JOURNAL OF EMERGING TECHNOLOGIES IN ACCOUNTING
Vol. 15, No. 2
Fall 2018
pp. 15–25

American Accounting Association
DOI: 10.2308/jeta-52311

How Much Automation Is Too Much? Keeping the Human Relevant in Knowledge Work

Steve G. Sutton
Vicky Arnold
NHH Norwegian School of Economics and University of Central Florida

Matthew Holt
University of Dayton

ABSTRACT: With the rapid advances in data analytics, machine learning, and continuous monitoring along with other related advances in artificial intelligence-based technologies, our solution as researchers to many of today's business problems increasingly becomes one of, "Can I fix the problem through automation?" However, as we find that artificial intelligence increasingly provides us with the power to replace knowledge workers with automated systems, rarely is the question asked, "Should we automate knowledge work?" There are a host of questions that should be addressed including (1) whether automation is the most effective solution, (2) if there are ethical dilemmas associated with replacing the human element, and (3) if there are societal implications of displacing large numbers of knowledge workers. The focus of this discussion is on understanding the impact of knowledge-based systems on human users' knowledge acquisition and retention and outlining an alternative research strategy that centers more on transferring knowledge to the user during the work production process in order to maintain human expertise and relevance in professional decision making. Contemporary research still argues that human-computer collaboration may outperform either on their own; but, to limit the deskilling effect of knowledge-based systems and alternatively promote skill development, we call upon academic researchers to seek better ways to keep the human relevant in a broad range of knowledge work fields. Further, we suggest that expanding the philosophical discussions of the ethics of artificial intelligence-based technologies and the corollary impact on the rapid decline of the professions is necessary.

Keywords: technology dominance; automation bias; deskilling; epistemology; artificial intelligence; accounting profession.

BACKGROUND

Arnold and Sutton (1998) began work on the theory of technology dominance over concerns about the unintended consequences of the use of knowledge-based systems—decision-aiding technologies with embedded artificial intelligence (AI) components. Today, not only are these unintended consequences becoming increasingly clear, but organizations' current strategies are likely to exacerbate the problem. The issues are not necessarily new, but the consequences are arguably more severe than expected. Forty years ago, Braverman (1974) wrote in detail about the conflict between labor and capital. He noted the scientific approaches that were used to deconstruct manual work into small pieces that could easily be replicated with minimal training and minimal skill. His concern was related to how skilled craft was being broken into small pieces that required little skill, taking away the pride of the craftsman and devaluing the worker because any small piece of the process was easily replaceable. In time, machines could be used to repeat the simple processes. Bryniolfsson and McAfee (2011), in *Race Against the Machine,* illustrate that knowledge work is proving easier to dissect and automate than much manual work. The widespread automation of knowledge work has many people questioning the role for the human who wants to pursue meaningful work in the future.

The authors thank participants at the 2017 International Conference on Enterprise Systems, Accounting and Logistics and, in particular, Constantinos Stefanou, Miklos Vasarhelyi, Bernhard Wieder, and two anonymous reviewers for their helpful comments.

Editor's note: Accepted by Miklos A. Vasarhelyi.

Submitted: September 2017
Accepted: September 2018
Published Online: November 2018

Some may question how relevant the race against the machine really is for the typical knowledge worker, but the challenge appears widespread. A recent *The New York Times* essay (Wakabayashi 2017) focused on how automation was transforming work, with a primary focus on workers who train robots to do their jobs. The essay included interviews with five diverse knowledge workers, whose job description was to teach an AI system to replace them in part or in whole. At roughly the same time, *The Economist* (2017) examined how more and more investment firms are experimenting with machine learning algorithms and how AI is beginning to displace highly paid financial analysts and fund managers who have not proved particularly effective over time. Many professionals are beginning to look a lot like pilots who sit back with little control or intervention as they watch the airplane fly (Martins and Soares 2012).

The New York Times interview with the travel agent was particularly telling. The agent recognized that the AI system was rapidly learning to do most of her job, and she viewed herself to be in a race against the system. However, her comfort came from the things that were unique in her knowledge that 30 years of experience brought; one example was how to schedule breakfast in Disney World at a time before the park opened so that you could get your children's picture at Cinderella Castle alone before the mobs arrived (Wakabayashi 2017). But, what about the new travel agents entering the workplace that do not have 30 years of work experience? How do they acquire knowledge to compete with AI? Do they have any chance of winning, or even keeping close, in the race with the machine? Extant research does not paint a very optimistic picture.

In a detailed case study of an accounting firm's audit practice support system, Dowling and Leech (2014) note how new auditors feel empowered by the system and feel like they know so much more than they would without the system. However, are they actually acquiring their own knowledge while using the system? On closer examination, the Dowling and Leech (2014) story suggests the auditors are guided not much differently than factory workers on an assembly line. Experienced auditors feel the system is very restrictive and limits their ability to focus where their knowledge tells them they should.[1] From that perspective, this story is much more consistent with a technology dominance effect (Arnold and Sutton 1998). Indeed, other research demonstrates how these restrictive systems appear to have hampered knowledge acquisition and integration by accounting professionals during the early years of their careers (Dowling, Leech, and Moroney 2008). Perhaps the role for researchers should be to focus on the design of systems that might promote knowledge acquisition and skill development and try to limit technology dominance and concurrent deskilling effects.

These issues, as well as the related ethical questions and implications for professions, are briefly explored in an effort to lay a broad framework for a new line of AI research in accounting. Our discussion is exploratory and does not yield answers; rather, we hope to foster a discourse that will motivate a new active research agenda on keeping the accounting knowledge worker relevant and that has a central focus of increasing the knowledge base of the users of AI-based systems.

TECHNOLOGY DOMINANCE AND AUTOMATION BIAS

Technology dominance relates to the concern over the dominating influence that technology may have over the user, which allows the user to take a more subservient position—in essence, the user deferring to the technology in the decision-making process. Another related concern is the potential deskilling effect that can occur through continued technology reliance. This deskilling effect can occur either through attrition of existing skills or lack of development of skills that would be a normal part of expertise development (Arnold and Sutton 1998). Technology dominance is a broad concern in many disciplines, but it is worth noting here that the medical research applying the theory of technology dominance has developed parallel terminology with a preference for "automation bias" as opposed to "technology dominance" (Goddard, Roudsari, and Wyatt 2011, 2014). Still, the effects are the same as concerns heighten over experts yielding to AI.

Most research on the theory of technology dominance has focused on the reliance component of the theory. Reliance is important; but, in the broader view of technology dominance, reliance is simply a necessary precondition for dominance to occur. The reliance model has been well tested and found to generally hold under a variety of conditions (Triki and Weisner 2014). The reliance model focuses on four aspects that drive reliance: user's experience level, problem complexity, user's familiarity with the system, and user's cognitive fit with the system's underlying decision process. If the user has low experience, the other factors are irrelevant because reliance is the only feasible way for the user to solve the problem. If the user has high experience, then the other three aspects of the reliance model are necessary for the experienced decision maker to warrant the effort to use the system (Arnold and Sutton 1998). Recent work has attempted to expand upon this model by suggesting several potential factors that may better refine the assessment of reliance (Triki and Weisner 2014). Further, Goddard et al. (2011, 2014) examine the cognitive fit factor and assess it via trust and confidence—two factors that technology dominance views as intermediary outcomes that result from cognitive fit between the user's and the system's decision processes.

[1] As an example, a primary issue highlighted by Dowling and Leech (2014) is the frustration that audit workpaper reviewers felt with their inability to focus on high-risk areas because the system would not let them proceed without electronically signing off on every minor workpaper first, and the system restricted the work processes they could follow.

While reliance is important for dominance to occur, the changes in the current professional decision-making environment suggest researchers' energies might be better spent on the second part of the model—the potential negative effects from technology dominance, the deskilling of users, and potential limitations on epistemological growth and innovation. Reliance is important, but increasingly users are given little choice by their organizations but to use the available AI systems and to justify any decision not to use AI (e.g., Dowling and Leech 2014). If, as the research evidence shows, users are not developing the skills and expertise that would normally come with experience (the deskilling effect of system use), then eventually all users basically become users of low experience (i.e., low expertise) and the other factors really are not relevant—the user has to rely on the system.

Multiple studies show that dominance can have a deleterious impact on decision making. Bias caused by the use of knowledge-based systems has been demonstrated in a variety of research settings: taxpayers exhibit lower confidence in tax compliance decisions leading to overpayments of taxes (Masselli, Ricketts, Arnold, and Sutton 2002), insolvency practitioners overreact to newly acquired evidence (Arnold, Collier, Leech, and Sutton 2004), and medical doctors ignore their own diagnoses (even when correct) because their diagnosis was not one presented by the AI system (Goddard et al. 2011, 2014). Deskilling effects have similarly been recorded in many environments including apparent decreases in confidence by doctors in their ability to diagnose patients (Goddard et al. 2011, 2014), inability of auditors with similar experience to perform audit tasks when their experience has been with highly restrictive (dominating) audit support systems (Dowling et al. 2008), less knowledge acquisition from using automated tax compliance systems (Noga and Arnold 2002), and knowledge management systems (McCall, Arnold, and Sutton 2008), and decreased levels of creativity by marketing managers whose work is supported through machine learning-based data analytic systems (Wortmann, Fischer, and Reinecke 2015).

The evidence appears fairly convincing across studies that technology dominance effects can result in poorer decision making as the user becomes dominated by the technology. The research also suggests that low experience users fail to learn from the systems, while the pattern among more experienced users is that they steadily lose confidence and deskilling effects are often present. To understand these phenomena, we need more research on this second half of the theory of technology dominance. While the patterns are predictable and supported by preliminary evidence, we have a rather limited understanding of how dominance takes hold, how technology affects expertise development, and why users fail to assimilate expertise while having similar experiences to experts before them. The how and why are critical pieces that are needed if we are to have a chance to counteract these effects and keep the human relevant in professional decision-making environments.

From a practical standpoint, the inhibition of expertise will almost certainly lead to a lack of innovation in the long run. Thus, automation at the expense of expertise seems a short-sighted solution. Perhaps research should take on a greater focus of pursuing technologies that promote the development of expertise in the user and systems that better leverage the collaboration of the user and the system where both contribute to the decision-making process. At a minimum, these concerns appear to warrant consideration when designing AI systems (Balasubramanian, Lee, Poon, Lim, and Yong 2017).

COUNTERACTING THE EFFECTS OF TECHNOLOGY DOMINANCE

Counteracting the effects of technology dominance is a multifaceted problem that researchers need to consider. Consider the travel agent in *The New York Times* essay; if the only way to win the race with AI is to have 30 years of prior experience, then where are the future experts going to develop?

Research in education has explored ways to instill experiential expertise in novices more rapidly. One of the more promising approaches has been the application of constructivist learning. In a traditional approach, students are taught the declarative knowledge (facts, rules, definitions) they need and are then provided the building blocks for using that declarative knowledge. In a constructivist learning environment, the focus is on immersing the student in real cases that provide professional experience faster (Hmelo 1998; Hmelo-Silver 2004; Jonassen, Howland, Marra, and Crismond 2008; Milne and McConnell 2001). Constructivist learning uses the reenactment of cases with visual delivery of the actual events taking place, and it involves the learner in synthesizing the information and making professional decisions. The environment is all simulated, but actual cases are delivered in a much shorter period to expedite the learning experience. The medical profession was the first to adopt constructivist learning in an effort to experiment with better preparing future doctors through these case experiences (Hmelo 1998; Schmidt, Rotgans, and Yew 2011).

The strategies have also received some attention in accounting. The INCASE project,[2] a design science project focused on developing a case delivery system for automated delivery of constructivist learning experiences, yielded software where

[2] INCASE is a distributed case delivery system that operates over the web to enable self-paced experiential learning. The testing of the systems was based on the use of 12 reenacted actual insolvency cases used in training with insolvency professionals having one to three years of experience. Preliminary results indicate the system is effective in significantly improving the participants' cognitive representation of key information in insolvency decision making (see Arnold, Collier, Leech, Sutton, and Vincent [2013] for more detail).

modules could be easily integrated into the system and the learner would be self-guided. The proof of concept was tested in several training sessions with junior insolvency professionals, and the assessments suggest the approaches were very effective at giving the junior-level accountants a better understanding of the overall engagement process and decision-making environment (Arnold et al. 2013). The results suggest junior professionals could develop a better and faster understanding of the decision-making environment, but that is short of expertise development. Still, the initial results hold some promise as a possible building block in the development of new experts. Simulations can play an important supplemental role, but realism is critical and the learner needs to feel the experience and have a sense of the engagement process and activities.

To meaningfully move forward, any effort to develop new experts in the emerging AI-dominated workplace will most likely require a better understanding of the psychological processes underlying expertise development. The expertise research in psychology is much broader than the expertise literature in auditing that arose in the 1980s and 1990s (e.g., Libby and Luft 1993; Bonner 2008). While auditing moved away from expertise as an area of interest, psychologists in various paradigms continued to explore how experts think and how expertise develops. These advances in psychology warrant further consideration in accounting and auditing as accounting researchers explore the relevancy of humans in a professional decision-making environment.

CAN EXPERTISE BE RAPIDLY DEVELOPED? DESIGNING SYSTEMS TO SKILL

Contemporary thoughts on expertise in complex decision environments associate it with deep structural domain knowledge that permits recognition of complex relational patterns below the superficial level of domain problems (Chi and VanLehn 2012; Goldwater and Schalk 2016). Emerging expertise literature in psychology focuses on various cognitive skills that facilitate the acquisition of such knowledge by novices. A primary example is analogical reasoning; a skill that allows an expert to identify appropriate similar events, despite superficial differences, in order to leverage past experience in recognizing patterns and selecting strategies (Gentner and Colhoun 2010; Holyoak and Richland 2014). Analogical reasoning is closely associated with expertise (Day and Goldstone 2012; Dumas, Alexander, and Grossnickle 2013). Both share the property of utilizing the deep structural knowledge of domains. Research on analogy in accounting is scant, and results have been mixed. However, a study by Magro and Nutter (2012) has provided some clarification regarding the mixed results, and it finds that tax experts utilize analogical reasoning to a greater extent than novices in making complex tax decisions. To date, no accounting research on methods to improve analogical reasoning by novices has been published. Finding effective methods to do so could speed up expertise development, and embedding such interventions in user systems would lead to more regular exposure, further increasing the speed of expertise development.

Another skill with potential to improve relational knowledge in novices is that of systems thinking.[3] Borrowing primarily from the systems dynamics paradigm, a small stream of research in accounting has begun to examine approaches to improve decision making by using systems thinking interventions (e.g., O'Donnell 2005; Brewster 2011). The preliminary results suggest this is an area AIS researchers should consider in terms of helping novices not only to understand domains better by thinking of them as systems, but to understand any system better by understanding its inner workings. Further work in knowledge representation paired with ideas from general systems thinking may allow researchers to help users map any entity as a system. Systems thinking interventions would then be all the more powerful.

The key is to dissect expertise into the core cognitive processes that are consistent among experts, integrate those processes when designing systems, and design systems in ways that reinforce the transfer of those cognitive processes to the user. Such a strategy should thereby offset the deskilling effects that occur when a greater proportion of domain tasks and subtasks are automated. The challenge is to determine how to build systems that meet firms' expectations for work production while at the same time developing users' cognitive processes.

Over 30 years ago, researchers began to explore designs for transferring knowledge from knowledge-based systems to the users of those systems (Eining 1991). These efforts largely revolved around the use of explanation systems borrowing predominantly on Anderson's (2000) ACT-R (and earlier versions) to tailor embedded explanation systems toward assisting novices' development of declarative and procedural knowledge (Smedley and Sutton 2004, 2007). This research finds only modest success in improving learner's knowledge acquisition, and in some cases shows no significant improvement. Gregor and Benbasat (1999) synthesized the work in this area and developed a conceptual model for how and when explanations should be delivered to the user. Subsequent work using their framework demonstrates strong differences in explanation use among novices and experts, suggesting that their prescribed delivery of explanations should be adjusted based on a moderating effect of expertise (Arnold, Clark, Collier, Leech, and Sutton 2006). The conceptual model and subsequent research combine to

[3] Systems thinking is a broad construct, examined within different paradigms. A simple bifurcation may be into general systems thinking and systems science. Systems thinking and systems science have been described as similar ways of thinking about the world, but with different applications (Churchman 1971; Cabrera, Colosi, and Lobdell 2008; Midgley 2003).

suggest that knowledge transfer from knowledge-based systems to users through explanations might be more successful if the type of knowledge explanation is tailored to the user and delivery is automatic but nonintrusive.

Other recent research explores the design of knowledge-based systems' interfaces, with a focus on arranging the inputs and menus on the screen in a less orthodox fashion that emphasizes groupings similar to how experts organize and relate information cues (J. Rose, McKay, Norman, and A. Rose 2012). Using fairly simple structures of information cues, J. Rose, A. Rose, and McKay (2007) were able to help novices subconsciously evolve their own knowledge structures closer to those of experts.

More recent work focuses on integrating the newer conceptualizations of explanation use with interface knowledge structures to examine effects on knowledge acquisition via system use during normal productive work. The preliminary results are encouraging; systematically providing explanations based on users' knowledge level, in combination with a knowledge structure based interface design, yields significant improvements in the development of expert-like knowledge structures (Arnold, Leech, Rose, and Sutton 2018).

The preceding work has made strides in helping users acquire knowledge and skill; and, the work also stands as a source of hope for future progress, even though the gains, thus far, have been incremental. One thing that researchers could consider is the information loss that exists in the knowledge representations presently incorporated into these systems. Better knowledge representations would allow for both greater flexibility in the system and availability of knowledge for acquisition by the user. One aspect of this is likely visual presentation of the relationship between knowledge components; but another part is making systems more collaborative so the user is involved in the decision processing and acts as a teammate with the system—where the user has a reason to reflect upon the representations. However, creating such representations will require significant innovation. More sophisticated representations that include additional information (e.g., the relations between elements within the domain) can be used not only in developing the engine of the system, but also in the interface, in order to make the knowledge more accessible to users. In addition to use in traditional knowledge-based systems, such interfaces could also be incorporated into more advanced applications that combine elements of expert systems and machine learning, such as those developed by Davis, Massey, and Lovell (1997) and Lombardi and Dull (2016). Further, exploring how cognitive skills important to expertise, such as analogical reasoning or systems thinking, can be enhanced through interventions or interfaces embedded within systems has the potential to make significant strides in stemming the deskilling effects of technology dominance, and may also be able to improve human-machine collaboration.

IMPROVING HUMAN-MACHINE COLLABORATION

In the long run, researchers may need to step away from traditional models of expertise development and reconsider what expertise is in an AI world. The human is not going to win the race against AI, so the solution is to alter the race from a competition into a collaboration. As more work becomes automated, as is predicted in the accounting profession (Frey and Osborne 2013; CEDA 2015), practitioners will require an augmented skillset to work effectively with AI-enabled systems. Additionally, determining the relative strengths of humans and machines and adopting a focus on designing systems that effectively take advantage of these strengths can improve human-machine collaboration as well as overall quality of the work. So, if the future is a collaboration between the human and AI, how does that affect the type of expertise needed in such a world?

At the most basic level, the answer is that as tasks for which professionals are responsible become increasingly outsourced to sophisticated algorithms, professionals will need to have at least some conceptual understanding of what the algorithms are doing. But to move beyond this basic concept to implementation is the challenge, as it is unclear what the best approach is to deliver this understanding. The deeper domain knowledge required of today's experts is not dispensable. Future professionals will still need to acquire such knowledge. What will be required is additional knowledge beyond that; development of effective, efficient delivery systems will be crucial.

Researchers can explore new methods of aiding professionals who do not have a significant background in data science or computer science in gaining a working understanding of the automating processes. A starting point may be to consider the significant overlap in human and machine cognition. Perhaps an explanation of how machines learn, couched within the similarities and differences to how humans learn, can make the processes more relatable. Returning to analogy, for example, similarity, which is at the core of analogical reasoning and some argue to all of human cognition (Hofstadter 2001; Gentner and Colhoun 2010), is also one of the central ideas across many paradigms in machine learning (Domingos 2015). A deeper understanding of how humans and machines learn and exploit knowledge through the fundamental idea of similarity could assist in generating solutions that aid both in professionals gaining an understanding of machine learning and in designing systems that take advantage of the relative strengths of professionals and machines working together to accomplish a task. As an example, the primary strength of machine learners is the use of various algorithms to find patterns (similarities) across datasets that are too large for humans to hold in working memory, let alone use for computations. However, humans are very good at analogizing across domains (Hofstadter 2001), whereas machines are not (Domingos 2015). With such high expectations on professionals, whether this knowledge acquisition can all be done in the course of productive work is questionable. Much of this will likely take place through some form of training (and innovating new methods of delivery for this training) through interactive systems, simulations, or intelligent tutors. But reinforcement by work production systems will still be necessary for success.

Potentially more difficult challenges lie beyond the answers regarding human-machine collaboration, however. One piece that raises questions on our traditional views of expertise is the Sparrow, Liu, and Wegner (2011) *Science* paper on the Google effect. The Google effect is based on evidence showing that human users of search engines (virtually everyone in a professional world) are losing their ability to store information in memory. In what appears to be a form of deskilling, users do not expend the energy to store information in memory that they perceive they can easily retrieve via a search engine at another time. More importantly, Sparrow et al. (2011) show that, in the process of developing a reliance on search engines, users are losing their mental capability to store information.

Sparrow et al. (2011) triggered a new stream of research examining the phenomenon that is known in psychology as transactive memory (Ward 2013). Transactive memory is offline memory—information the user has stored via other avenues than their own brain. Subsequent research findings suggest that while users are losing their ability to store information, they instead appear to develop an ability to store how to find the information: so the memory capability is reallocated from storing facts to storing search strategies. While the inability to store what would be generally thought of as declarative knowledge (facts, rules, definitions) raises concerns, the other side of the finding might be a stark reality check that our traditional views on expertise need to change in an AI environment. An open question is whether users can still proceduralize knowledge if it is stored offline. If knowledge can be proceduralized while using transactive memory, that ability would represent a promising evolutionary process that suggests the human is capable of evolving in interesting ways in an AI environment.

The bigger question may be whether procedural knowledge can be stored in some form in transactive memory. To date, the research has been focused on understanding the effects on declarative knowledge. However, it is not unreasonable to think that users could proceduralize knowledge while drawing on declarative knowledge that is stored in transactive memory, although that is an empirical question yet to be studied. What would really change the expertise equation is if somehow the user could learn to store procedural knowledge in transactive memory. AI is focused on solving problems that require procedural knowledge. Plausibly, the human decision maker might best play a role in the collaboration with AI if they became the selector of AI components rather than trying to actually compete with AI routines. In other words, could the human develop a distinct type of expertise with the necessary procedural knowledge to draw on the appropriate AI components as necessary to work through a decision process?

A good incubator for investigating this unknown may be through the focus on data analytics use in auditing. Vasarhelyi (2017) notes that his research group is working with the PCAOB to guide the use of data analytics in auditing, and suggests that this approach should focus on interactive data analysis. Vasarhelyi (2017) views the use of interactive data analysis in auditing as a sequential process: (1) the auditor uses data analytics to learn about the data; (2) once the analytics have been used to identify patterns, the auditor backs away from the data and develops appropriate filters; (3) the auditor reruns the analytics with a prioritization of the exceptions highlighted in Stage 1; and (4) the auditor weights the exceptions to focus on the "notable items." This process provides an interactive relationship between the auditor and AI in order to collaborate on identifying the real concerns in an audit. This provides an excellent environment to study how to make the professional more relevant in such a collaborative process, what skillset is needed to interact with AI technologies, and how to give professionals the relevant expertise needed to be a valued part of the collaboration. In this case, the AI components can reasonably become an extension of the auditor's expertise—for procedural knowledge to be stored in transactive memory.

REIGNITING THE PHILOSOPHICAL DISCOURSE

Multiple philosophical issues warrant greater discourse in light of the advances, proliferation, and power of contemporary AI. Within the philosophy community there is discourse being undertaken on the ethics underlying the use of AI, but much of the energy in this community has been absorbed by concerns over smart robots, and particularly military use of these robots. Another branch of philosophy is more focused on the general effects of AI-based technologies on the work conducted within the professions and the implications of decline of the professions in western society. These two branches of philosophy are briefly discussed in the following subsections, with a focus on the implications for accounting researchers.

Ethics Issues Surrounding AI

Although many accounting researchers may not see AI as an ethical issue, AI ethics has caught the attention of governmental bodies and researchers primarily concerned with robotics and smart machines. Cath, Wachter, Mittelstadt, Taddeo, and Floridi (2018) summarize three reports put out by the U.S. (Executive Office of the President 2016), U.K. (U.K. Government Office for Science 2016), and EU (European Parliament Committee on Legal Affairs 2016) related to "AI and the

Good Society." The European report is the narrowest and focuses primarily on ethical issues surrounding AI-based machines and robotics. The U.K. report is more focused on how the government can support AI development efforts to promote AI and robotics that make life better for individuals. They advocate forming a commission to oversee the legal and ethical issues surrounding AI and robotics. The U.S. report is the broadest based, yet Cath et al. (2018) describe the report as basically reflecting the optimism of the tech culture in Silicon Valley. The report advocates avoiding regulations that would stymie the industry and letting "a thousand flowers bloom." On the other hand, the report advocates a focus on developing AI technologies that augment human capabilities rather than replacing human capabilities. Further, it advocates monitoring future development to assure universal benefit and to assure that AI does not replace workers, without new opportunities, in a way that would increase income inequality. The breadth of issues covered in the U.S. report appear to have benefitted from the open discussion forums that allowed a broader constituency to partake in the discourse, representing in a small way an approach advocated in the early accounting AI ethics (Dillard and Yuthas 1997).

Overall, the reports are fairly reflective of the discourse that has taken place in recent times on AI and ethics. The current dialogue is largely dominated by a focus on robotics and other AI-based machines and how ethical values should be instilled in any learning machine (Yampolskiy 2012; Russell, Dewey, and Tegmark 2015). The other area receiving substantial attention based on current controversies is the area of information privacy and information validity (Helbing et al. 2017). These concerns have moved to the forefront of the discussion because of the manner in which information is collected from users online and the purported misuse of these data in recent elections.

While these two foci are undoubtedly critical to future society, researchers are increasingly calling for a near-term focus on how AI could reshape life. Russell et al. (2015) summarize the discourse from the previous conference on "The Future of AI: Opportunities and Challenges," where the consensus discussion focused on the need to no longer be neutral on the progress and development of AI, but rather the need for the research community to proactively focus on how AI can be used to the benefit of humanity as a whole. Russell et al. (2015) argue that the focus of the AI ethics discussion needed to become one that considered the effects of widespread automation of work, the displacement of workers, the resulting transformations in the labor market, and the potential effects on income inequality. The development of ethical robotics and machines has significant concerns for society, but these other effects could create substantial disruptions within society much sooner.

From an accounting research perspective, these near-term concerns relate to the basic question raised in this paper's discussion, "How much automation is too much?" Is it ethical to blindly pursue new methods that better enable us to automate knowledge work and remove the knowledge worker from the process? If our AI-based research in accounting information systems increasingly focuses on automated solutions to solve human failures in the decision-making process, and the by-product is a rapid decline in the number of human professionals required to deliver professional accounting services, is society better off?

As is fairly common in the AI research stream, these issues were the subject of discourse late in the 20th century,[4] but the discourse waned during a time when AI was being reconsidered for how best to implement—a time when standalone expert systems seemed less viable and embedded AI within broader systems became the focus. But the discourse from the 20th century is still relevant today, and re-igniting this discourse is important at a societal level.

One stream of this research focused on the extraction of individuals' expertise to develop expert systems and focused on the ethics around knowledge rights between employees and employers (Sutton, T. Arnold, and V. Arnold 1995). As this stream evolved, it took on more of a societal focus related to the widespread proliferation of AI-based systems. Various ethical reasoning approaches were considered, with the most applicable deemed as contractarian ethics, which argued from the "veil of ignorance." The "veil of ignorance" suggests that ethical choice is the choice that would be be made when one was unaware of whether they were the displaced knowledge worker or the benefactor of the system (T. Arnold, V. Arnold, and Sutton 1997). While this justice focus was viewed as most applicable in a business environment, teleological reasoning was also explored as a more accepted societal focus, applying a utilitarian type focus on the "greater good." The research here emphasized the disconnect between act-based teleological reasoning, which might sacrifice one person's expertise for the widespread availability of expertise, and rule-based teleology, which argues this evaluation must be made in light of all experts being susceptible to replacement with AI systems (Sutton, T. Arnold, and V. Arnold 1997–1998). It is this rule-based application that begins to get at the societal problems that are now emerging with widespread use of AI.[5]

Implications for the Accounting Profession

The other area of philosophical discussion currently taking place that should be of particular concern to accounting researchers is the effect of automation on the professions (R. Suskind, and D. Suskind 2016). The professions have held a

[4] See Dillard and Yuthas (2002) for a detailed review of this research.

[5] Parallel research at the time also focused on broadening the constituencies that were considered through stakeholder ethics (Dillard and Yuthas 2001) and responsibility ethics (Yuthas and Dillard 1996).

American Accounting Association

significant role in western society for centuries, with the specialized knowledge, certification, and licensing processes that are characteristics of a profession providing comfort and confidence to society's members through the services provided. The professions have been important to improving trust in medicine, public financial statement audits, quality of legal services, and competency in engineering-driven products, among others (Kultgen 1988). However, the professions are under attack on multiple fronts in western society. We see a waning of the respect that society has for professions in general, as well as an apparent decrease in professionals' satisfaction with simply holding respect in society unless that respect also comes with very high remuneration (Callahan 2007). Even the major accounting firms no longer speak of being a profession, but rather about being leaders in the accounting industry (Lampe, Garcia, and Tassin 2016). Still, perhaps the greatest attack on the professions is coming from technology. As professional work is increasingly automated, its mystique is dissolved, and the work is increasingly executed by paraprofessionals using smart technologies. As a result, the basic need for, and the continued existence of, many professions is fundamentally threatened (Suskind and Suskind 2016). We ask the question of whether the automation of professional work is for the greatest good. Do paraprofessionals using smart technologies provide the same comfort and trust to society? Or, are professions no longer valued nor needed in society?

Recent studies on the future of work suggest the automation of professional work will take a significant toll on accounting professionals. An Oxford study (Frey and Osborne 2013) predicts that as many as 94 percent of accountants could be displaced by technology within ten years. Accounting is not alone, of course, as a recent study by CEDA (2015) suggests that 40 percent of current overall Australian jobs could be displaced by technology within ten years. While the exact figures are debatable, the takeaway is that many accounting jobs are likely to be automated away in the not too distant future. Looking at these predictions, one should ask whether there will be an accounting profession in the future if we do not need the human decision makers. At a minimum, it seems that accounting researchers should study how the use of complex data analytics, machine learning, and the general automation of accounting and auditing tasks affect society's interests.

CONCLUDING THOUGHTS

The accounting information systems research community is at an important juncture in the transformation of professional accounting work. The prevailing forces suggest radical change in practice as efficiency and effectiveness demands necessitate the increased use of machine learning and other AI techniques to better (and presumably, more continuously) analyze data. To date, much of the research suggests that the human is the weak link in the decision model and to overcome the deficiencies in human training and ability, automation is the solution. Ethical issues aside, merely from a practical perspective this is a precarious assumption. As AIS researchers, we can either produce research that further reinforces those trends or we can step back and think about how the model might be changed—how we might help keep the human relevant in professional accounting decision making.

The theory of technology dominance is adopted as a lens for viewing and understanding how technology can deskill users and reinforce less desirable decision-making patterns by users. The theory provides a foundation for understanding the phenomena that are occurring with current technology designs. We argue that these aspects of the theory should be expanded upon to better understand how and why deleterious effects in human ability and decision processes occur in an effort to consider alternative system designs that could facilitate user knowledge acquisition and expertise development, and for system designs that leverage the strengths of both the human user and knowledge-based systems.

To achieve these objectives, researchers will need to rethink what expertise is in today's world. Expertise should probably not be assessed in terms of the human brain operating in isolation, but whether the human brain can efficiently import facts, information, and processing rules to effectively make decisions in complex decision domains. This may require a different type of expertise, or at least the strengthening of certain dimensions of expertise, that allow humans to more effectively use the AI embedded in emerging work systems. Essentially, human expertise needs to be fostered to allow the human to bring unique capability to decision environments, not to compete with AI, but to complement AI in a collaborative human-computer decision process. Humans do not appear to have much chance in winning the race against AI, but we may be able to better team with the machines in order to achieve even better decision-making outcomes.

REFERENCES

Anderson, J. R. 2000. *Learning and Memory*. New York, NY: John Wiley & Sons, Inc.

Arnold, T., V. Arnold, and S. G. Sutton. 1997. Toward a philosophical foundation for ethical development of audit expert systems: A contractarian approach. *Research on Accounting Ethics* 3: 211–232.

Arnold, V., and S. G. Sutton. 1998. The theory of technology dominance: Understanding the impact of intelligent decision aids on decision makers' judgments. *Advances in Accounting Behavioral Research* 1: 175–194.

Arnold, V., P. A. Collier, S. A. Leech, and S. G. Sutton. 2004. Impact of intelligent decision aids on expert and novice decision-makers' judgments. *Accounting & Finance* 44 (1): 1–26. https://doi.org/10.1111/j.1467-629x.2004.00099.x

Arnold, V., S. A. Leech, J. Rose, and S. G. Sutton. 2018. *Can Knowledge Based Systems be Designed to Counteract Deskilling Effects?* Working paper, The University of Melbourne.

Arnold, V., N. Clark, P. A. Collier, S. A. Leech, and S. G. Sutton. 2006. The differential use and effect of knowledge-based system explanations in novice and expert judgment decisions. *MIS Quarterly* 30 (1): 79–97. https://doi.org/10.2307/25148718

Arnold, V., P. A. Collier, S. A. Leech, S. G. Sutton, and A. Vincent. 2013. INCASE: Simulating experience to accelerate expertise development by knowledge workers. *Intelligent Systems in Accounting, Finance & Management* 20 (1): 1–21. https://doi.org/10.1002/isaf.1337

Balasubramanian, G., H. Lee, K. W. Poon, W. K. Lim, and W. K. Yong. 2017. Towards establishing design principles for balancing usability and maintaining cognitive abilities. In *Design, User Experience, and Usability: Theory, Methodology, and Management: DUXU 2017: Lecture Notes in Computer Science*, Volume 10288, edited by A. Marcus and W. Wang. Cham, Switzerland: Springer

Bonner, S. E. 2008. *Judgment and Decision Making in Accounting*. Upper Saddle River, NJ: Prentice Hall.

Braverman, H. 1974. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York, NY: New York University Press.

Brewster, B. E. 2011. How a systems perspective improves knowledge acquisition and performance in analytical procedures. *The Accounting Review* 86 (3): 915–943. https://doi.org/10.2308/accr.00000040

Brynjolfsson, E., and A. McAfee. 2011. *Race Against the Machine*. Lexington, MA: Digital Frontier Press.

Cabrera, D., L. Colosi, and C. Lobdell. 2008. Systems thinking. *Evaluation and Program Planning* 31 (3): 299–310. https://doi.org/10.1016/j.evalprogplan.2007.12.001

Callahan, D. 2007. *The Cheating Culture: Why More Americans Are Doing Wrong to Get Ahead*. New York, NY: Houghton Mifflin Harcourt.

Cath, C., S. Wachter, B. Mittelstadt, M. Taddeo, and L. Floridi. 2018. Artificial intelligence and the "good society": The U.S., EU, and U.K. approach. *Science and Engineering Ethics* 24 (2): 505–528. https://doi.org/10.1007/s11948-017-9901-7

Chi, M., and K. VanLehn. 2012. Seeing deep structure from the interactions of surface features. *Educational Psychologist* 47 (3): 177–188. https://doi.org/10.1080/00461520.2012.695709

Churchman, C. W. 1971. *The Design of Inquiring Systems: Basic Concepts of Systems and Organizations*. New York, NY: Basic Books.

Committee for Economic Development in Australia (CEDA). 2015. *Australia's Future Workforce?* Melbourne, Australia: Committee for Economic Development in Australia.

Davis, J., A. Massey, and R. Lovell II. 1997. Supporting a complex audit judgment task: An expert network approach. *European Journal of Operational Research* 103 (2): 350–372. https://doi.org/10.1016/S0377-2217(97)00125-2

Day, S., and R. Goldstone. 2012. The import of knowledge export: Connecting findings and theories of transfer of learning. *Educational Psychologist* 47 (3): 153–176. https://doi.org/10.1080/00461520.2012.696438

Dillard, J., and K. Yuthas. 1997. Fluid structures: A structuration approach to evaluating information technology. *Advances in Accounting Information Systems* 5: 247–271.

Dillard, J., and K. Yuthas. 2001. A responsibility ethic for audit systems. *Journal of Business Ethics* 30 (4): 337–359. https://doi.org/10.1023/A:1010720630914

Dillard, J., and K. Yuthas. 2002. Ethics research in AIS. In *Researching Accounting as an Information Systems Discipline*, edited by V. Arnold and S. G. Sutton. Sarasota, FL: American Accounting Association.

Domingos, P. 2015. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York, NY: Basic Books.

Dowling, C., and S. Leech. 2014. A Big 4 firm's use of information technology to control the audit process: How an audit support system is changing auditor behavior. *Contemporary Accounting Research* 31 (1): 230–252. https://doi.org/10.1111/1911-3846.12010

Dowling, C., S. Leech, and R. Moroney. 2008. Audit support system design and the declarative knowledge of long-term users. *Journal of Emerging Technologies in Accounting* 5 (1): 99–108. https://doi.org/10.2308/jeta.2008.5.1.99

Dumas, D., P. Alexander, and E. Grossnickle. 2013. Relational reasoning and its manifestations in the educational context: A systematic review of the literature. *Educational Psychology Review* 25 (3): 391–427. https://doi.org/10.1007/s10648-013-9224-4

Economist, The. 2017. Machine-learning in finance: Unshackled algorithms. (May 27): 68.

Eining, M. 1991. The impact of expert systems on experiential learning in an auditing setting. *Journal of Information Systems* (Spring): 1–16.

European Parliament Committee on Legal Affairs. 2016. *Civil Law Rules on Robotics* (2015/2103 (INL)). Brussels, Belgium: European Parliament.

Executive Office of the President. 2016. *Artificial Intelligence, Automation and the Economy*. Washington, DC: GPO.

Frey, C., and M. Osborne. 2013. *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Available at: https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

Gentner, D., and J. Colhoun. 2010. Analogical processes in human thinking and learning. In *Towards a Theory of Thinking*, edited by B. Glatzeder, V. Goel, and A. Muller. New York, NY: Springer-Verlag.

Goddard, K., A. Roudsari, and J. Wyatt. 2011. Automation bias—A hidden issue for clinical intelligence/decision support system use. *Studies in Health Technology and Informatics* 164: 17–22.

Goddard, K., A. Roudsari, and J. Wyatt. 2014. Automation bias: Empirical results assessing influencing factors. *International Journal of Medical Informatics* 83 (5): 368–375. https://doi.org/10.1016/j.ijmedinf.2014.01.001

Goldwater, M. B., and L. Schalk. 2016. Relational categories as a bridge between cognitive and educational research. *Psychological Bulletin* 142 (7): 729–757. https://doi.org/10.1037/bul0000043

Gregor, S., and I. Benbasat. 1999. Explanations from intelligent systems: Theoretical foundations and implications for practice. *Management Information Systems Quarterly* 23 (4): 497–530. https://doi.org/10.2307/249487

Helbing, D., B. S. Frey, G. Gigerenzer, E. Hafen, M. Hagner, Y. Hofstetter, J. van den Hoven, R. V. Zicari, and A. Zwitter. 2017. *Will Democracy Survive Big Data and Artificial Intelligence?* Available at: https://www.scientificamerican.com/article/will-democracy-survive-big-data-and-artificial-intelligence/

Hmelo, C. 1998. Problem-based learning: Effects on the early acquisition of cognitive skill in medicine. *Journal of the Learning Sciences* 7 (2): 173–208. https://doi.org/10.1207/s15327809jls0702_2

Hmelo-Silver, C. 2004. Problem-based learning: What and how do students learn? *Educational Psychology Review* 16 (3): 235–266. https://doi.org/10.1023/B:EDPR.0000034022.16470.f3

Hofstadter, D. R. 2001. Analogy as the core of cognition. In *The Analogical Mind: Perspectives from Cognitive Science*, edited by D. Gentner, K. J. Holyoak, and N. Boicho, 499–538. Cambridge, MA: The MIT Press.

Holyoak, K., and L. Richland. 2014. Using analogies as a basis for teaching cognitive readiness. In *Teaching and Measuring Cognitive Readiness*, edited by H. F. O'Neil, R. S. Perez, and E. L. Baker. New York, NY: Springer-Verlag.

Jonassen, D., J. Howland, J. Marra, and D. Crismond. 2008. *Meaningful Learning with Technology*. Upper Saddle River, NJ: Prentice Hall.

Kultgen, J. 1988. *Ethics and Professionalism*. Philadelphia, PA: University of Pennsylvania Press.

Lampe, J. C., A. Garcia, and K. L. Tassin. 2016. A post-SOX history of U.S. public accountancy. The history of deprofessionalization in U.S. public accountancy: Part III. *Research on Professional Responsibility and Ethics in Accounting* 20: 3–29. https://doi.org/10.1108/S1574-076520160000020001

Libby, R., and J. Luft. 1993. Determinants of judgment performance in accounting settings: Ability, knowledge, motivation, and environment. *Accounting, Organizations and Society* 18 (5): 425–450. https://doi.org/10.1016/0361-3682(93)90040-D

Lombardi, D. R., and R. B. Dull. 2016. The development of AudEx: An audit data assessment system. *Journal of Emerging Technologies in Accounting* 13 (1): 37–52. https://doi.org/10.2308/jeta-51445

Magro, A. M., and S. E. Nutter. 2012. Evaluating the strength of evidence: How experience affects the use of analogical reasoning and configural information processing in tax. *The Accounting Review* 87 (1): 291–312. https://doi.org/10.2308/accr-10161

Martins, E., and M. Soares. 2012. Automation under suspicion—Case: Flight AF-447 Air France. *Work (Reading, MA)* 41: 222–224. https://doi.org/10.3233/WOR-2012-0160-222

Masselli, J., R. Ricketts, V. Arnold, and S. G. Sutton. 2002. The impact of embedded intelligent agents on tax compliance decisions. *The Journal of the American Taxation Association* 24 (2): 60–78. https://doi.org/10.2308/jata.2002.24.2.60

McCall, H., V. Arnold, and S. G. Sutton. 2008. Use of knowledge management systems and the impact on the acquisition of explicit knowledge. *Journal of Information Systems* 22 (2): 77–101. https://doi.org/10.2308/jis.2008.22.2.77

Midgley, G. 2003. Science as systemic intervention: Some implications of systems thinking and complexity for the philosophy of science. *Systemic Practice and Action Research* 16 (2): 77–97. https://doi.org/10.1023/A:1022833409353

Milne, M., and P. McConnell. 2001. Problem-based learning: A pedagogy for using case material in accounting education. *Accounting Education* 10 (1): 61–82. https://doi.org/10.1080/09639280122712

Noga, T., and V. Arnold. 2002. Do tax decision support systems affect the accuracy of tax compliance decisions? *International Journal of Accounting Information Systems* 3 (3): 125–144. https://doi.org/10.1016/S1467-0895(02)00034-9

O'Donnell, E. 2005. Enterprise risk management: A systems-thinking framework for the event identification phase. *International Journal of Accounting Information Systems* 6 (3): 177–195. https://doi.org/10.1016/j.accinf.2005.05.002

Rose, J., A. Rose, and B. McKay. 2007. Measurement of knowledge structures acquired through instruction, experience, and decision aid use. *International Journal of Accounting Information Systems* 8 (2): 117–137. https://doi.org/10.1016/j.accinf.2007.04.002

Rose, J., B. McKay, C. Norman, and A. Rose. 2012. Designing decision aids to promote the development of expertise. *Journal of Information Systems* 26 (1): 7–34. https://doi.org/10.2308/isys-10188

Russell, S., D. Dewey, and M. Tegmark. 2015. Research priorities for robust and beneficial artificial intelligence. *AI Magazine* 36 (4): 105–114. https://doi.org/10.1609/aimag.v36i4.2577

Schmidt, H., J. Rotgans, and E. Yew. 2011. The process of problem-based learning: What works and why. *Medical Education* 45 (8): 792–806. https://doi.org/10.1111/j.1365-2923.2011.04035.x

Smedley, G., and S. G. Sutton. 2004. Explanation provision in knowledge-based systems: A theory driven approach for knowledge transfer designs. *Journal of Emerging Technologies in Accounting* 1 (1): 41–61. https://doi.org/10.2308/jeta.2004.1.1.41

Smedley, G., and S. G. Sutton. 2007. The effect of alternative procedural explanation types on procedural knowledge acquisition during knowledge-based systems use. *Journal of Information Systems* 21 (1): 27–51. https://doi.org/10.2308/jis.2007.21.1.27

Sparrow, B., J. Liu, and D. Wegner. 2011. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science* 333 (6043): 776–778. https://doi.org/10.1126/science.1207745

Suskind, R., and D. Suskind. 2016. *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford, U.K.: Oxford University Press.

Sutton, S. G., T. D. Arnold, and V. Arnold. 1995. Toward an understanding of the philosophical foundations for ethical development of audit expert systems. *Research on Accounting Ethics* 1: 61–74.

Sutton, S. G., T. D. Arnold, and V. Arnold. 1997–1998. Teleological foundations for the ethical implications of expert systems development: Act versus rule based reasoning. *Accounting Forum* 21 (3/4): 463–474.

Triki, A., and M. Weisner. 2014. Lessons learned from the literature on the theory of technology dominance: Possibilities for an extended research framework. *Journal of Emerging Technologies in Accounting* 11 (1): 41–69. https://doi.org/10.2308/jeta-51078

U.K. Government Office for Science. 2016. *Artificial Intelligence: An Overview for Policy-Makers*. London, U.K.: Government Office for Science.

Vasarhelyi, M. A. 2017. *Keynote Address: Audit Data Analytics*. International Symposium on Accounting Information Systems, Valencia, Spain.

Wakabayashi, D. 2017. Meet the people who train robots (to do their jobs). *The New York Times* (April 28): 4.

Ward, A. F. 2013. Supernormal: How the Internet is changing our memories and our minds. *Psychological Inquiry* 24 (4): 341–348. https://doi.org/10.1080/1047840X.2013.850148

Wortmann, C., P. M. Fischer, and S. Reinecke. 2015. *"Too Much of a Good Thing?" How Big Data Changes Managerial Decision Making in Marketing*. Society for Judgment and Decision Making Annual Meeting, Chicago, IL.

Yampolskiy, R. V. 2012. Artificial intelligence safety engineering: Why machine ethics is a wrong approach. In *Philosophy and Theory of Artificial Intelligence SAPERE 5*, edited by V. C. Miller, 389–396. Berlin, Germany: Springer.

Yuthas, K., and J. Dillard. 1996. An integrative model of accounting expert system design and implementation. *Advances in Accounting Information Systems* 4: 55–80.