

Beyond extraordinary: Theorizing artificial intelligence and the self in daily life

Keywords: Artificial intelligence, Human-machine communication, Symbolic interactionism, Theory of the self, Method, Narrow AI, Vocal Social Agent, Siri, Cortana

Citation: Guzman, A.L. (2018). Beyond extraordinary: Theorizing artificial intelligence and the self in daily life. In Z. Papacharissi (Ed.), *A Networked Self: Human Augmentics, Artificial Intelligence, Sentience* (pp. 83-96). New York, NY: Routledge.

Version: Submitted draft. (See book for final version).

Contact: alguzman@niu.edu

Andrea L. Guzman, PhD, is an assistant professor of communication at Northern Illinois University. She specializes in human-machine communication with an emphasis on artificial intelligence and voice-based agents.

Beyond extraordinary: Theorizing artificial intelligence and the self in daily life

Artificial intelligence is extraordinary – at least, that is how it has been culturally conceived. AI with its goal of imbuing machines with human-like intelligence has been portrayed as the pursuit of a “dream” (Ekbia, 2008) or, even more extraordinarily, described as humans stepping into the role of a god (e.g. McCorduck, Minsky, Selfridge, & Simon, 1977). Similarly, the implications of AI for self and society have been cast at a level well beyond that of ordinary technology. The promises and consequences of AI have not been so much about who we are as individuals: rather, at stake is who we are as humans and what it means to be human (e.g. Bostrom, 2016; Haugeland, 1985; Turing, 1950; Turkle, 1984). AI seemingly places our collective humanity up for grabs in ways that can elevate or decimate the human race, depending on how it is implemented (e.g. Leverhulme Centre for the Future of Intelligence, 2017).

For these reasons, I found myself perplexed a few years ago while conducting a research project regarding AI. The study’s focus was to better understand people’s conceptualizations of vocal social agents¹ – AI-enabled, voice-based technologies such as Apple’s Siri – and how these perceptions played back into people’s own understanding of the self in relation to AI (Guzman, 2015). But, I did not encounter people questioning human intelligence or their own abilities after chatting with Cortana. There was no wringing of hands or zealous optimism regarding what these programs meant for them as individuals, let alone humanity, now or in the future. Instead, people’s reflections on the self in light of agents contained themes similar to those associated with everyday technologies, such as the mobile phone. People thought of themselves as lazy for vocally asking an agent for directions instead of just typing an address into the device. They did

not want people to perceive them as being rude for talking to their phone while around others. Users were more interested in how a sassy Siri could make them the center of attention than in what a disembodied voice complete with personality meant for the future of the human race. Neither the technology nor its impact on the self from the perspective of users seemed extraordinary; rather, the self in relation to talking AI seemed to be, well, ordinary – just like any other technology.

Why such a disconnect between what scholars have long touted as the implications of AI for the self and what participants were telling me? The answer could be technological. Cortana and company routinely fail at their job, reminding us that they are nowhere near human. VSAs do not possess the full range of human-level intelligence and have limited application. They are not taking jobs or even performing tasks outside the context of a personal assistant. The answer also could be methodological. Maybe I was not asking the right questions or taking into account other ways that aspects of the self can manifest within conversations about AI and technology. The self and AI are not exactly subjects people routinely consider or talk about.

I have come to realize, however, that my struggle to make sense of my data was a theoretical problem that had intertwined technological and methodological components. I had grounded my research in the predominant conceptualization of AI as extraordinary: AI has been portrayed as standing apart from other technological endeavors because it seeks to recreate the human mind within the machine. By its very nature, then, AI has been said to challenge our own human nature (e.g. Dechert, 1966; Haugeland, 1985; Turing, 1950; Turkle, 1984). The result has been that for most of the history of AI, the self in relation to AI has been theorized primarily at the

metaphysical level, with emphasis placed on the implications of machine intelligence for our humanity. However, in my research, I was interested in studying the self at the individual level, as a way of theorizing back to the metaphysical level, not vice-versa. Furthermore, I was looking at a specific AI technology, not AI as a whole. My study also was taking place almost 70 years removed from when AI and its implications for humanity were first theorized. In the interim, particularly within the past decade, artificial intelligence has evolved significantly. In short, the predominant way of conceptualizing AI and theorizing AI in relation to the self could no longer account for what AI had become. The old theory did not fit the new reality.

The purpose of this chapter is to further challenge the predominant ways we have theorized AI in relation to the self while providing a new approach to this research that is better attuned to the evolving nature of artificial intelligence. I have specifically written this chapter for scholars like myself who study AI but are outside the field and face the difficult task of adapting and merging existing theories of AI with theories in our own areas. To that end, I begin by providing a brief overview of the historical, technological, and theoretical context in which the self came to be theorized in terms of AI's inherent disruption to human nature. Next, I explain ongoing and emerging problems with this predominant view of AI and the self. Once readers have had a chance to understand why and how existing theories of AI and the self are inadequate for the study of today's AI, I introduce a research approach that addresses these shortcomings by drawing from emerging theory in human-machine communication.

AI as extraordinary: The stakes for humanity

The first group to suggest that artificial intelligence would pose a challenge to human nature were the very scholars working to make AI a reality; although, their motivation for doing

so varied. For Turing, broaching people's reactions to the human-like qualities of thinking machines was a pragmatic issue (Gandy, 1999). He anticipated that categorizing machines as intelligent would disrupt long-held beliefs about human nature, and so when writing about machine intelligence, Turing (1948, 1950) confronted these objections head-on, refuting them. Where Turing saw a potential hurdle to scientific advancement, others saw an opportunity. Some scholars tapped into the human desire to control nature, including their own, as a way to promote AI. McCorduck describes the pursuit of AI as nothing short of a means through which humans may triumph over nature, adding "Artificial intelligence comes blessed with one of the richest and most diverting histories in science because it addresses itself to something so profound and pervasive in the human spirit" (McCorduck et al., 1977, p. 954). Other scholars warned that such challenges to human nature would upend people's sense of their place in the social and natural order. In advocating for scholars to consider the consequences of automation Theobald (1966) declares, "Man will no longer need to toil: he must find a new role in the cybernetics era which must emerge from a new goal of self-fulfillment" (p. 68-69). Since AI's founding, the implications of AI for our humanity have been an integral part of AI research.

Any interaction with technology has the potential to influence the self, and, in that regard, AI is not unique. However, volumes have not been written about how garage door openers will force us to rethink our humanity. Even technologies that have garnered significant scholarly attention regarding the self, such as social media (e.g. Papacharissi, 2011), have not generated a level of metaphysical inquiry anywhere near that of AI. Underpinning the pursuit of AI has been the theoretical conceptualization of the mind as machine (Haugeland, 1985). What has historically situated AI as extraordinary, setting it apart from garage doors, social media, and other

technologies is that AI seeks to recreate a part of us that we have conceived as being uniquely human – our minds. It is our minds that separate us from animals and objects and are key to our humanity, or so we have thought. As Turkle (1984) explains of AI scholars: “In the course of exercising their profession, they have made questions about human intelligence and human essence their stock and trade” (p. 20). AI’s challenge to humanity has been promulgated in the ways that researchers have theorized AI (Ensmenger, 2012) as well as in the scientific and cultural discourse surrounding AI (Ekbria, 2008).

The pitting of human against machine in some sort of intellectual competition has been an integral part of the development and testing of AI theory and technology (Ensmenger, 2012) and in setting the stakes for AI and the self. The precedent of theorizing AI through intellectual face-offs between human and machine begins with Turing’s (1950) “test” that set the bar for machine intelligence with a computer program’s ability to pass itself off as a human typist. Work toward developing automated chess programs that could beat human opponents also became an important testing ground for AI theory and technology (Ensmenger, 2012). The use of chess was significant at the time because chess has a long history of being conceptualized as *the* game requiring the utmost level of human intellect (Ensmenger, 2012). In using chess as the field’s theoretical testbed, AI scholars raised the ontological stakes for humans and machines to one of its highest levels, if not *the* highest level at the time. Computers could transcend their nature by beating humans at their “best” game, while humans were defending their unique ontological position. When IBM’s Deep Blue finally beat the world’s best chess player, the moment was hailed as a triumph for machines and a loss for humanity (Bloomfield & Vurdubakis, 2008). This scenario has played out repeatedly as new intellectual battles between human and machine –

such as a game of Jeopardy (Kroeker, 2011) or, more recently, AlphaGo (Anthes, 2017) – have continued to push the stakes of AI to higher and higher levels.

Our understanding of AI and its meaning for the self also are established through what Ekbia (2008) calls the “talk of AI.” This talk includes how scientists explain the concept of AI, frame its issues, and discuss its findings as well as media representations of AI, including science fiction plot lines. The idea of thinking machines and the connection between humans and machines has garnered significant media interest and speculation. According to Ekbia (2008) one of the reasons that AI has attracted such media attention is “that people sense that the issues raised by AI are directly related to their self-image as human beings” (p. 319). It is of little surprise, then, that the talk of AI regarding the self remains fixated on our human nature. The theory of mind as machine, around which AI was founded, is itself a metaphor, as is Simon’s (1999) later inversion, “machine as mind.” In both metaphors, the nature of the (human) mind and the nature of the machine are being compared for the sake of making a conceptual connection between the two (as well as what AI researchers hoped would eventually be a theoretical and technological one). In addition, AI technologies and their function are primarily described in human terms (Ekbia, 2008). The heavy use of metaphor and anthropomorphic language in the talk of AI reflects the theoretical claims of the connection between human and machine underpinning AI while, simultaneously, reinforcing them (Agre, 1997; Ekbia, 2008). In doing so, the talk of AI also perpetuates a focus on AI’s implications for our humanity.

The factors I have discussed so far – the theories underpinning AI, the methods of studying AI, and the talk of AI – are integral to forming and maintaining a focus on the meaning of AI for our

human nature. However, there is one meta-factor underlying all these others that needs to be discussed: the origin of AI in speculation based on technological reality. AI did not begin with the creation of a mind within a machine. Rather, it begins with an idea of the future potential of computers. Today, computers are taken for granted, but in the mid-twentieth century, their capabilities were mind-boggling. It is important to remember that it was Turing who played an integral role in the development of computing at this time and, thus, was one of the few people in the world who had access to these machines and had witnessed their capabilities. As Gandy (1999) has argued, Turing's (1950) famous essay on machine intelligence is better understood as "propaganda" aimed at convincing people to realize that "computers were not merely calculating engines but were capable of behavior which must be accounted as intelligent" (p. 123). For Turing, computers already had a degree of intelligence; therefore, it followed that they could be made to be more intelligent. Similarly, the 1956 Dartmouth workshop that is considered the beginning of formal AI research within the U.S. originated out of an idea of what could be possible with computers. Its proposal reads, "The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it" (McCarthy, Minsky, Rochester, & Shannon, 2006, p. 12). The result of this "conjecture" was a field formed entirely around an idea of what "could be." Graubard (1988) explains the power of speculation in shaping the unique character of AI research:

Had the term artificial intelligence never been created, with its implication that a machine might soon be able to replicate the intelligence of a human brain, there would have been less incentive to create a research enterprise of truly mythic proportions. In a very fundamental sense, AI is something of a myth (p. v).

Myth here carries the denotative meaning with a positive connotation of the pursuit of something of great consequence that has yet to be achieved.

Conceptualizations of who we are in relation to AI, then, have formed around the myth that is AI. The pursuit of a myth meant that the daily work of AI and the technologies developed were interpreted and evaluated not only within the context of how they furthered understanding of the mind and technology in the present but also based on how they contributed to the ultimate goal of AI that was yet to be realized. And so, advances in AI have served as a way of thinking about what could be possible in the future. This is one of the many reasons why AI has been not only a technological pursuit but also a philosophical endeavor (Dennett, 1988). As Dennett (1988) explains, “Most AI projects are explorations of ways things might be done and as such are more like thought experiments than empirical experiments” (p. 289). These thought experiments extend to the implications of AI for who we are. Analysis of the self in AI has taken place within the context of the future, one in which the ultimate goal of AI – which presents the highest stakes for humanity – is realized.

Reconciling the future of AI against its present

The predominant way we have theorized AI in relation to the self has remained fixated on what the future realization of AI would mean for our humanity. However, problems exist with this theoretical approach to AI and the self. The idea and study of artificial intelligence is not without its critics, and some of the key aspects underlying the study of AI that also serve as the basis for how we theorize the self, have been called into question. In addition, the discourse of AI that perpetuates the connection between human and machine, and thus reinforces the stakes of

AI for humanity, does not always accurately represent the technological reality of AI.

Furthermore, artificial intelligence – the concept, its study, and the resulting technologies – has continued to evolve while the predominant theory of the self in AI has largely remained static.

The current state of AI and its integration into our daily lives cannot be fully understood through a theoretical lens developed around mid-twentieth century technology and the possibilities of the future.

Metaphors drawing parallels between humans and machines have figured heavily into AI theory and discourse (Agre, 1997; Ekbis, 2008; Haugeland, 1985; Neisser, 1966). Some scholars have expressed concern with the central role of these metaphors in how we conceptualize and talk about AI (e.g. Agre, 1997; Dreyfus, 1999; Neisser, 1966). Their criticism is not so much in the general use of metaphor in AI research (see Agre, 1997); rather, scholars have been critical of how the metaphors connecting body and machine have been so thoroughly subsumed that they are no longer questioned (e.g. Neisser, 1966) and, in some cases, have stopped functioning as metaphors altogether (e.g. Agre, 1997; Ekbis, 2008). As metaphor, the meaning of “body as machine” draws a comparison between humans and machines, but when the comparison is no longer approached as a metaphor, the body becomes a type of machine. Within some veins of AI research the literal application of the metaphor has replaced the figurative (Haugeland, 1985). These same metaphors also have played an integral role in theorizing AI as intrinsically challenging our human nature. As the figurative aspect has receded, the relationship between human and machine is no longer one of comparison; instead, they are equated to one another. The ontological line that once stood between human and machine (Turkle, 1984) is erased, and, thus, the challenge of AI to human nature has seemingly been firmly established. The result is

not only a potentially erroneous way of viewing the self in relation to AI but also a quashing of the need to look for any alternative ways of theorizing the self and AI.

Researchers also have questioned the framing of AI's capabilities within the context of human traits. They argue that anthropomorphic portrayals of AI technology can be false or, at the very least, misleading because they focus on the similarities between human and machine without accounting for the differences (e.g. Agre, 1997; Ekbis, 2008). The high-stakes "intellectual battles" between human and machine can also misrepresent advances in AI and just how close AI is coming to achieving human-level intelligence. While the creation of a chess program that could beat humans seemed to be a sign of success for AI, Ensmenger (2012) explains that the technological reality was more complicated. Yes, chess programs could play chess, and play it well, but the technology behind these programs contributed less to the advancement of AI than work in other areas. IBM's Watson (Kroeker, 2011) and Google's DeepMind (Anthes, 2017) do have impressive AI capabilities and applications for AI well beyond the games they play against humans. Still, as with chess, the battles between humans and these technologies are largely symbolic, in that they do not represent the complexity with which humans and AI compare and contrast to one another. Overall, AI theory and the talk of AI do not just focus our attention toward one aspect of understanding the self – the implications of AI for humanity. They often create and reinforce a skewed picture of AI and its meaning for the self.

Like any physical portrait, this picture of AI has largely remained frozen, capturing what AI was thought to be as it was first emerging. Early efforts in AI coalesced around the creation of technology possessing general intelligence (Franklin, 2014). Now called artificial general

intelligence, this is the conceptualization of AI as extraordinary that is theorized as putting our human nature up for grabs. But whether this level of AI can be achieved has also been the subject of debate (Dreyfus, 1999; “Human-Level AI,” 2017). Recent advances in technology, new ways of theorizing knowledge, and a better understanding of the human brain have given some additional credence to the feasibility of realizing human-level AI (Bostrom, 2016; “Can We Copy the Brain?,” 2017; Müller, 2016). Some scholars have even gone as far as to consider the point when technology surpasses human intelligence (e.g. Bostrom, 2016; Kurzweil, 2005). As with human-level AI, the stakes for machine “superintelligence” are high, according to Bostrom: “This is quite possibly the most important and most daunting challenge humanity has ever faced. And – whether we succeed or fail – it is probably the last challenge we will ever face” (2016, p. v). Although aged, the portrait of AI formed in its early years has staying power.

But, as myth, AI also has nebulous aspects, so that when viewed from a different angle or without its prominent frame, the picture of AI can shift and so too can its meaning for the self. Artificial intelligence grew out of a set of ideas regarding what could be possible with computers, but these ideas were never codified into a universal definition of AI (Franklin, 2014). The Dartmouth workshop that established AI research within the U.S. did not operationalize AI, referring to it only as the “artificial intelligence problem” (McCarthy et al., 2006). Moor (2006) explains: “The field of AI was launched not by agreement on methodology or choice of problems or general theory, but by the shared vision that computers can be made to perform intelligent tasks” (p. 87). That “vision” did initially center around the pursuit of artificial general intelligence as a means to better understand the human mind, but even that work followed divergent theoretical and methodological paths (Dreyfus & Dreyfus, 1988). Eventually the

intertwined goals of understanding the mind through building technology became two research trajectories within AI: the study of the mind and the engineering of “intelligent” machines (Franklin, 2014). These trajectories can intersect, but, more often, proceed separately (Boden, 2016), drawing from and contributing back to varied conceptualizations of AI.

Work within the engineering of AI, in particular, has offered an alternative way of theorizing AI. As the field progressed, some scholars shifted their focus away from the long-term goal of creating a computerized mind toward the shorter-term objective of the commercialization of technology that could carry out certain tasks requiring a specific type of knowledge. Forty years after AI’s founding, Winston (1987) summarized the field’s different priorities: “The primary goal of Artificial Intelligence is to make machines smarter. The secondary goals of Artificial Intelligence are to understand what intelligence is (the Nobel laureate purpose) and to make machines more useful (the entrepreneurial purpose” (p. 1). While AI scholars initially sought to create autonomous machines to understand the human mind, researchers began to view the building of intelligent devices as an end unto itself. Today, the engineering of smart devices is the prominent area of research within AI (Boden, 2016). While efforts toward artificial general intelligence are ongoing, most of the AI programs found in commercial applications are categorized as narrow AI (Holdren & Smith, 2016), technology that operate “intelligently within some relatively narrow domain” (Franklin, 2014, p. 16).

Advances in AI as well as in computing generally also have further pushed AI out of research labs and into everyday spaces. Boden (2016) explains just how widespread AI has become: “AI’s practical applications are found in the home, the car (and the driverless car), the office, the bank,

the hospital, the sky ... and the Internet, including the Internet of Things (which connects the ever-multiplying physical sensors in our gadgets, clothes, and environments” (p. 1). The AI that we daily interact with – whether or not we realize it – are neither the fanciful sentient entities of science fiction nor the higher-intelligence technologies some AI researchers aspire to build. Instead, we encounter AI technologies that perform specific tasks throughout the various aspects of our lives. Yet the predominant lens available to us for theorizing the self within the context of AI was formed around the AI technologies of a fictional tomorrow, not today.

Returning to my opening anecdote, it is now easy to see why I was so baffled by my data. I had developed my original interview questionnaire grounded in the predominant view of the self in relation to AI, with AI being conceptualized within the context of artificial general intelligence; however, I was studying the self in relation to a very different type of AI (a voice-based agent) that operates within a specific framework, as an assistant, and is encountered within the mundane aspects of everyday life, not an experience projected onto the future. I was applying a theory of the self that was developed around AI as extraordinary to the study of AI that was situated within the ordinary. The theoretical lens was an inadequate match for my subject.

AI within the ordinary

Until recently, AI has been theorized as extraordinary and primarily existed outside the realm of the ordinary. AI has been physically restricted to research labs, theoretically projected onto the future, and culturally confined to the media, meaning that the average person had little-to-no opportunity to experience AI first-hand. Therefore, it also has been difficult for scholars to study the implications of AI for the self at the individual level (another reason why theory

regarding the self in AI remained stagnant). The evolution of AI that necessitates a new way of theorizing AI and the self has also loosened some of these former barriers to this research. Scholars have the opportunity to study how individuals conceptualize themselves in relation to AI that is increasingly part of their everyday lives. The challenge now is finding a theoretical and methodological way forward that is not restricted to what AI once was or was projected to be but that is responsive to what AI is and who we are.

Theory being developed in the emerging area of human-machine communication provides one such approach. In my research regarding artificial intelligence (Guzman, 2015, 2017) and technology more generally (Guzman, 2016, 2018), I have advocated for scholars to conceptualize interactions with devices and programs as a form of communication. From this theoretical perspective, a person does not simply use AI, as one would use a tool; rather, they make sense of AI in a dynamic process *with* these technologies. Both the human and machine play an active role in how the machine is understood (see also Neff & Nagy, 2016). The technology, in this case AI, is no longer positioned as only a medium, passing information among humans, but also is theorized as functioning as a type of communicator.

The utility of this theoretical conceptualization of our interactions with AI is that it opens up a way to simultaneously understand how we view a specific AI technology and how we see ourselves in light of it. The process of communication is at the heart of how we come to know others and ourselves (Blumer, 1969; Goffman, 1959; Mead, 1934); this includes both direct interaction with a communication partner and indirect communication with others about them. As I have adapted this epistemological position originally formed around humans to a human-

machine context, the machine becomes a communication partner. The focus then is on how we make sense of the AI technology through what unfolds in our interactions with it as well as what we learn about it from others, and how we then assess who we are in relation to a specific AI program or device. Intrinsic in this phenomenological approach to human-machine communication is that the conceptualizations of self and AI being formed are understood as temporal and contextual.

Following this approach, the study of who we are in relation to AI involves taking into consideration both the characteristics of the particular technology in question as well as the various elements that contribute to the self. As I tried to make apparent earlier in this chapter, simply referring to a technology as AI is not sufficient. Given the ambiguity surrounding the different types of AI, it is incumbent upon scholars to understand what type of AI they are studying and how its characteristics may potentially affect how people understand the technology and themselves. An in-depth example of assessing an individual AI technology from a communication perspective is available in Guzman (2017), so here I touch on only a few key elements.

Given the different ways AI is conceptualized, scholars should determine where a particular technology fits within AI as a whole: Who built the technology? What is its purpose? What does the technology do? What is the level of intelligence involved? Is the technology geared toward a highly specific (narrow) task or does it operate more generally? Attention also should be focused on aspects of people's direct interactions with the technology that inform their understanding of it: In what context(s) do people interact with or encounter the technology? How does the

interface present the technology to the user? How does interaction take place between someone and the technology, including the modes of communication and the messages (verbal and nonverbal) exchanged? Because what we know of others also is built through what we are told about them, researchers also may need to take into account the messages regarding a particular technology and AI more generally: How does the company market the technology? How do media depictions frame the technology? What similarities exist between the technology being studied and pop cultural depictions of AI?

Beyond their own assessment of the AI technology, scholars will have to account for users' own interpretations of the technology. As I learned in my research regarding vocal social agents, users have wide-ranging definitions of what constitutes AI. Some people thought of Cortana, Siri, and Google's voice-technology as AI because the programs could answer questions and provide information fairly well. Other people thought agents were a weaker form of AI or something that was close to but not quite AI, because the agents had flaws or did not possess the full-range of human intelligence. Still others did not consider agents to be AI because the programs were too simple to be considered intelligent, and a few people could not categorize these technologies because they did not know enough about AI to make a determination. In deciding whether Siri and similar programs were AI, people drew on knowledge about AI gained from personal or professional experience as well as media representations. They also compared and contrasted AI against other technologies. In some cases, people defined AI in similar ways but reached different conclusions as to whether agents qualified as AI.

Users' wide-ranging conceptualizations of AI further underscore the complexity of the study of AI and the self. If a person does not consider a device or program to be AI, when from a technological and research stand point it could be defined as AI, how are we to interpret their understanding of self in relation to AI? What if a participant's conceptualization of AI is based on a different but also accepted definition of AI? It is entirely possible that in a research project a participant and the researcher may hold divergent but externally valid definitions of AI.

Approaches like the one I have outlined here dictate that careful attention be paid to the etic and the emic and to the weighing of one against the other. In that, the study of AI is no different. But AI's nebulousness and relative newness, as an actual versus fictional technology, require extra care in formulating our research questions and in interpreting the data.

This is why it is so important for us as researchers to guard against our own fixation with the extraordinariness of AI that has been so entrenched within the predominant ways that AI has been theorized (and why I dedicated two-thirds of this chapter to explaining why this view is so pervasive and problematic). What led me to the theoretical breakthrough underlying this chapter was a decision to look beyond the dominant view of AI in relation to the self. Rather than focus on people's reactions to and interpretations of the human-like traits of agents – the characteristics of AI that are supposed to challenge our human nature – I shifted my focus to understanding what my participants told me again and again mattered the most to them, the “ordinary” aspects of AI as technology. People routinely compared agents to everyday technologies (mobile phones, navigation devices) and expectations surrounding technology generally (utility, usability, etc) to decide whether, how, and when to use an agent. This is not to say that people did not recognize human-like elements within these agents. These characteristics – such as the ability to speak, to

learn, and to enact a particular personality – played critical roles in how people made sense of agents. But most people viewed these human-like traits within the context of the machine, assessing whether these traits made agents a better technology than other programs and devices.

The technological aspects of voice-based agents also mattered the most for how people incorporated these technologies into performances and perceptions of the self. Some people who adopted Cortana or Siri regularly used the agents because they viewed themselves as being tech-forward. Others rejected agents because they considered the programs to be yet another example of unnecessary technology and prided themselves on not falling into these trends or becoming overly dependent on technology. Relatedly, some users described themselves or other people as being lazy for using a voice-based technology when typing a request worked well-enough. And so, a limited level of intelligence and the ability to communicate did not set up an ontological face-off between Cortana, a machine, and the user, a human. Nor did the human-like elements of agents invoke in most people an existential reflection on what it means to be human. In the minds of people using vocal social agents, what is at stake is their productivity, not their humanity.

Beyond extraordinary: AI as ordinary

As I was writing this chapter, the professional organization IEEE released a special magazine issue on AI (“Can We Copy the Brain?,” 2017) with a Q&A feature asking prominent AI scholars to weigh in on the future of artificial general intelligence. When asked “how brainlike [sic] computers would change the world,” Rodney Brooks replied, in part:

Since we won't have intelligent computers like humans for well over 100 years, we cannot make any sensible projections about how they will change the world, as we don't understand what the world will be like at all in 100 years ("Human-Level AI," 2017).

Although I do not always agree with Brooks, his point is an important one to remember in moving forward with the study of AI and its impact on the understanding of the self. Brooks's comment echoes an ongoing critique of how we discuss and theorize AI as if it suddenly appears out of nowhere. This is another example of how AI is extraordinary. AI is unlike any other technology not just because it possesses human-like elements but that by having these traits, AI is somehow magically exempt from the normal processes through which we make sense of our world. At least, that is how we have approached the study of AI and the self. For 70 years we have conceptualized AI as a technology that inherently challenges our very nature and our humanity. Yet, the AI that will allegedly lead to such ontological upheaval did not exist when this theoretical viewpoint was first introduced and still does not exist today. Our theory of what AI means for the self has been based on interpretations of contemporary technology projected onto the future.

But, as I have briefly demonstrated here, when we adopt an approach that situates the study of AI and the self within the context of our lived experiences, a different picture of AI appears. And this is what I mean by AI as ordinary. AI may be modeled on humans, thus, making it potentially different from other technologies, but it is not exempt from the ways we make sense of our world. Furthermore, no matter how human-like technology becomes, our own sense of self also remains rooted in our interactions with the world around us. And so, regardless of the degree of

the intelligence in the technology before us, to understand AI and ourselves, we must remember to go beyond the idea of AI as extraordinary, and approach it as ordinary.

Notes

1. Most of the empirical work mentioned within this chapter is derived from my dissertation (Guzman, 2015), and some of the terms used here and discussions of findings may differ slightly from how I originally presented them within the dissertation. For example, I now use the term vocal social agent, or VSA, to refer to the agents I studied. (see Guzman, 2017).

References

- Agre, P. E. (1997). *Computation and human experience*. Cambridge, UK: Cambridge University Press.
- Anthes, G. (2017). Artificial intelligence poised to ride a new wave. *Communications of the ACM*, 60, 19–21. <https://doi.org/10.1145/3088342>
- Bloomfield, B. P., & Vurdubakis, T. (2008). IBM's Chess Players: On AI and Its Supplements. *The Information Society*, 24, 69–82. <https://doi.org/10.1080/01972240701883922>
- Blumer, H. (1969). *Symbolic interactionism: Perspective and method*. Englewood Cliffs, NJ: Prentice-Hall.
- Boden, M. (2016). *AI: Its nature and future*. Oxford, UK: Oxford University Press.
- Bostrom, N. (2016). *Superintelligence: Paths, dangers, strategies*. Oxford, UK: Oxford University Press.
- Can We Copy the Brain? (2017, May 31). *IEEE Spectrum: Technology, Engineering, and Science News*. Retrieved June 1, 2017, from <http://spectrum.ieee.org/static/special-report-can-we-copy-the-brain>

- Dechert, C. R. (Ed.). (1966). *The social impact of cybernetics*. New York, NY: Simon and Schuster.
- Dennett, D. C. (1988). When philosophers encounter artificial intelligence. *Daedalus*, 117 (Artificial Intelligence), 283–295.
- Dreyfus, H. L. (1999). *What computers still can't do: A critique of artificial reason* (6th ed.). Cambridge, MA: MIT Press.
- Dreyfus, H. L., & Dreyfus, S. E. (1988). Making a mind versus modelling the brain: Artificial intelligence back at the branchpoint. *Daedalus*, 117 (Artificial Intelligence), 15–34.
- Ekbia, H. R. (2008). *Artificial dreams: The quest for non-biological intelligence*. New York, NY: Cambridge University Press.
- Ensmenger, N. (2012). Is chess the drosophila of artificial intelligence? A social history of an algorithm. *Social Studies of Science*, 42(1), 5–30.
<https://doi.org/10.1177/0306312711424596>
- Franklin, S. (2014). History, motivations, and core themes. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence*. Cambridge, UK: Cambridge University Press.
- Gandy, R. (1999). Humans versus mechanical intelligence. In P. Millican & A. Clark (Eds.), *Machines and thought: The legacy of Alan Turing* (Vol. 1, pp. 125–136). Oxford, UK: Oxford University Press.
- Goffman, E. (1959). *The Presentation of Self in Everyday Life*. New York, NY: Anchor Books.
- Graubard, S. R. (1988). Preface to the Issue, “Artificial Intelligence.” *Daedalus*, 117 (Artificial Intelligence), V–VIII.

- Guzman, A. L. (2015). *Imagining the Voice in the Machine: The Ontology of Digital Social Agents*. University of Illinois at Chicago, Chicago, IL.
- Guzman, A. L. (2016). The Messages of Mute Machines: Human-Machine Communication with Industrial Technologies. *Communication* 1, 5(1). Retrieved from <http://scholarworks.umass.edu/cpo/vol5/iss1/4>
- Guzman, A. L. (2017). Making AI Safe for Humans: A Conversation with Siri. In R. W. Gehl & M. Bakardjieva (Eds.), *Socialbots and Their Friends: Digital Media and the Automation of Sociality*. New York, NY: Routledge.
- Guzman, A. L. (2018). What is human-machine communication, anyway? In A. L. Guzman (Ed.), *Human-Machine Communication: Rethinking Communication, Technology, & Ourselves*. New York, NY: Peter Lang.
- Haugeland, J. (1985). *Artificial intelligence: The very idea*. Cambridge, Mass.: The MIT Press.
- Holdren, J. P., & Smith, M. (2016). *Preparing for the Future of Artificial Intelligence*. Washington, D.C.: National Science and Technology Council, Committee on Technology. Retrieved from <http://www.datascienceassn.org/sites/default/files/Preparing%20for%20the%20Future%20of%20Artificial%20Intelligence%20-%20National%20Science%20and%20Technology%20Council%20-%20October%202016.pdf>
- Human-Level AI Is Right Around the Corner—or Hundreds of Years Away. (2017, May). *IEEE Spectrum: Technology, Engineering, and Science News*. Retrieved from <http://spectrum.ieee.org/computing/software/humanlevel-ai-is-right-around-the-corner-or-hundreds-of-years-away>

- Kroeker, K. L. (2011). Weighing Watson's impact. *Communications of the ACM*, 54(7), 13.
<https://doi.org/10.1145/1965724.1965730>
- Kurzweil, R. (2005). *The singularity is near*. New York, NY: Penguin.
- Leverhulme Centre for the Future of Intelligence. (2017). About: Preparing for the age of intelligent machines. Retrieved August 12, 2017, from <http://lcfi.ac.uk/>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine*, 27(4), 12.
- McCorduck, P., Minsky, M., Selfridge, O. G., & Simon, H. A. (1977). History of Artificial Intelligence. In *IJCAI* (pp. 951–954). Retrieved from <http://ijcai.org/Proceedings/77-2/Papers/083.pdf>
- Mead, G. H. (1934). Obstacles and promises in the development of the society. In *Mind, Self and Society*. Chicago, IL: University of Chicago Press.
- Moor, J. (2006). The Dartmouth College artificial intelligence conference: The next fifty years. *AI Magazine*, 27(4), 87.
- Müller, V. C. (2016). New developments in the philosophy of AI. In V. C. Müller (Ed.), *Fundamental Issues of Artificial Intelligence* (Vol. 376, pp. 1–4). Switzerland: Springer.
- Neff, G., & Nagy, P. (2016). Automation, algorithms, and politics: Talking to bots: Symbiotic agency and the case of Tay. *International Journal of Communication*, 10, 17.
- Neisser, U. (1966). Computers as tools and as metaphors. In C. R. Dechert (Ed.), *The Social Impact of Cybernetics* (pp. 71–93). New York, NY: Simon and Schuster.
- Papacharissi, Z. (2011). *A networked self*. New York, NY: Routledge.

- Simon, H. A. (1999). Machine as mind. In P. Millican & A. Clark (Eds.), *Machines and thought: The legacy of Alan Turing* (Vol. 1, pp. 81–102). Oxford, UK: Oxford University Press.
- Theobald, R. (1966). Cybernetics and the problems of social reorganization. In C. R. Dechert (Ed.), *The Social Impact of Cybernetics* (pp. 39–70). New York, NY: Simon and Schuster.
- Turing, A. M. (1948). Intelligent machinery. In D. C. Ince (Ed.), *Mechanical Intelligence* (pp. 107–128). New York, NY: North-Holland.
- Turing, A. M. (1950). Computing machinery and intelligence. In D. C. Ince (Ed.), *Mechanical Intelligence* (pp. 133–160). New York, NY: North-Holland.
- Turkle, S. (1984). *The second self*. New York, NY: Simon & Schuster.
- Winston, P. H. (1987). Artificial intelligence: A perspective. In W. E. L. Grimson & R. S. Patil (Eds.), *AI in the 1980s and Beyond: An MIT Survey* (pp. 1–12). Cambridge, MA: The MIT Press.

