

# META-HUMAN SYSTEMS = HUMANS + MACHINES THAT LEARN

Kalle Lyytinen  
Case Western Reserve University

Jeffrey V. Nickerson  
Stevens Institute of Technology

John L. King  
University of Michigan

(Accepted version for the *Journal of Information Technology*<sup>1</sup>)

## Abstract

Metahuman systems are new, emergent, sociotechnical systems where machines that learn join human learning and create original systemic capabilities. Metahuman systems will change many facets of the way we think about organizations and work. They will push information systems research in new directions that may involve a revision of the field's research goals, methods and theorizing. Information systems researchers can look beyond the capabilities and constraints of human learning toward hybrid human/machine learning systems that exhibit major differences in scale, scope and speed. We review how these changes influence organization design and goals. We identify four organizational level generic functions critical to organize metahuman systems properly: delegating, monitoring, cultivating, and reflecting. We show how each function raises new research questions for the field. We conclude by noting that improved understanding of metahuman systems will primarily come from learning-by-doing as IS scholars try out new forms of hybrid learning in multiple settings to generate novel, generalizable, impactful designs. Such trials will result in improved understanding of metahuman systems. This need for large scale experimentation will push many scholars out from their comfort zone, because it calls for the revitalization of action research programs that informed the first wave of socio-technical research at the dawn of automating work systems.

**Keywords:** Machine Learning, Learning theory, Technology, Work groups, Job Design, Organizational Forms, Monitoring, Embodiment, Autonomy

---

<sup>1</sup> Lyytinen, K., Nickerson, J. V., King, J. (2020) Meta-human Systems = Humans + Machines That Learn, *Journal of Information Technology*, <https://doi.org/10.1177/0268396220915917>

# META-HUMAN SYSTEMS = HUMANS + MACHINES THAT LEARN

## INTRODUCTION

Machines that learn interact adaptively with their environment while increasing their capabilities through experience. Such machines are now being embedded in apps, news feeds, video streaming services, and email filters. Machines that learn do so in a manner similar to human trial-and-error learning. But because of digitally enabled transfer capability, they differ in learning speed, scope and scale. We use the term *machines that learn* here however as a distinct category from the common noun phrase in computing fields called *machine learning*. Machine learning is a specialized kind of computational process (Domingos, 2012) whereas machines that learn abstract from the details of the computation and focus on the process features, outcomes and emergent features of systems like their scale, scope, speed and architecture. This paper addresses machines that learn as parts of wider systems where *both humans and machines learn jointly* (Davenport and Kirby, 2016). These are called *meta-human* systems. The capabilities of machines to learn can change in such systems at a different rate than in how humans learn. Our focus is on the socio-technical consequences of introducing machines that learn to work settings and systems as organizations start deploying meta-human systems.

We use the Greek *meta* as a prefix for such systems because it fits with the nature of these systems etymologically, colloquially, and prescriptively. Etymologically *meta* means ‘after’ and ‘beyond,’ as in *post-human*. Metahuman systems invites us to explore things that go beyond what machines or humans *alone* possess. Such superhuman aspects of meta-human systems show now often up in popular books (Bostrom, 2017). They can also be oppressive (Zuboff, 2015, 2019). There is also a colloquial meaning of *meta* in denoting “about” and “above” as in

meta-studies that analyze and synthesize findings of other studies. In our case meta used colloquially relates to expanding and synthesizing ways in which humans and machines are connected in work systems and that study of metahuman systems encompass all systems using machines that learn in the context of human work systems. Our prescriptive use of *meta* is focused on “beyond” in the sense of models of models and related meta-models that humans use to reason about assessments made by machines come to fore in the examination of metahuman systems. Before humans delegate to a machine a task as they would to a colleague, for example when doctors who use AI to diagnose cancer and delegate that task to a machine they want to understand how the AI machine has been trained to calibrate their delegation. In this regard they either understand the process or engage in constant monitoring of it (Cai et al., 2019). In such settings higher meta-levels affect always design of systems at work. For example, a fleet of autonomous vehicles operates in an ecology of designers, passengers, controllers and regulators who need to work together and react to some situations that the (meta)system creates. They must always consider the next generation and understand also the design for design. Meta-design, if done well, allows to design better systems, and possibly paves way to systems that design themselves.

When machines learn, their capabilities change in unexpected ways and at a speed and scale different than we have encountered before. Currently, Information Systems (IS) research does not reason pragmatically and diligently about this. Some IS research presupposes that machines do not change or change very slowly, as seen in on use (Venkatesh et al., 2003) and affordance (Leonardi, 2011) related theorizing. Some advanced sociotechnical research presupposes symmetry between machine and human (Latour, 2005). The truth lies likely in between:

machines capability changes, sometimes rapidly, even if not rivaling the capabilities of humans because machine and human learning processes are embodied differently (Lawrence, 2017).

Organizational bureaucracies where metahuman systems now operate serve important purposes (du Gay, 2005). This is particularly true with respect to delegation and training of people that comes with the design of metahuman systems (Barnard, 1938; Simon, 1997). Meta-human systems are changing and will change these dramatically in future. For example, knowledge production in Wikipedia is already now shared between humans and bots which both approve and manage processes (Zheng et al., 2019). The proposed IS research agenda involves multiple levels of meta: first understanding differences in cognitive architecture and how this affects the design of organizational meta-processes for delegation and monitoring (operational processes), cultivation (capability-building processes), and reflection (systems improvement processes).

It is still early days of metahuman systems, but these systems already now receive growing attention under different monikers in multiple discourses. Simon (1997, p. 234ff) anticipated meta-human systems that build with their own initiative knowledge beyond the developments of the organization's internal research departments. Now such systems are used in automated drug discovery (Schneider, 2018) and other fields. The concept is creeping in management discourse when Uber drivers are managed by algorithms that delegate driver tasks irrespective of human resistance (Lee et al., 2015; Möhlmann and Henfridsson, 2019). The notion appears when doctors now use adaptive robotic surgery systems to delegate surgery tasks to machines, but find at the same time that they must protect patients from swinging robotic arms (Sergeeva et al., 2018). Meta-human systems are showing up in multiple work settings: industrial organization, market making, election results, transportation and urban mobility, education, and scientific

research, just to name a few. Machines that learn have been overhyped similar to the hype of 1930s Hugo Gernsback futuristic fanzines. They can produce negative, unanticipated effects (Marcus and Davis, 2019; Zuboff, 2015, 2019). Innocuous rectangles on traffic signs can fool autonomous vehicles leveraging the way machines are trained to cause accidents (Heaven, 2019). Given a certain brittleness of machines that learn, meta-human systems are susceptible to new types of attacks. IS research agenda has from early days blended social *and* technical (Hirschheim, 1985; Kling, 1980). IS researchers should now consider meta-human systems as unique kinds of socio-technical systems in their inquiries to remain relevant (Sarker et al., 2019).

An ongoing experiment in metahuman type learning provides an illustrative and powerful example of the immediate salience of the issues the field faces. The Never Ending Language Learner (NELL) learns substantially on its own and through humans about language and its uses. It reads web pages on its own to build a repository of contemporary beliefs, and then uses humans to explore the exact meaning of these beliefs (Pedro and Hruschka, 2012). By 2015 it had amassed over 80 million beliefs, with high confidence in over 2 million (Mitchell et al., 2018). When NELL fails to recognize a belief, it tweets human followers, parses the replies, and integrates the results to improve and expand its understanding (Carlson et al., 2010; Pedro and Hruschka, 2012). In learning terms NELL is a mixed-initiative, autonomous learning system (Horvitz, 1999; Parasuraman et al., 2000)(Horvitz, 1999; Parasuraman et al., 2000), and illustrates a growing trend in which machines use humans to further learning (Ekbja and Nardi, 2014). Any technology like NELL will go as part into meta-human systems. Such systems require novel research to address challenges like those identified by Zuboff (2015, 2019) regarding hidden or latent monetizing of consumer's 'behavioral surplus, by Aleksander (2017) that robotic systems far from human intelligence demonstrate character, capabilities and

motivations different from humans while lacking transparency required for organizational use. Meta-human systems will blur the boundaries between the socio (the human) and the technical (the machine) in unanticipated ways which calls for novel inquiry.

The IS field grew out from the need for experts who can skillfully work at and bridge the boundary of the social and the technical (Su et al., 2017). Meta-human systems create a new kind of need for experts who can work at a rapidly shifting learning boundary in organizations. Advances in machine learning technologies will change learning assumptions on which the foundations of modern management were built. Herbert Simon hinted at this several decades ago when he noted the difference in speed between human learning and learning by machines (Simon, 1983, p. 26-27). Even if it takes a while to get machines to learn specific tasks or content, once learned, the machine can disseminate the learned capability quickly (Kallinikos et al., 2013; Lyytinen et al., 2016). Meta-human systems will bring higher speed of machines that learn and alter the scale and scope in which meta-human systems learn. A myriad of consequences is likely to follow - many of them not intended. This calls for a revision of the IS research agenda how we think of sociotechnical systems where learning forms one of the main dimensions of their behavior.

## **THE CONCEPT OF META-HUMAN SYSTEM**

Meta-human systems are a hybrid of humans and machines that learn, complementing and amplifying capabilities that potentially make such systems better at learning than either humans or machines separately. What makes the story challenging and interesting is that machines and humans have different cognitive architectures and so far organizational design has operated under the assumption that all learning components in organizations have similar (though

somewhat varying in level and skill) of cognitive architectures and ergo learning speed, scope and embodiment. The emergence of different cognitive architectures explains differences in learning speed which affects the scale and scope of learning at the system level. This calls for a joint common definition of learning that covers learning as an outcome irrespective of cognitive architecture. In this essay we define learning as *a process of increasing capabilities in a configuration of agents*: human, machine, or mixed.

Past IS research assumes that humans learn and machines do not (Venkatesh et al., 2003). At the same time humans have used machines to improve the overall performance of a system given machine characteristics of accuracy, speed or nearly fault-free behavior. In consequence, research on learning in organizational contexts has focused on human cognition and related learning (Leonardi, 2012). That machines *can* learn is sort of implied by actor network theory (Latour 2005), critical realism (Mingers et al., 2013), and DeLanda's (2019) theory of assemblages. However, they do not recognize the emerging learning capability of hybrids where both machines and humans learn. As the discussion of cognitive architectures below will show, it does not ultimately matter whether humans and machines learn through similar or dissimilar processes. Nevertheless, much of the research, including ours, in the context of meta-human systems proceeds on assuming the presence of mechanisms for learning in both humans and machines (Lake et al., 2015). The dynamics that surface when these types of learning agents with different architectures and skills interact are still unknown. Using a machine in a cognitive loop will for example causes humans also to anthropomorphize the machine expert (Cai et al., 2019), or only selectively follow machine advice when it reinforces human's own biases (Green and Chen, 2019).

A thorough review of learning literature is beyond the scope of this essay. Learning spans many disciplines with many denotations and connotations. We constrain our discussion to learning relevant to meta-human systems. We differentiate between trial-and-error learning, when skills might be learned for the first time by humans and machines, and diffusion-based learning when skill are transferred (Argall et al., 2009). Given this distinction learning can be first simplified to the observation that any agent's learning involves some sort of trial and error process (Thorndike, 1932) with intermittent rewards and punishments regarding goals (Sutton and Barto, 2018). An agent acts in an environment, and based on environmental reaction, decides what to do next. As it moves through this process it acquires new capabilities that make it better fit operating in that environment. The increased capabilities that accrue from learning can be measured by testing what an agent can do after a learning event in that setting. Testing might involve tasks to infer cognitively what agents can do when they are actively functioning as a learning system. Machines that learn receive reward or punishment based on the environment's response to the machine's action (Sutton and Barto, 2018). Even with their current limitations, this simple learning process allows machines to move into domains hitherto reserved to human learning. There are trade-offs: reinforcement learning in machines requires massive amounts of data, while humans learn through similar process using a small number of instances (Lake et al., 2017).

Trial-and-error learning encompasses also meta- level: learning to learn. Trial and error single loop learning builds specific capabilities to do something as well as generic capabilities to assimilate such specific capabilities. Double loop learning builds an ability to achieve higher learning capability which covers also changing goals and assumptions of learning (Argyris, 2004). Trial and error learning is often costly and slow for both humans and machines so specific



solutions have been developed to circumvent it which enable *diffusion-based learning*. In such learning mode humans and animals learn through imitation based learning (i.e. through observation or social participation). By watching others do something successfully or engaging in such activity with others, allows one to infer and repeat expected patterns of behavior (Frith and Frith, 2012). Multiple other mechanisms to diffuse knowledge might work in humans (and in many animals), with the fastest conjectured mechanism a kind of mirroring at the neurological level (Rizzolatti and Fogassi, 2014). Machines cannot do any of the above yet to diffuse knowledge. But this advantage of humans is met by a machine advantage in that they possess very fast mechanism of direct copying of learned capability accomplished by transferring lines of code, data, or matrix weights. Analogous tasks can therefore be transferred in machines by a process akin to human transfer learning (Mestre, 2006), wherein certain layers of a matrix are copied and recombined (Tan et al., 2018).

Four aspects influence generally how meta-human systems will learn either through trial-and-error learning or diffusion based learning

1. *Trial and error learning is context-dependent and requires language and cultural knowledge.* Moral judgments that separate right from wrong fall here. Humans tend to be better at this than machines (Marcus and Davis, 2019; Smith, 2019). Sometimes humans can do it and machines cannot. Even when machines can do it, humans tend to do this more robustly. This might change over time, but for now humans have the advantage. Even basic language skills are elusive for machines: a prize has been offered for a machine agent that can sustain a 20 minute conversation with a human, but by 2018 the median conversation lasted less than 2 minutes (Ram et al., 2018).

2. *Trial and error learning takes time and requires significant attention.* Performance time in humans might be less than in machines, but machines can operate relentlessly, 24 hours a day in dedicated mode. Elapsed time can be shorter for machines. Recent deep learning variants have made machines to progress in learning within realms where simulations can be used to produce infinite amounts of data.
3. *Diffusion based learning from one agent to another is fast if the agents are machines.* Learning by knowledge transfer in machines speed up learning and response in meta-human systems. In contrast learning between humans and machines is slow due to significant differences in diffusion learning skills and mechanisms. This can only be improved by significant improvements in human-computer interfaces.
4. *In diffusion based learning the use of sensory learning that extends human senses with those of machines can empower human learning* (Sinz et al., 2019). Machine augmentation in vision alone includes an arsenal of microscopes, telescopes, night-vision equipment, and radar and such sensory capabilities have been a component of meta-human systems for some time. There are concerns about atrophication resulting from such advances for humans including the wide use of GPS tools reducing human unassisted navigation skills (Robbins, 2010) or how calculators have lowered arithmetic skills in humans.

A unique aspect of human learning is that it always happens in a cultural context (Engeström et al., 2016). Distributed cognition posits that any human knowledge, the consequence of learning, is embedded in the material organization of humans and machines in a cultural setting (Clark, 2017). Knowledge can reside in a human team, it can be embedded in their tools, or it can

be inscribed in the physical environment through maps, signs, switches, and other kinds of interfaces. No one yet knows how meta-human systems distribute (or should distribute) learning objectives and capabilities as part of larger organizational or institutional processes. IS research involving meta-human systems requires deepening understanding of learning in ways that goes beyond notions of anthropocentric learning. Scholars should not project onto machines the capabilities and processes of humans, and vice versa. But speed, scope of learning matters, and differences in degree in any of them can and will become differences in kind at the meta-human system level.

### **Cognitive Architecture**

Embodiment in human learning suggests that humans' cognitive processing related to learning depends on the architecture of human bodies (Lawrence 2017). Inbound and outbound information used in cognitive processing associated with learning in humans is communicated through sensory experience embodied in the human body (this includes also spoken or written language). The human brain may have 100 million times the processing power of a desktop computer, but the computer can exchange information with another computer 100 million times as fast as humans. Meta-human system designs can and should take advantage of these differences in cognitive architectures. A human takes perhaps 10,000 hours of learning to become an expert due to the biological foundation of human learning (Ericsson et al., 1993). Machines operate at different rates and the underlying technologies are improving. Moreover, as noted, machines *already* transfer information very fast. This was observed early on by Herbert Simon: "The first obvious fact about human learning is that it's horribly slow..." (Simon, 1983, p. 26-27). Simon was not sanguine that machines could match human learning, though it was

worth trying, because machines could replicate acquired capability quickly with other machines. Few doubt the organizational, social and economic consequences, if machines rival or surpass human learning and significant cognitive tasks observed in organizational settings. Even if it takes years to get there (such as conducting tax planning or trading equities with even 95% effectiveness), the machine's ultimate advantage is in the fast transfer and replication of such capability which would have revolutionary consequences of related organizational tasks (Kallinikos et al., 2013; Lyytinen et al., 2016). But there is no need for machines to rival or surpass human learning in all aspects of complex knowledge tasks. The tasks can be distributed in specific ways in meta-human systems and such meta-human systems are already faster than biological-only- human based systems learning organizations. Acceleration of the scope of transfer learning and diffusion based learning between machines will affect how information systems, cognitive psychology, computer science, biology, and other fields approach learning in social settings (Kallinikos et al., 2013; Lake et al., 2017; Lawrence, 2017; Rosen, 1991).

There are many sensational, science fiction-like claims about machines that learn without limits (Vinge, 1993). Most predictions are simple extrapolations of exponential improvements in raw 'machine power' like those applying Moore's law where a given amount of money buys roughly twice as much computing capability every two years (e.g., Kurzweil (2000)). It is not clear that machines can rival human learning or design themselves to learn and replicate because of just their ever greater computing power. Much depends on the definition of intelligence (Aleksander, 2017); it also depends on its origin. Human intelligence reflects human needs: surviving and adapting as biological and social organisms in evolving, complex environments. Machines possess no such needs. In consequence, machine intelligence might not serve all human needs, because teleology and meaning are part of human intelligence but not that of

machine (Ackoff and Emery, 2005). Not surprisingly machine survival in the current discourse is treated primarily in technical terms and posed as a problem of security or fault tolerance that are high priorities for human existence. Prior to machines that learn it was not necessary to care about this in a similar way. That may be changing.

## **THE EMERGENT FIELD OF META-HUMAN SYSTEMS**

Machines that can think and their effects have excited and informed IS research for more than half a century. We posit machines that learn constitute a new threshold event that will impact the future social organization and all disciplines that deal with organizing (Marcus and Davis, 2019; Smith, 2019). For example the NELL system shows that there is a new direction where meta-human systems are evolving which will apply and shape organizational practices involving machines in the years to come. Whether still in the lab or deployed, it is time to anticipate research challenges that come with meta-human systems. We will note some of them below by illustrating features of meta-human systems operating in the laboratories and reviewing uses of meta-human systems in select industries.

### **Meta-human Systems in the Lab**

Machines that can rapidly edit and copy knowledge (Yoo et al., 2010) make it possible for machine trainers to quickly pass new skills to machine trainees (Buciluă et al., 2006; Hinton et al., 2015; Rusu et al., 2016). Parallel processing can in addition speed this up by a factor of ten (Nair et al., 2015). Finally, such knowledge can be compressed. As a result data sets based on new learning models aid significantly both trial-and-error learning and diffusion learning:

training machines can correct trainee machines. Then machines can distill their knowledge into smaller and more compact neural networks. Because of this, trainee machines can quickly match or surpass trainer machines.

There are trade-offs with these advances. Machines will learn faster and better than humans in specific knowledge tasks that rely on pattern recognition with clear goals (e.g., recognizing tumor patterns in radiographs). But in settings with high costs of failure this machine advantage might not be reliable enough. For example, autonomous vehicles that malfunction can cause significant damage. While high speed driving games have been built and simulated to advance machine learning (Johnson-Roberson et al., 2016), these simulations don't train for a wide enough range of potential driving hazards (Santana and Hotz, 2016). New techniques are now emerging to address this weakness such as using empirical data extracted from cameras in vehicles to ground the simulations. More generally, new models of agent behavior and learning patterns will be necessary to build machines that learn for tasks where learning takes place in a context of many human and machine agents interacting, cooperating, and competing, as with driving (Panait and Luke, 2005; Lowe et al., 2017).

Human intervention is still expected to set up initial machine capabilities and embedding them into the physical infrastructure, and for their subsequent modification and maintenance. Most machines that learn are not able to learn on their own (Blue and Andoh-Baidoo, 2010; Sutherland, 2008). However, this is now changing. After a learning harness is set up, a neural network system Deep Q uses reinforcement learning algorithms to take action while following a learned policy (Sutton and Barto, 2018). It can also learn on its own and pass its knowledge to other machines. For example, Deep-Q has been made to play video games by controlling a joystick, and then discovering strategies for playing the game that require complex cognitive

processing (Figure 1). To accomplish this Deep Q considers multiple appropriate computational methods and evaluates how they meet its needs to maximize the consequences of its actions expressed in game points. As of 2015 Deep Q had learned 49 types of games, with skill on 29 comparable to a professional human game expert (Mnih et al., 2015).

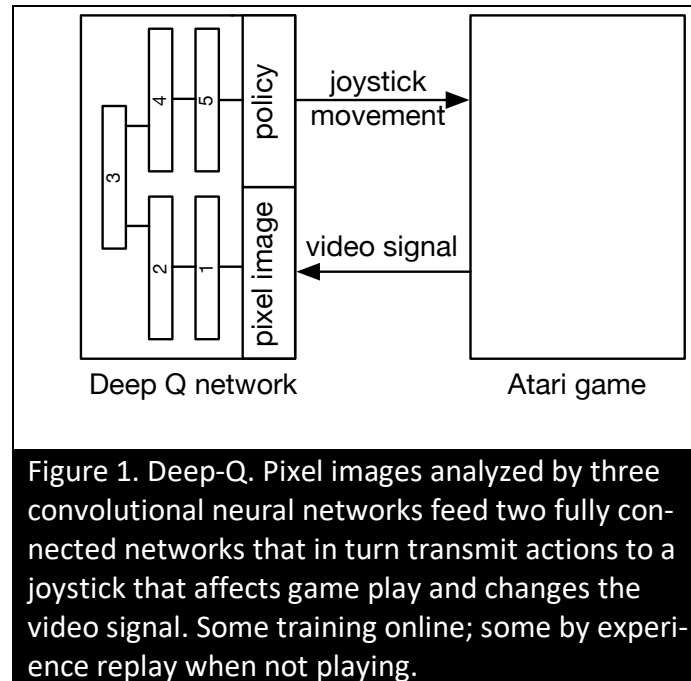


Figure 1. Deep-Q. Pixel images analyzed by three convolutional neural networks feed two fully connected networks that in turn transmit actions to a joystick that affects game play and changes the video signal. Some training online; some by experience replay when not playing.

Alphabet's DeepMind has a Go game player that learns from humans (Moyer, 2016; Wang et al., 2016). It can also play by itself to learn game strategies no human invented (Silver et al., 2018). It can start from scratch, continue learning, and work with other machines to build related game knowledge. Similar machines play now capture-the-flag games and discover strategies that humans have not used (Jaderberg et al., 2019). These successes, however, are all in closed worlds and artificial domains. These are machines that do not understand cultural settings and their strategies therefore can have effects that are anti-social (Aleksander, 2017). For example, autonomous cars using machines that learn can sometimes brake sporadically after accelerating

and cause accidents. For machines this is acceptable behavior to learn to drive but such behavior lies outside expected norms in any normal traffic situation.

Early models of the human mind in artificial intelligence sought to approximate thought through the manipulation of symbols: good old fashioned AI (Lieto et al., 2018). Subsequent models have been probabilistic. They used Bayesian modeling techniques to try to infer backward to causes from phenomenon (Lake et al., 2017; Tran et al., 2017). Indeed, recent AI based workflow research treats humans as ‘computational nodes’ in larger networks of probabilistic hybrid agents. Machines recruit humans onto teams based on predicted compatibility with other group members and their behaviors (Retelny et al., 2014). An important topic for IS research in the meta-human systems is machines selecting humans for joint tasks. This problem already shows up in ride sharing algorithms and decisions about hiring employees.

Because of advances in machine learning the current understanding in the design and use of meta-human systems is that humans pick problems (knowledge tasks) and then create a learning environment for machines to learn about the tasks. After the initial choice the machines can teach themselves. But there are caveats with this scenario: when the design of the scope is initially done by humans some important knowledge tasks might not be amenable to this approach (Marcus and Davis, 2019). Moreover, everyday tasks are not like a video game; furthermore, a person’s options in most tasks greatly exceed the possibilities presented by game controls to a machine. More generally, there is a larger system surrounding machines that learn which consists of people in companies and universities who train machine creators and the related models of building the machine software (Mackenzie, 2017). Overall, laboratory examples and results in machine learning suggest that machines will continue to improve in learning without close



human monitoring. Possibly they may learn to design parts of themselves (Dean, 2019) which will pose new challenges and implications for IS research.

### **Meta-human Systems in Operational Use**

Meta-human systems are already in use in a plethora of industries and settings covering: chip design (Bricaud, 2012), 3D printing metamodels (Kyriakou et al., 2017), travel and tourism (Orlikowski and Scott, 2015), among others. The machines in these systems incorporate various algorithms (Faraj et al., 2018), and are not constrained to the classical organizational containers of earlier sociotechnical systems theory (Winter et al., 2014). They often cross boundaries of market and infrastructure in emergent, poorly understood, complex, and distributed forms. We selected three illustrative examples of meta-human systems which currently shape organization of work and have produced many unanticipated outcomes: high-frequency trading (HFT, sometimes called algorithmic trading), elections; and autonomous vehicles.

HFT shows how speed and scale effects arise from meta-human systems (Lawrence, 2017; MacKenzie, 2019). Trading machines rely now on complex learning algorithms to trade equities quickly and at very high volumes. HFT systems use timely, accurate and extensive market information and draw on that information to learn continuously. Older HFT without learning capability has become over the last 20 years a standard and have brought significantly lower prices per trade, narrowed spreads, created higher price fluctuations, fragmented markets (Arnoldi, 2016; Kearns and Nevmyvaka, 2013; Mattli 2019). Machines that learn are bringing new capabilities and problems to how financial market operate and how investors and investment banks operate (Hendershott and Riordan, 2013; Lenglet, 2011).

Meta-human system HFT can transact equity trades within a few milliseconds – about a million times faster than human traders. It can also adjust to environmental changes including its own actions or those of competing meta-human systems (Cartea et al., 2016). Humans cannot react in time to prevent frequent, small, short-duration anomalies (e.g., crashes and spikes) (Mattli 2019). Johnson et al. (2013) found 18,520 anomalies of less than 1.5 seconds duration in a five-year period. This is much faster rate than traditional market crashes produced by human exuberance, bias and/or herd behavior. Meta-human systems HFT can now produce unexpected incidents like the flash crash of 2010 (Kirilenko et al., 2017) when humans become secondary to the machine and make decisions that come too late. Comparatively slow humans cannot act fast enough when meta-human HFT runs amok (MacKenzie, 2018). The fact that the U.S. Securities and Exchange Commission (SEC) has mandated circuit breakers to stop runaway HFT suggests the potential for unintended consequences of such machine based behavior (Kim and Yang, 2004).

Meta-human system HFT is also becoming common and pervasive. Many large investment banks (e.g., Blackrock) no longer use human traders in several asset classes or markets (Thomas, 2017). The SEC estimated in 2013 that meta-human system HFT was used for more than 50% of all U.S. stock trades (SEC Staff, 2014) and it has since grown to over 60% (Economist 2019). The SEC has mandated large scale monitoring and data collection to capture forensic information, new market making rules, real-time monitoring systems (e.g., MIDAS), and consolidated audit trails (e.g., CATS) (O’Hara, 2015). Meta-human HFT can do a great deal of damage before humans even know what is going on, and blaming human owners for prohibited trading practices makes little sense, if the prohibited behavior was executed by machines that do not ‘know’ better due to lack of understanding of cultural setting. Tight coupling that come with

electronic trading can precipitate “normal accidents” in adjacent meta-human system HFT (Perrow, 2010). For example, gas stations whose meta-human systems watch neighboring stations in real time and alter prices accordingly have been accused of collusion and price-fixing, (Schlechner, 2017). These systems can now be so complicated that consumers might need their own meta-human systems to cope (e.g., an app to avoid neighborhood fuel price fixing). Alternative monitoring and forensic investigation might be required to hold rule breakers accountable.

Manipulations of voter behavior illustrates the importance of scale in how these systems learn. In the 2016 U.S. presidential election, social media manipulated voter attitudes in critical swing states (Valentino et al., 2017). Adjusting the algorithms or adding humans to cure problems created by machines might not solve this problem. Similarly, highly assistive (nearly autonomous) vehicles are an example of meta-human systems already deployed that redefine the relationship between human and machine. They illustrate the importance of scope. Drivers are expected to react correctly to unanticipated situations (Shapiro, 2016). To address this contemporary vehicles carry hundreds of millions of lines of code and 25-30% of the value of a new car can be in software enabled functions (Charette, 2009). Collision avoidance aids have helped bad drivers get better, but they still work through drivers (Gage et al., 2015). Nearly autonomous vehicles that send data to other cars, raise questions of who is in control. Insurance companies entice drivers with less expensive insurance for continuous monitoring with, despite privacy implications (Ohlsson et al., 2015). Liability shifts from drivers to software designers, while rapid progress in sensors, data communications, and related software functions are met by lags in complements such as roads, signals, laws, training, licensing, insurance, and repair (King and Lyytinen, 2005). The significance and cost of social convention to drive on one side of the

road was illustrated by Sweden's right-hand drive switch (*Högertrafikomläggningen*) of September 3, 1967. Meta-human systems will affect such long-standing conventions. For example, in the transportation ecosystem Uber's computer-assisted dispatch uses data originating from customer mobile phones to assign drivers and facilitate pickups (Chen et al., 2015), changing the convention of communication between passenger and dispatcher to communication between phone and dispatch, both meta-human systems.

## **META-HUMAN SYSTEMS: TOWARDS AN IS RESEARCH AGENDA**

### **The Evolution of Technology and Organizing toward Socio-Technical Thinking and Meta-Human Systems**

Machines that learn are rapidly penetrating work systems. It does not matter whether they are as intelligent as humans: they still provide an opportunity for design increasingly powerful meta-human systems. Rapid machine based accumulation, distribution and dissemination of knowledge will alter organizing and organizations for ever. The period of organizing where only humans learned will end. This section addresses how an IS research agenda might recognize and respond to this inevitable trend and technological juggernaut.

The IS field's earlier focus on management information systems will prevail. The field still focuses on the prerogatives and challenges of management and organizing in the context of using novel technologies. The best way to understand the new challenges meta-human systems pose for the IS field is by reviewing the evolving role of technology in the management of organizations and organizing of work over a long period. This helps put the recent change into a broader historical perspective and understand its significance in a context. Different scholars

will frame the evolution differently, but for our purposes we will organize the grand arc of human organizing into four epochal changes which were punctuated by shifts in the mechanisms and technologies of organizational learning.

The first shift took place about 10 millennia ago from hunter-gatherer to agrarian production (Day and Walter, 1989). This introduced change in organization due to changes in human mobility (hunter-gatherers are nomadic, farmers are settled), specialization (job skills differentiated), and governance (given the surplus of agrarian production and related needs for coordination settlements could now grow from tribes into empires). The outcome was an emerging state that became a testbed for organization, along with new agriculture practice and knowledge (such as predicting seasons), rudimentary manufacturing (tools), and trade. Other than a few examples of machines (the plow, possibly the wheel) played role in this shift.

The second epochal change, the industrial revolution that began in Europe in the 15th century in Italy and then spread, made machines and associated production factors (e.g., electricity) central to everyday human life. The speed of production, transportation, communication, etc., enabled by machines changed radically in degree which became changes in kind. Records kept at that time demonstrate the importance of speed on scale and scope to improve production (Chandler and Hikino, 2009). Production governance dependent on scale and scope became essential to modern management (Hughes, 1994). Through affecting scale and scope managers complemented the invisible hand of the marketplace with the visible hand of direction (Chandler, 1993; Smith, 1937). Meta-human systems in management made machines as central to organizing as new forms of governance was in the shift from hunter-gatherer to agriculture. Machines that learn are central to meta-human systems that learn fast, thereby affecting scale and scope that requires rethinking of work systems design, as well as how to build and evolve them.

The third epochal change was synthetic, bringing together insights from centuries of governance, agriculture, trade, and industrialization into the functional view of executives as documented in Chester Barnard's classic *The Functions of the Executive* (Barnard 1938). He built on insights from Henri Fayol's study of the French state (Parker and Ritson, 2005) "scientific management" (Taylor, 1911) and the sociology of organization (e.g, Max Weber's studies of bureaucracy and organizing values such as the Protestant Ethic) (Kilcullen, 1996). This functionalist work formed a foundation for distinct management functions and functional separation and their related knowledge bases as still reflected in terms like "the marketing function", or "the information systems function." Meta-human systems involving information technology are likely to be affected and shaped by the IS function, reinforcing the centrality of IS research in meta-human systems. It should also draw IS researchers into shaping the design of future work systems. The functional view is still the foundation of the IS research agenda below that concentrates on four managerial functions related to meta-human systems.

The fourth epochal change elevates the machine to a closer and more intimate status with humans in work systems. The socio-technical systems movement of the mid 20th Century punctuated this trend when the British Government nationalized coal mining at the same time that fast, new long wall coal mining machines appeared. These machines brought about work system re-design needs as mining work had to refocus from getting coal out of the seam to getting coal out of the mine. To inform the re-design research took place in real coal mines and not in laboratories. The action research involved management and unions and mixed knowledge creation with real improvements in coal mining processes (Trist, 1981). Technologies and social systems need to be designed jointly and cover mutual relationships between people, technology, tasks, and organization. This perspective later fundamentally shaped the IS research goals and

framing (Sarker et al., 2019). In the early wave of socio-technical research machines did not learn while teams of humans learned to help achieve system-level goals of efficiency and work satisfaction (Sarker et al 2019; Jaderberg et al., 2019). Poorly designed sociotechnical systems with inadequate concern with mutual relationships were shown to fail and produce unintended or unwanted outcomes (such as alienation or vulnerability).

As computing machines moved in the 60's and 70's beyond performing simple arithmetic, sorting, and matching carried out by earlier punch-card systems to increasingly complex real time transaction recording where machine exhibit more autonomous and complex agency the shift called also for better learning among system users. In consequence, sociotechnical thinking had to change again (Alvesson and Sandberg, 2011; Ekbja and Nardi, 2014; Kallinikos, 2011; Sutherland, 2008; Seidel et al., 2018, 2019). Several extensions to understand the relationships between machines and people were put forward. These included Latour's concept of Actor-Network Theory (ANT) (Leonardi, 2012) where machines join humans as 'actants' to 'perform' networks (Latour, 2005). Now machines can act autonomously (Parasuraman et al., 2000) as parts of 'networks of humans and machines', and ANT provides a "...powerful tool to regain the sense of heterogeneity, but is a "...bad tool for differentiating associations" (Latour 1996, p 380). Because of this ANT still does not fully open to the possibility of what happens when machines learn by themselves and from other machines in association with humans. But it brings machines closer to the human level in the work system analysis. The relationship between human and machine is more nuanced but human agency is primary. As a result, humans and machines 'imbricate' one another (Leonardi, 2011). For example, different capability levels in machines that learn will bring about significantly different allocations of tasks and related work sequences.

As these systems learn and adapt based on human action, the uses, goals and ways of using machines will also change.

Generally, the sociotechnical tradition embraces mutually enforcing, sustainable relations between people, technology, tasks, and organization but also expects causal effects to emerge from features of technology – which are mediated by specific structures and features of task, people and organization. In such causal analysis humans and machines enter during work system design into relationships that need to be rendered analytically separable for causal analysis but which still remain ontologically intertwined. Markus and Rowe's (2018) recent review of sociotechnical thinking in IS research identifies different notions of causality when machines of varying kinds are integrated into work systems with specific effects. Agency in networks will involve machines which have causal powers that differ from those that are subjected to clean human control. Therefore, sociotechnical theory and how it conceives the causal effects between the machines and humans must change. Machines that learn will have new kinds of emergent system level effects on work systems. For example, the concern for human autonomy, integrity and value in design needs to be contrasted to F.W. Taylor's scientific management (1911), which modeled humans as if they were machines, components of work systems. Now, if machines behave increasingly more like humans or beyond what humans can do cognitively should the principles of the work system design also change?

One important element in meta-human system design is that human learning makes humans versatile. When viewed as machines, humans can be seen to act like truly general purpose machines applicable to many, even so far unconceived, tasks (Helpman, 1998). Computer and communications technologies which underlie the capabilities of machines that learn can also be characterized as general purpose technologies (David, 1990). Machine based learning expands



the general purpose character of such technologies making built machines closer to the human. In this sense, meta-human systems will create general purpose systems with new types of system level scale and scope economies that will shape the future of organizing.

Both effects are important in understanding the evolution of management thought so far and how they will play out in the future. Machines draw closer to humans as interpreted in the socio-technical and their variant ANT traditions. The origins of managerial thought suggest that when machines get better the joint capabilities with humans in meta-human systems will greatly improve organization's ability to carry out a wider range of tasks in new ways and introduce new tasks. IS research into meta-human systems can be understood by applying a functional view into how to organize meta-human systems even though the growth of meta-human systems in organizations will shift their management away from the classic functional view.

### **A Functional Frame for IS Research Involving Meta-Human Systems**

Four functions – delegation, monitoring, cultivating and reflecting – capture an initial IS research agenda involving metahuman systems with the objective of designing and integrating meta-human systems into work systems. As with most functional views of socio-technical systems, the agenda seeks to achieve a level of control that keeps the systems operational and helps achieve organization level goals that derive from new scale and scope economies. This type of control captures essentially the idea of management as it is now thought of (Yates, 1988)(p. xiv). Controlling meta-human systems will in the coming decade and beyond become part of management practice and essential element of work systems design. This control might accommodate levels of machine autonomy to acquire capabilities, but follow also human goals as part of cooperative effort (Lawrence, 2017), or in setting goals (Aleksander, 2017). At the

same time machines will use humans for goal-setting the way Google's search engine uses human input to discover and present content deemed useful when measured by click-through-rates. With these new capabilities come new trade-offs. Narrow goals will constrain system level learning, while broader goals is likely to produce multiple unintended consequences (e.g., content manipulation, search results biased by stereotype). Controls that work for humans in a meta-human system will this not work for machines, and vice versa.

IS research into new IT capability, changing properties of decision making, and organizational effectiveness, will remain relevant for IT based design and implementation generally, but are likely insufficient in guiding how to think of control of meta-human systems in future organizations. In future, managers are still likely to operate a visible hand that shakes the invisible hand of the market, but many hand-shaking tasks will be carried out by meta-human systems where the hand and the shake are based on machine based learning (Chandler, 1993; Smith, 1937). The functionalist view will remain relevant in this new setting. Simon's 1948 classic *Administrative Behavior* (Simon 1977) anticipated new kinds of digital technologies that would affect all the functions but he never questioned their value and role in organizing. This has also remained a core assumption in IS research: reap system-level rewards by exploiting technology use via effective adjustments of the social across the functions. Humans will set goals, refine plans, explain decisions, train other humans, and build or buy technology. It is human resistance and failure to learn that causes implementation failure. But machines that learn will make possible meta-human system where machines that learn operate with higher speed and autonomy, and this will affect the scale and scope of organizational learning ways of implementing systems. In consequence, organizing needs to be re-thought to ensure that technology and people remain re-aligned.

Four functional areas are suggested as key levers to manage meta-human systems (Table 1). The first two, delegating (letting machines that learn and humans do their job), and monitoring (anticipatory watching for consequences) relate to meta-human system operations. Delegating assumes agency, hierarchy, authority, and decision rights and responsibilities that are part of any meta-human system. Monitoring involves attention costs, appropriate levels of psychological monitoring to balance indifference versus micromanagement with adequate subordinate learning (Simon, 1997). Both functional processes defined management in the 20th century (Holmstrom and Milgrom, 1991). The next two, cultivating (mindful enhancement of learning by setting up the right conditions) and reflecting (learning at the meta-human systems level) are new functions that relate to the autonomy and speed that characterizes all elements of meta-human systems. Cultivating is about system enhancement that touches both machine and human training, related management and organization. Reflecting captures the idea of double-loop learning that characterizes meta-human systems and allows critically evaluate their learning processes and outcomes (Argyris, 2004; Schön, 1983).

The suggested framework is similar to many early IS research frameworks that helped frame pertinent topics for a nascent research field (Mason and Mitroff, 1973). The framework is just one possible sensitizing approach- not the only approach. No one can predict accurately what will happen with meta-human systems or know how to best organize them. Much depends on framing and theorizing whereby IS researchers can learn to what differences meta-human systems make in different settings and how the management can learn to adjust machine and human relationships in ways that improve how all socio-technical elements fit together for established system goals (Markus and Rowe, 2018).

<b>Function</b>	<b>Research Questions</b>	<b>Theory base</b>	<b>Potential Empirical Sites</b>	<b>Design/experiment</b>
<i>Delegating</i>	What types of governance processes are needed to map responsibility to humans, and how best to provide this?	Forms and sources of governance relevant to meta-human systems (Zheng et al., 2019)	Medical devices (Therac-25) Ride sharing companies (Uber, Lyft) Autonomous vehicle trials (Tesla, Waymo)	Experimental engineering Exploratory field studies Online gaming experiments
	What skills need to evolve to design and test responsibility-based processes?	Sociotechnical system design Business process models (Sergeeva et al., 2018)	Process re-engineering Study of effects and non-intended side effects Programming task evolution Data scientist career growth	New curricula and relevant evaluation Study of design curricula and effects
<i>Monitoring</i>	How should machine-based decisions operating at high speed scale and scope be monitored? How to achieve this sustainably?	(Real time) accounting/auditing High reliability organizing (Rozario and Vasarhelyi, 2018)	High frequency trading (flash crashes) Electric grid monitoring (failure and successes) Robotic surgery (Intuitive Surgical)	Studies of system scaling and interactions Simulation experiments of emergent system behaviors, Trading floor analyses and experiments
	How should humans interact with machines augmented with new functions? How to achieve this sustainably?	HCI theory, Agency theory involving groups and technology, Studies of digital transformation and digital ontologies (Cai et al., 2019)	Design of machines (e.g. Synopsys software), games (Ubisoft, Activision), autonomous vehicles (Tesla, Waymo, Uber), etc.	Interface theories and experiments Studies of work system change and behaviors
<i>Cultivating</i>	How best to achieve the right selection of training data for learning? How to achieve this sustainably?	Statistical learning theory (Wexler et al., 2019)	Bots (Microsoft, Amazon) Newsfeeds(Facebook, Twitter) Face recognition (UK, China, Las Vegas)	Adversarial networks (generative and otherwise)
	What criteria are useful for evaluating learning regarding meta-human systems? How to affect learning?	Cognitive science Learning science Education (Lake et al., 2017)	Organizational learning in manufacturing and service industries	Cognitive modelling Simulation of meta-human systems Behavioral economic experiments
<i>Reflecting</i>	What processes help evaluate meta-human system progress? How to create such processes?	Human- Computer Interaction IT governance models (Buchwald et al., 2014)	Diffusion studies of past disruptive technologies: elevators, traffic signals, office automation.	Longitudinal analyses of technological change and evolution Metadesign experiments
	What meta-human system features can be anticipated? How best to accomplish this?	STS theory Cybernetics Technology history and policy (Kallinikos et al. 2013)	Large projects (Manhattan, SAGE, space exploration, the Internet, etc.) Standardization efforts.	Technology-focused field studies Public policy interventions Regulatory actions

## **Delegating**

Delegating grants authority and resources to agents to accomplish tasks, freeing resources for other work. In meta-human systems agents can be machines, humans or hybrid agents.

Delegation is about *discovering* what will happen if we shift the agency in the system. It is

almost impossible to know *ex ante* how to configure the overall system if multiple possibilities are available for delegation (Colombo and Delmastro, 2004; Vickers, 1985). Delegation in meta-human systems integrates machines that learn into tasks formerly restricted to human learning. For example, delegation of driving control requires new kind of autonomy in traffic control, road systems with sensors, and changes in insurance systems (King and Lyytinen, 2005). Delegation to automated traders invites illegal or unpredictable trading (Scopino, 2015). Delegation of editorial decisions to newsfeeds can lead to the manipulation of voter behavior (Lazer et al., 2018).

The result of this delegation is new kinds of human/machine hybrids. Currently, our understanding of delegation advances through learning-by-doing and mostly in the field given the limitations of laboratory experiments. Delegation comes with new specific agency problems, and different policies can be used to evade agent responsibility in the event of failure (Bartling and Fischbacher, 2012). Research is needed on how humans can best delegate to machines, the parameters of delegation, undesired behaviors to be avoided, and desired behaviors to be hoped for. A good start might be to look at research into interactive algorithms for ways human/machine interaction can be improved (Goldin and Wegner, 2006; Wegner, 1998). Fine-grained individual interactions might be collected into larger chunks of capability that allow higher-level delegation similar to the way sociologists use grounded interpersonal actions as a way of building higher level abstractions related to identity and organization (Abbott, 2016).

Machine delegation to humans is a nascent and important topic as machine capabilities continue to improve. Social media companies rely now on learning-by-doing to watch humans accept and place advertising, then later delegate this role to machines thereby for example generating election fairness concerns (Berghel, 2017; Lazer et al., 2018). Machines often do not

understand what humans find important. Systems like NELL help reveal what works for the collective as a whole (including machines and humans), and enable machines to supervise humans (Pedro and Huschka, 2012; Retelny et al., 2014). There are trade-offs in new forms of delegation. Legal constraints may prevent some forms of delegation, demanding responsibility from humans (Solaiman, 2017). New work quality issues will emerge: human deskilling can become a significant source of work dissatisfaction (Trist, 1981). Machine substitution for humans in pure profit-seeking can risk reputational or emotional loss. Machines cannot understand feelings that guide human systems (Hubbard, 2010). Observational studies will help assess how and when to shift agency (List and Pettit, 2011), and how roles are already shifting (Mortensen and Haas, 2018).

## **Monitoring**

Agency requires monitoring (Jensen and Meckling, 1976), but monitoring mechanisms timed for humans will be too slow for machines. For example HFT proceeds so quickly that monitoring needs to be performed by second level meta-human systems. Such monitoring is also social, as when algorithms act on changes in liquidity (Hendershott and Riordan, 2013). These systems are already driving research into how monitor and regulate market activity (Lenglet, 2011; Treleven and Batrinca, 2017; Siering et al., 2017). Likewise, election monitoring systems can be audited to figure out what truly happened (Masterson, 2019). IS researchers have raised concerns for meta-human system surveillance for hire (Clarke, 2019). Monitoring autonomous driving systems is currently a complex topic and requires interdisciplinary research (Koopman and

Wagner, 2017). For example, the complexity of handovers and takeovers in semiautonomous vehicles already taxes human cognitive load, and itself needs monitoring (Sibi et al., 2016).

During monitoring continuous auditing might help flag anomalous outputs (Kiesow et al., 2016). or invoke rule-based triggers inside meta-human processes (Rozario and Vasarhelyi, 2018; Singh et al., 2013). In institutionally regulated settings such as finance or transportation authorities can and need to be notified of fraudulent behavior or emergent changes in goals (Abbasi et al., 2012). There is need for observational studies to show which architectures work best under different circumstances for monitoring goals. Election meddling is often a consequence of myopic architectural choice (Berghel, 2017) and can only be countered by taking a broader view. The recently introduced SEC's forensic and real time monitors may lead to a better understanding of HFT (SEC staff, 2014) and better mechanisms.

Overall, IS research is needed to better understand the consequences of using different strategies of monitoring. Three approaches in particular deserve attention. Disclosive archaeology help show what to look for – this is where humans run experiments on machines to determine how they work (Ananny, 2016; Barocas et al., 2013). However, algorithms like Facebook's news feed are changing constantly, so the utility of disclosive archaeology is limited. Transparent design involves stakeholders and users, and shows them explicitly what is driving the system (Introna, 2007; Lyons et al., 2017). But few organizations are willing to commit to such extensive transparency. Moreover, even with a goal of transparency, creating machines that divulge and explain their behavior in way humans can interpret is difficult. Data disclosure by the machine as it operates can provide some means for monitoring (Lenglet, 2019), but fast machines require fast monitoring to keep up : this in turn calls for more systems! In principle, monitoring computers can better report to humans as now seen in smart watches, phones, and

clickstream analyzers. However, few companies disclose enough data to enable such monitoring of their own server-based capabilities.

## **Cultivating**

The cultivation metaphor comes from agriculture (trimming, correct spacing, fertilizing and removal of insect pests and weeds). It is a stage in mentoring (Chao, 1997; Humberd and Rouse, 2016). Cultivating systems or technology has been discussed in the IS context for some time (Bergqvist and Dahlberg, 1999; Hanseth, 2010; Henfridsson and Bygstad, 2013; Markus and Benjamin, 1996). Systems are prone to evolve organically, so cultivation focuses attention on meta-human systems learning and knowledge allocation to different system components (Martinez, 2014). HFT for example has a long history of research that has focused on cultivation. Algorithmic traders interact and change the way to think about machine training (MacKenzie, 2019). Autonomous driving systems, in contrast, focus on space rather than time. In this regard cars can be trained to anticipate running animals, but we do not have models for all animals in any landscape. Kangaroos jump differently than deer (O'Rourke, 2017), and so a car may need to be trained differently to match the fauna in the environment. Election systems are complicated by serious social as well as technical challenges. In particular, social media has political impact (Kreiss and McGregor, 2018), and currently machine learning systems can be trained to recognize political ads (Qi et al., 2016). But controversies over political ads and how far we can go point to the challenges of cultivation. For example, can the AI discriminate between antisocial political ads and eusocial political ads? Should those who run the AI systems censor ads? There are no easy answers to these questions (Kreiss and McGregor, 2019; *The Washington Post*, 2019).



Organizations might not be ready now to cultivate their meta-human systems as suggested by multiple failures, examples of inherent bias, and systems running amok (Kulshrestha et al., 2017; Lambrecht and Tucker, 2019). However, work on these areas has proceeded and is receiving more attention. For bias, one can visualize causes and adjust the system accordingly (Wexler et al., 2019). As organizations acquire this capability, meta-human systems will come to address social processes with multiple viewpoints. An immediate need is to cultivate systems that serve human-defined and human-monitored purposes (Salimans et al., 2017), but eventually, machines might cultivate humans. Electronic tutoring might be an early exemplar of this (Wenger, 2014).

## **Reflecting**

Reflective practice helps human experts to build novel knowledge in any professional domain (Schön, 1983). Reflection is also important to aesthetics and ethics of metahuman system design (Bostrom and Ord, 2006). For example, reflection about HFT should anticipate large scale changes in the financial industry (Gomber et al., 2018). Metahuman systems have now been accepted for increased liquidity and more efficient investing. But they also precipitate new approaches to regulation (Arner et al., 2016). Likewise, election systems are interwoven into media and other complex systems and call for reflection. Questioning algorithms that control news feeds and control advertising is far from affecting change. There are also disciplinary differences: sociology focuses on collective action (Coleman, 2017), while IS researchers concentrate on false news in social media (Aral and Eckles, 2019; Vosoughi et al., 2018). Autonomous vehicles may also pose a category error, attributing greater cognitive ability to robots than they warrant (Aleksander, 2017). Perhaps the road to full vehicle autonomy is very long though some say that benefits such as safety can be achieved without full automation

(Bailey and Erickson, 2019). The argument for full automation calls for changes in economic systems to handle labor displacement (Srnicek and Williams, 2015). Other analyses of labor displacement are more nuanced, considering which jobs in which orders might be displaced (Brynjolfsson and McAfee, 2014). From the IS research perspective, reflection here is about any case of metahuman systems can lead to a far-ranging analysis of societal impacts and values.

Asimov's famous laws for robots proved circumventable (Asimov, 1942), perhaps inevitably, as they can be viewed as an early attempt to introduce reflection into automation studies. Clearly machines can harm humans because of poor design, careless programming, or malicious action. Machines that reflect on the behavior of other machines to detect problems seem far-off (Goertzel, 2012). We arguably don't know how to build machines that understand and conform to human values (Etzioni and Etzioni, 2016). Thus far, grounding in human values requires reflection by humans. Reinforcement and reward engineering might be starting points for future research on how to design machines that reflect human values (Dewey, 2014). But the question of whether ethics for meta-human systems can be learned top-down, bottom-up, or both has not even been answered (Yampolskiy, 2016).

IS research has applied the idea of reflective practice to understand how system design works (cf. (Córdoba, 2007; Mathiassen, 1998; Redmiles and Nakakoji, 2004; Stroulia and Goel, 1995). Reviews, consulting meetings, and coaching are but a few of the mechanisms used for such reflection. Reflection, however, has remained a uniquely human endeavor; and at present, machines do not engage in reflection. On the other hand, meta-human systems already outperform some human experts in law, finance, and medicine. But, as noted above, they have not demonstrated an ability to invent mechanisms for their own governance or to understand

what they have learned and consequently, change their parameters for delegation, monitoring, or cultivation.

IS researchers can do case studies of reflective practice for meta-human systems founded on action research of the sociotechnical systems tradition. However, many IS researchers within professional schools have seen such action research as methodologically flawed due to the within-system effects of the researcher. Moreover, there is concern about such research being supported by those who might have a particular goal related to the outcome (e.g., owners, executives, and workers). Conflicts of interest, confounds, cognitive dissonance and confirmation bias threaten action research. Meta-human systems might require IS researchers to move into the unfamiliar and get out of their comfort zone. Two recent studies are signposts in this regard: Zuboff's look at Google and Microsoft (2019) and Google's look at its own AI systems as used by doctors (Cai et al. 2019).

## **CONCLUSION**

Information systems increasingly incorporate machines that learn. Together with humans, these machines that learn combine to form meta-human systems which exhibit sociotechnical systems of new ilk. On the one hand they embody classic machine advantages of calculating, recording and transferring information in simple intellectual tasks. On the other hand they come with higher level cognitive skills that affect the speed of organizational learning, and shape the scale and scope economies of organizations in new ways. Meta-human systems afford IS research in four areas of work system organization: delegating, monitoring, cultivating, and reflecting. Differences in the ways humans and machines learn will affect how meta-human systems are organized.

This paper is a call to arms for IS research in the era of meta-human systems. It asks IS researchers to address issues of human goals and values in settings where meta-human systems evolve or are applied. Achieving benefits and avoiding problems will require better understanding of systems level learning, and how learning emanating from meta-human systems affects large ecosystems in manufacturing, agriculture, transportation, finance, medicine, and other fields. Design science researchers can explore and test alternative designs for meta-human systems similar to the sociotechnical systems tradition that confronted disruptive technology and was successfully adopted by IS research. IS researchers should join with practitioners and scholars from administration, policy, computing, law, and other fields relevant to understand, design and implement meta-human systems.

## Funding

This material is based upon work supported by the National Science Foundation under Grants 1909803, 1717473, 1745463, 1442840, and 1422066.

## References

- Abbasi A, Albrecht C, Vance A, et al. (2012) Metafraud: A meta-learning framework for detecting financial fraud. *MIS Quarterly* 36(4). pdfs.semanticscholar.org: 1293–1327.
- Abbott A (2016) *Processual Sociology*. University of Chicago Press.
- Ackoff RL and Emery FE (2005) On Purposeful Systems: An Interdisciplinary Analysis of Individual And Social Behavior As a System of Purposeful Events. Transaction Publishers.
- Aleksander I (2017) Partners of Humans: A Realistic Assessment of the Role of Robots in the Foreseeable Future. *Journal of Information Technology* 32(1). SAGE Publications Ltd: 1–9.
- Ananny M (2016) Toward an Ethics of Algorithms: Convening, Observation, Probability, and Timeliness. *Science, technology & human values* 41(1). SAGE Publications Inc: 93–117.
- Aral S and Eckles D (2019) Protecting elections from social media manipulation. *Science* 365(6456). science.sciencemag.org: 858–861.
- Argall BD, Chernova S, Veloso M, et al. (2009) A survey of robot learning from demonstration. *Robotics and autonomous systems* 57(5). Elsevier: 469–483.
- Argyris C (2004) Double-Loop Learning and Implementable Validity. In: Tsoukas H and Mylonopoulos N (eds.) *Organizations as Knowledge Systems: Knowledge, Learning and Dynamic Capabilities*. London: Palgrave Macmillan UK, pp. 29–45.
- Arner DW, Barberis J and Buckley RP (2016) FinTech, RegTech, and the reconceptualization of financial regulation. *Nw. J. Int'l L. & Bus.* 37. HeinOnline: 371.
- Arnoldi J (2016) Computer Algorithms, Market Manipulation and the Institutionalization of High Frequency Trading. *Theory, Culture & Society* 33(1). SAGE Publications Ltd: 29–52.
- Asimov I (1942) Runaround. *Astounding Science Fiction* 29(1): 94–103.
- Bailey DE and Erickson I (2019) Selling AI: The Case of Fully Autonomous Vehicles. *Issues in science and technology* 35(3). Issues in Science and Technology: 57–61.
- Barnard CI (1938) *The Functions of the Executive*. Reprint, Harvard University Press, 1968.
- Barocas S, Hood S and Ziewitz M (2013) Governing Algorithms: A Provocation Piece. Available at SSRN 2245322. DOI: 10.2139/ssrn.2245322.
- Bartling B and Fischbacher U (2012) Shifting the Blame: On Delegation and Responsibility. *The Review of economic studies* 79(1). Narnia: 67–87.
- Berghel H (2017) Oh, What a Tangled Web: Russian Hacking, Fake News, and the 2016 US Presidential Election. *Computer* 50(9). ieeexplore.ieee.org: 87–91.
- Bergqvist J and Dahlberg P (1999) Scalability through cultivation. *Scandinavian Journal of Information Systems* 11(1). aisel.aisnet.org: 1.

- Blue J and Andoh-Baidoo FK (2010) Directive Decision Devices: Extending the reach of automation into the finance domain. *Expert systems with applications* 37(1). Elsevier: 45–54.
- Bostrom N and Ord T (2006) The reversal test: eliminating status quo bias in applied ethics. *Ethics* 116(4). journals.uchicago.edu: 656–679.
- Bricaud P (2012) Reuse Methodology Manual: For System-on-a-Chip Designs. Springer Science & Business Media.
- Brynjolfsson E and McAfee A (2014) The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. W. W. Norton & Company.
- Buchwald A, Urbach N and Ahlemann F (2014) Business Value through Controlled IT: Toward An Integrated Model of IT Governance Success and Its Impact. *Journal of Information Technology* 29(2). SAGE Publications Ltd: 128–147.
- Bucilua C, Caruana R and Niculescu-Mizil A (2006) Model compression. *Proceedings of the 12th ACM*. dl.acm.org. Available at: <https://dl.acm.org/citation.cfm?id=1150464>.
- Cai CJ, Winter S, Steiner D, et al. (2019) Hello AI: Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW). ACM: 104.
- Carlson A, Betteridge J, Kisiel B, et al. (2010) Toward an architecture for never-ending language learning. In: *Twenty-Fourth AAAI Conference on Artificial Intelligence, 2010*. aaii.org. Available at: <https://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/viewPaper/1879>.
- Cartea Á, Jaimungal S and Kinzebulatov D (2016) ALGORITHMIC TRADING WITH LEARNING. *International Journal of Theoretical and Applied Finance* 19(04). World Scientific Publishing Co.: 1650028.
- Chandler AD and Hikino T (2009) *Scale and Scope: The Dynamics of Industrial Capitalism*. Harvard University Press.
- Chandler AD Jr (1993) *The Visible Hand*. Harvard University Press.
- Chao GT (1997) Mentoring Phases and Outcomes. *Journal of vocational behavior* 51(1). Elsevier: 15–28.
- Charette RN (2009) This car runs on code. *IEEE Spectrum* 46(3). IEEE: 3.
- Chen L, Mislove A and Wilson C (2015) Peeking beneath the hood of uber. In: *Proceedings of the 2015 Internet Measurement Conference, 2015*, pp. 495–508. ACM.
- Clark A (2017) Embodied, Situated, and Distributed Cognition. In: Bechtel W and Graham G (eds.) *A Companion to Cognitive Science*. Oxford, UK: Blackwell Publishing Ltd, pp. 506–517.
- Clarke R (2019) Risks inherent in the digital surveillance economy: A research agenda. *Journal of Information Technology* 34(1). SAGE Publications Ltd: 59–80.
- Coleman J (2017) The mathematics of collective action. content.taylorfrancis.com. Available at: <https://content.taylorfrancis.com/books/download?dac=C2017-0-53071-0&isbn=9781351479714&format=googlePreviewPdf>.
- Colombo MG and Delmastro M (2004) Delegation of Authority In Business Organizations: An Empirical Test. *The Journal of industrial economics* 52(1). Wiley Online Library: 53–80.
- Córdoba J-R (2007) Developing Inclusion and Critical Reflection in Information Systems Planning. *Organization* 14(6). SAGE Publications Ltd: 909–927.
- Davenport TH and Kirby J (2016) Just how smart are smart machines? *MIT Sloan Management Review* 57(3). Massachusetts Institute of Technology, Cambridge, MA: 21.

- David PA (1990) The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *The American economic review* 80(2). American Economic Association: 355–361.
- Day RH and Walter J-L (1989) Economic growth in the very long run: On the multiple-phase interaction of population, technology, and social infrastructure. In: Barnett WA, Geweke J, and Shell K (eds.) *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity*. books.google.com, pp. 253–289.
- Dean J (2019) The Deep Learning Revolution and Its Implications for Computer Architecture and Chip Design. *arXiv [cs.LG]*. Available at: <http://arxiv.org/abs/1911.05289>.
- DeLanda M (2019) *A New Philosophy of Society: Assemblage Theory and Social Complexity*. Bloomsbury Publishing.
- Dewey D (2014) Reinforcement learning and the reward engineering principle. In: *2014 AAAI Spring Symposium Series*, 2014. aaii.org. Available at: <https://www.aaii.org/ocs/index.php/SSS/SSS14/paper/viewPaper/7704>.
- Domingos PM (2012) A few useful things to know about machine learning. *Communications of the ACM* 55(10). pdfs.semanticscholar.org: 78–87.
- du Gay P (2005) *The Values of Bureaucracy*. OUP Oxford.
- Ekbia H and Nardi B (2014) Heteromation and its (dis) contents: The invisible division of labor between humans and machines. *First Monday* 19(6). firstmonday.org. Available at: <http://firstmonday.org/article/view/5331/4090>.
- Engeström Y, Lompscher J and Rückriem G (2016) *Putting Activity Theory to Work: Contributions from Developmental Work Research*. Lehmanns Media.
- Ericsson KA, Krampe RT and Tesch-Römer C (1993) The role of deliberate practice in the acquisition of expert performance. *Psychological review* 100(3). American Psychological Association: 363.
- Etzioni A and Etzioni O (2016) AI assisted ethics. *Ethics and information technology* 18(2). Springer: 149–156.
- Faraj S, Pachidi S and Sayegh K (2018) Working and organizing in the age of the learning algorithm. *Information and Organization* 28(1). Elsevier: 62–70.
- Frith CD and Frith U (2012) Mechanisms of social cognition. *Annual review of psychology* 63. annualreviews.org: 287–313.
- Gage T, Bishop R and Morris J (2015) The Increasing Importance of Vehicle-Based Risk Assessment for the Vehicle Insurance Industry. *Minn. JL Sci. & Tech.* 16. HeinOnline: 771.
- Goertzel B (2012) *Creating Internet Intelligence: Wild Computing, Distributed Digital Consciousness, and the Emerging Global Brain*. Springer Science & Business Media.
- Goldin D and Wegner P (2006) Principles of Interactive Computation. In: Goldin D, Smolka SA, and Wegner P (eds.) *Interactive Computation: The New Paradigm*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 25–37.
- Gomber P, Kauffman RJ, Parker C, et al. (2018) On the fintech revolution: interpreting the forces of innovation, disruption, and transformation in financial services. *Journal of management* 35(1). Taylor & Francis: 220–265.
- Green B and Chen Y (2019) The Principles and Limits of Algorithm-in-the-Loop Decision Making. *Proc. ACM Hum. -Comput. Interact.* 3(CSCW). New York, NY, USA: ACM: 50:1-50:24.

- Hanseth O (2010) From Systems and Tools to Networks and Infrastructures—from Design to Cultivation: Towards a Design Theory of Information Infrastructures. In: *Industrial Informatics Design, Use and Innovation: Perspectives and Services*. IGI Global, pp. 122–156.
- Heaven D (2019) Why deep-learning AIs are so easy to fool. *Nature* 574(7777): 163–166.
- Helpman E (1998) General Purpose Technologies and Economic Growth. MIT Press.
- Hendershott T and Riordan R (2013) Algorithmic Trading and the Market for Liquidity. *Journal of Financial and Quantitative Analysis* 48(4). Cambridge University Press: 1001–1024.
- Henfridsson O and Bygstad B (2013) The Generative Mechanisms of Digital Infrastructure Evolution. *MIS Quarterly* 37(3). Management Information Systems Research Center, University of Minnesota: 907–931.
- Hinton G, Vinyals O and Dean J (2015) Distilling the Knowledge in a Neural Network. *arXiv [stat.ML]*. Available at: <http://arxiv.org/abs/1503.02531>.
- Hirschheim R (1985) Information systems epistemology: An historical perspective. *Research methods in information systems*. North-Holland, Amsterdam: 13–35.
- Holmstrom B and Milgrom P (1991) Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.* 7: 24.
- Horvitz E (1999) Principles of Mixed-initiative User Interfaces. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 1999, pp. 159–166. CHI '99. ACM.
- Hubbard FP (2010) Do androids dream: personhood and intelligent artifacts. *Temple law review* 83. HeinOnline: 405.
- Hughes TP (1994) Technological momentum. *Does technology drive history* 101. Cambridge, MA: MIT Press.
- Humberd BK and Rouse ED (2016) Seeing You in Me and Me in You: Personal Identification in the Phases of Mentoring Relationships. *AMRO* 41(3). Academy of Management: 435–455.
- Introna LD (2007) Maintaining the reversibility of foldings: Making the ethics (politics) of information technology visible. *Ethics and information technology* 9(1). Kluwer Academic Publishers: 11–25.
- Jaderberg M, Czarnecki WM, Dunning I, et al. (2019) Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* 364(6443). [science.sciencemag.org](http://science.sciencemag.org): 859–865.
- Jensen MC and Meckling WH (1976) Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of financial economics* 3(4): 305–360.
- Johnson Neil, Zhao G, Hunsader E, et al. (2013) Abrupt rise of new machine ecology beyond human response time. *Scientific reports* 3. [nature.com](http://nature.com): 2627.
- Johnson-Roberson M, Barto C, Mehta R, et al. (2016) Driving in the Matrix: Can Virtual Worlds Replace Human-Generated Annotations for Real World Tasks? *arXiv [cs.CV]*. Available at: <http://arxiv.org/abs/1610.01983>.
- Kallinikos J, Aaltonen A and Marton A (2013) The Ambivalent Ontology of Digital Artifacts. *MIS Quarterly* 37(2). Management Information Systems Research Center, University of Minnesota: 357–370.
- Kearns M and Nevmyvaka Y (2013) Machine learning for market microstructure and high frequency trading. In: David Easley ML de PAMO (ed.) *High Frequency Trading - New Realities for Traders, Markets and Regulators*. Risk Books.



- Kiesow A, Schomaker T and Thomas O (2016) Transferring continuous auditing to the Digital Age--The Knowledge Base after three decades of research. *aisel.aisnet.org*. Available at: [https://aisel.aisnet.org/ecis2016\\_rp/42/](https://aisel.aisnet.org/ecis2016_rp/42/).
- Kilcullen J (1996) Max Weber: On Bureaucracy. Available at: <https://philpapers.org/rec/kilmwo>.
- Kim YH and Yang JJ (2004) What Makes Circuit Breakers Attractive to Financial Markets? A Survey. *Financial Markets, Institutions and Instruments* 13(3). Wiley Online Library: 109–146.
- King JL and Lyytinen K (2005) Automotive informatics: Information technology and enterprise transformation in the automobile industry. *Transforming enterprise: The economic and social implications of information technology*. Cambridge, MA: MIT Press: 283–333.
- Kirilenko A, Kyle AS, Samadi M, et al. (2017) The flash crash: High-frequency trading in an electronic market. *The Journal of finance* 72(3). Wiley Online Library: 967–998.
- Kling R (1980) Social analyses of computing: Theoretical perspectives in recent empirical research. *ACM Computing Surveys (CSUR)*. *dl.acm.org*. Available at: <https://dl.acm.org/citation.cfm?id=356806>.
- Koopman P and Wagner M (2017) Autonomous Vehicle Safety: An Interdisciplinary Challenge. *IEEE Intelligent Transportation Systems Magazine* 9(1). *ieeexplore.ieee.org*: 90–96.
- Kreiss D and McGregor SC (2018) Technology Firms Shape Political Communication: The Work of Microsoft, Facebook, Twitter, and Google With Campaigns During the 2016 U.S. Presidential Cycle. *Political Communication* 35(2). Routledge: 155–177.
- Kreiss D and McGregor SC (2019) The “Arbiters of What Our Voters See”: Facebook and Google’s Struggle with Policy, Process, and Enforcement around Political Advertising. *Political Communication*. Taylor & Francis: 1–24.
- Kulshrestha J, Eslami M, Messias J, et al. (2017) Quantifying Search Bias: Investigating Sources of Bias for Political Searches in Social Media. *arXiv [cs.SI]*. Available at: <http://arxiv.org/abs/1704.01347>.
- Kurzweil R (2000) *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. Penguin.
- Kyriakou H, Nickerson JV and Sabnis G (2017) Knowledge Reuse for Customization: Metamodels in an Open Design Community for 3D Printing. *MIS Quarterly* 41(1): 315–332.
- Lake BM, Salakhutdinov R and Tenenbaum JB (2015) Human-level concept learning through probabilistic program induction. *Science* 350(6266). *science.sciencemag.org*: 1332–1338.
- Lake BM, Ullman TD, Tenenbaum JB, et al. (2017) Building machines that learn and think like people. *The Behavioral and brain sciences* 40. *cambridge.org*: e253.
- Lambrecht A and Tucker C (2019) Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management science* 65(7). INFORMS: 2966–2981.
- Latour B (2005) *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford.
- Lawrence ND (2017) Living Together: Mind and Machine Intelligence. *arXiv [cs.AI]*. Available at: <http://arxiv.org/abs/1705.07996>.
- Lazer DMJ, Baum MA, Benkler Y, et al. (2018) The science of fake news. *Science* 359(6380). *science.sciencemag.org*: 1094–1096.
- Lee MK, Kusbit D, Metsky E, et al. (2015) Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers. In: *Proceedings of the 33rd Annual*

- ACM Conference on Human Factors in Computing Systems*, New York, NY, USA, 2015, pp. 1603–1612. CHI '15. ACM.
- Lenglet M (2011) Conflicting Codes and Codings: How Algorithmic Trading Is Reshaping Financial Regulation. *Theory, Culture & Society* 28(6). SAGE Publications Ltd: 44–66.
- Lenglet M (2019) Algorithmic Finance, Its Regulation, and Deleuzian Jurisprudence: A Few Remarks on a Necessary Paradigm Shift. *Topoi. An International Review of Philosophy*. Springer Netherlands: 1–9.
- Leonardi PM (2011) When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Quarterly*. JSTOR: 147–167.
- Leonardi PM (2012) Materiality, sociomateriality, and socio-technical systems: What do these terms mean? How are they different? Do we need them. In: Leonardi P. M., Nardi BA, and Kallinikos J (eds.) *Materiality and Organizing: Social Interaction in a Technological World*. Oxford University Press Oxford.
- Lieto A, Bhatt M, Oltramari A, et al. (2018) The role of cognitive architectures in general artificial intelligence. *Cognitive systems research* 48. Elsevier: 1–3.
- List C and Pettit P (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Lowe R, Wu YI, Tamar A, et al. (2017) Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In: Guyon I, Luxburg UV, Bengio S, et al. (eds.) *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc., pp. 6379–6390.
- Lyons JB, Sadler GG, Koltai K, et al. (2017) Shaping Trust Through Transparent Design: Theoretical and Experimental Guidelines. In: *Advances in Human Factors in Robots and Unmanned Systems*, 2017, pp. 127–136. Springer International Publishing.
- Lyytinen K, Yoo Y and Boland RJ Jr (2016) Digital product innovation within four classes of innovation networks. *Information Systems Journal* 26(1). Wiley Online Library: 47–75.
- Mackenzie A (2017) *Machine Learners: Archaeology of a Data Practice*. MIT Press.
- MacKenzie D (2018) Material Signals: A Historical Sociology of High-Frequency Trading. *The American journal of sociology* 123(6). The University of Chicago Press: 1635–1683.
- MacKenzie D (2019) How algorithms interact: Goffman's 'interaction order' in automated trading. *Theory, Culture & Society* 36(2). SAGE Publications Sage UK: London, England: 39–59.
- Marcus G and Davis E (2019) *Rebooting AI: Building Artificial Intelligence We Can Trust*. Knopf Doubleday Publishing Group.
- Markus ML and Benjamin RI (1996) Change Agency - The Next IS Frontier. *MIS Quarterly* 20(4). Management Information Systems Research Center, University of Minnesota: 385–407.
- Markus ML and Rowe F (2018) Is IT changing the world? Conceptions of causality for information systems theorizing. *MIS Quarterly* 42(4). researchgate.net: 1255-1280;
- Martinez ME (2014) *Education as the Cultivation of Intelligence*. Routledge.
- Mason RO and Mitroff II (1973) A Program for Research on Management Information Systems. *Management science* 19(5). INFORMS: 475–487.
- Masterson M (2019) Protecting Election Infrastructure: A View from the Federal Level. In: Brown M, Hale K, and King BA (eds.) *The Future of Election Administration: Cases and Conversations*. Cham: Springer International Publishing, pp. 55–60.
- Mathiassen L (1998) Reflective systems development. *Scandinavian Journal of Information Systems* 10(1). aisel.aisnet.org: 12.

- Mestre JP (2006) Transfer of Learning from a Modern Multidisciplinary Perspective. IAP.
- Mingers J, Mutch A and Willcocks L (2013) Critical Realism in Information Systems Research. *MIS Quarterly* 37(3). Management Information Systems Research Center, University of Minnesota: 795–802.
- Mitchell T, Cohen W, Hruschka E, et al. (2018) Never-ending Learning. *Communications of the ACM* 61(5). New York, NY, USA: ACM: 103–115.
- Mnih V, Kavukcuoglu K, Silver D, et al. (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540). nature.com: 529–533.
- Möhlmann M and Henfridsson O (2019) What People Hate About Being Managed by Algorithms, According to a Study of Uber Drivers. *Harvard Business Review*, 30 August. Available at: <https://hbr.org/2019/08/what-people-hate-about-being-managed-by-algorithms-according-to-a-study-of-uber-drivers> (accessed 1 December 2019).
- Mortensen M and Haas MR (2018) Perspective—Rethinking Teams: From Bounded Membership to Dynamic Participation. *Organization Science* 29(2). INFORMS: 341–355.
- Moyer C (2016) How Google’s AlphaGo beat a Go world champion. *Atlantic* 28.
- Nair A, Srinivasan P, Blackwell S, et al. (2015) Massively Parallel Methods for Deep Reinforcement Learning. *arXiv [cs.LG]*. Available at: <http://arxiv.org/abs/1507.04296>.
- O’Hara M (2015) High frequency market microstructure. *Journal of financial economics* 116(2). Elsevier: 257–270.
- Ohlsson J, Händel P, Han S, et al. (2015) Process Innovation with Disruptive Technology in Auto Insurance: Lessons Learned from a Smartphone-Based Insurance Telematics Initiative. In: vom Brocke J and Schmiedel T (eds.) *BPM - Driving Innovation in a Digital World*. Cham: Springer International Publishing, pp. 85–101.
- Orlikowski WJ and Scott SV (2015) Exploring material-discursive practices. *Journal of management studies* 52(5). Wiley Online Library: 697–705.
- O’Rourke M (2017) The Kangaroo Conundrum. *Risk Management: An International Journal* 64(7). Risk and Insurance Management Society, Inc.: 40.
- Panait L and Luke S (2005) Cooperative Multi-Agent Learning: The State of the Art. *Autonomous agents and multi-agent systems* 11(3). Hingham, MA, USA: Kluwer Academic Publishers: 387–434.
- Parasuraman R, Sheridan TB and Wickens CD (2000) A model for types and levels of human interaction with automation. *IEEE transactions on systems, man, and cybernetics. Part A, Systems and humans: a publication of the IEEE Systems, Man, and Cybernetics Society* 30(3). [ieeexplore.ieee.org](http://ieeexplore.ieee.org): 286–297.
- Parker LD and Ritson PA (2005) Revisiting Fayol: Anticipating Contemporary Management. *British Journal of Management* 16(3). Wiley Online Library: 175–194.
- Pedro SDS and Hruschka ER (2012) Conversing Learning: Active Learning and Active Social Interaction for Human Supervision in Never-Ending Learning Systems. In: *Advances in Artificial Intelligence – IBERAMIA 2012*, 2012, pp. 231–240. Springer Berlin Heidelberg.
- Perrow C (2010) The meltdown was not an accident. In: Michael L and Paul MH (eds.) *Markets on Trial: The Economic Sociology of the U.S. Financial Crisis: Part A*. Research in the Sociology of Organizations. Emerald Group Publishing Limited, pp. 309–330.

- Qi L, Zhang C, Sukul A, et al. (2016) Automated Coding of Political Video Ads for Political Science Research. In: *2016 IEEE International Symposium on Multimedia (ISM)*, December 2016, pp. 7–13. [ieeexplore.ieee.org](http://ieeexplore.ieee.org).
- Ram A, Prasad R, Khatri C, et al. (2018) Conversational AI: The Science Behind the Alexa Prize. *arXiv [cs.AI]*. Available at: <http://arxiv.org/abs/1801.03604>.
- Redmiles D and Nakakoji K (2004) Supporting reflective practitioners. In: *Proceedings of the 26th International Conference on Software Engineering*, 2004, pp. 688–690. IEEE Computer Society.
- Rizzolatti G and Fogassi L (2014) The mirror mechanism: recent findings and perspectives. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 369(1644): 20130420.
- Robbins J (2010) GPS navigation... but what is it doing to us? In: *2010 IEEE International Symposium on Technology and Society*, 2010, pp. 309–318. IEEE.
- Rosen R (1991) *Life Itself: A Comprehensive Inquiry Into the Nature, Origin, and Fabrication of Life*. Columbia University Press.
- Rozario AM and Vasarhelyi MA (2018) How Robotic Process Automation Is Transforming Accounting and Auditing. *The CPA Journal*. [search.proquest.com](http://search.proquest.com). Available at: <http://search.proquest.com/openview/e07d840d837b53ca2d62cab251bb07ba/1?pq-origsite=gscholar&cbl=41798>.
- Rusu AA, Rabinowitz NC, Desjardins G, et al. (2016) Progressive Neural Networks. *arXiv [cs.LG]*. Available at: <http://arxiv.org/abs/1606.04671>.
- Salimans T, Ho J, Chen X, et al. (2017) Evolution Strategies as a Scalable Alternative to Reinforcement Learning. *arXiv [stat.ML]*. Available at: <http://arxiv.org/abs/1703.03864>.
- Santana E and Hotz G (2016) Learning a Driving Simulator. *arXiv [cs.LG]*. Available at: <http://arxiv.org/abs/1608.01230>.
- Sarker S, Chatterjee S, Xiao X, et al. (2019) The Sociotechnical Axis of Cohesion for the IS Discipline: Its Historical Legacy and its Continued Relevance. *MIS Quarterly* 43(3). [misq.org](http://misq.org): 695–719.
- Schlechner S (2017) Why do gas station prices constantly change? Blame the algorithm. *Wall Street Journal vom 8: 2017*.
- Schneider G (2018) Automating drug discovery. *Nature reviews. Drug discovery* 17(2). [nature.com](http://nature.com): 97–113.
- Schön DA (1983) *The Reflective Practitioner: How Professionals Think in Action*. Basic books.
- Scopino G (2015) Do Automated Trading Systems Dream of manipulating the Price of Futures Contracts-Policing Markets for Improper Trading Practices by Algorithmic Robots. *Florida law review* 67. HeinOnline: 221.
- SEC Staff (2014) Equity market structure literature review Part II: High frequency trading.
- Sergeeva A, Huysman M and Faraj S (2018) Losing Touch: How Robots Transform the Practice of Surgery. *Proceedings: a conference of the American Medical Informatics Association / ... AMIA Annual Fall Symposium. AMIA Fall Symposium 2018(1)*. Academy of Management: 11429.
- Shapiro D (2016) Accelerating the race to autonomous cars. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 415–415. ACM.

- Sibi S, Ayaz H, Kuhns DP, et al. (2016) Monitoring driver cognitive load using functional near infrared spectroscopy in partially autonomous cars. In: *2016 IEEE Intelligent Vehicles Symposium (IV)*, June 2016, pp. 419–425. [ieeexplore.ieee.org](http://ieeexplore.ieee.org).
- Siering M, Clapham B, Engel O, et al. (2017) A taxonomy of financial market manipulations: establishing trust and market integrity in the financialized economy through automated fraud detection. *Journal of Information Technology* 32(3): 251–269.
- Silver D, Hubert T, Schrittwieser J, et al. (2018) AlphaZero: Shedding new light on the grand games of chess, shogi and Go. *DeepMind.com*.
- Simon HA (1997) *Administrative Behavior* (4th expanded edition; 1947). The Free Press, NY.
- Singh K, Best PJ, Bojilov M, et al. (2013) Continuous auditing and continuous monitoring in ERP environments: Case studies of application implementations. *Journal of Information Systems* 28(1). American Accounting Association: 287–310.
- Sinz FH, Pitkow X, Reimer J, et al. (2019) Engineering a Less Artificial Intelligence. *Neuron* 103(6). Elsevier: 967–979.
- Smith A (1937) *The wealth of nations* [1776]. na. Available at: <https://mi01000971.schoolwires.net/cms/lib05/MI01000971/Centricity/Domain/440/Primary%20Source%20Articles%20Smith%20and%20Marx.pdf>.
- Smith BC (2019) *The Promise of Artificial Intelligence: Reckoning and Judgment*. MIT Press.
- Solaiman SM (2017) Legal personality of robots, corporations, idols and chimpanzees: a quest for legitimacy. *Artificial Intelligence and Law* 25(2). Springer: 155–179.
- Srnicek N and Williams A (2015) *Inventing the Future: Postcapitalism and a World Without Work*. Verso Books.
- Stroulia E and Goel AK (1995) FUNCTIONAL REPRESENTATION AND REASONING FOR REFLECTIVE SYSTEMS. *Applied artificial intelligence: AAI* 9(1). Taylor & Francis: 101–124.
- Su N, King JL and Grudin J (2017) Staying alive: The IS field at the half century mark. In: *The Routledge Companion to Management Information Systems*. Routledge, pp. 490–503.
- Sutherland JW (2008) Directive decision devices: Reversing the locus of authority in human–computer associations. *Technological forecasting and social change* 75(7). Elsevier: 1068–1089.
- Sutton RS and Barto AG (2018) *Reinforcement Learning: An Introduction*. MIT Press.
- Tan C, Sun F, Kong T, et al. (2018) A Survey on Deep Transfer Learning. In: *Artificial Neural Networks and Machine Learning – ICANN 2018*, 2018, pp. 270–279. Springer International Publishing.
- Taylor FW (1911) *The Principles of Scientific Management*. New York: Harper & Brothers.
- The Washington Post* (2019) Twitter’s new rules ban political ads from candidates but not ads about causes. 15 November. The Washington Post. Available at: <https://www.washingtonpost.com/technology/2019/11/15/twitters-new-political-ads-rules-bans-them-candidates-not-about-causes/> (accessed 25 December 2019).
- Thomas L (2017) At BlackRock, machines are rising over managers to pick stocks. New York Times. <https://www.nytimes.com/2017/03/28/business/dealbook/blackrock-actively-managed-funds-computer-models.html>. Accessed 28. [advtechconsultants.com](http://advtechconsultants.com). Available at: [https://advtechconsultants.com/Machines%20are%20Rising%20Over%20Managers%20to%20Pick%20Stocks\\_%20NYTimes%20Mar%202017.pdf](https://advtechconsultants.com/Machines%20are%20Rising%20Over%20Managers%20to%20Pick%20Stocks_%20NYTimes%20Mar%202017.pdf).

- Thorndike EL (1932) The fundamentals of learning. Teachers College Bureau of Publications. Available at: <https://psycnet.apa.org/record/2006-04535-000>.
- Tran D, Hoffman MD, Saurous RA, et al. (2017) Deep Probabilistic Programming. *arXiv [stat.ML]*. Available at: <http://arxiv.org/abs/1701.03757>.
- Treleven P and Batrinca B (2017) Algorithmic Regulation: Automating Financial Compliance Monitoring and Regulation Using AI and Blockchain. *Journal of Financial Transformation* 45. Capco Institute: 14–21.
- Trist EL (1981) The evolution of socio-technical systems: A conceptual framework and an action research program. Occasional Paper No. 2. Ontario Quality of Working Life Center.
- Valentino NA, King JL and Hill WW (2017) Polling and Prediction in the 2016 Presidential Election. *Computer* 50(5). [ieeexplore.ieee.org](http://ieeexplore.ieee.org): 110–115.
- Venkatesh V, Morris MG, Davis GB, et al. (2003) User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly* 27(3). Management Information Systems Research Center, University of Minnesota: 425–478.
- Vickers J (1985) Delegation and the Theory of the Firm. *The Economic journal of Nepal* 95(Supplement). Oxford University Press Oxford, UK: 138–147.
- Vinge V (1993) Technological singularity. In: VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute, 1993, pp. 30–31. [cs.umu.se](http://cs.umu.se).
- Vosoughi S, Roy D and Aral S (2018) The spread of true and false news online. *Science* 359(6380). [science.sciencemag.org](http://science.sciencemag.org): 1146–1151.
- Wang F, Zhang JJ, Zheng X, et al. (2016) Where does AlphaGo go: from church-turing thesis to AlphaGo thesis and beyond. *IEEE/CAA Journal of Automatica Sinica* 3(2). [ieeexplore.ieee.org](http://ieeexplore.ieee.org): 113–120.
- Wegner P (1998) Interactive foundations of computing. *Theoretical computer science* 192(2). Elsevier: 315–351.
- Wenger E (2014) Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge. Morgan Kaufmann.
- Wexler J, Pushkarna M, Bolukbasi T, et al. (2019) The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE transactions on visualization and computer graphics*. [ieeexplore.ieee.org](http://ieeexplore.ieee.org): 1–1.
- Winter S, Berente N, Howison J, et al. (2014) Beyond the organizational ‘container’: Conceptualizing 21st century sociotechnical work. *Information and Organization* 24(4). Elsevier: 250–269.
- Yampolskiy RV (2016) Taxonomy of pathways to dangerous artificial intelligence. In: *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, 2016. [aaai.org](http://aaai.org). Available at: <https://www.aaai.org/ocs/index.php/WS/AAAIW16/paper/viewPaper/12566>.
- Yates J (1988) Creating organizational memory: systematic management and internal communication in manufacturing firms, 1880-1920. Cambridge, Mass.: Sloan School of Management, Massachusetts Institute of .... Available at: <https://dspace.mit.edu/bitstream/handle/1721.1/47068/creatingorganiza00yate.pdf>.
- Yoo Y, Henfridsson O and Lyytinen K (2010) Research Commentary—The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. *Information Systems Research* 21(4). INFORMS: 724–735.
- Zheng L (nico), Albano CM, Vora NM, et al. (2019) The Roles Bots Play in Wikipedia. *Proc. ACM Hum. -Comput. Interact.* 3(CSCW). New York, NY, USA: ACM: 215:1-215:20.

Zuboff S (2015) Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology* 30(1). Springer: 75–89.

Zuboff S (2019) *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Profile Books.