AIS Transactions on Human-Computer Interaction

Volume 13 | Issue 2 Article 5

6-30-2021

Intelligence Augmentation: Towards Building Human-Machine Symbiotic Relationship

Lina Zhou
University of North Carolina at Charlotte, lzhou@uncc.edu

Souren Paul
Northern Kenntucky University, souren.paul@gmail.com

Haluk Demirkan
University of Washington-Tacoma, haluk@uw.edu

Lingyao Yuan
Iowa State University, lyuan@iastate.edu

Jim Spohrer
IBM Almaden Research Center, spohrer@us.ibm.com

See next page for additional authors

Recommended Citation

Zhou, L., Paul, S., Demirkan, H., Yuan, L., Spohrer, J., Zhou, M., & Basu, J. (2021). Intelligence Augmentation: Towards Building Human-Machine Symbiotic Relationship. *AIS Transactions on Human-Computer Interaction, 13*(2), 243-264. https://doi.org/10.17705/1thci.00149
DOI: 10.17705/1thci.00149

Intelligence Augmentation: Towards Building Human-Machine Symbiotic Relationship

Authors
Lina Zhou, Souren Paul, Haluk Demirkan, Lingyao Yuan, Jim Spohrer, Michelle Zhou, and Julie Basu

Transactions on Human-Computer Interaction

Volume 13 Issue 2

6-2021

Intelligence Augmentation: Towards Building Human-Machine Symbiotic Relationship

Lina Zhou

Department of Business Information Systems and Operations Management, Belk College of Business, University of North Carolina at Charlotte, lzhou@uncc.edu

Souren Paul

Department of Business Informatics, Northern Kentucky University, souren.paul@gmail.com

Haluk Demirkan

Milgard School of Business, University of Washington Tacoma, haluk@uw.edu

Lingyao (Ivy) Yuan

Department of Information Systems and Business Analytics, Debbie & Jerry Ivy College of Business, Iowa State University, lyuan@iastate.edu

Jim Spohrer

Cognitive Opentech Group, IBM, spohrer@us.ibm.com

Michelle Zhou

Juji, mzhou@juji-inc.com

Julie Basu

smartQED, julie@smartqed.com

Recommended Citation

Zhou, L., Paul, S., Demirkan, H., Yuan, L., Spohrer, J., Zhou, M., & Basu, J. (2021). Intelligence augmentation: Towards building human-machine symbiotic relationship. *AIS Transactions on Human-Computer Interaction*, *13*(2), pp. 243-264.

DOI: 10.17705/1thci.00149

Available at http://aisel.aisnet.org/thci/vol13/iss2/5

Transactions on Human-Computer Interaction

Research Commentary DOI: 10.17705/1thci.00149 ISSN: 1944-3900

Intelligence Augmentation: Towards Building Human-Machine Symbiotic Relationship

Lina Zhou
Department of Business Information Systems and Operations Management
Belk College of Business, University of North Carolina at Charlotte
lzhou8@uncc.edu

Souren Paul
Department of Business Informatics
Northern Kentucky University
souren.paul@gmail.com

Haluk Demirkan
Milgard School of Business
University of Washington Tacoma
haluk@uw.edu

Lingyao (Ivy) Yuan
Department of Information Systems & Business Analytics
Debbie & Jerry Ivy College of Business, Iowa State University
lyuan@iastate.edu

Jim Spohrer
Cognitive Opentech Group, IBM
spohrer@us.ibm.com

Michelle Zhou
Juji
mzhou@juji-inc.com

Julie Basu
smartQED
julie@smartqed.com

Abstract:

Artificial intelligence, which people originally modeled after human intelligence, has made significant advances in recent years. These advances have caused many to fear that machines will surpass human intelligence and dominate humans. Intelligence augmentation (IA) has the potential to turn the tension between the two intelligence types into a symbiotic one. Although IA has not gained momentum until recent years, the idea that machines can amplify human abilities has existed for many decades. Expanded from a panel discussion on Intelligence Augmentation at the 2020 International Conference of Information Systems (ICIS), we define IA in light of its history and evolution and classify IA based on its capabilities, roles, and responsibilities. Based on reviewing the IA literature in terms of research themes, enabling technology, and applications, we identify key research issues, challenges, and future opportunities.

Keywords: Intelligence Augmentation, Human-Machine Symbiotic, Intelligence, Artificial Intelligence.

Fiona Nah was the accepting senior editor for this paper.

1 Introduction

Today, humanity has turned visions about theoretical computer science into reality through advancements in technologies, especially in artificial intelligence (AI), and the availability of big data across many different domains (Shneiderman, 2020). Businesses now invest heavily in AI to improve their operational efficiency and value creation. They deploy chatbots as customer service agents and robots in fulfillment centers. People interact with AI technologies in their daily lives in order to seek quick answers or assistance from virtual assistants, such as Alexa, Siri, and Cortana, through voice commands. Cars have self-driving and auto pilot capabilities, collision-avoidance systems, and more to help human drive safely and responsively. Artificial intelligence's widespread success in many different domains has stirred much excitement. However, we have also witnessed increasing concerns about the possibility that machines could control or dominate humanity and, thus, the tension between human intelligence and machine intelligence.

In the intelligence augmentation (IA) paradigm, the two intelligence types—namely, human and artificial intelligence—work together as a symbiotic system. IA can help organizations gain a competitive advantage and even survive. According to Aron and Sicular (2019), IA is expected to "create $2.9 trillion of business value and 6.2 billion hours of worker productivity" globally in 2021. Decision support/augmentation will even "surpass all other types of AI initiatives to account for 44 percent of the global AI-derived business value" by 2030. Specifically, 70 percent of enterprises will augment employees' productivity and tasks by implementing AI such as virtual assistants or other NLP-based conversational agents and robots by 2021 (Maita et al., 2018). Currently, the coronavirus disease of 2019 (COVID-19) pandemic has accelerated society's digital transformation such that many more organizations and individuals now adopt technologies such as online tools and AI. Therefore, it has become both timely and critical to investigate IA from both theory and practice perspectives, and identify research opportunities and open issues (Carter, Liu, & Cantrell, 2020; Robert, Bansal, Melville, & Stafford, 2020).

In this research commentary, we draw and significantly expand on a panel discussion on IA at the 2020 International Conference on Information Systems (ICIS). The panel, which comprised entrepreneurs and researchers in areas related to IA, generated great conversations on some key issues with IA, such as "what is IA?", "how does it differ from AI?", "what implications does IA have for future of work?", and "What kind of culture shift does IA bring to the workplace?". To further enlighten IA research, in this paper, we draw from intelligence theories to elucidate the symbiotic relationship between human and machine intelligence. In addition, we categorize research themes, enabling technology, and IA applications from the literature to guide future research and identify key research issues, opportunities, and challenges. Furthermore, we discuss best practices of adopting IA for companies and ethical and governance issues related to IA to inform IA practice.

2 Intelligence Augmentation

2.1 Definition

According to dictionary.com, augmentation refers to the act or process of augmenting something ("Augmentation", n.d.)—adding to it in a way that makes it bigger or better. We define IA as enhancing and elevating human's ability, intelligence, and performance with the help of information technology. IA stresses human-machine collaboration or human-machine symbiosis where machines perform what they do best (e.g., computing, recording, and doing routine, repetitive work) to aid humans in doing what humans do best (e.g., abstract reasoning, creating, and making in-depth discoveries about people and the world).

IA does not simply represent AI rebranded. Even though the underlying technology empowering IA and AI overlap, they have fundamentally different goals and foci. IA focuses on making people smarter, whereas AI focuses on making machines smarter. Unlike the traditional view that sees AI as autonomous systems that can fully automate tasks, workflows, and/or business processes and operate without human involvement, IA focuses on AI systems that work with humans to outperform either one alone. In particular, IA can augment human users' decision-making process and capabilities by providing otherwise hidden or inaccessible data-driven insights.

Humans' central role in IA not only helps improve the trust and reciprocity of human and machine intelligence but also helps allay the fears and concerns associated AI systems' proliferation (Mohanty & Vyas, 2018). Advances in AI can undoubtedly help advance IA. However, things that make people smarter (augmentations) include tools and organizations (Norman, 2014).

In defining IA, we emphasize the relationship between human and machines in general. IA constitutes just one aspect of augmentation: for strength (steam engine), perception (telescope, microscope), memory (writing, reading), thinking (logic), problem-solving and decision making (mathematics), investment (Black-Sholes), and so on. Kline (2020) calls augmentation via technology the techno-extension factor. While using the technology, humans need to be aware of the role machines (or algorithms) plays and how much they should rely on them. Taking GPS navigation for instance, should a driver follow the machine even though it directs the driver to a lake? To what extent should users rely on machines depends on not only the tasks but also how they understand the human self. Humans have (whether conscious or unconscious) biases. Using machine learning algorithms can reduce or even avoid potential human biases. Moreover, this relationship not only exists when adopting or using IA technology but also in creating it. When designing IA technology, designers need to be aware of users. After all, IA technology focuses on enhancing human capabilities.

2.2 Theory on Intelligence

We draw from theories on intelligence to more deeply understand the human intelligence-machine intelligence symbiotic relationship. Researchers have proposed different theories on intelligence, which we can group into four major types: psychometric theories, cognitive theories, cognitive-contextual theories, and biological theories (Bray & Kehle, 2013; Dweck & Henderson, 1989; Mackintosh, 2011). Psychometric theories come from work that has studied individual differences in test performance on cognitive tests. Questions about the structure of human intelligence, which includes general intelligence's importance, have dominated the psychometric theories. Cognitive theories come from work that studies the processes involved in intelligent performance. These processes range from the simple (e.g., inspection time) to the complex (e.g., working memory). Different theorists have focused on different processes (or aspects of these processes, such as processing speed). Cognitive-contextual theories emphasize processes that demonstrate intelligence in a particular context (such as a cultural environment). Major theories include triarchic theory of intelligence (Sternberg, 1985), the theory of multiple intelligences (Gardner, 2011), and theory of cognitive development (Piaget, 1977). Biological theories emphasize the relationship between intelligence, and the brain, and its functions.

Despite the different schools of thoughts on intelligence, researchers generally accept that machines can have intelligence in specific areas and even emotional intelligence but that they may not possess the general intelligence. The multiple intelligence theories (Gardner, 2011) help understand the strengths of human and machine intelligence. For instance, humans generally surpass machines in linguistic, interpersonal, creative (or experiential), and contextual (or practical) intelligence. On the other hand, machines can surpass humans in some logical/mathematical and analytical intelligence areas. The relative strengths of humans and machines will elucidate their relationship and lay the theoretical foundation for IA.

With guidance from intelligence theories, we characterize human intelligence and machine intelligence in spectra along multiple concrete dimensions, such as structured versus unstructured decisions, specialized versus general intelligence, computational depth versus breadth, repetitive versus non-routine/creative decisions, static/certain versus dynamic/uncertain tasks, knowledge versus wisdom, and experiential versus reflective intelligence (see Figure 1). Generally, machine intelligence has a strong potential for the categories on the left, whereas human intelligence for the categories on the right. Specifically, machine intelligence has the capabilities to address reasonably well-defined problems that have a narrow scope and a repetitive, experiential, non-creative, and static nature, while humans excel at defining and solving unfamiliar problems or making creative decisions and judgment that require flexibility and skills to adapt dynamically. Some artificially intelligent machines have outperformed humans in certain types of specialized intelligence (Malone, 2018) or some specific tasks that have clear boundary conditions. For instance, machines' efficiency in computation and ingestion, retrieval, and linking information far outpace humans' efficiency in doing so. On the other hand, humans have general intelligence (Malone, 2018), which allows them to perform a wide range of tasks. It is common sense, intuition, moral judgment, conceptual understanding, (bounded) rationality, reflective cognition, feeling, empathy, and sensation that epitomize humans' superiority and enables them to use information to reason, strategize, and handle uncertainties in addressing complex problems or new situations. In addition, reflective intelligence or thought is "the critical component of modern civilization: it is where new ideas come from" (Norman, 2014, p. 27).

The findings from the comparisons suggest that machine and human intelligence complement each other and that integrating them can potentially make a better world. As Norman (2014, p. 225) describes: "the automation works best when conditions are normal. When conditions become difficult...then the automation is also likely to fail. In other words, the automation takes over when it is least needed, gives up when it is most needed.”.

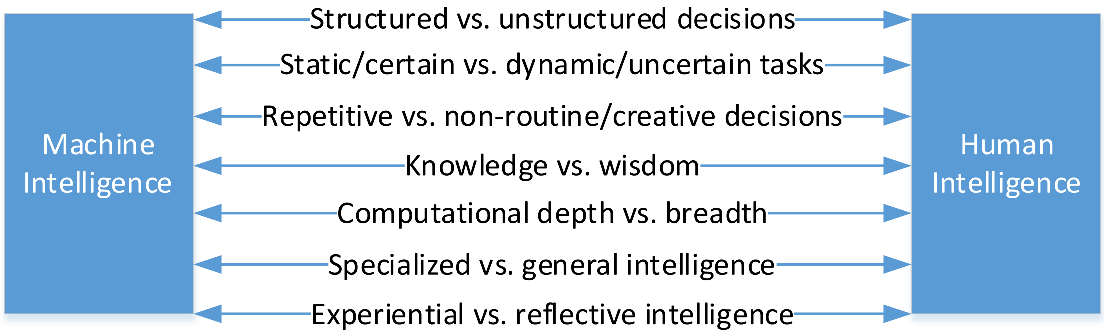

Figure 1. Multi-dimensional Comparison between Human and Machine Intelligence

2.3 History and Evolution of Intelligence Augmentation

Intelligence augmentation refers to a partnership between person and machine in which both contribute their strengths. The idea that machines can amplify human abilities has existed for many years and has a long lineage that leads to the present day. In 1945, Bush talked about how machines were evolving and how they had started to make people’s lives easier and then urged the men of science to record information and make the knowledge accessible to others. In 1960, Licklider (1960, p. 4) proposed cooperation between people and computers:

> *Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers. It will involve very close coupling between the human and the electronic members of the partnership. The main aims are 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In the anticipated symbiotic partnership, men will set the goals, formulate the hypothesis, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions in technical and scientific thinking. Preliminary analyses indicate that the symbiotic partnership will perform intellectual operations much more effectively than man alone can perform them.*

Technology pioneer Douglas Engelbart (1962, p. 1) also advocated for IA in writing:

> *By “augmenting human intellect” we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems. Increased capability in this respect is taken to mean a mixture of the following: more-rapid comprehension, better comprehension, the possibility of gaining a useful degree of comprehension in a situation that previously was too complex, speedier solutions, better solutions, and the possibility of finding solutions to problems that before seemed insoluble.*

Bush and Licklider envisioned and funded programs that benefited Engelbart building working systems. Douglas Engelbart went on to give what researchers have since called “the mother of all demos” (Engelbart & English, 1968) at a conference that the Association for Computing Machinery/Institute of Electrical and Electronics Engineers (ACM/IEEE) sponsored in 1968. He introduced a select group to practically implement all elements of what would later become personal computing.

Recent advancements in technology have started to drive people and society rather than people driving them. Norman (2014, p. xi) is one of the first authors who raised the issue of designing people-centered machines that will augment people:

> *Society has unwittingly fallen into a machine-centered orientation to life, one that emphasizes the needs of technology over those of people, thereby forcing people into a supporting role, one for which we are most unsuited.*

Kline (2020) discusses how cybernetic, information-using, feedback-controlled processes constitute a fundamentally important aspect of the living world in general and especially of human behavior and human-designed systems. These ideas have continued to develop and now emphasize the need for organizations and society to evolve to develop better human-tool capabilities to address problems and/or opportunities. Many people realize that adaptions to the environment have driven human evolution. The current stage of human evolution constitutes an information-rich human-made environment that leads to different goals and selection pressures for different capabilities (Spohrer & Engelbart, 2004). Bush and Licklider envisioned that advanced information technologies would accelerate the co-evolution and human-machine symbiosis for individuals and organizations and accelerate how much information they could store, process, and replicate and the speed at which they did so. Figure 2 illustrates a basic way for how the human system and tool system may co-evolve to result in an enhanced (cheaper, faster, better, etc.) capability infrastructure. Individuals and organizations tap into the societal capability infrastructure to enhance their performance and achieve goals.

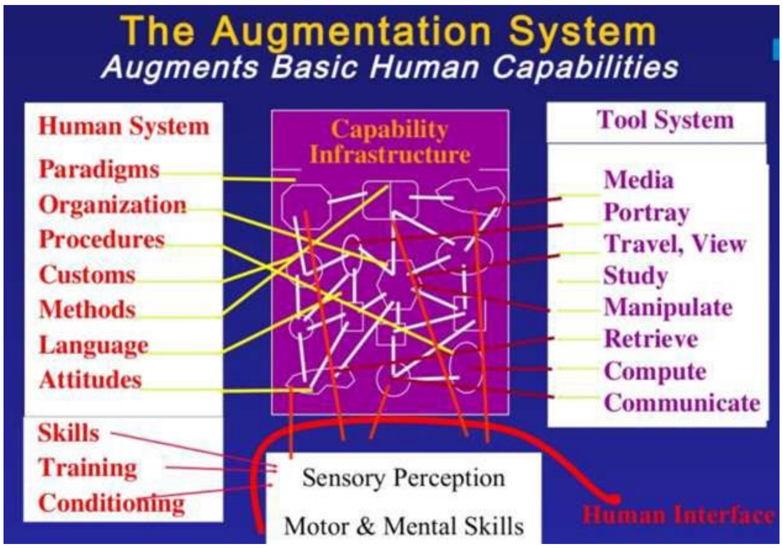

Figure 2. How Machines can Augment Basic Human Capabilities

Boumphrey (2016) sees systems as driving the AI revolution—these systems make us more productive, mobile, connected, and able to compete in the global world economy. They also improve safety in hazardous environments and in tedious jobs. Malone (2018) explores the different ways groups of people make decisions and how new forms of AI, especially machine learning, can help them do so. He predicts that AI, robotics, and automation will destroy many jobs—including jobs that require highly skilled knowledge—while also creating new ones.

2.4 IA Roles

As we summarize in Section 2.3, AI and IA have fundamentally different goals and foci. AI focuses on developing and advancing technology to think or act like humans, whereas IA focuses on expanding and extending humans' abilities. In addition to technology efficiency, effectiveness, and performance, user satisfaction, perception, and other experience also constitute important factors in designing and evaluating IA artifacts.

AI helps enable IA. Computers undoubtedly continue to increase in their ability to do things that only humans could once do (Demirkan, Spohrer, & Welser, 2016). Today, smart machines have become like humans by recognizing voices, processing natural language, learning, and interacting with the physical world through their vision, smell, touch and other senses, mobility, and motor control. In some cases, they do a much faster and better job than humans at recognizing patterns, performing rule-based analysis on very large amounts of data, and solving both structured and unstructured problems. Significant publications have talked about how smart machines with AI may take jobs from humans by replacing them. In contrast, people do not perceive IA as threatening since it does not replace humans but enhances their capabilities.

Automating (automation with AI) versus augmenting intelligence (augmentation with AI) represent two sides the same coin, and many jobs undoubtedly involve 100 percent routine, highly repeatable tasks that AI may completely automate. On the contrary, many partially routine and partially non-routine jobs will be amenable to automation that augments humans who look after their non-routine aspects (Rouse & Spohrer, 2018). We can view the automation-augmentation continuum as involving a mix between two different types of cognitive systems: biological and digital. Each cognitive system can play a range of roles: tool, assistant, collaborator, coach, and mediator. The progression from cognition tool to cognitive mediator requires cognitive systems with increasingly sophisticated task, world, self, user, and institutional context models (Siddike, Spohrer, Demirkan, & Kohda, 2018). Table 1 shows the different cognitive systems (tool, assistant, collaborator, coach, mediator) in relation to their roles and the increasingly complex self, other, and world models that they require.

Many people today use IA not just in their work but also their daily lives, which has democratized AI-enabled IA. Today, virtual assistants such as Apple's Siri, Microsoft's Cortana, Google's Now, Amazon's Echo, IBM's Watson, and other cognitive tools have begun to reach a level of utility that will provide a foundation for a new generation of cognitive collaborators and cognitive coaches (see Figure 3). Nonetheless, cognitive mediators require sufficient capability and trust to augment and scale expertise that is not only higher but also built from cognition-as-a-service building blocks that do not yet fully exist (Spohrer & Banavar, 2015). These capabilities will help people to achieve deeper insights into huge amounts of structured and unstructured data, which will boost both creativity and productivity. In the coming decade, we predict that most people will have many types of cognitive tools, assistants, collaborators, coaches, and mediators as a form of IA in their smartphones or equivalent technologies such as wearables and the environment. Cognitive mediators represent an evolution in both technology capability and social trust.

	Task & World Model/ Planning & Decisions	Self Model/ Self Capacity & Limits	User Model/ User's Episodic Memory and Identity	Institutions Model/ Trust & Social Acts
Tool	+	-	-	-
Assistant	++	+	-	-
Collaborator	+++	++	+	-
Coach	++++	+++	++	+
Mediator	+++++	++++	+++	++

Figure 3. IA Progression in Terms of Roles and Interaction Context

3 Literature Review on IA Research

To gain insights into the state of IA research, we conducted a literature search in digital libraries. We chose to use Web of Science (WoS) because it not only provides access to multiple databases but also provides comprehensive citation information across many academic disciplines, such as engineering, social sciences, computer science, life sciences, biomedical sciences, and the arts and humanities. We selected the following search terms: "augment* intelligent*", "intelligence augment*", "augment* human intelligence" where "*" denotes a wild card to retrieve papers published since 1960. The search resulted in 80 papers in total. Even though we obtained relatively few papers, they serve as lens to observe IA research themes and trends.

3.1 Research Trend

We plotted the papers' distribution by year in Figure 4. The figure shows that the first publication appeared in 1993 (i.e., Skagestad, 1993). The paper focused on ways to adapt computers to improve human thinking. As Skagestad (1993, p. 157) puts it: "The automation of intellectual housekeeping tasks was intended to bring about qualitative changes in our thinking, not just to enable us to do more of the same kind of thinking we had been doing all along".

The second paper (i.e., Borges & Baranauskas, 1998) did not appear until five years later. It proposed a user-centered approach to designing an expert system for training in the manufacturing context. One year later, another paper discussed the principles of wearable computers that can help humans manage, sort, and filter information to become intimately connected to their daily lives (Billinghurst & Starner, 1999).

Only one paper on IA appeared over the first decade of the 21st century. It examined the neural network augmented intelligent control of a turbo-fan engine to minimize a performance measure online (Kulkarni &

KrishnaKumar, 2003). From 2010 to 2017, the number of publications on IA remained small yet relatively steady. Researchers have developed IA in many different application domains, such as neurosurgery, distance learning, precision medicine, imaging reading, disaster mitigation, law, and design. Interestingly, Bauer (2010) raised concerns about a transhumanism future. Transhumanism advocates for transforming human nature via technologies such as pharmacology and nanotechnology. This issue seems to have attracted even greater concern than the concern around AI replacing humans.

Research on IA has gained momentum and has grown significantly since 2018 (see Figure 4). The distribution of citation count (see Figure 5) shows a similar trend to the publication count distribution but at a much larger scale. The number of citations that IA publications received, which indicates their research impact, started to grow exponentially since 2014. The drop in 2021 comes from the fact that we collected data in 2021 and, thus, the year had not yet completed. The statistics demonstrate IA's increasing research impact in recent years.

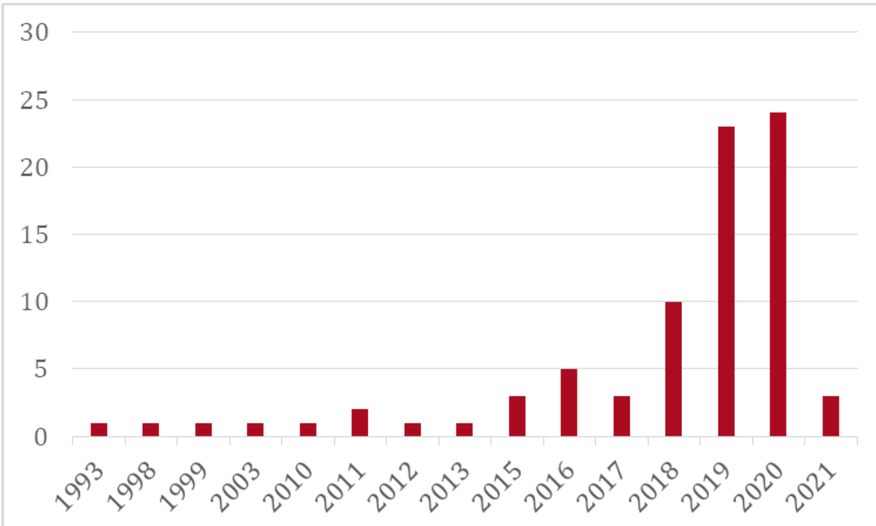

Figure 4. Number of IA Publications

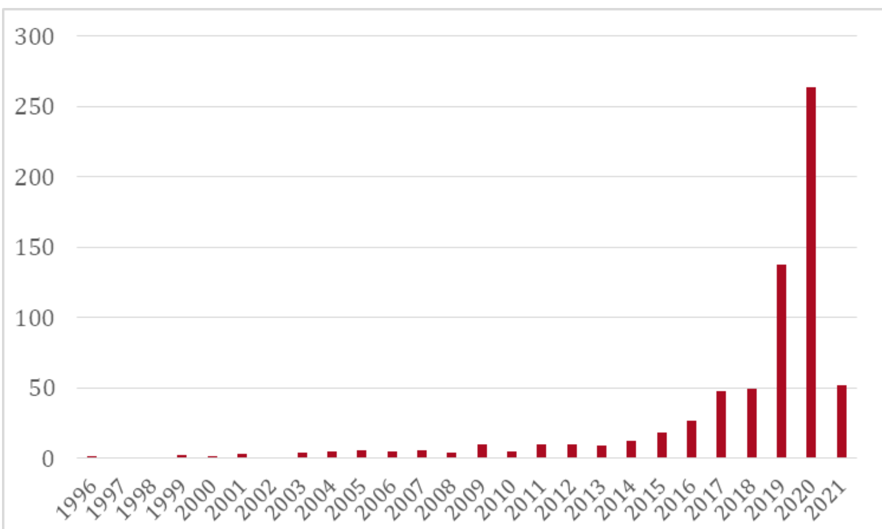

Figure 5. Number of Citations that IA Publications Received

3.2 Research Disciplines

To understand the breadth of IA research, we summarized the (sub-)disciplines that the IA publications came from, which we sorted in descending order of their publication counts (≥ 2) in Table 1. In the table, we made a few minor modifications to the original discipline classification from the WoS such as treating "information systems" as a separate discipline and adding "medicine" and "business" as new broad disciplines to group several (sub-)disciplines.

We can see from Table 1 that the IS discipline generated the largest number of publications. Among the top ten (sub-)disciplines, four (hardware architecture, artificial intelligence, software engineering, theory methods) belong to computer science. In addition, medicine had the largest number of (sub)-disciplines and total publication count from the table. Other disciplines that had significant IA presence include business (general, finance, management, etc.), telecommunication, educational research, architecture, information science, law, and materials science. In addition to the above-mentioned diverse disciplines, other (sub-)disciplines such as multidisciplinary sciences, interdisciplinary applications, and neuroscience clearly demonstrate the interdisciplinary nature of IA research.

Table 1. Number of IA Publications by Discipline

Discipline / (sub-)discipline	Count
Information systems	10
Electrical electronic engineering	7
Computer science / hardware architecture	5
Medicine / medicine medical imaging	5
Computer science / artificial intelligence	4
Medicine / neurosciences	4
Medicine / pharmacy	4
Telecommunications	4
Computer science / software engineering	3
Computer science / theory methods	3
Education / educational research	3
Medicine / medical informatics	3
Multidisciplinary sciences	3
Architecture	2
Business	2
Business / finance	2
Medicine / cardiovascular systems	2
Medicine / clinical neurology	2
Computer science / interdisciplinary applications	2
Medicine / dermatology	2
Ethics	2
Medicine / healthcare sciences services	2
Information science / library science	2
Law	2
Business / management	2
Materials science / characterization testing	2

3.3 Research Areas and Topics

Moving beyond research disciplines, we further looked into the specific areas that IA research has addressed. To this end, we first extracted author-defined keywords that reflect specific research topics and then manually clustered them into areas based on their similarities. Table 2 lists some of the research areas and related research topics that we summarized. Unsurprisingly, the research areas largely overlap with the research disciplines (see Table 1). Additionally, we also identified several new research areas, such as disaster management, workforce management, regulation, human cognition, and environmental and urban studies.

The IA research topics in each research area not only suggest the related research goals of IA but also illustrate how organizations and researchers have applied IA and evidence of IA in the real world. For instance, researchers have discussed IA in educational contexts to support blended learning, distance/electronic learning, interactive learning, intelligent tutoring, social learning, and flipped classrooms. In medicine, IA has been applied to improve clinical decisions and patient safety via personalized medicine, robotic-assisted surgery, and formulary management. In the business context, IA has helped address goals in relation to consumers, marketing, user profiling, negotiation, financial technology (fintech), and hyper-personalization, and so on. In particular, hyper-personalization takes personalized marketing one step further by employing AI to deliver more relevant content and information to individual users. Interestingly, organizations have also used IA to enhance their regulatory technology (regtech), particularly in the finance sector.

Table 2. Sample IA Research Areas and Topics

Research areas	Sample research topics
Education (learning)	Blended learning Interactive learning environment Intelligent tutoring system Social learning
Medicine	Clinical decision tool Personalized medicine Point-of-care systems Proactive health management Robotic assisted surgery Formulary management Risk stratification Learning health systems
Business	Risk analysis User profiling Negotiation support Financial technology Hyper-personalization Customer ROI and lifetime value assessment
Disaster management	Disaster mitigation Disaster prevention
Environmental and urban studies	Corrosion rate prediction Urban space restructuring and transition Energy optimization
Human cognition	Attention control Intersemiotic translation Perception
Regulation	Regulatory technology decentralized governance financial regulation

3.4 Enabling Technologies

To better illustrate the technologies that enable IA, or machine intelligence, we extracted technology-related terms from the retrieved publications' keywords. Note that, among the top 10 keywords based on publication count, seven related to technology (artificial intelligence, machine learning, deep learning, convolutional neural networks, (artificial) neural networks, big data, cloud computing, cognitive computing, natural language processing, and industry 4.0). The fourth industry revolution (industry 4.0) describes the current trend toward an inclusive, human-centered future via adopting converging technologies. In addition to the above fundamental technology enablers, other technologies that support IA include chatbots, cyborgs, financial technology, regulatory technology, blockchain, social computing, optimization algorithms, augmented reality, data analytics, data science, decision modeling, design supports, edge computing, emotion recognition, face detection, image recognition, knowledge based systems, predictive models, reasoning, visualization, virtual machines, simulation, the industrial Internet of things (IIoT), and cognitive technologies. Finally, we generated a word cloud of the keywords extracted from the IA publications (see

Figure 6) using MonkeyLearn. The figure highlights IA's central themes and its relationship with AI and other enabling technologies. Interestingly, IA frequently co-occurs with real-world evidence, organizational decision making, clinical decision making, and so on, suggesting the real-world impact of IA.

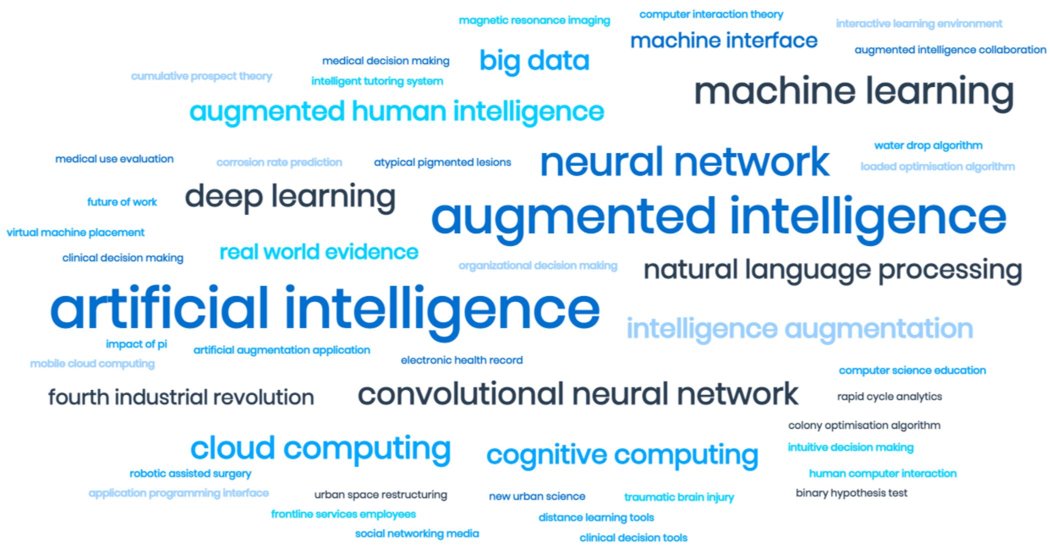

Figure 6. A Word Cloud of Keywords in the IA Publications

4 A Basic IA Framework

Based on conceptualizing IA; theoretically analyzing machine intelligence and human intelligence, IA roles, and capabilities; and reviewing the IA literature, we propose a basic IA ecosystem framework. The framework comprises five key components: goals, humans, machines (technologies), governance, and environment (see Figure 7).

4.1 Goals

Intelligence involves the ability to achieve goals (Malone 2018), and performing tasks is the pathway moving toward the goals. IA aims to help humans achieve their goals by performing related tasks. The forces that drive humans' goals may include certain organizational contexts, business objectives, individual users' needs or preferences, changes in environmental and technology factors, and so on. Additionally, machines can also assist humans with identifying and refining the goals. Accomplishing different goals may require different types of tasks, which can range from generation (planning and creativity), to choice (problem solving and decision making, negotiation (cognitive conflict and mixed motives), to execution (contests and performances) (McGrath, 1984, p. 61). For the purpose of IA, we assume that humans and machines work collaboratively; and thus the negotiation and contests tasks refer to the context of humans versus other humans or IA systems instead of humans versus machines within an IA system. Different tasks may rely on different types of intelligence. For instance, creativity tasks tend to focus on creative, unstructured, dynamic, reflective, and contextual intelligence; cognitive conflict tasks emphasize interpersonal and contextual intelligence; and problem solving tasks typically require specialized, experiential, structured, logical and analytical intelligence. Further, the goals are subject to continuous refinement based on the monitoring and evaluation of the outcomes from performing the related tasks.

4.2 Humans

When designing IA, designers must consider human users to understand when, what, and how to augment human abilities. This component not only draws on human intelligence theories but also reflects the human system that feeds the capability infrastructure in the augmentation system (see Figure 2). Human factors not only include their psychological, physiological, demographic, and behavioral characteristics but also other dynamic features such as related experience, literacy, and education. The information about human users helps designers adapt and even personalize IA solutions to specific users.

4.3 Technologies

Drawing on the strengths of machine intelligence, this discussion focuses on enabling technologies as tools, the other cornerstone of the capability infrastructure in the IA system (see Figure 2). Based on the specificity of technologies that enable IA, we classify the technologies from the literature into two categories: general-purpose technologies and domain-specific technologies. *General-purpose technologies* cover AI and smart technologies such as machine learning, natural language processing, image recognition, and the IoT since one can potentially apply them to any specific domain or application context. In contrast, *domain-specific technologies* address domain-specific tasks that can result from customizing or adapting general-purpose technologies. For instance, financial technologies, regulatory technologies, and clustered regularly interspaced short palindromic repeats (CRISPR) address financial, regulatory body, and genomics issues, respectively. Starting with tools, the design of smart machines can progress toward assistants, collaborators, coaches, and mediators with different capabilities, roles, and responsibilities to augment human intelligence by providing assistance in achieving goals. It is worth noting that the role of enabling technologies is not limited to goals but includes interfacing with humans and machines in their interactions. Such examples may include avatars, chatbots, cognitive computing, and brain-computer interfaces.

4.4 Governance

Without a governance structure, humans would not be able to interact with machines in predictable ways. Both machines and humans need to abide by laws that the governance form dictates to prevent them from taking drastic measures that may lead to disruptive actions. In addition to control, governance structure also implies responsibility, accountability, and ethics, which make machines and humans take either autonomous or interdependent actions toward common goals. Malone (2018) discusses five governance systems: hierarchies, democracies, markets, communities, and ecosystems. Governance systems gather and interpret information (or sensing), improve memory (e.g., by structuring symptoms of medical databases for better diagnosis), and increase learning through pattern recognition. While humans will always be the final decision makers in each system, their decisions will be systematically informed through knowledge bases with analytical capabilities (Arndt, 2020; Malone, 2018). Except for one paper we reviewed that used decentralized governance, governance issues have received little attention in the current IA literature.

4.5 Environment

A variety of internal and external environment factors can affect the development and adoption of IA. These factors include institutional structure, financial resources, human resources, technology infrastructure, economic and social environments, policies, and global environment, and key stakeholders such as customers, competitors, regulatory agencies, and IA research and development. For instance, the technology infrastructure comprises hardware and software platforms, networking technologies, and data-management technologies. An IA system must constantly monitor and respond to changes in these environment factors. Indeed, IA environments constantly change due to new developments in technology, political shifts, evolving consumer preferences, new regulations, and emerging international events. For instance, the ongoing COVID-19 pandemic and government intervention policies have helped misinformation spread. IA needs to respond quickly to changes in their environments in order to sustain and survive. Considering these factors will require more sophisticated cognitive models.

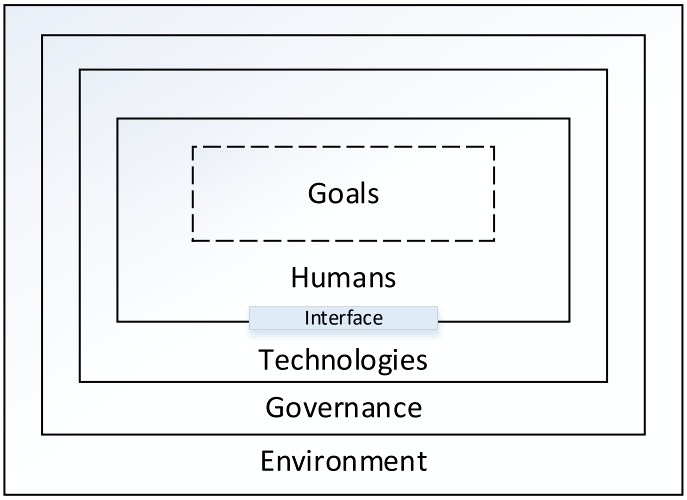

Figure 7. A Basic Framework of Intelligence Augmentation

5 Research Opportunities, Issues, and Challenges

Based on the literature search and categorization that we discuss in Section 4, we found that 1) overall, researchers have conducted little IA research despite the increase starting from 2019, 2) researchers from many disciplines have conducted IA research (including some interdisciplinary IA research), 3) IA research has significant implications in many different domains, 4) most IA research comes from the IS discipline. Based on our observations and the panel's expertise, we emphasize the importance of human-machine relationship in IA related research rather than viewing the human intelligence and machine intelligence as two isolated entities. We believe the relationship represents the fundamental core of the intelligence augmentation. In this section, we focus on the human-machine symbiotic relationship and discuss the opportunities, issues, and challenges in IA research with guidance from the IA framework.

5.1 Governance with Trusted AI and Responsible People

Among the proposed IA framework's components, researchers have not sufficiently studied governance. We expect decision governance to be one area that will witness significant changes in the future. Governance ensures the human-machine symbiotic relationship can work by establishing trust.

AI-enabled IA already exists in our cars, homes, and jobs. IA also increasingly supports human decision making. While AI holds the promise of delivering valuable insights and knowledge across many different applications, whether society broadly adopts AI systems will rely heavily on whether we can trust their output. When we use more tools, assistants, collaborators, coach, and mediators for our personal lives and work, we need make sure that we trust them. Humans trust technology based on understanding how it works and assessing its safety and reliability. To trust a technology and recommend an algorithm's decision, we need to learn and understand its reliability and fairness. Many organizations have begun working to build trusted AI principles. For example, IBM's commitment to make AI more trustworthy drove the organization to join the LF AI Foundation that supports and sustains open source innovation in AI, machine learning, and deep learning. The Trusted AI committee that the LF AI Foundation established has identified eight trustworthy AI principles: reproducibility, robustness, equitability, privacy, explainability, accountability, transparency and security ((R)REPEATS) (Cardoso, 2021). We describe each principle in Table 3.

As people gain more powers (e.g., augmented capabilities with intelligence, physical strength, social interaction, perception), we will need further research on developing trusted and responsible IA solutions and new theories, models and approaches to implement (R)REPEATS.

Table 3. Trusted AI Principles (adopted from LF AI & Data) (Cardoso, 2021)

• ***Reproducibility*** *is the ability of an independent team to replicate in an equivalent AI environment, domain or area, the same experiences or results using the same AI methods, data, software, codes, algorithms, models, and documentation, to reach the same conclusions as the original research or activity. Adhering to this principle will ensure the reliability of the results or experiences produced by any AI.* • ***Robustness*** *refers to the stability, resilience, and performance of the systems and machines dealing with changing ecosystems. AI must function robustly throughout its life cycle and potential risks should be continually assessed and managed.* • ***Equitability*** *for AI and the people behind AI should take deliberate steps—in the AI life-cycle—to avoid intended or unintended bias and unfairness that would inadvertently cause harm.* • ***Privacy*** *requires AI systems to guarantee privacy and data protection throughout a system's entire lifecycle. The lifecycle activities include the information initially collected from users, as well as information generated about users throughout their interaction with the system e.g., outputs that are AI-generated for specific users or how users responded to recommendations. Any AI must ensure that data collected or inferred about individuals will not be used to unlawfully or unfairly discriminate against them. Privacy and transparency are especially needed when dealing with digital records that allow inferences such as identity, preferences, and future behavior.* • ***Explainability*** *is the ability to describe how AI works, i.e., makes decisions. Explanations should be produced regarding both the procedures followed by the AI (i.e., its inputs, methods, models, and outputs) and the specific decisions that are made. These explanations should be accessible to people with varying degrees of expertise and capabilities including the public. For the explainability principle to take effect, the AI engineering discipline should be sufficiently advanced such that technical experts possess an appropriate understanding of the technology, development processes, and operational methods of its AI systems, including the ability to explain the sources and triggers for decisions through transparent, traceable processes and auditable methodologies, data sources, and design procedure and documentation.* • ***Accountability*** *requires AI and people behind the AI to explain, justify, and take responsibility for any decision and action made by the AI. Mechanisms, such as governance and tools, are necessary to achieve accountability.* • ***Transparency*** *entails the disclosure around AI systems to ensure that people understand AI-based outcomes, especially in high-risk AI domains. When relevant and not immediately obvious, users should be clearly informed when and how they are interacting with an AI and not a human being. For transparency, ensuring that clear information is provided about the AI's capabilities and limitations, in particular the purpose for which the systems are intended, is necessary. Information about training and testing data sets where feasible, the conditions under which AI can be expected to function as intended and the expected level of accuracy in achieving the specified purpose, should also be supplied.* • ***Security and safety*** *of AI should be tested and assured across the entire life cycle within an explicit and well-defined domain of use. In addition, any AI should be designed to also safeguard the people who are impacted.*

5.2 Mitigating Bias

Bias violates the equability principle of trusted AI. Humans have well-documented cognitive biases such as confirmation bias, which means their emotions may cloud their judgments, they may overgeneralize from personal experience (Mohanty & Vyas, 2018), they may show mental set and functional fixedness, and so on (Sternberg & Sternberg 2008). Researchers have increasingly realized that human biases have made their way into AI systems and even produced harmful results (Manyika, Silkberg, & Presten, 2019). For instance, artificial intelligence can reflect and amplify social bias (e.g., gender and racial bias (Benjamin, 2019)) in dangerous ways (Smith, 2020), which determines who gets a job interview or bank loan. We can attribute this bias to the data and algorithms that AI uses and to human biases. Many AI principles, such as understanding and measuring fairness (see Table 3), can serve as potential solutions to tackling bias in AI. For instance, counterfactual fairness focuses on keeping AI model's decisions unchanged in situations where sensitive personal attributes change. Establishing responsible process in deploying AI and making designers' choice more transparent can help mitigate bias (Manyika et al., 2019). Nevertheless, bias could have a much broader scope, and many may still go unnoticed. The ultimate solution will likely be a holistic one, which requires people, technology, and organization.

5.3 Ethical Issues

People have augmented their biological capacity through artifacts for many years (e.g., running shoes to improve physical capacity, calculators to boost mathematical skills, and eye glasses to enhance vision). As the technologies we use rapidly become more sophisticated, more personalized, and smarter, our ability to interact and use these technologies will also expand. These changes and adaptations create a new wave of ethical issues that need further research. Some ethical issues around AI-enabled IA include:

- **Customer data privacy**: digital platforms already collect and analyze people's behavioral data and build models to predict their shopping behavior (Amazon, Google, Facebook etc.). They can perform mind and social interaction modeling to anticipate others' decisions, actions, and interests. Organizations try to use as much customer data for their gain. The clear line between using data and invading individual privacy lacks a clear definition. Should we rely on companies' "good conscious"? Or should government regulate how organizations can use customer data? As we still need to address issues on the ownership and use of customer data still, more questions emerge such as whether or to what extent people can monetize their own data (such as DataSwift).
- **Digital nudging**: nowadays, people's daily activities increasingly depend on numerous mobile applications. While collecting data constantly, mobile applications on our smartphones and/or other digital Web-enabled devices also push targeted advertisements into our social media feeds, communication channels and webpages. While such applications expedite the decision-making process for many people, they also influence their decisions. Consciously and unconsciously, those applications "nudge" people to make various decisions by taking advantage of our cognitive biases, heuristics, and habits. As such, these applications raise ethical questions about data privacy, autonomy, consent, and how AI and data should (or should not) be used to nudge individuals and groups.
- **Digital addiction**: technology has undoubtedly enriched our lives and enhanced our abilities. At the same time, we, as humans, face digital addiction. As technologies assist our lives, we also found ourselves incapable from separating from our devices. Due to individual differences, for some people, digital addiction can be harmful as other types of addictions. Electronic commerce no longer constitutes a competitive advantage but rather a surviving necessity. Keeping customers "hooked" on various digital platforms and establishing company virtual presence on those platforms at all times become the goal.
- **Cognitive extension**: the way we view the mind can affect our norms and values, such as what psychological disorders we recognize, what kinds of treatments we provide, and how we assess people's cognitive capacities. We may see the direct effects of AI when cognitive assistants such as Alexa start making our personal appointments for meetings. More people now discuss how smart machines can augment our cognitive capabilities. These machines may have an impact on long-term memory or individual experiences.
- **Brain-computer interface**: a brain-computer interface, sometimes called a neural control interface or mind-machine interface, refers to a direct communication pathway between an enhanced or wired brain and an external device (Krucoff, Rahimpour, Slutzky, Edgerton, & Turner, 2016). These devices have the potential to change our cognitive profiles and even challenge how we understand what it means to be human. As such, they may begin to influence our decision-making processes. Ethical issues arise whenever only some people use performance-enhancing technologies in competitive situations, such as sports, academics, and businesses.

5.4 Human-AI Interaction

We undoubtedly need to improve the algorithms and techniques that underlie IA and AI to make better IA and AI. However, researchers also need to pay more attention to the human factors in the "collaborative" relationship. For example, one study shows that human trust on intelligent agents could depend on human users' personality traits (Zhou, Mark, Li, & Yang, 2019). Some key questions remain unanswered. For instance, how do humans consciously or unconsciously feel and respond to AI? How should we incorporate user experience in evaluating AI technology? What real human-technology dynamics does this relationship contain? What human factors does one need to consider in designing and evaluating an IA system? How can we provide methods to better support human-AI interaction?

IA can also provide unique opportunities for researchers to better investigate the existing research domain that involves complex and dynamic human interactions. Methodologically, using IA technologies facilitating group collaboration has the potential to shed new light into understanding the human interaction in a group setting. Since one can manipulate IA easier than humans, research can simulate different scenarios to observe and measure the changes in the group dynamic by using IA to study humans.

Many IA technologies have expanded from being merely a supportive tool, such as calculators, to having a more collaborative role, such as providing analytics for decision making. Those high intelligent IA technologies will keep emerging in the upcoming years. Existing human-computer interaction theories may still apply to this emerging area, but we may need to adapt or modify them to fit IA research. Of course, it may be that no current overarching theory that one can use to understand IA use exists. Existing theories that view technology as a tool may not apply in this new information era with users co-existing with intelligent technology. We may face a situation without an applicable overarching theory for more "intelligent" IA technology due to the evolving and complex nature of the interaction between IA technology and users. Many researchers may recognize the great opportunities associated with an emerging technology such as IA. On the other hand, IA has great practical implications to industry. Answers to all the research questions raised above will ultimately lead society to better design and improve IA technology. How can we improve the interaction between human and machine? Can users shape the direction in which technology develops? Or, perhaps a better question, should they?

5.5 Future of Workplace Culture

As our lives increasingly require more machine-human collaborations, people will need to become used to collaborating with machines in their personal and professional lives. Just like how people today ubiquitously adopt computers and smartphones, people will have to adopt AI/intelligent computer systems in their lives. Thus, we need to democratize AI education, AI tool use, and AI adoption. Workplace culture constitutes one type of environmental factor in IA.

People matter the most in any business. With the help of IA, every person in the workforce has to imagine a world in which they have 100 digital workers working for them as our smartphone applications "grow up" and obtain speech interfaces. In short, everyone becomes a manager of 100 digital workers. Knowing how to manage IA and AI would be crucial to run business successfully. Helping every employee adapt and grow into responsible "managers" of digital workers will be critical for businesses.

This emphasis on digital workers will lead to the evolution of a new form of workplace culture. Everyone in the workplace will face higher expectations. Digital workers will augment human intelligence to the extent that they can effectively use the digital workers to achieve their goals and satisfy all the "higher demands" that others place on them. It will be interesting to see how the reliance on digital workers will shape organizational culture's traditional dimensions, such as people orientation, outcome orientation, easygoingness, detail orientation, and team orientation.

5.6 Skill Shift in Workforce Training

In IA, each human in a centralized control position must assess the distributed situation in which many machines evolve. Generally, the situation assessment involves acquiring, organizing, and abstracting information about the environment. Workers also need to overcome resistance to change and master the skills they need to effectively partner with these IA systems. In doing so, they will find it easier to build and try out AI applications, which will lower the barrier to entry.

Work in the future depends on collaborative intelligence between people and AI. Collaborative intelligence has begun to create an environment for people and AI to co-create value and enable IA. As a result, work and jobs in the future will change. The augmentation system (see Figure 2) suggests that the radical changes that the IA perspective introduces requires new professional skills, knowledge, language, attitude, customs, methods, procedures, organization, and so on. We have recently heard discussions about I-shaped, T-shaped, M-shaped, pi-shaped, and dash-shaped professionals. These shapes describe whether a professional specializes deeply in one area (I shaped), specializes in just one area but has good knowledge and communication skills across many other areas (T shaped), specializes in two areas (pi shaped), specializes in several fields (M-shaped), or has broad rather than deep knowledge (i.e., a generalist) (dash shaped).

We need more T-shaped people in order to co-create value with AI. A need for T-shaped people first arose in conjunction with computing professionals. Today, we refer to the exponential rate of change in computing as Moore's Law after Gordon Moore who observed that the number of transistors on a chip doubled every 18 to 24 months. However, we have seen growth in knowledge, technology, and organizations outside computing as well. While the traditional Moore's Law depends on organizations investing in creating smaller and cheaper transistors, a generalized Moore's Law would depend on organizations investing to create more T-shaped people whom technology augments/enables and creating and harnessing new knowledge's value better.

T-shaped professionals and citizens continue to learn over their entire lives, and a next-generation education system has prepared them to compete for collaborators on projects that matter most to them (Moghaddam, Yurko, Demirkan, Tymann, & Rayes, 2020). They remain ready for team work and collaboration across local and global networks and to build smarter service systems. They can communicate broadly and empathetically, seek challenges, engage deeply with problems, and think critically. They have an entrepreneurial mind, find opportunities, and have much curiosity and imagination. Being T-shaped is mostly about how people and AI provide service to each other and co-create value. They provide significant opportunities for renovating our education systems for prepare students for work in the future. A promising direction for education involves training non-STEM students to understand the IA area, to learn to use AI tools, and to even join efforts to teach and train IA.

5.7 Multidisciplinary Perspective

By its nature, IA constitutes a multidisciplinary research area. Even though researchers have conceived of IA for many years, it has not taken off until recently. Existing theories may still apply to this nascent area, or we may have yet to establish an overarching theory for IA technology due to its new and unique nature. We need to understand technology in a more multi-faceted way before we can establish theories. To establish such understanding in order to build theory, we need different research methods (qualitative and quantitative), different philosophical views (positivism, interpretivism), and collaboration with different disciplines (computer science, engineering, social science, psychology, etc.). Multidisciplinarity also concurs with the call for preparing T-shaped people (Demirkan & Spohrer, 2018). The more diverse viewpoints we can obtain, the more holistic understanding we can develop towards IA. Our understandings will guide us to better improve IA or, even better, understand humans better through it. As we state in Section 5.4, AI has a broader role beyond extending humans' social and cognitive capabilities. It can also create value by helping humans better understand themselves. For instance, IA may provide new opportunities for people to observe and understand their cognitive biases and behavioral and emotional patterns.

6 Conclusions

With a focus on combining the strengths of human and machine intelligence, IA has the possibility to fundamentally change the role that technology plays in our work and life. The significant advancement and accessibility of AI technologies and related technology infrastructure has made the present a historical moment for IA. As IA enters into our daily lives and workplace, it creates tremendous opportunities to observe how humans interact with IA, which will provide ample evidence and help researchers and scientists develop new theories to explain and guide efforts to cultivate human-machine synergy for human goods, improve enabling technologies, and even prepare and train the future workforce. The relationship between human and machine will likely change as technology evolves. Ultimately, IA focuses on creating value for humans, be it business outcomes, cognition enhancement, or innovations. Thus, as a research field, IA has strong real-world relevance and requires a multidisciplinary perspective that considers goals, human factors/context, effective and efficient technologies, the ways in which humans and technologies interact, governance structures, and environmental constraints and their feedback. To advance the field, academics, industry, and regulatory agencies will need to make a concerted effort. The AI area has abundant research opportunities and issues that invite cross-disciplinary collaboration.

Acknowledgments

We thank the Editor-in-Chief, Fiona Nah, and the Managing Editor, Gregory D. Moody, for their support and guidance throughout the process. We also thank Dongsong Zhang for his constructive suggestions and Adam LeBrocq for his thoughtful edits and comments.

References

Arndt, F. (2020). *Book review: Superminds: The surprising power of people and computers thinking together.* Thousand Oaks, CA: Sage.

Aron, D., & Sicular, S. (2019). Leverage augmented intelligence to win with AI. *Gartner*.

Augmentation. (n.d.). In *Dictionary.com.* Retrieved from https://www.dictionary.com/browse/augmentation

Bauer, K. A. (2010). Transhumanism and its critics: Five arguments against a posthuman future. *International Journal of Technoethics*, *1*(3), 1-10.

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code.* Medford, MA: Polity Press.

Billinghurst, M., & Starner, T. (1999). Wearable devices: new ways to manage information. *Computer*, *32*(1), 57-64.

Borges, M. A. F., & Baranauskas, M. C. C. (1998). A user-centred approach to the design of an expert system for training. *British Journal of Educational Technology*, *29*(1), 25-34.

Bray, M. A., & Kehle, T. J. (2013). *The Oxford handbook of school psychology*. Oxford, UK: Oxford University Press.

Bush, V. (1945). As we may think. *The Atlantic Monthly*, *176*(1), 101-108.

Cardoso, J. Z. (2021). LF AI & data announces principles for trusted AI. *The Linux Foundation*. Retrieved from https://lfaidata.foundation/blog/2021/02/08/lf-ai-data-announces-principles-for-trusted-ai/

Carter, L., Liu, D., & Cantrell, C. (2020). Exploring the intersection of the digital divide and artificial intelligence: A hermeneutic literature review. *AIS Transactions on Human-Computer Interaction*, *12*(4), 253-275.

Demirkan, H., & Spohrer, J. C. (2018). Commentary—cultivating T-shaped professionals in the era of digital transformation. *Service Science*, *10*(1), 98-109.

Demirkan, H., Spohrer, J. C., & Welser, J. J. (2016). Digital innovation and strategic transformation. IT *Professional*, *18*(6), 14-18.

Dweck, C. S., & Henderson, V. L. (1989). *Theories of intelligence: Background and measures.* ERIC.

Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework.* Menlo Park, CA: Stanford Research Institute.

Engelbart, D. C., & English, W. K. (1968). A research center for augmenting human intellect. In Proceedings of the Joint Computer Conference.

Gardner, H. E. (2011). Frames of mind: The theory of multiple intelligences. Hachette, UK: Cambridge University Press.

Kline, S. J. (2020). Conceptual foundations for multidisciplinary thinking. Palo Alto, CA: Stanford University Press.

Krucoff, M. O., Rahimpour, S., Slutzky, M. W., Edgerton, V. R., & Turner, D. A. (2016). Enhancing nervous system recovery through neurobiologics, neural interface training, and neurorehabilitation. *Frontiers in Neuroscience*, *10*.

Kulkarni, N. V., & KrishnaKumar, K. (2003). Intelligent engine control using an adaptive critic. *IEEE Transactions on Control Systems Technology*, *11*(2), 164-173.

Licklider, J. C. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, 1, 4-11.

Boumphrey, R. (2016). *Foresight review of robotics and autonomous systems: Serving a safer world*. Lloyd's Register Foundation.

Mackintosh, N. (2011). *IQ and human intelligence*. Oxford, UK: Oxford University Press.

Maita, K., Poitevin, H., Howson, C., Sau, M., Hunter, E., & Sicular, S. (2018). Predicts 2019: AI and the future of work. *Gartner*.

Malone, T. W. (2018). How human-computer "superminds" are redefining the future of work. *MIT Sloan Management Review*, *59*(4), 33-41.

Manyika, J., Silberg, J., & Presten, B. (2019). What do we do about the biases in AI. *Harvard Business Review*. Retrieved from https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai

McGarth, E. J. (1984). *Groups: Interaction and performance*. Englewood Cliffs, NJ: Prentice-Hall.

Moghaddam, Y., Yurko, H., Demirkan, H., Tymann, N., & Rayes, A. (2020). The future of work: How artificial intelligence can augment human capabilities. New York, NY: Business Expert Press.

Mohanty, S., & Vyas, S. (2018). How to compete in the age of artificial intelligence: Implementing a collaborative human-machine strategy for your business. New York, NY: Apress.

Norman, D. (2014). Things that make us smart: Defending human attributes in the age of the machine. New York, NY: Diversion Books.

Piaget, J. (1977). *The development of thought: Equilibration of cognitive structures* (A. Rosin, Trans.). New York, NY: Viking.

Robert, L. P., Jr., Bansal, G., Melville, N., & Stafford, T. (2020). Introduction to the special issue on AI fairness, trust, and ethics. *AIS Transactions on Human-Computer Interaction*, *12*(4), 172-178.

Rouse, W. B., & Spohrer, J. C. (2018). Automating versus augmenting intelligence. *Journal of Enterprise Transformation*, *8*(1-2), 1-21.

Shneiderman, B. (2020). Human-centered artificial intelligence: Three fresh ideas. *AIS Transactions on Human-Computer Interaction*, *12*(3), 109-124.

Siddike, M. A. K., Spohrer, J., Demirkan, H., & Kohda, Y. (2018). A framework of enhanced performance: People's interactions with cognitive assistants. *International Journal of Systems and Service-Oriented Engineering*, *8*(3), 1-17.

Smith, C. S. (2020). Dealing with bias in artificial intelligence: Three women with extensive experience in A.I. spoke on the topic and how to confront it. *The New York Times.* Retrieved from https://www.nytimes.com/2019/11/19/technology/artificial-intelligence-bias.html

Skagestad, P. (1993). Thinking with machines: Intelligence augmentation, evolutionary epistemology, and semiotic. *Journal of Social and Evolutionary Systems*, *16*(2), 157-180.

Spohrer, J., & Banavar, G. (2015). Cognition as a service: an industry perspective. *AI Magazine*, *36*(4), 71-86.

Spohrer, J. C., & Engelbart, D. C. (2004). Converging technologies for enhancing human performance: Science and business perspectives. *Annals of the New York Academy of Sciences*, *1013*(1), 50-82.

Sternberg, R., & Sternberg, K. (2008). *Cognitive psychology*. Belmont, CA: Wadsworth.

Sternberg, R. J. (1985). *Beyond IQ: A triarchic theory of human intelligence.* Cambridge, UK: Cambridge University Press.

Zhou, M. X., Mark, G., Li, J., & Yang, H. (2019). Trusting virtual agents: The effect of personality. *ACM Transactions on Interactive Intelligent Systems*, *9*(2-3), 1-36.

About the Authors

Lina Zhou is a Professor of Management Information Systems at the University of North Carolina at Charlotte. Her research focuses on improving human decision making and knowledge management through both the design and development of intelligent systems and understanding of human behavior. She has published in journals such as *MIS Quarterly*, *Journal of Management Information Systems*, various ACM and IEEE Transactions, *Information & Management*, and *Decision Support Systems*. Her research has been funded by the National Science Foundation.

Souren Paul is a Professor of Information Systems and the Chair of Business Informatics department at Northern Kentucky University. His research interests are in areas of virtual teams, collaboration systems, behavioral information security, and augmented intelligence. He has published research articles in *Journal of Management Information Systems*, *Decision Support Systems*, and *Information & Management*. He has served as Conference Co-Chair for the 2015 Americas Conference on Information Systems (AMCIS) and the 2020 International Conference on Information Systems (ICIS).

Haluk Demirkan is an Assistant Dean of Analytics Innovations and Milgard Endowed Professor of Service Innovation and Business Analytics at the Milgard School of Business, University of Washington-Tacoma. He has a Ph.D. in information systems and operations management from the University of Florida. His research interests include driving digital business innovation with analytics, smart service systems, cloud computing, and cognition for value co creation and outcomes. He has five books and more than 170 publications.

Lingyao (Ivy) Yuan is an assistant professor of information systems in the Debbie and Jerry Ivy College of Business at Iowa State University. Her research interests include topics on the impact of non-cognition behavior and decision-making, digital human, and artifical intelligence. She has been published in information systems journals, including *Journal of Management Information Systems*, *Journal of the Association for Information Systems*, and *Decision Sciences*. She also serves on the editorial board of the *International Journal of Electronic Commerce*.

Jim Spohrer directs IBM's open source Artificial Intelligence developer ecosystem effort. He led IBM Global University Programs, co-founded Almaden Service Research, and was CTO Venture Capital Group. After his MIT BS in Physics, he developed speech recognition systems at Verbex (Exxon) before receiving his Yale PhD in Computer Science/AI. In the 1990's, he attained Apple Computers' Distinguished Engineer Scientist and Technologist role for next generation learning platforms. With over ninety publications and nine patents, he received the Gummesson Service Research award, Vargo and Lusch Service-Dominant Logic award, Daniel Berg Service Systems award, and a PICMET Fellow for advancing service science.

Michelle Zhou is a Co-founder and CEO of Juji, Inc., a California-based company that powers Cognitive Artificial Intelligence (AI) Assistants in the form of chatbots. She is an expert in the field of Human-Centered AI, an interdisciplinary area that intersects AI and Human-Computer Interaction (HCI). Zhou has authored more than 100 scientific publications and 45 patent applications on subjects including conversational AI, personality analytics, and interactive visual analytics of big data. Earlier in her career, she spent 15 years at IBM Research and the Watson Group, where she managed the research and development of Human-Centered AI technologies and solutions, including IBM Watson Personality Insights. She serves as Editor-in-Chief of *ACM Transactions on Interactive Intelligent Systems* (*TiiS*) and an Associate Editor of *ACM Transactions on Intelligent Systems and Technology* (*TIST*). She was formerly the Steering Committee Chair for the ACM International Conference Series on Intelligent User Interfaces. She is an ACM Distinguished Member and received a PhD in Computer Science from Columbia University.

Julie Basu is a technology industry veteran in Silicon Valley with over 20 years of experience in leading-edge enterprise software. She was a Director of Engineering at Oracle and currently heads her own startup smartQED. Her interests include AI and big data, and any technology that improves everyday life. Julie holds an MS and a PhD in Computer Science from Stanford University.