provided their outlook for future research in these emerging fields.

Originality/values:

This paper aggregates the panelists' viewpoints, as expressed at the DOTI workshop. Crucial ethical and theoretical issues related to AI and robots in both work and personal life domains are addressed. Promising research directions to these cutting-edge research fields are also proposed.

Keywords:

Artificial intelligence, Robots, Digitization of the Individual (DOTI), Ethical issues, Panel

Paper type:

Viewpoint

1. Introduction

Artificial intelligence (AI) is a term used for describing information systems that allow machines to present intelligence that resembles some aspects of human intelligence, such as making decisions based on ongoing learning from inputs (Nilsson, 2014). It has emerged as the software engine that drives the Fourth Industrial Revolution, a technological force that influences all disciplines, economies, and industries (World Economic Forum 2020a). Robots, in many cases build on such AI capabilities, to resemble and mimic human behaviour and decisions or to improve such decisions (Richardson, 2015). Both AI and robots, are important aspects of the digitization of individuals, as they can interface with many work and non-work life domains, and satisfy autonomy, relatedness and competence needs of humans (Matt et al. 2019; Turel et al. 2020). They are already prevalent in produces such as autonomous cars, recommendation systems, production lines, smart homes, and social media (Dirican, 2015).

The explosive development of AI technologies and robots have already given rise to a host of challenging moral and ethical issues, including unemployment, inequality, humanity, artificial stupidity, racism, security, evil genies, singularity, and robot rights (Sun and Medaglia 2019). With the advancement of AI and robots, tech giants such as Alphabet, Amazon, Facebook, IBM and Microsoft have urged more discussions and conversations addressing ethics and risk assessment for the emerging technologies (Seeber et al. 2020; World Economic Forum 2016). The use of AI-based agents may also hinder productivity by creating a negative perception of the user (Prakash and Das, 2020)

To explore and address these timely and significant issues (which can be viewed as "grand challenges"), we organized a panel discussion with three experts. Our objective was to discuss and evolve viewpoints on ethical issues of AI and robots in individuals' lives. This discussion was held as part of the *Digitization of the Individual (DOTI)* workshop, which was held in conjunction with the International Conference on Information Systems in 2019. The three panelists in alphabetical order included: Ting-Peng (TP) Liang, College of Management, National Sun Yat-Sen University; Lionel Robert, School of Information, University of Michigan; and Suprateek (Supra) Sarker, School of Commerce, University of Virginia's McIntire. The panel was moderated by Manual Trenz from the University of Göttingen.

In the workshop, the three panelists shared their research experiences related to AI and robots and explained their motivation to focus on such issues. They also debated the underlying nature, potential, and effects of AI across work and personal life domains.

Further, they discussed the ethical dilemmas with regards to AI research and practice. Finally, the panelists elaborated on potential areas for future research. The following sections present their responses to the four questions that were discussed during the workshop.

2. What aspect of AI do you study; and what motivated you to examine such issues?

The panelists' research experiences with AI and robots were diverse, and their motivations and research foci were also very different. For example, TP's AI-related research started from his thesis on developing knowledge-based systems. Lionel focused on developing safeguards to ensure so-called black-box AI-algorithms adhere to core values: fairness, trust and ethics. Supra began to explore the phenomenon of big data and data-driven AI with the advent of big data in recent years. He emphasized the use of the sociotechnical lens in understanding the AI phenomenon. Their detailed conversations are summarized here:

TP: Artificial intelligence (AI) is a discipline that intends to build powerful models for solving complex problems that are typically handled by human beings, such as stock investment and game playing. Instead of being a "garbage in garbage out" technology, AI allows a mechanical system to behave more flexible and thus more useful in a dynamic environment. My AI-related research started from my thesis when developing knowledge-based systems (or called "expert systems") gained much attention in mid-1980. My main research areas include expert system development, knowledge acquisition and inductive learning, artificial neural networks (ANN), and, more recently, data mining and text mining. It's interesting to see that ANN in 1990 has evolved into today's deep learning and to become much more powerful modelling techniques. My early work in Management Science (Liang, 1992) was one pioneer work on integrating statistical models (e.g., Bayesian theory) and tree induction. My recent work focuses on applying AI techniques to build better business models and providing interpretation of complex ANN models.

Lionel: I study fair AI systems for managing employees in organizations. AI systems have been found at times to be unfair to workers. Unfairness toward workers leads to decreased worker effort and increased worker turnover. It is becoming clear that the use of AI in organizations is having transformative effects, and this is particularly the case in relation to fairness, trust and ethics. Despite this, little attention has been paid to proposing a theoretical and systematic approach to designing fairer AI systems. This is particularly problematic when we consider that in a recent survey of 1,770 managers from 14 counties, 86% of managers stated that they planned to use AI systems for managing their workers and 78% of the managers said they trust decisions made by AI systems (Kolbjørnsrud, Amico, &

Thomas, 2016). Therefore, it is vital for us as IS scholars to understand this and to develop safeguards to ensure that these so-called black-box algorithms adhere to core values: fairness, trust and ethics.

To be fair, AI is rapidly changing every aspect of our society, not only organizations. Unfortunately, the use of AI has also amplified biases. As such, it is important for academia, business and government leaders to begin to work together. This is particularly important if we hope to both highlight challenges and identify potential solutions to address such challenges. For example, it is easy for IS scholars to focus on addressing the problems associated with managers who ignore the limitations of AI models, which results in biased decisions that are hard to detect (Robert, Pierce, Marquis, Kim, & Alahmad, 2020). But there are also deeper societal issues that cannot be ignored simply because they fall outside the scope of a single organization. The use of AI is leading to displaced workers and affordable-housing shortages. These types of challenges are leading to further inequalities in our society. These challenges are also global in nature (Robert, Gaurav & Lütge, 2020). This means the solutions have to be scalable across borders, which is much more problematic than we imagine. For bias, AI is an important topic across many domains. But when we have discussions on the topic they are very North American oriented. This means, for example, that we discuss AI biases in relation to their impacts on women, minorities and people of color. But these are protected groups in North America, so it is not clear that, let's say, China is concerned about AI bias toward women or other minority groups. Protected groups are not universal and often do not translate well across borders.

Supra: I am interested in studying AI application design, implementation, use, and management through a sociotechnical lens, as articulated by us in a recent paper (e.g., Sarker et al. 2019). Very simply stated, as a member of the IS discipline, I wish to understand how and why AI applications that are intricately tied to human activity succeed or fail. To be clear, I am not particularly interested in AI being used to solve problems that are difficult puzzles where human volition, emotions, reflexivity, creativity, improvisation, sensemaking, identity, power and politics, and so on, are not in play. The second aspect of my interest reflects the simultaneous and possibly recursive focus of the sociotechnical perspective on instrumental goals (such as monetary value, and making processes more efficient) and humanistic goals (such as human well-being that include ensuring privacy, dignity, and fairness).

As a graduate student, I was intrigued with many concepts and tools related to AI; however, AI seemed to almost disappear into the background for many years, and then it

reappeared in the mainstream Information Systems stage recently with the advent of Big Data. As I have sought to explore the phenomenon of Big Data (Abbasi et al. 2016), and data-driven AI, I have been intrigued with AI applications related to humans in business and society, and the potential of AI technologies/techniques hold to do good and likewise to do harm, I have been a bit disappointed on the fact that as members of the IS discipline, many of us have been overly fascinated with the tales about AI technology without being sufficiently critical. In this regard, I would like to characterize my research stance toward AI applications as cautious but optimistic.

3. How does AI affect individuals in their work and personal life domains? (Probe in question: Are there differences between the nature, potential, and effects of AI across these two domains?)

The panel discussion addressed the issues related to the impact of AI on work and personal life domains. TP and Supra discussed opportunities and threats of AI to both work and personal life domains. Lionel emphasized how AI blurred the boundaries between work and personal life domains. Their detailed conversations are summarized here:

TP: AI has impact on several aspects of our work and personal life. For instance, the fact that Alpha Go can defeat top human players totally changed our understanding of the Go game and attitude toward AI. Now, human experts often consult computer systems after they lose to another human player. AI becomes a super player and God of the game. It also changes game rules to include "no access to digital device during competition to prevent cheating." Certain products such as face recognition is everywhere from streets to office entrance. Auto-driving may also change the way vehicles are used.

The use of AI has pros and cons. One the one hand, AI-embedded software makes our life and work easier in many ways. AI may also provide insightful knowledge that may take centuries for humans to learn (e.g. AlphaGo). On the other hand, AI also brings in a number of challenges such as the ethical issues involved in autonomous driving - when an autonomous vehicle is malfunctioning and needs to detour, should it keep the original direction to hurt the person on the vehicle or switch to a safe direction but hit an innocent pedestrian. This kind of hard moral decisions may have existed for a long time, but AI makes it necessary to be explicit in order to code into the software. AI can also invade into many areas sensitive to our life, such as privacy and safety. AI algorithms along with big data collected from monitoring our daily activities are capable of revealing a lot of private

information that we would not want to be revealed otherwise. Hence, regulations regarding data collection and use are becoming critical to protect our traditional lifestyle.

Lionel: First we need to acknowledge that the lines between work and personal life are gone. This is especially true in the gig economy, where AI is being rapidly deployed to manage workers. Therefore, it is not clear that classifications of work and personal life domains are still useful going forward. That being said, we can probably divide work into paid and unpaid efforts. Paid efforts are any efforts that are financially compensated, whereas unpaid efforts are any efforts that are not financially compensated. So someone who employs the AI-enabled Uber application is using AI to engage in financially compensated efforts.

Typically, we can classify the impacts of AI as making activities more efficient, more effective, or enabling completely new activities altogether. In my classes, I refer to this as the 3E model of technology implications. If we apply the 3E model across both financially compensated and un-compensated efforts, the biggest differences are the incentives. AI used for financially compensated efforts will be adopted to the degree that it increases compensation or extrinsic value. AI used for financially uncompensated efforts will be adopted to the degree that it increases intrinsic value. In many cases the same AI application can be used for both.

This means that the nature, potential, and effects of AI across these two domains can at times vary drastically but may not. For example, an AI application designed to support paid efforts might focus on reducing the need for co-workers to interact directly. The AI application might be designed to literally replace a human co-worker. However, an AI application designed to support unpaid efforts might focus on establishing and encouraging direct interactions between individuals, for example by finding and linking gamers of like interest who might find commonality based on their shared interest. Likewise, an AI application designed to support paid efforts might link a worker who is in need of knowledge to another worker who has that knowledge. Although the incentives are different, the intended use of the AI will not always be different.

Supra: I believe AI has potential to transform work as well as personal life. It can make tasks easier and us smarter. It can make the unthinkable feasible because of the amount of data and computing power on our fingertips, and advances in techniques such as deep learning. However, at work, increasingly people are questioning what will be left for humans to do (in terms of employment to be able to make a living as well as pertaining to self-actualization) as AI is able to undertake most of our tasks? Most experts believe AI technologies are nowhere close to the stage when they can seamlessly take over all of the

intellectual work that most of us do – AI can replace us in patches. There is thus talk about task augmentation and task assemblage, not just task substitution (Rai et al. 2019). At home, AI can turn from a facilitator to a manipulator of smart devices and even our interpersonal relationships, and AI can also be harnessed to use our data strategically against our own interests. Even scarier is the vision of smart self-learning algorithms getting out of control. Needless to say, an AI system based on machine learning that becomes dysfunctional and out of control can be very dangerous for a company, an individual, or the society.

4. What are some of the ethical dilemmas you see with regard to AI in practice and AI research? (Probe in: Will humans eventually rise against AI machines? Why?)

In discussing the ethical dilemmas, Lionel and Supra raised their concerns on AI biases, while TP believed that AI can be a good complementary partner for humans in many areas. Their detailed conversations are summarized here:

TP: Another issue is that "will AI make human look stupid?" A good example is the AlphaGo case. Most players were confident that humans would beat AlphaGo before the game, but this confidence was totally destroyed after human players lost three consecutive games. Even though AlphaGo lost the fourth game due to a system bug, human players have lost their confidence for good and they tried hard to mimic or learn tricks from computer software. Many strategies human professional players have used for hundreds of years are dropped and now many AI strategies become popular. Hence, I do not see a situation where humans will compete with machines. Instead, I think AI can be a good complementary partner in many areas. Of course, we also see that, in some areas, AI-enabled robots are replacing human workers. For example, an intelligent teller program can help answering customer questions more efficiently, especially for handling repetitive and structured tasks. This is particularly significant in manufacturing and certain low-level service jobs. Hence, AI research can improve task performance and productivity in many areas, but we also need to design a social system that can accommodate this trend of declining job opportunities (unemployment) for structured work.

Lionel: The biggest dilemmas are related to AI biases, AI accountability and privacy. Although AI bias is receiving a lot of attention, it is still not always clear how we can identify bias. Many of the same biases that exist in the broader community are finding their way into AI systems that are managing employees. This is most prominently highlighted in the recent cases involving human resource (HR) practices where AI systems are screening applicants based on, in some cases, biased attributes potentially disadvantaging women and other

minorities. But if these AI systems are mimicking the biases of the managers employing them, it is not clear that those same managers will recognize such biases. If we can identify the biases, then what can we do to mitigate the negative consequences of AI unfairness or biases on employees? This is another problem that is garnering more attention than actual answers.

AI accountability is getting less attention but is vital to addressing any questions related to ethical dilemmas (Robert et al. 2020a; Robert et al. 2020b). AI accountability helps to answer the question: Who should be held accountable for the actions of an AI? This is much more problematic than it initially seems. For example, let's take the case of an AI HR application, which is often deployed by a third-party company on behalf of an organization located in the same or a different country. Who should be held accountable when the AI HR application engages in unethical behavior like hiring biases? We also have to recognize that there are different ethical frameworks among different nations. The problem of AI accountability, especially as a global problem, is slowly being acknowledged (World Economic Forum, 2020b).

Data are the lifeblood of AI, but it is not always clear whether the data being used were collected ethically. Privacy can be viewed as the ability to self-select oneself from being observed. The expectations of privacy are not absolute but rather relative across the globe. For example, U.S. citizens should have very little expectation of privacy with regard to data collected as the result of work on behalf of an employer. This is not the case in the European Union, where European citizens can have much higher expectations of privacy even with working on behalf of an employer. Citizens in other countries like China can have even less expectation of privacy than U.S. citizens. Is it unethical for a U.S. company to deploy an AI to collect data in China in a way that would be illegal or inappropriate in the U.S.? Is it unethical for a European company to employ an AI built on data collected in the U.S. in a way that would be illegal or inappropriate in the European Union?

Supra: Outlining the ethical dilemmas in a general way within a short answer is difficult. However, balancing the instrumental goals with humanistic goals is one way to think of the dilemma (Sarker et al. 2019). There are numerous well-known examples of AI and other intelligent technologies harming stakeholders who are already marginalized part of the society (Noble 2018). These include racist search engines, racist image recognition algorithms, sexist recruitment, opaque financial e-scores, and social credit systems (e.g., O'Neal 2016; Noble 2018).

AI researchers and practitioners must ask themselves, why must an AI technique be applied in a certain area? Is it necessary? Along with a justification based on expected benefits, they must ask: What are possible dangers or pitfalls to the society? Today, the dominant pattern is that of AI developers and company executives enjoy the upside while the downside, when they emerge, is to be borne by the society. Unless we develop and enforce standards of algorithmic accountability, formulate principles of data collection, storage, use, correction, and retirement that build on Mason's PAPA, and develop managerial and design code of ethical behaviors that are linked to professional credibility and advancement, we could see ourselves in the dystopian world that currently might seem alien or implausible.

5. What is next for AI research, or what can we do better as an area of research? Comment on emerging AI areas and research techniques.

TP suggested two methodology-based directions for future research: (1) To develop more powerful and domain-dependent AI techniques, and (2) to develop methods to make AI models understandable by human users. Lionel proposed research studies on systematic and theoretical approaches to audit and evaluate AI systems. Supra urged IS researchers to pursue approaches and perspectives to questions on AI that are somewhat unique to and distinctive to the IS discipline. Their detailed conversations are summarized here:

TP: In my opinion, this third wave of AI has its value in solving a few problems such as image recognition, but is still limited by not having enough understanding of human intelligence and generalizable learning algorithms. This restriction is primarily due to limited explainability of the ANN-based deep learning algorithms. In most cases, researchers may prove that a deep learning model is better in modelling a dataset to obtain higher fit or predictability, but the model is a black box that the human user is unable to know why the model performs better or when the model does not perform well. This kind of low human interpretable nature restricts the value of AI in serious projects where the decision maker has to take full responsibility. An example is that the developer told media that he did not know why AlphaGo made a silly mistake to lose the fourth game to the human player. The damage may not be so bad when AI plays a game, but would be disastrous if it happened in a nuclear power plant or an autopilot vehicle. Hence, two directions are important for the future: one is to develop more powerful and probably domain-dependent AI techniques and the other is to develop methods to make AI models understandable by human users. The second area of making AI understandable is called interpretable AI (XAI) is gaining much attention recently.

Only when AI algorithms are understandable can humans examine more ethical and social issues involved in the increased use of AI.

Lionel: Going forward, two areas will become increasingly important: AI autonomy and AI audits. AI autonomy is the degree to which the AI can operate independently in making decisions and taking actions. An important question to ask ourselves is: How much autonomy should the AI have? For example, an AI HR application's autonomy can range from low (e.g., making recommendations to interview), to moderate (e.g., making recommendations to hire) to high (e.g., actually hiring the employee). Questions like: When should an HR human manager be contacted? Should the AI just identify and recommend applications? Should the AI wait for permission from a human manager to approve any action? Another question not to be overlooked is whether or not an AI should be involved. This is much more of a philosophical question, but there may be areas that we simply do not want AI to be involved. A future design agenda could be undertaken to study when AI use is appropriate or inappropriate.

A small but ongoing area of research is focusing on auditing AI systems. An AI audit is an inspection of the AI's underlying logic and data sources to clearly understand and validate the AI (Robert et al., 2020a). AI audits are becoming increasingly popular because of new EU laws that are requiring some level of auditing to be performed before an AI is deployed. Therefore, we are now seeing more discussion around AI audits as an attempt to demonstrate that the AI complies with laws, regulations, or policies. As IS scholars we need to focus more of our efforts toward designing AI auditability. AI auditability is an affordance directed at designing AI that allows or supports the ability to be audited (Robert et al., 2020a). AI audits can be conducted manually through logical walk-throughs, automatically via code checks, or even by large-scale computer simulations. AI auditability is important because it may hold the key to addressing issues such as AI bias, AI accountability and issues related to privacy. Therefore, going forward IS scholars should seek to develop a design agenda to promote AI auditability.

Supra: One of the issues that has bothered me a great deal is how some AI enthusiasts, wealth-maximizing corporate managers, and self-engrossed researchers downplay potential concerns raised about Big Data and AI applications.

First, they use the rhetoric of "Unintended Consequences", implying the intention was all good but something unanticipated happened. For me, this would imply that there is a need to anticipate different types of consequences as a mandatory and critical part of a design and implementation project. Second, they use the rhetoric of "The Dark Side of Technology";

after all, they maintain, if there is the Bright side we enjoy, do we not need to accept the Dark side? Third, the offer the well-known utilitarian argument that the technology is beneficial overall, that is, it does "the greatest amount of good for the greatest number of people". However, precisely who gains the advantages and who pays the price (and to what degree) is largely left to imagination. Fourth, it is argued that knowledgeable subjects have authorized the use of data underlying the AI system, even though the presumed volition of subjects in signing away their rights has been a result of certain questionable practices. Fifth, there is a tendency of technical staff pleading immunity from the responsibility of negative consequences (say, unfair discrimination) resulting from their work. They maintain, "Our job is technical – we implement what is asked of us by our managers." I wonder if the staff members would be as comfortable if they discriminated an individual physically in the same manner as their brainchild algorithm-driven-by-data tends to do. And finally, there is the promise of "Almost there!": "With some more and better training data, all the issues will be fixed – all biases will go way."

With respect to future AI research in IS, I feel every topic in AI is a potential area for research for some in our community. For example, some may choose to deal with philosophical questions such as can machines think? Others may try to come up with more efficient algorithms for learning. Yet others may study the experience of humans with some team-members being robots (You and Robert 2018). Whatever their topic, I urge the community to pursue approaches and perspectives to questions on AI that are somewhat unique to and distinctive of the IS discipline, in addition to trying to emulate (and hopefully to surpass) what philosophers, organization scientists, psychologists, statisticians, management scientists, and computer scientists are very adept at and in approaches that they have ownership of.

The promise of AI in changing our world for the better in unquestionable. However, by reflecting on our societal responsibilities and carefully implementing some adjustments as suggested, we can potentially have a profound effect on the evolution of AI technologies, such that it becomes something that humanity can wholeheartedly embrace and live in harmony with.

6. Conclusions

Overall, the panelists shed light onto their own foci on the AI and robotics field and charted possible ways for future research in these domains. It is noteworthy that even though

each panelist has quite different research foci on AI and robot issues, they presented quite a consensus in terms of the need to further study the ethical dilemmas with AI research and practice. TP put more emphasis on the relationship between AI and humans (i.e., whether AI and robots will replace humans). Lionel highlighted the problems of AI biases and suggested that future research should attend to the auditing and evaluation of AI systems. Supra adopted the sociotechnical lens to view the AI phenomenon and concluded his discussion by highlighting the need for unique and distinctive contributions by the IS discipline. Based on such discussions, we call for treating issues related to the effects of AI and robots as grand challenges, and to further study then.

References

Abbasi, A., Sarker, S., and Chiang, R. (2016). "Big Data Research in IS: Toward an Inclusive Research Agenda," Journal of the AIS, Vol. 17, Issue 2, pp. i – xxxii.

Dirican, Cüneyt. "The impacts of robotics, artificial intelligence on business and economics." Procedia-Social and Behavioral Sciences 195 (2015): 564-573.

Kolbjørnsrud, V., Amico, R., & Thomas, R. J. (2016, November 2). How artificial intelligence will redefine management. Harvard Business Review. Retrieved from https://hbr.org/2016/11/how-artificial-intelligence-will-redefine-management

Mason, R. O. (1986). Four ethical issues of the information age. MIS Quarterly, 10(1), 5-12.

Matt, C., Trenz, M., Cheung, C.M.K., and Turel, O. 2019. "The Digitization of the Individual: Conceptual Foundations and Opportunities for Research," *Electronic Markets* (29:3), May 18, pp. 315–322.

Nilsson, N. J. (2014). Principles of artificial intelligence. Morgan Kaufmann. Palo Alto, CA

Noble, S. U. (2018). Algorithms of Oppression: How Search Engines Reinforce Racism. New York University Press, 2018.

O'Neal, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy, New York: Crown Books.

Rai, A., Constantinides, P. and Sarker, S. (2019) "Next generation digital platforms: toward human-AI hybrids," MIS Quarterly , 44 (1). iii-ix

Richardson, K. (2015). An anthropology of robots and AI: Annihilation anxiety and machines. Routledge. Abingdon-on-Thames, UK

Robert, L. P., Pierce, C., Marquis, E., Kim, S., Alahmad, R. (2020a). Designing Fair AI for Managing Employees in Organizations: A Review, Critique, and Design Agenda, Human-Computer Interaction, https://doi.org/10.1080/07370024.2020.1735391.

Robert, L. P., Gaurav, B., & Lütge, C. (2020b). ICIS 2019 SIGHCI Workshop Panel Report: Human–Computer Interaction Challenges and Opportunities for Fair, Trustworthy and Ethical Artificial Intelligence. AIS Transactions on Human-Computer Interaction, 12(2), pp. 96-108. DOI: 10.17705/1thci.00130.

Sarker, S., Chatterjee, S., Xiao, X., and Elbanna, A. (2019),"The Sociotechnical Perspective as an 'Axis of Cohesion' for the IS discipline: Recognizing its Historical Legacy and Ensuring its Continued Relevance," MIS Quarterly, Vol. 43, No. 3, pp. 695-719.

Seeber, I., Waizenegger, L., Seidel, S., Morana, S., Benbasat, I. and Lowry, P.B. (2020), "Collaborating with technology-based autonomous agents: Issues and research opportunities", Internet Research, Vol. 30, No. 1, pp. 1-18

Sun, T.Q., and Medaglia, R. 2019. "Mapping the Challenges of Artificial Intelligence in the Public Sector: Evidence from Public Healthcare," *Government Information Quarterly* (36:2), pp. 368-383.

Turel, O., Matt, C., Trenz, M., and Cheung, C.M.K. 2020. "An Intertwined Perspective on Technology and Digitised Individuals: Linkages, Needs and Outcomes," *Information Systems Journal* (30:6), pp. 929-939.

World Economic Forum (2016). Top 9 ethical issues in artificial intelligence. World Economic Forum. Retrieved from https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial-intelligence/

World Economic Forum (2020a). How to put AI ethics into practice: a 12-step guide. World Economic Forum. Retrieved from https://www.weforum.org/agenda/2020/09/how-to-put-ai-ethics-into-practice-in-12-steps/

World Economic Forum (2020b). Model artificial intelligence governance framework and assessment guide. World Economic Forum. Retrieved from https://www.weforum.org/projects/model-ai-governance-framework

You, S. and Robert, L. P. Jr. (2018) "Emotional Attachment, Performance, and Viability in Teams Collaborating with Embodied Physical Action (EPA) Robots," Journal of the Association for Information Systems: Vol. 19 : Iss. 5 , Article 2.

Pralash, A.V. and Das, S. (2020) "Intelligent Conversational Agents in Mental Healthcare Services: A Thematic Analysis of User Perceptions," Pacific Asia Journal of the Association for Information Systems: Vol. 12: Iss. 2, Article 1. DOI: 10.17705/1pais.12201.

Liang, T.P. (1992), "A Composite Approach to Inducing Knowledge for Expert Systems Design," Management Science, 38:1, pp. 1-17, 1992.