

## *Research Perspectives:* **The Rise of Human Machines: How Cognitive Computing Systems Challenge Assumptions of User-System Interaction**

Sebastian Schuetz<sup>1</sup>, Viswanath Venkatesh<sup>2</sup>

<sup>1</sup>University of Arkansas, USA, [swschuetz@outlook.com](mailto:swschuetz@outlook.com)

<sup>2</sup>University of Arkansas, USA, [vvenkatesh@vvenkatesh.us](mailto:vvenkatesh@vvenkatesh.us)

### Abstract

Cognitive computing systems (CCS) are a new class of computing systems that implement more human-like cognitive abilities. CCS are not a typical technological advancement but an unprecedented advance toward human-like systems fueled by artificial intelligence. Such systems can adapt to situations, perceive their environments, and interact with humans and other technologies. Due to these properties, CCS are already disrupting established industries, such as retail, insurance, and healthcare. As we make the case in this paper, the increasingly human-like capabilities of CCS challenge five fundamental assumptions that we as IS researchers have held about how users interact with IT artifacts. These assumptions pertain to (1) the direction of the user-artifact relationship, (2) the artifact's awareness of its environment, (3) functional transparency, (4) reliability, and (5) the user's awareness of artifact use. We argue that the disruption of these five assumptions limits the applicability of our extant body of knowledge to CCS. Consequently, CCS present a unique opportunity for novel theory development and associated contributions. We argue that IS is well positioned to take this opportunity and present research questions that, if answered, will lead to interesting, influential, and original theories.

**Keywords:** Cognitive Computing Systems (CCS), Intelligent Agents (IA), Artificial Intelligence (AI), Expert Systems (ES), Assumptions, Research Agenda

Allen Lee was the accepting senior editor. This research perspectives article was submitted on July 31, 2018, and underwent two revisions.

## 1 Introduction

*Cognitive computing systems* (CCS) are new types of systems that mimic human cognitive abilities (Maresca & Stanganelli, 2016). The capabilities are impressive—from IBM Watson beating the world's best Jeopardy! player in 2011<sup>1</sup> to powering services that allow businesses to digitally replace customer service agents in 2016, CCS are clearly capable of unprecedented feats. The consumer market for CCS has also been booming. For example, the global market

for smart speakers has recently exploded into a US\$4.5 billion dollar industry and is expected to reach US\$40 billion by 2024 (Global Market Insights Inc., 2018). Smart speakers are just one example of CCS with scores of other CCS applications being developed that are slowly but surely disrupting the manifold spaces of our daily lives (Marr, 2016). Whereas previous technological advancements made systems more powerful, more connected, and/or more mobile, CCS are a uniquely disruptive advancement that aims at making machines *more human*.

<sup>1</sup> A recording of the show can be found at: [https://www.youtube.com/watch?v=WFR3IOm\\_xhE](https://www.youtube.com/watch?v=WFR3IOm_xhE)

CCS fundamentally challenge our long-held beliefs about what falls into the realm of human ability and what is machine capability. Fueled by advances in AI, CCS are capable of various human feats such as perception and learning with significant implications. The new cognitive capabilities allow CCS to enter domains that have remained exclusive to humans. For instance, Amazon Alexa already allows customers to order from Amazon's e-commerce platform using natural speech. No longer do users need to rely on artificial interfaces (e.g., monitor, mouse, and keyboard), but instead may interact with machines as they would with other humans. This dramatically blurs the lines between the (thus far) clear-cut fronts of human abilities and computer capabilities.

Meanwhile, CCS disrupt our beliefs about what machines can and cannot do; the IS literature still maintains the traditional notion that systems are tools with some consistent functionality that can be used by humans to generate some outcome (Benbasat & Zmud, 2003). This tool perspective is associated with many assumptions that dictate how we as IS researchers think about how humans use IT artifacts and how IT artifacts generate outcomes. For example, we generally assume that humans use IT artifacts through an artificial interface (e.g., a touch-enabled display), or we assume roles that define humans as users and IT artifacts as tools. This view, and these assumptions, has largely not changed (Demetis & Lee, 2018). Maybe as a consequence, an incremental research paradigm has developed and plagued IS research (Grover & Lyytinen, 2015).

As we make the case in this paper, CCS fundamentally challenge these and other assumptions. We believe that these assumptions no longer hold for systems with more human-like capabilities, such as CCS. Humans are not only inherently prone to anthropomorphism, the attribution of human-like characteristics to inanimate objects and animals, but human-CCS interaction can *actually* resemble human interactions to an as yet unprecedented degree. In fact, research has already shown that humans relate to these machines more like humans than objects (Aleksander, 2004; Lankton, McKnight, & Tripp, 2015; Schroeder & Epley, 2016; Waytz, Haefner, & Epley, 2014). As we argue here, the development of CCS is the result of a clear progression toward more human-like capabilities. CCS can thus not be classified as a technological fad, but as machines capable of human-like interactions. Consequently, CCS have ushered in a paradigm change regarding human-machine interactions, thereby rendering the artifact-based paradigm of IS obsolete (Alter, 2015; Demetis & Lee, 2018). The emergence of CCS opens an entire domain of research questions that cannot yet be answered with our existing theories. As research on CCS enters uncharted territory, the potential boundary conditions of our existing knowledge, tied to

key underlying assumptions we have made, present an opportunity to develop novel theories that are influential and interesting (Alvesson & Sandberg, 2011; Weick, 1989). We believe that IS research is well-positioned to exploit this opportunity to generate original theories that ultimately make a difference in our daily interactions with CCS.

Despite increasing public and commercial interest in CCS, research on cognitive computing remains absent from the IS Basket journals (see Appendix A). If at all, extant IS research has demonstrated an interest in the underlying technologies of CCS, such as machine learning (Greenwald, Kannan, & Krishnan, 2010; Li et al., 2009; Mayer et al., 2014). Aside from two studies (Aleksander, 2004; Lankton et al., 2015), the actual implications of a system's cognitive capabilities remain largely unexplored. Recently, related topics have gained some traction at our most prestigious conferences (e.g., Rzepka & Berger, 2018; Wuenderlich & Paluch, 2017). However, these studies are still in the early stages of inquiry into CCS and thus treat CCS from traditional perspectives or as an isolated technological instantiation (e.g., as conversational agents). A fundamental understanding of how CCS differ from preceding systems and how CCS' human-like capabilities question the applicability of our existing knowledge base is missing but needed if IS scholars want to take advantage of this unique opportunity to develop novel, impactful theories.

Against this backdrop, we discuss the singular opportunity that CCS presents to IS research. Specifically, we discuss the novel capabilities and characteristics of CCS and investigate why CCS represent a permanent, progressive development. Then, we discuss how the unique capabilities of CCS challenge five traditional IS assumptions, rooted in discussions about assumptions underlying research (Alvesson & Sandberg, 2011) about the user-IT artifact interaction and subsequently illustrate how challenging these assumptions requires the development of new theories that render existing ones inapplicable. To aid the development of novel theories on CCS-specific phenomena, we propose several research questions that we believe will be of interest in the future. We thus hope to break ground for IS research to leverage this unique opportunity for conducting research that will ultimately impact the lives of individuals, organizations, and society.

## 2 Background

### 2.1 The Emergence of CCS

CCS represent the culmination of a long tradition devoted to creating machine capabilities (Figure 1) equivalent to or better than human abilities in certain

areas (Rich & Knight, 1991). Early efforts exploited computers' superior processing speed and memory systems to create machines that would be better at retaining and aggregating data. As such, *decision support systems* (DSS) were developed, referring to systems that employ "decision rules and models, coupled with an extensive database" (Turban & Watkins, 1986, p. 122). DSS allowed decision makers to query systems to produce factual information in the form of aggregated data, reports, or even charts (Turban & Watkins, 1986). However, it was still up to the decision makers to draw the inferences from those data. Thus, the next step in the development of machine capabilities was to devise reasoning capabilities. Consequently, *expert systems* (ES) were developed, which are systems that allow for "propagating inferences over the knowledge base" (Turban & Watkins, 1986, p. 122). The reasoning of expert systems allowed ES to mimic human experts (Turban & Watson, 1988) by providing explanations for given recommendations.

Research on DSS and ES dominated early IS research up until 1996 (Nevo, Nevo, & Ein-Dor, 2008) at which point several problematic issues with these systems became clear. Despite their technological capabilities and economic success, many of these systems were quickly abandoned by users (Gill, 1995). One of the key challenges was that these systems required structured information to interface with human users (Sviokla, 1990). Thus, humans had to adapt to the systems to formulate information and problems in ways that the computer would understand (Paradice & Courtney, 1987). As it turns out, this was problematic because users often did not provide adequate data (Kopsco, Pipino, & Rybolt, 1988) and thus the systems

often arrived at different conclusions than their human users (Paradice & Courtney, 1987). Consequently, these systems were gradually abandoned because they relied on uncooperative users.

The next development of machine capabilities empowered systems to operate autonomously from users. To that end, *intelligent agents* (IA) were developed, referring to a software that "acts 'intelligently' and 'in the place of' a human to perform a given task" (March, Hevner, & Ram, 2000, p. 334). With the power to autonomously react to and stimulate their environments (Russell & Norvig, 2010), intelligent agents were no longer reliant on human decision makers. Rather, intelligent agents could now autonomously serve human purposes. For example, some research investigated the utility of IA for placing bids in auction markets (Adomavicius, Gupta, & Zhdanov, 2009), facilitating interorganizational meetings (Glezer, 2003), and identifying the malicious intentions of border-crossing individuals (Nunamaker et al., 2011).

The aggregation of all of these capabilities meant that machines were capable of knowing, reasoning, and autonomously (re)acting. With these capabilities, machines possess significant human-like abilities. However, machines still suffered from the caveat that they were inherently reliant on structured data input, making it difficult for users to interact with them. Thus, logically, in the pursuit to build human-like computers, it became evident that machines needed capabilities that would allow them to make sense of their unstructured environments. In other words, they required cognitive capabilities.

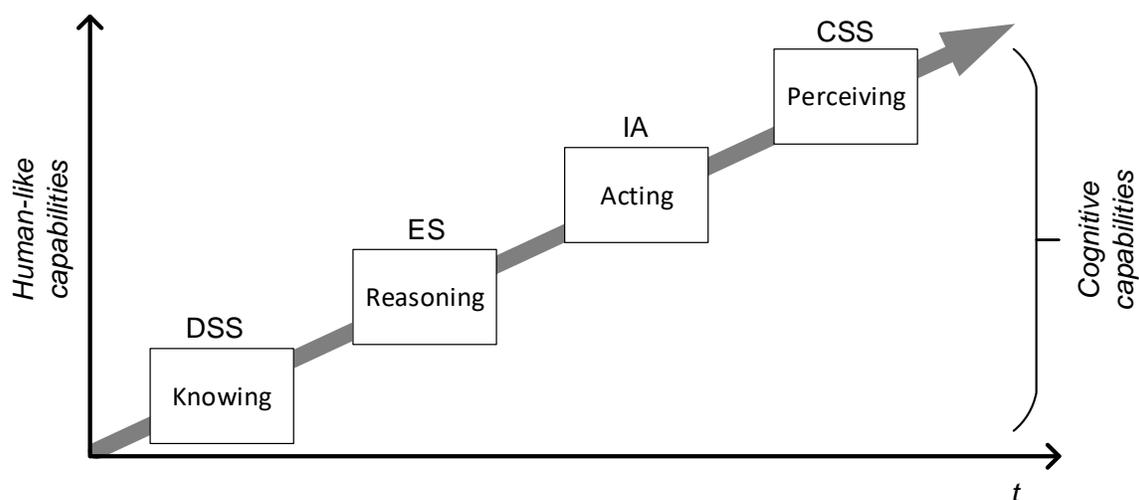


Figure 1. Progression of Machine Capabilities

**Table 1. Characteristics of CCS**

Characteristic	Description from the CCS Consortium	Example of a CCS
(1) Adaptive	They must learn as information changes and as goals and requirements evolve. They must resolve ambiguity and tolerate unpredictability. They must be engineered to feed on dynamic data in real time or near real time.	Google Maps changes its best route recommendations based on real-time traffic information.
(2) Interactive	They must interact easily with users so that users can define their needs comfortably. They may also interact with other processors, devices, and cloud services.	Amazon Alexa interacts using natural language.
(3) Iterative and stateful	They must aid in defining a problem by asking questions or finding additional source input if a problem statement is ambiguous or incomplete. They must “remember” previous interactions in a process and return information that is suitable for the specific application at that point in time.	Microsoft Cortana can identify problems when creating a new event, Apple’s Siri will ask for missing information.
(4) Contextual	They must understand, identify, and extract contextual elements, such as meaning, syntax, time, location, appropriate domain, regulations, user’s profile, process, task, and goal. They may draw on multiple sources of information, including both structured and unstructured digital information, as well as sensory inputs (visual, gestural, auditory, or sensor-provided).	Apple’s Siri incorporates contextual information; for example, when users search for restaurants, the result will be dependent on the users’ location.

Cognitive capabilities were achieved with recent advances in artificial intelligence that allow machines to perceive their environments. Traditionally, processing unstructured data such as text documents and audio-visual inputs has been understood as an exclusively human ability. However, with the development of more powerful machine learning techniques, machines have finally become capable of clustering, classifying, and making sense of the unstructured data that describe the world in which we live. To CCS, pictures, speech, and texts are comprehensible. These cognitive capabilities of CCS rely on a combination of new and existing capabilities (Figure 1). Theories of cognitive architecture, such as Soar, prescribe that for systems to have cognitive capabilities, they must have components that provide memory, reasoning, action, and perceptive capabilities (Laird, Newell, & Rosenbloom, 1987). CSS are the first systems to possess all of these and are thus the first generation of machines with cognitive capabilities.

The cognitive capabilities of CCS present a clear progressive development that has resulted in continuous additions to an existing capability base. As such, the current capability base renders machines now capable of cognition. This level of capability is permanent; as such, CCS represent enduring phenomena. At this point, there is no reason to believe that human beings will abandon their ability to create machines with cognitive capabilities enabling the capacity to process unstructured data and interact in a

more human-like fashion. On the contrary, these unique capabilities will allow CCS to enter even more domains of human life. Although CCS may not replace each and every technology that we currently employ, it is certain that CCS are here to stay.

## 2.2 The Interactive Characteristics of CCS

With these cognitive capabilities, CCS mimic human-like abilities at an unprecedented level. As defined by the Cognitive Computing Consortium (2014),<sup>2</sup> CCS are systems that are (1) *adaptive*, as they must learn from changing information, goals, and requirements; (2) *interactive*, as they must interact with humans and other systems easily; (3) *iterative and stateful*, as they must be able to further narrow down a problem until understood and must remember previous interactions; and (4) *contextual*, as they must consider contextual elements (Table 1).

## 3 Challenging Assumptions

It is precisely these characteristics that enable CCS to engage in new types of user-system-environment interactions. We argue that these new types of interactions break with the traditional assumptions that we as IS researchers have held about user-system interactions. Because CCS challenge these assumptions, they afford IS the opportunity to create influential and, ultimately, interesting original theories

<sup>2</sup> Member organizations include BA-Insight, Babson College, Basis Technology, Cognitive Scale, CustomerMatrix, Decision Resources, Ektron, Google, HP Autonomy, IBM,

Microsoft/Bing, Next Era Research, Oracle, Pivotal, SAS, Saxena Foundation, Synthexis, and Textwise/IP.com.

(cf. Alvesson & Sandberg, 2011; Bartunek, Rynes, & Duane Ireland, 2006). To that end, it is necessary to understand how the new interactions challenge existing assumptions that have guided the inquiry into IS phenomena. However, delineating existing assumptions is difficult because existing assumptions are rarely formulated in the literature; consequently, they are rarely disputed or actively discussed (Alvesson & Sandberg, 2011). To attempt identifying, articulating, and challenging the preexisting assumptions that have governed IS research, we follow the recommendations of Alvesson and Sandberg (2011). We first begin by describing the traditional research paradigm on user-system interactions and then contrast it against an IS-external perspective to identify and articulate the existing assumptions.

### 3.1 Traditional View and Assumptions

For decades, technologies such as DSS, ES, and IA have been viewed as tools. The tool perspective is embodied in the terms “IT artifact” or “technology artifact,” defining technology as “a human-created *tool* whose *raison d’être* is to be used to solve a problem, achieve a goal or serve a purpose that is human defined, human perceived or human felt” (Lee et al., 2015, p. 8, emphasis added). Figure 2 illustrates this perspective. According to this tool perspective, IS scholars were compelled to study how the *use* of IT artifacts would lead to *impacts* (Benbasat & Zmud, 2003). Thus, a user would *use* (A) artifacts to *impact* (B) some outcome, similar to how a gardener’s use of garden shears would lead to trimmed hedges. Notwithstanding that this perspective has recently been challenged (Alter, 2015; Demetis & Lee, 2018; Lee, Thomas, & Baskerville, 2015), its basic assumptions about how IS scholars think about user interactions with IT artifacts remain intact.

To discern these assumptions, we turn to the basic model of human-computer interaction (Norman, 1986; 1988). Figure 3 shows an adapted version of the basic human-computer interaction model using IS terminology. In this model, users interact with IT artifacts through an artificial interface. Through this interface, the interaction develops through the user providing some input to the IT artifact and the IT artifact returning some output to the user.

Over decades of studying user-artifact interaction using this model, researchers have developed several implicit assumptions based on expectations that relate to the user, artifact, and the interaction between the two, as defined by interfaces that facilitate the exchange of inputs and outputs (Table 2). Specifically, we believe that various IS research areas have made assumptions with regard to *who* typically holds the role

of the user (Assumption 1), *who* defines the inputs to a system (Assumption 2), *how* a system produces outcomes (Assumption 3), *whether* humans can understand how systems arrive at these outcomes (Assumption 4), and *whether* there is always a computer interface between humans and systems (Assumption 5). Alvesson and Sandberg (2011) suggest that assumptions can range in scope from in-house assumptions to field assumptions. While in-house assumptions are shared by a specific school of thought, field assumptions are shared across multiple schools of thought and sometimes even across paradigms and academic disciplines. In the following sections, we will show how each assumption is shared across multiple IS research streams, in which case they constitute field assumptions. Moreover, for each of these assumptions, we will also make the case that the unique capabilities of CCS render these assumptions obsolete.

### 3.2 Assumption 1: Unilateral Relationship between User and Artifact

The first assumption of IS research in this model of user-artifact interaction is that there is a human user and that there is a clear unilateral relationship between the user and the artifact, as evident in the term “user.” This assumption is prevalent in the IS field. For example, in system use research, scholars explore the use of IT by individuals, groups, and organizations (Straub, Limayem, & Karahanna, 1995) and in privacy research, scholars study when and how users disclose information to systems (e.g., Chen & Sharma, 2013; Lowry, Cao, & Everard, 2011; Smith, Dinev, & Xu, 2011).

CCS break this assumption because they are interactive. The capability of CCS to easily interact with users allows CCS to use them for their own purposes. Thus, CCS are no longer simple tools and users are no longer simple users. Rather, CCS and users form complex systems in which artifacts use users to achieve their objectives (Alter, 2015; Demetis & Lee, 2018). This does not mean that CCS are no longer tools that serve a given purpose, but that CCS are no longer tools to the users but active agents that may act in their own interests. Examples of how users can be used by machines to achieve their objectives include Twitter bots that autonomously create content and shape public opinion (Markoff, 2017) and robot callers that call citizens to defraud them of their savings (Kiro, 2018). Thus, in contrast to the widely held assumption that artifacts are tools to users, users might also be tools to CCS (Demetis & Lee, 2018).

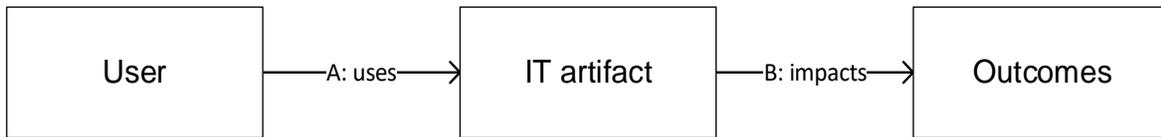


Figure 2. Traditional “Tool” View on User-Artifact Interaction

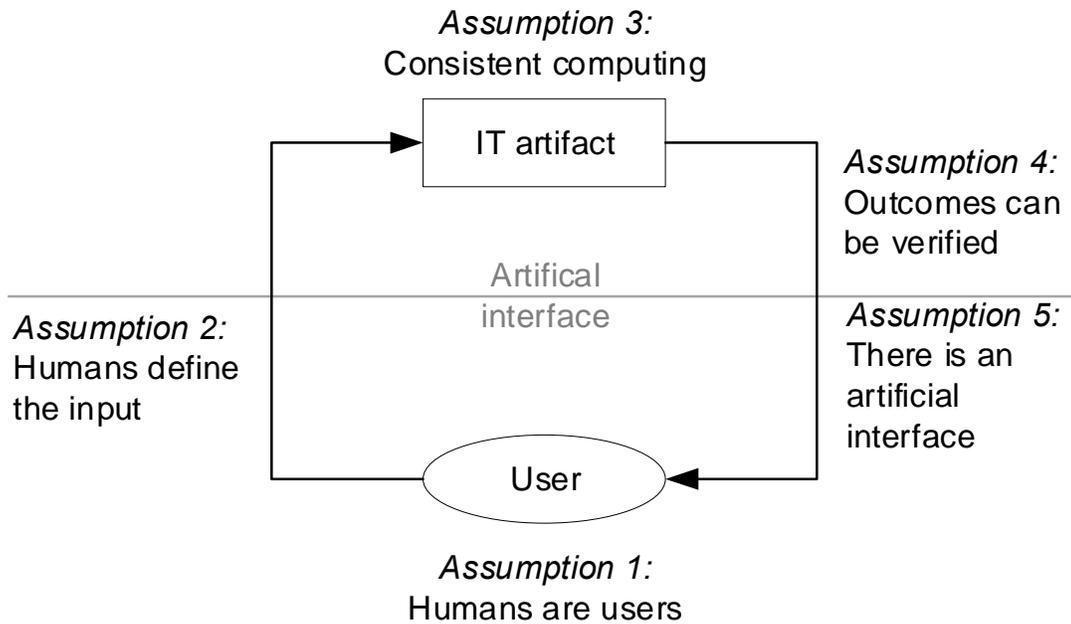


Figure 3. Assumptions of User-Artifact Interaction

Table 2. How CCS Challenge IS Assumptions

Element	IS Assumption	Assumption label	Scope	Challenged by CCS characteristic
User	1. Humans are users	Unliteral relationship	Field assumption (e.g., system adoption, privacy)	Interactive
Input	2. The developer defines the inputs	Ignorance of environment	Field assumption (e.g., system adoption, privacy, communication)	Adaptive, interactive
Computation	3. IT artifact use leads to consistent outcomes	Functional consistency	Field assumption (e.g., IS success, IT governance, IS development)	Adaptive
Output	4. The way the tool derives its outcomes is comprehensible and can be verified	Functional transparency	Field assumption (e.g., recommendation systems, IS development)	Contextual, adaptive
Interface	5. There is an artificial interface	Awareness of use	Field assumption (e.g., privacy, service science)	Interactive, iterative and stateful, and context aware

### **3.3 Assumption 2: Artifacts' Ignorance of the Environment**

The second assumption of IS research is that machines are generally isolated from their environments. This is evident in that most machines only generate outcomes in machine-specific environments (Demetis & Lee, 2018). For example, the direct outcomes of using an ERP system are mainly the manipulation of data and sometimes the generation of instructions capable of manipulating shop floor machinery. Another example would be the use of SPSS, the direct outputs of which would be constrained to the laptop or desktop environment. Further, machines are also generally unable to receive inputs from their environments unless explicitly enabled to do so. The only inputs that systems typically receive are from human-computer interfaces, sensors, or other computer interfaces. These interfaces are highly specific in terms of what information is received. Overall, the general notion is that systems operate in ignorance of their environments.

CCS challenge this assumption because they are adaptive and interactive. By being capable of adapting to changing information or by reacting to interactions, CCS can be stimulated by their environments. This argument is best illustrated by Demetis and Lee (2018), who describe how autonomous selling algorithms interacted with each other to lead to the 2010 Dow Jones Flash Crash. Demetis and Lee (2018) thus argue that some systems can generate outcomes and recursively react to themselves. Further, CCS can process unspecific input—that is, any kind of video or audio data. These rich, unspecific input streams can provide CCS with information that was not previously specified by its developers. A prominent example is a recent case in which Amazon Alexa served as a witness to a murder (Whittaker, 2018). Other examples include Alexa ordering goods from Amazon after hearing its name on TV (Liptak, 2017). Thus, CCS break with this assumption because of their general capability to react to their environments.

### **3.4 Assumption 3: Functional Consistency**

The third assumption of IS research is that the functionality of systems is directed by software owners and vendors. Using that logic, research has studied the external events that trigger organizations to expand the functional base of their digital infrastructures (e.g., Henfridsson & Bygstad, 2013) and how platform vendors control the evolution of their software products (Tiwana, Konsynski, & Bush, 2010). Because IS research sees systems as exhibiting stable functionality, functional deficiencies are mostly attributed to IT governance weaknesses (e.g., Benaroch & Chernobai, 2017). Thus, whether systems function as anticipated is largely a product of functional fit (Goodhue & Thompson, 1995) and utilization (Devaraj & Kohli,

2003). Due to the assumed stability of function, studies have placed a single system in different (e.g., cultural) contexts to explore context effects (e.g., Lowry et al., 2011).

In contrast to this assumption, CCS are adaptive. This trait, enabled by machine learning, obscures the assessment of their functional reliability. Through normal operation alone (i.e., learning from increasing amounts of data and receiving feedback), CCS acquire new functionalities—functionalities that are enabled by connections between thousands of artificial neurons spanning a network of dependencies that humans cannot truly understand. A DARPA director has hence labeled current advances statistically impressive but individually unreliable (Launchbury, 2017). An example of such adaptive behavior of a CCS would be Tay, a chat bot developed and deployed by Microsoft, that turned racist in less than 24 hours, based only on interactions it had with human counterparts (Vincent, 2016). Although this problem already persists with systems that operate in controlled environments (e.g., only feed on selected training data), the issue is likely to be more significant as CCS become exposed to unprecedented data streams that enable new and unforeseen behaviors.

### **3.5 Assumption 4: Functional Transparency**

The fourth assumption of IS research relates to the transparency of the tool's functionality. The common assumption is that users are biased (i.e., cognitively biased) but that artifacts are largely objective under the conditions of their programming (e.g., Paradice & Courtney, 1986). As such, artifacts derive their outcomes from a combination of logical rules and mathematical operations. The capability of machines to calculate correctly when given accurate and unbiased information is thus a key driver of trust in technology (Lankton, McKnight, & Thatcher, 2014). Researchers in the decision support systems (DSS) literature anticipate that the logical inferences drawn by systems such as expert systems create unbiased recommendations for users (Paradice & Courtney, 1987; Remus & Kottemann, 1986). Along those lines, they have even proposed a system to logically validate the biased knowledge of expert managers (Paradice & Courtney, 1986).

More importantly, CCS break with this assumption because they are contextual and adaptive. In order to achieve that adaptability, their protocol of calculation has shifted from deterministic (i.e., being programmed to calculate if-then clauses) to probabilistic (i.e., training neural networks to choose the most likely accurate answer). Thus, the results derived from a CCS system are derived from complex statistical models. Consequentially, CCS can incorporate many contextual factors without the knowledge of developers and users.

This implies that understanding how a CCS arrives at its outcomes is often not easily comprehensible. There have been extreme cases in which substantive biases in training data have caused surprising yet outrageous outcomes that are not in line with the concept of correct computing. For example, a Google algorithm consistently classified black people as gorillas, an effect based on the biased training set of white engineers (Barr, 2015). Other algorithms drawn from training data have developed a sexist view of women, as pictures depict women in kitchens more often than men (Simonite, 2017). The *New York Times* speaks of AI's "white guy problem."<sup>3</sup> The issue is not trivial. The related scientific discipline speaks of a "black box" problem (Castelvecchi, 2016; Russell & Norvig, 2010) and illustrates that such systems may behave in unexpected ways. Indeed, the unpredictability of AI-based systems is a key problem that DARPA and other research institutions are trying to solve (Launchbury, 2017; Robertson, 2017). Thus, unlike deterministic computing, CCS are inherently challenged in delivering the intended results. Although there is currently a lull in DSS research, it might experience a revival as CCS become increasingly important for health diagnostic purposes such as cancer detection and treatment (Metz, 2017).

### 3.6 Assumption 5: Users' Awareness of Artifact Use

The fifth assumption of IS research relates to user awareness. To date, research has assumed that users are aware that they are using an artifact because they are interacting with a machine-specific, artificial interface. This assumption is foundational to several IS research streams. Consider, for example, the technology acceptance stream (for an overview, see Venkatesh, Thong, & Xu, 2016) that builds on users' perceptions of technological characteristics, or SERVQUAL (Devaraj, Fan, & Kohli, 2002; Jiang, Klein, & Carr, 2002; Pitt, Watson, & Kavan, 1997) that builds on the foundational assumption that users are aware of their service use to form perceptions of service quality. Similarly, in privacy research, information disclosure is often studied from a rational choice perspective, i.e., privacy calculus (Smith et al., 2011), through examining what users choose to disclose (Anderson & Agarwal, 2011; Krasnova et al., 2010). This assumes that users know what they are using and disclosing. Overall, the current research paradigm assumes that users use of artifacts is intentional and deliberate.

This assumption is challenged by the CCS ability to authentically mimic human-like interactions. By being interactive, CCS can engage in human-like back-and-forth conversations. These conversations seem even

more authentic because CCS are iterative and stateful, capabilities that allow CCS to ask questions to specify a problem and remember previous answers. Finally, by considering context factors (e.g., time, location), CCS may render it substantially more difficult for human users to differentiate a CCS from a human agent—a fact that has recently been unwittingly demonstrated by Georgia Tech students who interacted with an online bot believing it was a teaching assistant (Etzioni, 2017). Beyond chat, other service providers already sell voice-enabled solutions (e.g., IBM Watson Virtual Agent or Nuance Conversational Interactive Voice Response). Although this scenario remains rare, since Alan Turing, computer scientists have long held the goal of creating capabilities that would make it impossible to differentiate an intelligent machine from a human being. With a projected uptake in CCS capabilities, we expect that such cases will become increasingly common. Especially for computer-mediated channels, such as service chat bots or voice calls, CCS can operate hidden—without users' knowledge of their presence. Thus, the prevailing assumption that users are aware and know when they are using artifacts is becoming increasingly obsolete.

## 4 Implications for IS Research

The five challenged assumptions are of great importance to IS research, as they pertain to the very nature of IS inquiry: the study of the development, use, and management of IS. As such, these field assumptions have been the keystones to many of the most influential scientific advancements in IS: for example, whether scholars theorized about system success (e.g., Delone & McLean, 2003), adoption (e.g., Venkatesh, Morris, Davis, & Davis, 2003) or computer-mediated communication (e.g., Dennis, Fuller, & Valacich, 2008), their work has been based on the conventional wisdom that systems are tools and humans are users (Assumption 1), that systems are generally unaware of their environments (Assumption 2), that systems are functionally consistent (Assumption 3), that systems operate as expected (Assumption 4), and that human users are aware of their interactions with systems (Assumption 5). Thus, these field assumptions have, over time, established themselves as *truths* because the artifacts under study did not change dramatically (see Nevo et al., 2008).

Generally, whenever one of these assumptions has been previously challenged, the resulting papers have been deemed interesting and influential and published in our most prestigious journals. Examples include Orlikowski and Scott (2008), Riemer and Johnston (2017), Carter and Grover (2015), and Demetis and Lee (2018), who challenged the unilateral relationship between users and

<sup>3</sup> [https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html?\\_r=0](https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html?_r=0)

systems; and Liu, Santhanam, & Webster (2017) and Polites and Karahanna (2013), who challenged the assumption that users are necessarily aware of their interactions with systems using gamified immersive experiences or habituated use. However, these papers all challenge one or more of the delineated assumptions based on specific contexts (e.g., organizational systems use or hedonic video games) but never because of a major technological advancement.

#### 4.1 Implications for Existing Theories

CCS challenge all of the aforementioned assumptions at once: CCS present a break with the incremental mode of technological advancement by advancing machines with a new kind of capability, cognition. It is beyond the scope of this paper to identify all the theories and research streams that are affected by this paradigm change. Nevertheless, the implications are dramatic. We discuss an example as a means of sparking further dialog in different steams: Consider the concept of computer self-efficacy (CSE), which is a key construct in IS research, and also an adaptation of reference theory (Bandura, 1977). The paper contextualizing general CSE by Compeau and Higgins (1995) is among the most cited papers in IS research and CSE has been shown to determine perceived ease of use (Agarwal, Sambamurthy, & Stair, 2000; Venkatesh, 2000), a core component of one of the most influential IS theories (Moody, Iacob, & Amrit, 2010). In fact, many influential theories use some sort of contextualized self-efficacy construct to capture an individual's ability to engage with and use systems—for example, Agarwal and Karahanna (2000) use self-efficacy in their study that underpins the current emerging stream of gamification, and Liang and Xue (2009) build on self-efficacy in their security threat avoidance model.

Yet for CCS, this core concept is irrelevant: self-efficacy is grounded in the assumption that individuals face systems and make judgments about their ability to perform a specific behavior within a specific system (i.e., to complete a task). However, for CCS, the user might not even be aware of their system use and their system use might be driven by the system's efficacy to facilitate an interaction. For instance, the Georgia Tech students did not consider their abilities to use a self-service system even for a moment when they were chatting with what they thought to be their teaching assistant (Etzioni, 2017). Moreover, if they were aware, using a CCS requires only the (minimal) skill of using natural language—written or spoken. How can self-efficacy be of primary concern when system use becomes as simple as ordering a pizza over the phone? Evidently, if system use for CCS often just means the ability to speak, then CSE is irrelevant to CCS research.

The downstream implications of changes in field assumptions can be dramatic: when core concepts such

as CSE become obsolete, then related theories lose their predictive power or even their validity. To expose this radical shift in perspective, consider the implications of challenging this assumption for some of the core IS theories: the technology acceptance model (TAM) (Venkatesh et al., 2016) or SERVQUAL (a theory of service quality borrowed from marketing) (Jiang et al., 2002; Pitt et al., 1997). Without users needing to consciously think about their system use, how useful is the concept of ease of use for TAM? How valid is the contextualization of SERVQUAL to information systems? Moreover, how can we even study CCS use intentions?

Evidently, the implications of challenging just one assumption are vast. However, the disruption of the institutionalized line of reasoning holds wide-ranging implications (Sandberg & Alvesson, 2010): any theory, concept, or method developed based on these legacy assumptions may not be applicable to CCS and, at the very least, require reevaluation. Thus, CCS provide an abundance of opportunities for IS research to reconceptualize and reevaluate existing phenomena in light of this new technology. An example is Lankton et al.'s study (2015), which recontextualized the trust-in-technology construct to more human-like systems. With five basic assumptions of human-system interaction being challenged, ample further research is needed to understand the specific implications for existing knowledge in various domains.

#### 4.2 CCS: A Unique Opportunity for Novel Theorizing

More importantly, the new modes of CCS-user interaction create a space for the creation of novel, original theories. While reevaluating existing theories will only help to answer *old* questions in light of this new technology, CCS pose an entirely *new* set of problems to the scientific community that, with our current knowledge, cannot be answered. One could argue that CCS is a phenomenon in which normal science ends and a paradigm shift occurs (Kuhn, 2012). While much of our extant knowledge was developed in a paradigm that focused on how IT can add value to organizations, we now live in an age in which such problems are well understood. Virtually any business recognizes the value of IT. In contrast, CCS ring in a new age of more human-like systems that bring about an unprecedented set of challenges. Never before have machines spoken like humans, driven cars, or identified pedestrians jaywalking. In recognizing the unique characteristics of CCS, IS research has the opportunity to expand its focus beyond questions like “How can IT increase the bottom line of an organization?” toward questions that probe how more human-like systems change the way we work, drive, collaborate, create, innovate, and live.

**Table 3. Directions for CCS-Specific Theorizing**

CCS assumption	Theoretical gap	Research question
1. Bilateral relationships	Bilateral relationships open new avenues for CCS-based <i>persuasion</i> and <i>collaboration</i> research.	RQ1: How can CCS effectively persuade people to follow system advice and orders? RQ2: What are the effects of CCS advice on individuals? RQ3: How can individuals and groups effectively collaborate with CCS? RQ4: What are the prerequisites of successful user-CCS collaboration?
2. Awareness of the environment	Awareness of the physical and digital environment creates new possibilities for CCS-based <i>augmentation</i> and <i>surveillance</i> .	RQ5: How can CCS effectively augment human decision-making? RQ6: What are the effects of subjection to CCS-based surveillance on individuals and society?
3. Functional adaptivity	Functional adaptivity creates new opportunities for CCS to <i>adapt their behavior</i> .	RQ7: What will be the key outcomes (positive and negative) of functional adaptivity in different contexts? (e.g., increase in user satisfaction, productivity, reliance) RQ8: What extent of functional adaptivity will be desirable for CCS to satisfy the dual outcomes of system and user success?
4. Functional opacity	Functional opacity affects the <i>economics</i> of developing CCS and how humans <i>rely on</i> and <i>trust</i> CCS.	RQ9: What are the economics of CCS development? RQ10: When can humans safely rely on CCS? RQ11: How can CCS foster humans' trust?
5. Unawareness of use	Humans' potential unawareness of use creates new phenomena around CCS <i>deception</i> and <i>substitution</i> of humans.	RQ12: How can CCS effectively deceive humans? RQ13: When will it be beneficial for CCS to keep users unaware of their system use, and when will it be detrimental? RQ14: What are the effects of substituting humans with CCS on individuals, organizations, and society?

CCS present unique opportunities for IS research to escape the dogma of gap-spotting (Alvesson & Sandberg, 2011) and scripted research (Grover & Lyytinen, 2015) that consumes more from reference theories than it contributes back to other disciplines (Polites & Watson, 2009). Indeed, its cognitive capabilities situate CCS in uncharted territory between scientific inquiry into humans (e.g., psychology, sociology, medicine) and inquiry into technology (e.g., computer science, engineering). We believe that IS is ideally positioned to take a lead into the inquiry of CCS because it draws from both worlds. By answering basic questions, such as when humans will start trusting CCS recommendations, IS could make unprecedented contributions to high-impact fields such as medicine (e.g., CCS as diagnosis and treatment systems), finance, and politics (e.g., CCS as advisors). Since practice is already engaging in realizing such endeavors, we call upon IS researchers

to embrace this new technology, critically scrutinize the underlying assumptions of what is new, and engage in paradigm-breaking research.

## 5 CCS Research Directions

We have challenged five assumptions on the grounds that we previously assumed that systems *function in a certain way*; however, because of the new capabilities of CCS, they *may now (also) function in other ways*. Consequently, it is necessary to update the traditional assumptions to more adequately reflect the new capabilities and characteristics of CCS. This means that it may be necessary to relax the assumptions that are, in the case of CCS, too restrictive (e.g., that users are humans, that interaction is facilitated via an artificial interface, or that users are always aware of their system use) or modify the assumptions that are based on presuppositions that do not apply to CCS

(e.g., that a systems' functionality is consistent and transparent). Relaxing and modifying such assumptions not only changes the boundary conditions of our existing theories (Busse, Kach, & Wagner, 2017), as discussed in the previous section, but also widens the landscape of phenomena that can be studied within the context of CCS. This creates theoretical gaps that present unique opportunities for future research. The following discussion illustrates some of the theoretical gaps and research opportunities that arise based on the inadequacy of these five assumptions for CCS and on updated assumptions coming into play (Table 3).

### 5.1 Bilateral Relationships

There is a need for new theories that allow us to understand how humans form relationships with CCS and how these relationships allow CCS to persuade individuals and collaborate with them. In terms of persuasion, IS research has mainly borrowed from the communication literature to examine the characteristics of persuasive messages. For example, IS research often uses the elaboration likelihood model (Petty & Cacioppo, 1986) that explains when arguments or peripheral message cues are more persuasive. Another example is the often used protection motivation theory (Rogers, 1983) that prescribes which arguments should be included in persuasive messages. However, these theories were developed in the context of persuasion in which single, isolated messages came from a single entity and were targeted to the general public (e.g., public service announcements). In contrast, CCS can persuade individuals in the more powerful context of a personal interaction. Unlike messages, CCS can engage individuals in a bilateral dialog using human language. Further, CCS could use personally relevant information to tailor these messages to the individual, just as an actual human would. Instead of generic messages like "smoking kills," a CCS could make in-time interventions with personalized arguments, such as "if you get lung cancer, your daughter Molly will grow up without a mother." There is an unprecedented potential for CCS to persuade users to change their behaviors in ways that could save or improve the lives of millions. However, the extent to which CCS can be successful with such approaches is also unknown, as research has demonstrated the potential for user backlash in response to strongly manipulative messages (Shen, 2015). Since existing theories are not equipped to guide the development of persuasive CCS applications (e.g., for healthcare purposes), we see a need to build new, CCS-specific theories of persuasion. Along these lines, we suggest that future research explore the following broad research questions:

**RQ1:** How can CCS effectively persuade people to follow system advice and orders?

**RQ2:** What are the effects of CCS advice on individuals?

Beyond persuasion, the new bidirectionality allows CCS to collaborate with humans. This is a novel notion for IS research because the previous paradigm has seen systems as tools that users would *use* (Benbasat & Zmud, 2003) but not collaborate *with*. Consequently, IS researchers have developed models that explain when IT use would lead to the desired outcomes. Examples of such research are the task-technology fit model (Goodhue & Thompson, 1995) and the IS success model (Delone & McLean, 2003). However, CCS can operate autonomously and thus *with* humans. For example, various bots already digitally coproduce content on Wikipedia (Young, Wigdor, & Kane, 2018). These bots are not explicitly used by humans in the sense of a tool but are separate entities that edit, update, and delete user-generated content. As such, Wikipedia content is cocreated and managed collaboratively between users *and* bots. The existing theory base on IT use cannot explain when and how these users and bots efficiently collaborate together. With the uptake of CCS, we expect such collaborations to become more common as it appears effective to divide tasks between humans and machines. Hence, there is a pressing need for new theories that can explain how users and CCS can effectively collaborate with each other. Initial conceptual work suggests that collaborations will pan out in weaker or stronger forms of symbioses (Veres, 2017). Various modes are conceivable: from the human being in charge and CCS being in a supporting and/or consulting role (like current smart assistants) to some of the human decisions being outsourced to CCS (e.g., scheduling meetings) to the CCS being in charge and the human being in a supporting role. Demetis and Lee (2018) suspect that the form of collaboration will be determined by the needs and requirements of each technological system. Some related research has shown that the viability of such collaborations also depends on the emotional bond between the human and the system (You & Lionel, 2018). We concur that future research will need to explore these collaborative forms as they emerge and answer research questions such as:

**RQ3:** How can individuals and groups collaborate effectively with CCS?

**RQ4:** What are the prerequisites of successful user-CCS collaboration?

### 5.2 Awareness of the Environment

We further see a need for theories on augmentation and surveillance. A key capability of CCS relates to perceiving its physical and digital environment through

unstructured data streams (e.g., video, audio, web content). This capability enables unprecedented opportunities for augmentation. Augmenting human decision-making through the use of systems is a core research stream in the IS field (Lyytinen & Grover, 2017; Nevo et al., 2008). To that end, IS literature has mostly explored how systems can improve decision-making through supporting communications, data availability and analysis, documentation, and knowledge storage and provisioning (Power, 2002). In contrast, CCS open up new possibilities to support decision-making based on augmenting human *perceptions*. For example, in the healthcare industry, CCS can help radiologists quickly find anomalies in MRI scans (Ahmed et al., 2017); in the banking industry, CCS can help identify complaint patterns in voice recordings of customer service calls. The sheer volume of unstructured data that can be processed by CCS greatly exceeds human abilities; thus, CCS can identify patterns that humans cannot. Yet it remains unknown how human decision makers will respond to findings that they themselves cannot see. Will they ignore them in sheer disbelief or resistance or blindly rely on them in lieu of understanding how these patterns emerge? Related research on expert systems suggests that high-skilled users in particular are likely to perceive such systems as threats (Gill, 1995). Even if such concerns were not an issue, research on antiphishing tools suggests that users are often reluctant to rely on tools even when they have much greater capabilities than the individual user (Abbasi et al., 2015). Hence, for CCS to successfully augment human decision-making, substantial challenges need to be overcome. Thus, we agree with Lyytinen and Grover (2017), who recently called for more research on how emerging information technology can augment decision makers, and we suggest that future research explore the following research question:

**RQ5:** How can CCS effectively augment human decision-making?

The perceptual capabilities of CCS further enable new levels of subjection to surveillance. The phenomenon of subjection to surveillance arises because governments (e.g., UK, Singapore, China) have begun to couple CCS with their surveillance systems. While previous research has always used the working assumption that people can choose whether to disclose their private information, as evident in the privacy calculus model (Smith et al., 2011), this is no longer true. For example, in Shenzhen, pedestrians are subject to constant surveillance through CCS that automatically recognize and punish (i.e., fine) individuals for jaywalking (Li, 2018). Although societies have been subjected to holistic surveillance before (e.g., the German Democratic Republic), never before have citizens been unwillingly subjected to holistic computer surveillance to the extent possible

with CCS. The implications of this phenomenon remain unexplored. In organizational information security research, there is evidence that extensive monitoring can increase human compliance with rules and policies (Vance, Lowry, & Eggett, 2015); however, there is also evidence that extensive monitoring can backfire and lead to even more violations of rules and policies (Lowry & Moody, 2015; Lowry et al., 2015; Posey et al., 2011). To what extent these findings are transferable to a societal context is still unknown, especially when citizens' freedom and lives are at stake. We see a pressing need for research to explore the effects of subjection to surveillance on individuals and society at large and thus propose the following research question:

**RQ6:** What are the effects of subjection to CCS-based surveillance on individuals and society?

### 5.3 Functional Adaptivity

Because CCS can gradually learn over time from their interactions, we see a need for new theories on CCS adaptivity. Although previous systems could, at best, adapt the *content* of their responses to meet user preferences (Lee, Ahn, & Bang, 2011; Liang, Lai, & Ku, 2006; Wattal et al., 2012), CCS can adapt their *functionality*. This means that CCS can change their behavior over time to achieve better outcomes. This is a novel capability that opens up new possibilities for creating more effective systems. Consider the potential of functional adaptivity for healthcare: over 50% of epilepsy patients in the UK do not take their medication regularly or at the correct times (Epilepsy Research UK, 2017) despite the widespread availability of medication reminder apps and alarms. In the absence of knowledge of how to create effective reminders, a CCS system could provide a solution by learning the best ways for delivering effective reminders. However, the degree to which adaptivity is desired and effective remains unknown. It is conceivable that too much adaptivity may render systems inherently ineffective. Consider, for example, a financial advisory CCS that gradually adapts its recommendations toward the recommendations that are most likely to be followed. To what extent would the user benefit from hearing recommendations that he *wants* to hear (e.g., you can save money by buying more fast food) vs. recommendations that he *needs* to hear (e.g., do you really need a new TV)? Which system would be adopted and abandoned and which system would be successful? We expect that answering these types of questions will have wide-ranging implications for the development of adaptive CCS. We propose the following research questions to provide guidance:

**RQ7:** What will be the key outcomes (positive and negative) of functional adaptivity in different contexts? (e.g., increase in user satisfaction, productivity, reliance)

**RQ8:** What extent of functional adaptivity will be desirable for CCS to satisfy the dual outcomes of system and user success?

## 5.4 Functional Opacity

The functional opacity of CCS requires new theories of development and reliance. For the development of previous systems and regardless of the development methodology used (e.g., waterfall model, spiral model, or agile), software development has generally followed the stages of (1) development, (2) testing (of work completed), and (3) deployment. This sequence works well with previous technologies: because of their deterministic nature (i.e., logic and rules), software engineers can conceive tests for functions and features that should always work, regardless of whether the software is tested in development or in the productive environment. Consequently, testing can precede the deployment stage because one would expect the software to pass the same tests in both environments. However, this does not work with CCS that are probabilistic in nature or on trained models that are highly contextualized to their training data sets. This creates two interrelated problems: First, deploying a (1) trained and (2) tested CCS in a new environment will likely lead to entirely different results because the environment provides different inputs (e.g., consider deploying the US version of Siri in Germany). Hence, the development methodology of CCS requires deployment to happen before testing. Second, because the CCS functionality arises from complex probabilistic functions, software engineers have no way of validating the function using logic. Consequently, the only way of successfully validating a CCS is through extensive testing. This is the reason why self-driving car manufacturers deploy their autonomous vehicles on the streets to test them, as this would be the original productive environment; this also explains why these cars are being tested over many years and tens of millions of miles. For example, Waymo recently reached 10 million miles on public roads in 25 American cities (Krafcik, 2018). However, because these 25 cities are the “production environment” of Waymo, their cars cannot be deployed to new cities or countries and expected to function with the same reliability, as each new city and country is idiosyncratic.

These two problems hold implications for how organizations develop software systems. What appears is that the development of CCS is (1) risky because it requires early deployment, and (2) costly because it requires tremendously expansive testing. Thus, the circumstances under which the development of CCS is viable are unclear. Further, it appears that CCS products can lend themselves to winner-take-all markets in which the first to produce a viable CCS will be able to license its software to clients and even

competitors. Hence, beyond the development of improved development and testing methods, we see a specific need for theory that guides managers in their understanding of when CCS development is risky and costly and when licensing is more desirable. Thus, we believe it to be crucial for future CCS research to address the following research question:

**RQ9:** What are the economics of CCS development?

Another key opportunity for CCS research is to explore when humans are willing to rely on CCS. Existing research has suggested that transparency is an important factor in users’ willingness to trust in and rely on systems (Wang & Benbasat, 2016), as users want to understand why and how a system makes its recommendations (Komiak & Benbasat, 2008; Wang & Benbasat, 2007). However, as we have argued, CCS are, because of their probabilistic nature, inherently incapable of explaining why they arrive at specific conclusions. Due to the complex probabilistic techniques employed, even developers have a difficult time explaining why a CCS returns a specific result. Consider, for example, autonomous cars that might suddenly not recognize a street barrier and hence cause a fatal accident (Thompson, 2018). With the safety of humans at risk, there is a continued need for theory that examines when it is safe for humans to rely on CCS and how CCS can foster trust in their functionality. Hence, we suggest that future research address the following questions:

**RQ10:** When can humans safely rely on CCS?

**RQ11:** How can CCS foster humans’ trust?

## 5.5 Unawareness of Use

The more human-like capabilities of CCS also create a need for new theories of deception and substitution. In terms of deception, existing IS research has mostly explored how criminals can deceive and defraud users via computer-mediated communication channels or on e-commerce websites (e.g., Xiao & Benbasat, 2011; Zhou, Burgoon, & Twitchell, 2004). Previous research has thus studied how systems (as the medium) aid the deceptive efforts of humans (the deceiver). With CCS, systems now have the ability to *also* take on the role of the deceiver. For example, a customer service CCS might imitate a human customer agent so effectively that customers may not realize that they are interacting with a CCS. With the uptake of CCS as service agents, such scenarios will become increasingly common. There is evidence suggesting that users interact differently when they are aware that they are interacting with a CCS versus a human being (e.g., Pickard, Roster, & Chen, 2016). Consequently, a need arises to understand how CCS successfully deceive humans and when such deceit can be beneficial or detrimental to user experience and business outcomes.

Along these lines, we suggest that future research address the following questions:

**RQ12:** How can CCS effectively deceive humans?

**RQ13:** When will it be beneficial for CCS to keep users unaware of their system use, and when will it be detrimental?

When CCS can effectively act like human agents, an opportunity arises to study when CCS could or should substitute for human agents. Although the topic of substitution has already attracted much attention in the popular press (e.g., “Will AI take over jobs?”), it has received scant scholarly attention. What is needed are theories that explain which tasks and roles CCS could and should substitute for to benefit individuals, organizations, and society at large. While the substitution of CCS for some tasks would clearly be of benefit for human individuals (e.g., consider the individuals tasked with filtering demeaning, violent, and abusive content from Facebook), CCS are often seen as a threat to many professions (e.g., self-driving cars are a threat to taxi drivers). However, widespread substitution could also create new opportunities for skilled individuals. For example, Wilson, Daugherty, & Morini-Bioanzino (2017) suggest that we will see new jobs devoted to training, explaining, and sustaining the functionality of CCS. Overall, it remains poorly understood where substitution is feasible and where it can have positive or negative effects on individuals, organizations, and society at large. Thus, we believe it to be crucial for future CCS research to address the following research question:

**RQ14:** What are the effects of substituting humans with CCS on individuals, organizations and society?

## 5.6 Methodological Considerations with CCS research

Challenging assumptions not only holds implications for what phenomena can be studied, but also for how the interactions between users and CCS can be studied. We now discuss some of the threats to reliability and internal and external validity that arise from using updated versions of the challenged assumptions.

**CCS Assumption 1** states that the user-system interaction can be bilateral and that participants may receive and react to stimuli from systems (e.g., a voice response from Alexa). We suspect that participants’ responses to system stimuli are especially prone to suffer from what is called the Hawthorne or observer effect, which suggests that participants may alter their behavior based on their knowledge of being observed (Adair, 1984). The Hawthorne effect might be particularly prevalent when studying the interaction between CCS and users, as participants may comply with CCS requests not because the requests are

persuasive or reasonable, but because they feel compelled to do so as a participant in an experimental study. This threatens the external validity of research findings because the observed effects may not reproduce in other (nonexperimental) environments. We thus recommend that future research use methods that are high in external validity, such as surveys, case studies, and field experiments (Karahanna et al., 2018).

**CCS Assumption 2** states that systems may be aware of their environment. This is the case because the rich, natural data that can be processed by CCS may carry information about environmental factors. For example, a voice command response can feature background noises, or a video feed for visual recognition may contain information about the weather. If not carefully controlled, these factors may pose threats to internal validity, as they might influence the responses of CCS to input streams. A classical parable tells a story of an algorithm that was trained to differentiate American tanks from Russian tanks based on the appearance of the tanks in pictures (see Murphy, 2017). But because the American tanks in the training data set were photographed on sunny days and the Russian on cloudy days, the algorithm started to make predictions based on brightness rather than on the appearance of the tanks. This parable thus illustrates that internal validity may be compromised if that which researchers thought would lead to predictions (i.e., appearance) was not the actual driver of the predictions. To cope with this threat, we recommend that researchers replicate their research inquiries in varying environments. If findings replicate to various environments, it may be reasonable to assume that environmental factors do not exert substantial effects on research findings.

**CCS Assumption 3** states that system functionality may not be functionally consistent. If this is the case, reliability is threatened because the interaction between the user and the CCS may change over the course of (1) a research inquiry and, more certainly, between (2) research inquiries, as the system learns based on preceding interactions. Consequently, results obtained from singular cross-sectional and especially longitudinal research inquiries may suffer from biases that arise from changing functionality and results may not be replicable even when using the same system in the same context with the same participants because of changes in the system’s functionality. To cope with this challenge, we suggest that future research should restrict learning within a single research inquiry and test and report differences in functionality between research inquiries.

**CCS Assumption 4** states that CCS may not be functionally transparent. This gives rise to threats to reliability when studying such systems. Specifically, threats to reliability arise when system inputs are not perfectly identical, that is, because inputs through

naturalistic interfaces (i.e., audio or visual) are rich in noise. As a consequence, even small deviations between the same input may result in different results. For example, a participant may use the same instruction (e.g., “Show me the time”) with different accents or tonalities. Many real-world examples of such reliability issues have recently been circulated.<sup>4</sup> Thus, to test interactions between users and CCS, we suggest that researchers pretest the reliability of the system with a set of standard responses in a powerful pilot study with the targeted demographic.

**CCS Assumption 5** states that users may not be aware of their system use. In research settings in which use-related attitudes are queried, instruments that probe such questions may introduce bias in that they make participants aware of their use in the first place. It is not yet understood whether users’ awareness of their system interactions and, especially, their becoming aware of such interactions influence their attitudes. However, we see a potential threat to external validity in that users might react differently in the real world in which some users may not be aware of the system use versus in research conditions in which users are likely to become aware through the instrumentation process. Consequently, we recommend that researchers carefully examine their research instruments so that they do not influence users’ natural beliefs through, for example, hints implicit in questions. If this is not possible, they should make the awareness of system use an explicit boundary condition of their research.

Taken together, challenging assumptions creates substantial new opportunities for exciting research but also holds implications for how we study CCS. We do not claim or aim to have identified all research opportunities or methodological implications that pertain to the study of CCS. Nevertheless, we hope to have inspired some researchers to study some of the fascinating new opportunities that arise through the

emergence of CCS, contribute to current practice through prescient contributions, and do so in a rigorous manner by carefully considering the methodological implications of the modified assumptions.

## 6 Conclusion

As CCS evolve and become an integral part of everyday life—just as the internet is today, compared to what it was two decades or even a decade ago—what will be important avenues for future research? Will it be another study about yet another construct added to our extant, often borrowed, theory base, or will it be unique and fundamental research exploring the novel user-system interactions enabled by the advent of CCS? As we have argued, CCS are the result of a gradual progression toward more human-like capabilities that break with many extant assumptions that have guided our research inquiry thus far. Consequently, many of our concepts, theories, and even methodologies are limited in their applicability to CCS research and opportunities for new, original theorizing arise. This is an excellent opportunity for IS researchers to make novel, high-impact contributions that may extend even beyond the field of IS. We have identified several research avenues that IS can follow to take advantage of this unique opportunity unfolding before us.

## Acknowledgments

The authors are grateful to Professor Allen Lee and the two anonymous reviewers who provided constructive guidance throughout the entire review process. Given the nature of the paper, we are especially thankful that they included reactions to the effect of “although I disagree, the authors’ view is reasonable.”

---

<sup>4</sup> For example the following video of a Scottish woman talking to Alexa in English ([https://www.youtube.com/watch?v=orD-e\\_W6Pic](https://www.youtube.com/watch?v=orD-e_W6Pic)) or the video of HP face tracking that

only works on white users (<http://www.youtube.com/watch?v=t4DT3tQqgRM>).

## References

- Abbasi, A., Zahedi, F. M., Zeng, D., Chen, Y., Chen, H., & Nunamaker, J. F. (2015). Enhancing predictive analytics for anti-phishing by exploiting website genre information. *Journal of Management Information Systems*, 31(4), 109-157.
- Adair, J. G. (1984). The Hawthorne effect: A reconsideration of the methodological artifact. *Journal of Applied Psychology*, 69(2), 334-345.
- Adomavicius, G., Gupta, A., & Zhdanov, D. (2009). Designing intelligent software agents for auctions with limited information feedback. *Information Systems Research*, 20(4), 507-526.
- Agarwal, R. & Karahanna, E. (2000). Time flies when you're having fun: Cognitive absorption and beliefs about information technology usage. *MIS Quarterly*, 24(4), 665-694.
- Agarwal, R., Sambamurthy, V., & Stair, R. M. (2000). The Evolving Relationship Between General and Specific Computer Self-Efficacy—An Empirical Assessment. *Information Systems Research*, 11(4), 418-430.
- Ahmed, M. N., Toor, A. S., O'Neil, K., & Friedland, D. (2017). Cognitive computing and the future of health care. <https://pulse.embs.org/may-2017/cognitive-computing-and-the-future-of-health-care/>
- Aleksander, I. (2004). Advances in intelligent information technology: Re-branding or progress towards conscious machines? *Journal of Information Technology*, 19(1), 21-27.
- Alter, S. (2015). The concept of "IT artifact" has outlived its usefulness and should be retired now. *Information Systems Journal*, 25(1), 47-60.
- Alvesson, M. & Sandberg, J. (2011). Generating research questions through problematization. *Academy of Management Review*, 36(2), 247-271.
- Anderson, C. L. & Agarwal, R. (2011). The digitization of healthcare: Boundary risks, emotion, and consumer willingness to disclose personal health information. *Information Systems Research*, 22(3), 469-490.
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191-215.
- Barr, A. (2015). Google mistakenly tags Black people as "gorillas": Showing limits of algorithms. <https://blogs.wsj.com/digits/2015/07/01/google-mistakenly-tags-black-people-as-gorillas-showing-limits-of-algorithms/>
- Bartunek, J. M., Rynes, S. L., & Duane Ireland, R. (2006). What makes management research interesting, and why does it matter? *Academy of Management Journal*, 49(1), 9-15.
- Benaroch, M. & Chernobai, A. (2017). Operational IT failures, IT value destruction, and board-level IT governance changes. *MIS Quarterly*, 41(3), 729-762.
- Benbasat, I. & Zmud, R. W. (2003). The identity crisis within the IS discipline: Defining and communicating the discipline's core properties. *MIS Quarterly*, 27(2), 183-194.
- Bordetsky, A. & Mark, G. (2000). Memory-based feedback controls to support groupware coordination. *Information Systems Research*, 11(4), 366-385.
- Busse, C., Kach, A. P., & Wagner, S. M. (2017). Boundary conditions: What they are, how to explore them, why we need them, and when to consider them. *Organizational Research Methods*, 20(4), 574-609.
- Carter, M. & Grover, V. (2015). Me, myself, and I(T): Conceptualizing information technology identity and its implications. *MIS Quarterly*, 39(4), 931-957.
- Castelvecchi, D. (2016). The blackbox of AI. *Nature*. [https://www.nature.com/polopoly\\_fs/1.20731!/menu/main/topColumns/topLeftColumn/pdf/538020a.pdf](https://www.nature.com/polopoly_fs/1.20731!/menu/main/topColumns/topLeftColumn/pdf/538020a.pdf)
- Chen, R. & Sharma, S. K. (2013). Learning and self-disclosure behavior on social networking sites: The case of Facebook users. *European Journal of Information System*, 24(1), 1-14.
- Compeau, D. R. & Higgins, C. A. (1995). Computer self-efficacy: Development of a measure and initial test. *MIS Quarterly*, 19(2), 189-211.
- Delone, W. H. & McLean, E. R. (2003). The DeLone and McLean model of information systems success: A ten-year update. *Journal of Management Information Systems*, 19(4), 9-30.
- Demetis, D. S. & Lee, A. S. (2018). When humans using the IT artifact becomes IT using the human artifact. *Journal of the Association for Information Systems*, 19(10), 929-952.
- Dennis, A. R., Fuller, R. M., & Valacich, J. S. (2008). Media, tasks, and communication processes: A theory of media synchronicity. *MIS Quarterly*, 32(3), 575-600.
- Devaraj, S., Fan, M., & Kohli, R. (2002). Antecedents of B2C channel satisfaction and preference:

- Validation e-commerce metrics. *Information Systems Research*, 13(3), 316-333.
- Devaraj, S. & Kohli, R. (2003). Performance impacts of information technology: Is actual usage the missing link? *Management Science*, 49(3), 273-289.
- Elkins, A. C., Dunbar, N. E., Adame, B., & Nunamaker, J. F. (2013). Are users threatened by credibility assessment systems? *Journal of Management Information Systems*, 29(4), 249-262.
- Epilepsy Research UK. (2017). Poll shows that almost 50% of people forget to take their medication at least once a month. <https://www.epilepsyresearch.org.uk/poll-shows-that-almost-50-of-people-forget-to-take-their-medication-at-least-once-a-month/>
- Etzioni, O. (2017). How to regulate artificial intelligence. <https://www.nytimes.com/2017/09/01/opinion/artificial-intelligence-regulations-rules.html?mcubz=0>
- Gill, T. G. (1995). Early expert systems: Where are they now? *MIS Quarterly*, 19(1), 51-51.
- Gill, T. G. (1996). Expert systems usage: Task change and intrinsic motivation. *MIS Quarterly*, 20(3), 301-301.
- Glezer, C. (2003). A conceptual model of an interorganizational intelligent meeting-scheduler (IIMS). *Journal of Strategic Information Systems*, 12(1), 47-70.
- Global Market Insights Inc. (2018). Global smart speaker market size worth over \$30bn by 2024. <https://www.gminsights.com/pressrelease/smart-speaker-market>
- Goodhue, D. L. & Thompson, R. L. (1995). Task-technology fit and individual performance. *MIS Quarterly*, 19(2), 213-236.
- Greenwald, A., Kannan, K., & Krishnan, R. (2010). On evaluating information revelation policies in procurement auctions: A Markov decision process approach. *Information Systems Research*, 21(1), 15-36.
- Gregor, S. & Benbasat, I. (1999). Explanations from intelligent systems: Theoretical foundations and implications for practice. *MIS Quarterly*, 23(4), 497-530.
- Grover, V. & Lyytinen, K. (2015). New state of play in information systems research: The push to the edges. *MIS Quarterly*, 39(2), 271-A275.
- Henfridsson, O. & Bygstad, B. (2013). The generative mechanisms of digital infrastructure evolution. *MIS Quarterly*, 37(3), 907-931.
- Jiang, J. J., Klein, G., & Carr, C. L. (2002). Measuring information system service quality: SERVQUAL from the other side. *MIS Quarterly*, 26(2), 145-166.
- Karahanna, E., Benbasat, I., Bapna, R., & Rai, A. (2018). Editor's comments: Opportunities and challenges for different types of online experiments. *MIS Quarterly*, 42(4), iii-x.
- Kiro. (2018). If you get this call, hang up immediately. <https://www.kiro7.com/news/local/if-you-get-this-call-hang-up-immediately/699237877>
- Komiak, S. Y. X. & Benbasat, I. (2008). A Two-Process view of trust and distrust building in recommendation agents: A process-tracing study. *Journal of the Association for Information Systems*, 9(12), 727-747.
- Kopsco, D., Pipino, L., & Rybolt, W. (1988). A Comparison of the manipulation of certainty factors by individuals and expert system shells. *Journal of Management Information Systems*, 5(1), 66-81.
- Krafcik, J. (2018). Where the next 10 million miles will take us. <https://medium.com/waymo/where-the-next-10-million-miles-will-take-us-de51bebb67d3>
- Krasnova, H., Spiekermann, S., Koroleva, K., & Hildebrand, T. (2010). Online social networks: Why we disclose. *Journal of Information Technology*, 25(2), 109-125.
- Kuhn, T. S. (2012). *The structure of scientific revolution*. University of Chicago Press.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1), 1064-1064.
- Lankton, N., McKnight, D. H., & Thatcher, J. B. (2014). Incorporating trust-in-technology into expectation disconfirmation theory. *Journal of Strategic Information Systems*, 23(2), 128-145.
- Lankton, N. K., McKnight, D. H., & Tripp, J. (2015). Technology, humanness, and trust: rethinking trust in technology. *Journal of the Association for Information Systems*, 16(10), 880-918.
- Launchbury, J. (2017). A DARPA perspective on artificial intelligence. <https://www.darpa.mil/attachments/AIFull.pdf>
- Lee, A. S., Thomas, M., & Baskerville, R. L. (2015). Going back to basics in design science: From the information technology artifact to the information systems artifact. *Information Systems Journal*, 25(1), 5-21.
- Lee, D.-J., Ahn, J.-H., & Bang, Y. (2011). Managing consumer privacy concerns in personalization: A

- strategic analysis of privacy protection. *MIS Quarterly*, 35(2), 423-A428.
- Li, T. (2018). Jaywalkers under surveillance in Shenzhen soon to be punished via text messages. <https://www.scmp.com/tech/china-tech/article/2138960/jaywalkers-under-surveillance-shenzhen-soon-be-punished-text>
- Li, X., Chen, H., Zhang, Z., Li, J., & Nunamaker, J. F. (2009). Managing knowledge in light of its evolution process: An empirical study on citation network-based patent classification. *Journal of Management Information Systems*, 26(1), 129-154.
- Liang, H. & Xue, Y. (2009). Avoidance of information technology threats: A theoretical perspective. *MIS Quarterly*, 33(1), 71-90.
- Liang, T.-P., Lai, H.-J., & Ku, Y.-C. (2006). Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management Information Systems*, 23(3), 45-70.
- Liptak, A. (2017). Amazon's Alexa started ordering people dollhouses after hearing its name on TV. <https://www.theverge.com/2017/1/7/14200210/amazon-alexa-tech-news-anchor-order-dollhouse>
- Liu, D., Santhanam, R., & Webster, J. (2017). Towards meaningful engagement: A framework for design and research of gamified information systems. *MIS Quarterly*, 41(4), 1011-1034.
- Lowry, P. B., Cao, J., & Everard, A. (2011). Privacy concerns versus desire for interpersonal awareness in driving the use of self-disclosure technologies: The case of instant messaging in two cultures. *Journal of Management Information Systems*, 27(4), 163-200.
- Lowry, P. B. & Moody, G. D. (2015). Proposing the control-reactance compliance model (CRCM) to explain opposing motivations to comply with organisational information security policies. *Information Systems Journal*, 25(5), 433-463.
- Lowry, P. B., Posey, C., Bennett, R. B. J., & Roberts, T. L. (2015). Leveraging fairness and reactance theories to deter reactive computer abuse following enhanced organisational information security policies: An empirical study of the influence of counterfactual reasoning and organisational trust. *Information Systems Journal*, 25(3), 193-230.
- Lyytinen, K. & Grover, V. (2017). Management misinformation systems: A time to revisit? *Journal of the Association for Information Systems*, 18(3), 206-230.
- March, S., Hevner, A., & Ram, S. (2000). Research commentary: an agenda for information technology research in heterogeneous and distributed environments. *Information Systems Research*, 11(4), 327-341.
- Markoff, J. (2017). Automated pro-Trump bots overwhelmed pro-Clinton messages, researchers say. <https://www.nytimes.com/2016/11/18/technology/automated-pro-trump-bots-overwhelmed-pro-clinton-messages-researchers-say.html>
- Marr, B. (2016). What everyone should know about cognitive computing. <https://www.forbes.com/sites/bernardmarr/2016/03/23/what-everyone-should-know-about-cognitive-computing/#3a6a91965088>
- Mayer, G., Adomavicius, G., Johnson, P. E., Elidrisi, M., Rush, W. A., Sperl-Hillen, J. M., et al. (2014). A machine learning approach to improving dynamic decision making. *Information Systems Research*, 25(2), 239-263.
- McLeod, R., Jr. & Jones, J. W. (1987). A Framework for Office Automation. *MIS Quarterly*, 11(1), 87-104.
- Metz, C. (2017). Using AI to detect cancer, not just cats. <https://www.wired.com/2017/05/using-ai-detect-cancer-not-just-cats/>
- Moody, D. L., Jacob, M. E., & Amrit, C. A. (2010). *In search of paradigms: Identifying the theoretical foundations of the IS field*. Paper presented at the European Conference on Information Systems.
- Mookerjee, V. S. & Dos Santos, B. L. (1993). Inductive expert system design: Maximizing system value. *Information Systems Research*, 4(2), 111-140.
- Murphy, H. (2017). Why Stanford researchers tried to create a "Gaydar" machine. <https://www.nytimes.com/2017/10/09/science/stanford-sexual-orientation-study.html>
- Nevo, S., Nevo, D., & Ein-Dor, P. (2008). Thirty years of IS research: Core artifacts and academic identity. *Communications of the Association for Information Systems*, 25(24), 221-242.
- Nissen, M. E. & Segupta, K. (2006). Incorporating software agents into supply chains: experimental investigation with a procurement task. *MIS Quarterly*, 30(1), 145-166.
- Norman, D. A. (1986). Cognitive engineering. In D. A. Norman & S. W. Draper (Eds.), *User centered system design* (pp. 31-61). Erlbaum.
- Norman, D. A. (1988). *The design of everyday things*. Basic Books.

- Nunamaker, J. F., Derrick, D. C., Elkins, A. C., Burgoon, J. K., & Patton, M. W. (2011). Embodied conversational agent-based kiosk for automated interviewing. *Journal of Management Information Systems*, 28(1), 17-48.
- Nunamaker, J. F., Jr., Konsynski, B. R., Minder, C., Vinze, A. S., Chen, Y. I. L., & Heltne, M. M. (1988). Knowledge-based systems support for information centers. *Journal of Management Information Systems*, 5(1), 6-24.
- Orlikowski, W. J. & Scott, S. V. (2008). Sociomateriality: Challenging the separation of technology, work and organization. *Academy of Management Annals*, 2(1), 433-474.
- Paradice, D. B. & Courtney, J. F. (1986). Controlling bias in user assertions in expert decision support systems for problem formulation. *Journal of Management Information Systems*, 3(1), 52-64.
- Paradice, D. B. & Courtney, J. F., Jr. (1987). Casual and non-casual relationships and dynamic model construction in a managerial advisory system. *Journal of Management Information Systems*, 3(4), 39-53.
- Petty, R. E. & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123-205.
- Pickard, M. D., Roster, C. A., & Chen, Y. (2016). Revealing sensitive information in personal interviews: Is self-disclosure easier with humans or avatars and under what conditions. *Computers in Human Behavior*, 65, 23-30.
- Pitt, L. F., Watson, R. T., & Kavan, C. B. (1997). Measuring information systems service quality: Concerns for a complete canvas. *MIS Quarterly*, 21(2), 209-209.
- Polites, G. L. & Karahanna, E. (2013). The Embeddedness of information systems habits in organizational and individual level routines: development and disruption. *MIS Quarterly*, 37(1), 221-246.
- Polites, G. L. & Watson, R. T. (2009). Using social network analysis to analyze relationships among IS journals. *Journal of the Association for Information Systems*, 10(8), 595-636.
- Posey, C., Bennett, R. J., Roberts, T. L., & Lowry, P. B. (2011). When computer monitoring backfires: Invasion of privacy and organizational injustice as precursors to computer abuse. *Journal of Information System Security*, 7(1), 24-47.
- Power, D. J. (2002). *Decision support systems: Concepts and resources for managers*. Quorum Books.
- Ransbotham, S., Fichman, R. G., Gopal, R., & Gupta, A. (2016). Special section introduction: Ubiquitous IT and digital vulnerabilities. *Information Systems Research*, 27(4), 834-847.
- Remus, W. & Kottemann, J. (1986). Toward intelligent decision support systems: an artificially intelligent statistician. *MIS Quarterly*, 10(4), 403-418.
- Rich, E. & Knight, K. (1991). *Artificial intelligence*. McGraw-Hill.
- Riemer, K. & Johnston, R. B. (2017). Clarifying ontological inseparability with Heidegger's analysis of equipment. *MIS Quarterly*, 41(4), 1059-1081.
- Robertson, R. (2017). Research aims to make artificial intelligence explain itself. <https://today.oregonstate.edu/archives/2017/jun/research-aims-make-artificial-intelligence-explain-itself>
- Rogers, R. W. (1983). Cognitive and physiological processes in fear appeals and attitude change: A revised theory of protection motivation. In *Social psychophysiology: A sourcebook* (Vol. 19, pp. 153-176). Guilford.
- Russell, S. & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Prentice Hall.
- Rzepka, C. & Berger, B. (2018). User interaction with AI-enabled systems: a systematic review of IS research. *Proceedings of the International Conference of Information Systems*.
- Sandberg, J. & Alvesson, M. (2010). Ways of constructing research questions: gap-spotting or problematization. *Organization*, 18(1), 23-44.
- Schroeder, J. & Epley, N. (2016). Mistaking minds and machines: how speech affects dehumanization and anthropomorphism. *Journal of Experimental Psychology: General*, 145(11), 1427-1437.
- Shen, L. (2015). Antecedents to psychological reactance: The impact of threat, message frame, and choice. *Health Communication*, 30(10), 1-11.
- Simonite, T. (2017). Machines taught by photos learn a sexist view of woman. <https://www.wired.com/story/machines-taught-by-photos-learn-a-sexist-view-of-women/>
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989-1015.
- Straub, D., Limayem, M., & Karahanna, E. (1995). Measuring system usage: Implications for is theory testing. *Management Science*, 41(8), 1328-1342.

- Sviokla, J. J. (1990). An examination of the impact of expert systems on the firm: The case of XCON. *MIS Quarterly*, 14(2), 127-127.
- Thompson, A. (2018). NTSB: Tesla's autopilot steered into barrier in fatal crash. <https://www.popularmechanics.com/technology/infrastructure/a21205127/ntsb-teslas-autopilot-steered-into-barrier-in-fatal-crash/>
- Tiwana, A., Konsynski, B., & Bush, A. A. (2010). Platform evolution: Coevolution of platform architecture, governance, and environmental dynamics. *Information Systems Research*, 21(4), 675-687.
- Turban, E. & Watkins, P. (1986). Integrating expert systems and decision support systems. *MIS Quarterly*, 10(2), 121-136.
- Turban, E. & Watson, H. J. (1988). Integrating expert systems, executive information systems, and decision support systems. *MIS Quarterly*, 10(2), 121-136.
- Vance, A., Lowry, P. B., & Eggett, D. (2015). A new approach to the problem of access policy violations: Increasing perceptions of accountability through the user interface. *MIS Quarterly*, 39(2), 345-366.
- Venkatesh, V. (2000). Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model. *Information Systems Research*, 11(4), 342-365.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478.
- Venkatesh, V., Thong, J. Y. L., & Xu, X. (2016). Unified theory of acceptance and use of technology: A synthesis and the road ahead. *Journal of the Association for Information Systems*, 17(5), 328-376.
- Veres, C. (2017). Strong cognitive symbiosis: Cognitive computing for humans. *Big Data Cognitive Computing*, 1(6), 1-17.
- Vincent, J. (2016). Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day. <http://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>
- Wang, W. & Benbasat, I. (2007). Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems*, 23(4), 217-246.
- Wang, W. & Benbasat, I. (2016). Empirical assessment of alternative designs for enhancing different types of trusting beliefs in online recommendation agents. *Journal of Management Information Systems*, 33(3), 744-775.
- Wattal, S., Telang, R., Mukhopadhyay, T., & Boatwright, P. (2012). What's in a "Name"? Impact of use of customer information in e-mail advertisements. *Information Systems Research*, 23(3), 679-697.
- Waytz, A., Haefner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113-117.
- Weick, K. E. (1989). Theory construction as disciplined imagination. *The Academy of Management Review*, 14(4), 516-531.
- Whittaker, Z. (2018). Judge orders Amazon to turn over Echo recordings in double murder case. <https://techcrunch.com/2018/11/14/amazon-echo-recordings-judge-murder-case/>
- Wilson, J. H., Daugherty, P. R., & Morini-Bioanzino, N. (2017). The jobs that artificial intelligence will create. *MIT Sloan Management Review*, 58(4), 14-16.
- Wuenderlich, N. & Paluch, S. (2017). A nice and friendly chat with a bot: User perceptions of AI-based service agents. *Proceedings of the International Conference on Information Systems*.
- Xiao, B. & Benbasat, I. (2011). Product-related deception in e-commerce: A theoretical perspective. *MIS Quarterly*, 35(1), 169-195.
- You, S. & Lionel, P. R. J. (2018). Emotional attachment, performance, and viability in team collaborating with embodied physical action (EPA) robots. *Journal of the Association for Information Systems*, 19(5), 1-56.
- Young, A., Wigdor, A. D., & Kane, G. C. (2018). Theorizing human and bot co-production effects on information quality. *Proceedings of the International Conference on Information Systems*.
- Zhou, L., Burgoon, J., & Twitchell, D. (2004). A comparison of classification methods for predicting deception in computer-mediated communication. *Journal of Management Information Systems*, 20(4), 139-165.

## Appendix

### Literature

#### IS Research on CCS

We could not find any articles referring to cognitive computing in the Senior Scholars' Basket of Journals<sup>5</sup> when searching Web of Science in February 2019.

#### IS Research on Machine Learning and Artificial Intelligence

Table A1 shows the results from a database search “Basket of 8” AND (“AI” OR “artificial intelligence” OR “machine learning” OR “cognitive computing” OR “intelligent agent”) from the past 17 years, conducted on Web of Science in May 2018.

**Table A1. Definitions and Measures of Managerial Actions**

Study, outlet	Type of study	Phenomenon of interest	Role of AI or ML
Ransbotham, Fichman, Gopal, & Gupta (2016), <i>ISR</i>	Special section introduction	How ubiquitous IT makes people more vulnerable	To motivate new research streams on algorithmic ethics and algorithmic bias
Mayer et al. (2014), <i>ISR</i>	Design science	Dynamic decision-making	To develop an algorithm that aids decision-making in complex, ill-structured contexts
Elkins, Dunbar, Adame, & Nunamaker (2013), <i>JMIS</i>	Empirical study: experiment, $n = 178$	Credibility assessment systems (rely on AI)	To test how users feel when expert systems give contradictory recommendations
Nunamaker et al. (2011), <i>JMIS</i>	Empirical studies: experiment 1 ( $n = 88$ ), experiment 2 ( $n = 202$ ), field experiment 3 ( $n = 29$ )	Value of an automated interviewing agent for border control	Interact with human beings and identify whether they carry a bomb
Adomavicius et al. (2009), <i>ISR</i>	Economical study, simulation	Intelligent agents for auctions	As an agent to place bids in auction markets
Li et al. (2009), <i>JMIS</i>	Design science	Knowledge evolution	To develop a classification algorithm that considers knowledge evolution
Nissen & Segupta (2006), <i>MISQ</i>	Empirical study, experiment ( $n = 84$ )	Effect of procurement agent on users' procurement performance	Recommendation system to support users
Glezer (2003), <i>JSIS</i>	Design science	Facilitating interorganizational meetings with software agents	Agent as a broker
Bordetsky & Mark (2000), <i>ISR</i>	Design science	Facilitating groupware collaboration with an intelligent agent	Agent as a broker
March, Hevner, & Ram (2000), <i>ISR</i>	Research commentary	Research avenues on intelligent agents	As an agent in place of a human
Gregor & Benbasat (1999), <i>MISQ</i>	Review article	Explanations of behaviors of intelligent agents	Intelligent agents as a technological aid to users

<sup>5</sup> The Basket includes the following eight journals: *European Journal of Information Systems*, *Information Systems Journal*, *Information Systems Research*, *Journal of Association for Information Systems*, *Journal of Information Technology*, *Journal of Management Information Systems*, *Journal of Strategic Information Systems*, *MIS Quarterly*.

## IS Research on Expert Systems

Table A2. Definitions and Measures of Managerial Actions

Study, outlet	Type of study	Phenomenon of interest	Theory used	Relevant key insights
McLeod & Jones (1987), <i>MISQ</i>	Case study	Office automation	Managerial role model	Utility of automation depends on the task.
Gill (1995), <i>MISQ</i>	Research commentary	Expert system usage	n/a	Technical success or economic success don't guarantee adoption or use.
Gill (1996), <i>MISQ</i>	Empirical study, survey	Expert system usage	Job design theory	Task (discretion, complexity, speed, and quality) and job (identity) factors positively affect current usage.
Kopsco et al. (1988), <i>JMIS</i>	Empirical study, experiment	Expert's certainty factor estimation	n/a	User's usage of certainty factors differs from that of expert systems.
Mookerjee & Dos Santos (1993), <i>ISR</i>	Design science study	Maximizing expert system's value to an organization	Induction algorithms	n/a
Nunamaker, Konsynski, Minder, Vinze, Chen, & Heltne (1988), <i>JMIS</i>	Design science study	Design of an information center expert system	n/a	Knowledge acquisition is a key concern for the successful development of an expert system.
Paradice & Courtney (1986), <i>JMIS</i>	Design science study	Debiasing expert systems	n/a	Expert systems suffer from biases of experts.
Paradice & Courtney (1987), <i>JMIS</i>	Design science study	Expert systems supporting managerial tasks	n/a	Humans arrive at different conclusions than expert systems.
Remus & Kottemann (1986), <i>MISQ</i>	Future directions	Design of an artificially intelligent statistician to avoid bias in decision makers	n/a	Humans are biased and DSS should thus also support a decision concerning which decision criteria to choose.
Sviokla (1990), <i>MISQ</i>	Case study	Impact of expert systems on task performance	n/a	Task performance significantly increased: quicker decisions and decisions of higher quality by distributing expertise.  Jobs need to be redesigned to fit expert system support.  Organizations might become overdependent on a system.
Turban & Watkins (1986), <i>MISQ</i>	Design science study	Integration of DSS and expert systems	n/a	Integration is beneficial because systems need to not just recommend what action to take, but also why. The authors thus argue that ES and DDS capabilities are needed.

## About the Authors

**Sebastian W. Schuetz** is an assistant professor at the University of Arkansas, Sam M. Walton College of Business. He received his PhD in information systems from the City University of Hong Kong in 2017. His work focuses on behavioral information security as well as the societal implications of artificial intelligence applications and has been published in the *Information Systems Journal*, *International Journal of Information Management*, and several international conferences.

**Viswanath Venkatesh**, who completed his PhD at the University of Minnesota, is a Distinguished Professor and Billingsley Chair at the University of Arkansas. He is widely regarded as one of the most influential scholars in business and economics, both in terms of premier journal publications and citations (e.g., Thomson Reuters' highlycited.com, Emerald Citations, SSRN). His research focuses on understanding the diffusion of technologies in organizations and society. His favorite project focuses on improving the quality of life of the poorest of the poor in India—which he has presented in various forums including at the United Nations. The sponsorship of his research has been about US\$10M. His work has appeared in leading journals in human-computer interaction, information systems, organizational behavior, psychology, marketing, medical informatics, and operations management, and included best paper awards (e.g., *Academy of Management Journal*). His works have been cited over 101,000 times (Google Scholar) and about 33,000 times (Web of Science), with an h-index of 73 and i-10 index of 110. He developed and maintains an IS research rankings web site that has received many accolades from the academic community including AIS' Technology Legacy Award. He has served in editorial roles in various journals. He is a Fellow of the Association for Information Systems and the Information Systems Society, INFORMS.

Copyright © 2020 by the Association for Information Systems. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints, or via email from [publications@aisnet.org](mailto:publications@aisnet.org).