Design Lessons From AI's Two Grand Goals: Human Emulation and Useful Applications

Ben Shneiderman

Abstract—**Researchers' goals shape the questions they raise, collaborators they choose, methods they use, and outcomes of their work. This article offers a fresh vision of artificial intelligence (AI) research by suggesting a simplification to two goals: 1) emulation to understand human abilities to build systems that perform tasks as well as or better than humans and 2) application of AI methods to build widely used products and services. Researchers and developers for each goal can fruitfully work along their desired paths, but this article is intended to limit the problems that arise when assumptions from one goal are used to drive work on the other goal. For example, autonomous humanoid robots are prominent with emulation researchers, but application developers avoid them, in favor of tool-like appliances or teleoperated devices for widely used commercial products and services. This article covers four such mismatches in goals that affect AI-guided application development: 1) intelligent agent or powerful tool; 2) simulated teammate or teleoperated device; 3) autonomous system or supervisory control; and 4) humanoid robot or mechanoid appliance. This article clarifies these mismatches to facilitate the discovery of workable compromise designs that will accelerate human-centered AI applications research. A greater emphasis on human-centered AI could reduce AI's existential threats and increase benefits for users and society, such as in business, education, healthcare, environmental preservation, and community safety.**

Index Terms—**Artificial intelligence (AI), autonomous systems, design, design lessons, human–computer interaction (HCI), humanoid robots, mechanoid appliances, simulated teammate, social impact, supervisory control, teleoperated devices.**

I. Introduction: What Is the Goal of AI Research? Emulation or Application

Goals of artificial intelligence (AI) research were proposed at least 60 years ago, when early conferences brought together those who believed in pursuing Alan Turing's question "can machines think?" [64]. A simplified summary might be that AI is getting computers to do what humans do, especially by emulating (some would say simulating) their perceptual, cognitive, and motor abilities.

This summary encompasses goals such as satisfying the classic Turing test, based on a keyboard and teletype (or screen) conversation in which users cannot tell if they are connected to a human or a machine. Other forms of Turing test include visually representing a human with computer-generated imagery (CGI) in a computer game or Hollywood film and making humanoid robots that think, see, speak, act, and move like a human. Russell [51] energetically argued that AI was "one of the principal avenues for understanding how human intelligence works but also a golden opportunity to improve the human condition—to make a far better civilization."

Emulation research on human perceptual, cognitive, and motor abilities includes pattern recognition (images, speech, facial, signal, etc.), natural language processing, translation of natural language, bipedal robots, emotionally responsive human faces, and game playing (checkers, chess, go, etc.). As the early emulation research evolved, useful applications became possible, but the emulation research that emphasized symbolic manipulation, gave way to statistical approaches, based on deep learning and machine learning, which functioned differently from humans.

The rich history of AI research has many methods and many voices advocating diverse goals. The visionary aspirations of AI researchers led to a range of inspiring projects. Proponents claim that AI is a historical turning point for humanity with great promise and existential dangers. Critics point out that many projects failed, as is common with ambitious new research directions, but others led to widely used applications, such as optical character recognition, speech recognition, and natural language translation. While critics say that AI applications remain imperfect, many applications are impressive and commercially successful.

While bold aspirations can be helpful, another line of criticism is that the AI emulation methods failed, giving way to more traditional engineering solutions, which succeeded. For example, IBM's famed Deep Blue chess-playing program, which defeated world champion, Garry Kasparov, in 1997, is claimed as an AI success. However, IBM's researcher who built Deep Blue, Feng-Hsiung Hsu, makes an explicit statement that the brute-force hardware solution did not use AI methods [22]. Another example is that AI-guided knowledge-based expert systems failed, but carefully engineered rule-based systems with human-curated rule sets succeeded in many business applications [28].

These debates about AI research goals dramatically influence government research funding, major commercial projects, academic research and teaching, and public impressions. This article simplifies the many goals of AI researchers into these two: 1) emulation and 2) application, and then describes four pairs of contrasting subgoals (Fig. 1). The sharply defined emulation and application goals help clarify important distinctions, but individuals are likely to have more complex beliefs, which fall in between these extremes.

These four contrasting subgoals provide this article's structure, as well as a guide to compromise designs. This article is intended to accelerate research on human-centered AI that produces useful applications with widespread benefits for users and society, such as in business, education,

Manuscript received April 26, 2020; accepted April 30, 2020. Date of publication May 5, 2020; date of current version June 2, 2020. This article was recommended for publication by Associate Editor K. Michael.

The author is with the Department of Computer Science, University of Maryland at College Park, College Park, MD 20742 USA, and also with the Human–Computer Interaction Laboratory, University of Maryland at College Park, College Park, MD 20742 USA (e-mail: ben@cs.umd.edu).

Digital Object Identifier 10.1109/TTS.2020.2992669

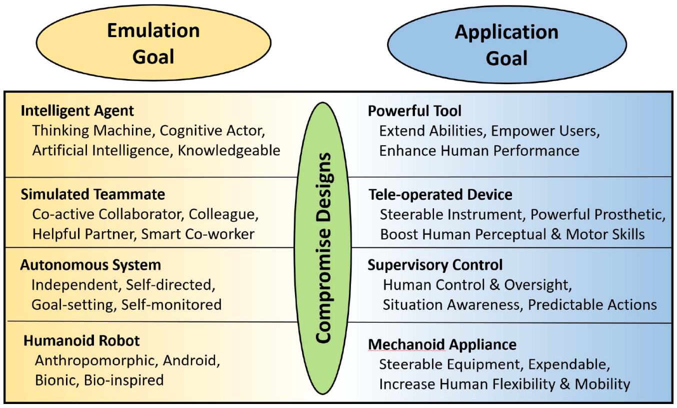

Fig. 1. Terminology and metaphors for the emulation goal and application goal, with the possibility of compromise designs that combine features from each goal.

healthcare, environmental preservation, and community safety.

Section II describes the emulation goal of understanding human perceptual, cognitive, and motor abilities to build computers that autonomously perform tasks as well as or better than humans. It summarizes the application goal of developing widely used products and services by using AI methods, which ensure human control. Both goals require science, engineering, and design research.

Section III focuses on conflicts that arise when product designers follow the emulation goal of building cognitive computers that they describe as smart, intelligent, knowledgeable, and capable of thinking. The resulting human-like products can be entertaining and appealing to some users, but these designs can exacerbate the distrust, fears, and anxiety that many users have of computers. The application goal community believes that computers are best designed to be powerful tools that amplify, augment, empower, and enhance humans. The compromise strategy could be to design tool-like user interfaces with AI-driven technologies for services, such as text messaging word suggestions and internal operations that transmit optimally across complex networks.

Section IV raises these questions: do designers benefit from thinking of computers as being teammates, partners, and collaborators? When is it helpful and when is there a danger in assuming human–human interaction is a good model for human–robot interaction? Application goal researchers and developers want to build teleoperated devices that extend human capabilities while providing superhuman perceptual and motor support, thereby boosting human performance while allowing human–human teamwork to continue. The compromise strategy could be to use emulation goal algorithms to implement automatic internal services that support the application goal of human control. This approach is implemented in many car driving technologies, such as lane following, parking assist, and collision avoidance.

Section V discusses the conflicts that arise when the emulation goal of autonomous systems leads designers of products and services astray. Rather than autonomous system acting alone, application goal researchers want to support supervisory control (sometimes called human in the loop), in which humans operate highly automated devices and systems. The compromise strategy could be to have emulation goal algorithms for highly automated features, with user interface designs that support human control and oversight by the way of comprehensible, predictable, and controllable user interfaces. This compromise strategy is in use in many NASA, industrial, utility, military, and air traffic control rooms, where rich forms of AI are used to optimize performance, but the operators have a clear mental model of what will happen next.

Section VI covers the many attempts by emulation goal that advocates to build humanoid robots over hundreds of years, which have attracted widespread interest, but limited commercial success. At the same time, mechanoid computers in the form of home appliances, mobile devices, and kiosks are great successes. Application goal champions prefer mechanoid (or mechanical) robots that are seen as steerable instruments, designed to increase flexibility or mobility while being expendable in rescue, disaster, and military situations. The compromise design could be to use limited humanoid services, which have proven acceptance, such as voice-operated virtual assistants embedded in mechanoid designs.

Awareness of how the different goals can produce avoidable conflicts lays the foundation for clearer thinking that leads to reliable, safe, and trustworthy systems [58]. Section VII offers conclusions and possible paths to constructive collaboration between the emulation goal and application goal communities.

II. Two Goals for AI Researchers and Developers

AI researchers and developers have offered many goals such as this one in a major textbook: 1) think like a human; 2) act like a human; 3) think rationally; and 4) act rationally [52]. Others see AI as a set of tools to augment human abilities or extend their creativity. For simplicity, this article focuses on two goals: 1) human emulation and 2) useful applications. Of course, some researchers and developers will be sympathetic to goals that fall in both communities or even other goals that fall in between. I repeat the caution that the sharply defined emulation and application goals are meant to clarify important distinctions, but individuals are likely to have more complex beliefs.

A. Emulation Goal

The emulation goal is to understand human perceptual, cognitive, and motor abilities to build computers that perform tasks as well as or better than humans. This goal includes the aspiration for humanoid robots, natural language and image understanding, commonsense reasoning, and artificial general intelligence (AGI).

Those who pursue the emulation goal have grand scientific ambitions and understand that it may take 100 or 1000 years, but they tend to believe that researchers will be able to understand and model humans faithfully [33], [48], [51]. Many researchers in this AI community believe that humans are machines, may be very sophisticated, but eventually building exact emulations of humans is a realistic and worthwhile grand challenge. They are cynical about claims of human exceptionalism or that humans are a separate category from computers. The famed *AI 100 Report* [60] states that "the difference between an arithmetic calculator and a human brain is not one of kind, but of scale, speed, degree of autonomy, and generality," which assumes that human and computer thinking are in the same category.

The desire to build computers that match human abilities is an ancient and deep commitment. Broadbent [7] states that: "humans have a fundamental tendency to create, and the ultimate creation is another human." This belief influences the terminology and metaphors that the emulation goal community feels strongly about. They often think of computers as smart machines, intelligent agents, and knowledgeable actors, and are attracted to the idea that computers are learning and require training. Their work often includes performance comparisons between humans and computers, such as the capability of oncologists versus AI programs to identify breast cancer tumors.

Many emulation goal researchers and developers believe that robots can be teammates, partners, and collaborators and that computers can be autonomous systems that are independent, capable of setting goals, self-directed, and self-monitoring. They see automation as merely carrying out requirements anticipated by the programmers/designs, while autonomy is a step beyond automation to support unanticipated goals and emergent behaviors based on new sensor data. Emulation goal protagonists promote embodied intelligence through humanoid (or anthropomorphic) robots, which are bioinspired (or bionic) to resemble human forms. Finally, the emulation goal work may be influenced by real problems, but their work is often on toy problems or based on synthetic data and testing is often in laboratory conditions.

Some researchers, legal scholars, and ethicists envision a future in which computers will have the responsibility and legal protection of their rights, much as individual humans and corporations [8]. They believe that computers and humanoid robots can be moral and ethical actors and that these qualities can be built into algorithms. This controversial topic is beyond the scope of this article, which focuses on design issues to guide the near-future research and the development of the next generation of technology.

B. Application Goal

The application goal drives researchers to develop widely used products and services by applying AI methods. This goal typically favors tool-based metaphors, teleoperated devices, supervised operation, and mechanoid appliances. These applications are described as instruments, apps, orthotics, prosthetics, utensils, or implements. These AI-guided products and services are built into the cloud, websites, laptops, mobile devices, home automation, kiosks, flexible manufacturing, and virtual assistants, and tailored for diverse application domains. An emulation goal airport assistant might be a mobile human-like robot that engaged in natural language conversation while an application goal airport kiosk would be a fixed device with touchscreen instructions to guide travelers.

Researchers and developers who pursue the application goal study aspects of human behavior and social dynamics, which relate to user acceptance of products and services. These researchers are typically enthusiastic about serving human needs, so they often partner with professionals to work on authentic problems and take pride in widely adopted innovations. They regularly begin by clarifying what the tasks are and thinking about the diverse stakeholders, their values, and societal/environmental impacts [14].

The application goal community frequently supports high levels of human control and high levels of computer automation, but they understand that there are applications that require rapid fully automatic operation (airbag deployment, pacemakers, etc.) and there are applications in which users prefer full human control (bicycle riding, piano playing, etc.). Between these extremes lie a rich design space that combines high levels of human control and high levels of automation. These researchers normally recognize the dangers of excessive automation and excessive human control, so they introduce interlocks that prevent failures while striving to find a balance that produces reliable, safe, and trustworthy systems [58].

The desire to make commercially successful products and services means that human–computer interaction (HCI) methods, such as design thinking, observation of users, usability testing, market research, and continuous monitoring of usage, are frequent processes employed by the application goal community. They recognize that users often prefer designs that are comprehensible, predictable, and controllable and that users want to increase their own mastery, self-efficacy, and creative control. They accept that humans need to be "in the loop," respect that users deserve explainable systems, and recognize that humans and organizations are the holders of responsibility, liability, and accountability [46]. They are sympathetic to audit trails, product logs, or flight data recorders to support retrospective forensic analysis of failures to improve reliability and safety, especially for life-critical applications, such as pacemakers, self-driving cars, and aviation [57].

Sometimes those pursuing the application goal start with emulation goal ideas and then do what is necessary to create successful products and services. For example, speech recognition research was an important foundation of successful virtual assistants, such as Apple's Siri, Amazon's Alexa, Google's Home, or Microsoft's Cortana, but user interface design methods were important complements. Similarly, natural language translation research was integrated into well-designed user interfaces for imperfect, but successful websites and services. The third example is that image understanding research-enabled automatic creation of alt tags, which are short descriptions of images that enable users with disabilities and others to know what is in a website image.

Autonomous humanoid robots with bipedal locomotion and emotionally responsive faces, inspired by the emulation goal make appealing demonstrations and videos. However, these designs often give way to four-wheeled boxes, tread-based mobility, or teleoperated drones without faces, which are needed for success in the application goal [44]. The language of robots may remain, as in "surgical robots," but these are teleoperated devices that allow surgeons to do precise work in tight spaces inside the human body.

Many emulation goal researchers believe that a general purpose humanoid robot can be made, which can serve tea to elders, deliver packages, and perform rescue work. In contrast, application goal researchers recognize that they have to tune their solutions for each context of use. Nimble hand movements, heavy lifting, or movement in confined spaces require specialized designs, which are not at all like a generic multipurpose human hand.

Lewis Mumford's book [43] *Technics and Civilization* has a chapter titled "The Obstacle of Animism," in which he describes how first attempts at new technologies are misled by human and animal models. He uses the awkward term "dissociation" to describe the shift from human forms to more useful designs, such as recognizing that four wheels have large advantages over two feet in transporting heavy loads over long distances. Similarly, airplanes have wings, but they do not flap like bird wings. Mumford stressed that "the most ineffective kind of machine is the *realistic* mechanical imitation of a man or other animal." He continues with this observation "circular motion, one of the most useful and frequent attributes of a fully developed machine is, curiously, one of the least observable motions in nature" and concludes "for thousands of years animism has stood in the way of development."

Another important topic for application goal researchers is supporting human connections, e.g., with social media and collaborative software. For example, teleconferencing services expanded dramatically during the COVID crisis as universities shifted to using online instruction with live instructors and AI-guided automated services such as massive online open courses (MOOCs), such as those from Khan Academy, edX, or Coursera. Businesses quickly expanded working from home (WFH) options for their employees and family, friends, businesses, and communities adopted services, such as zoom, to support lectures, discussions, and much more.

Many forms of collaboration are supported by social media platforms, such as Facebook, Twitter, and Weibo, which employ AI-guided services. These platforms have attracted billions of users, who enjoy the social connections, benefit from the business opportunities, connect communities in productive ways, and support teamwork as in citizen science projects. However, many users have strong concerns about privacy and security, as well as the misuse of social media by political operatives, criminals, terrorists, and hate groups to spread fake news, scams, and dangerous messages. AI algorithms and user interface designs have contributed to these abuses, but they also contribute to the solutions.

The brief definitions of AI and the simple emulation and application goals are incomplete and controversial, with many people having more complex beliefs. Still, I believe that they provide a foundation for clearer thinking about the science, engineering, and design of AI research and systems so that societal benefits can be delivered. A common problem occurs when the assumptions tied to the emulation goal are put to work in the application goal. However, there are often compromise designs that benefit each goal's practices, effectively combining AI and HCI methods.

III. Intelligent Agent or Powerful Tool?

By the 1940s, as modern electronic digital computers emerged, the descriptions included "awesome thinking machines" and "electronic brains." Dianne Martin's extensive review [36] includes her concern that "the attitude research conducted over the past 25 years suggests that the "awesome thinking machine" myth may have in fact retarded public acceptance of computers in the work environment, at the same time that it raised unrealistic expectations for easy solutions to difficult social problems."

In 1950, Alan Turing provoked huge interest with his essay "Computing Machinery and Intelligence," in which he raises the question: "Can Machines Think?" He proposed what has come to be known as the Turing test or the imitation game [64]. This thoughtful analysis catalogs objections, but he ends with "we may hope that machines will eventually compete with men in all purely intellectual fields." Many AI researchers who pursue the emulation goal have taken up Turing's challenge by developing machines that are capable of carrying out human tasks, such as playing chess, understanding images, and delivering customer support. The January 2016 issue of *AI Magazine* was devoted to articles with many new forms of Turing tests [34].

A related early, but more nuanced, vision came in J. C. R. Licklider's 1960 description of "man–computer symbiosis," which acknowledged differences between humans and computers, but stated that they would be cooperative interaction partners with computers doing the routine work and humans having insights and making decisions [30].

The widespread use of terms, such as smart, intelligent, knowledgeable, and thinking helped propagated terminology, such as machine learning, deep learning, and the idea that computers were being trained. Neuroscience descriptions of human brains as neural networks were taken up enthusiastically as a metaphor for describing AI methods, further spreading the idea that computers were like people.

IBM latched on to the term "cognitive computing" to describe their work on the Watson system. However, IBM's Design Director recently reported it "was just too confusing for people to understand" and added that "we say AI, but even that, we clarify as augmented intelligence." Google has long branded itself as strong on AI, but their current effort emphasizes "people and AI research" (PAIR, https://pair.withgoogle.com/). It appears that those pursuing the application goal increasingly recognize that the suggestion of computer intelligence should be tempered with a human-centered approach for commercial products and services.

Journalists have often been eager proponents of the idea that computers were thinking and that robots would be taking our jobs. Cover stories with computer-based characters were featured in popular magazines, such as *Newsweek* in 1980, which reported on "Machines that think," and *Time* magazine in 1996, which asked "Can Machines Think?"

Graphic artists have been all too eager to show thinking machines, especially featuring human-like heads and hands, which reinforce the idea that humans and computers are similar. A common theme was a robot hand reaching out to grasp a human hand. Popular culture in the form of Hollywood movies offered sentient computer characters, such as HAL in the 1968 film *2001: A Space Odyssey* and C3PO in the 1977 *Star Wars*. Human-like robots also played central roles in films, such as *The Terminator*, *The Matrix*, *Wall-E*, *Robot and Frank*, *Her*, and *Ex Machina* [63].

Computers were increasingly portrayed as independent actors or agents that ("who") could think, create, discover, and communicate. Journalists and headline writers were attracted to these notions producing headlines such as:

> "Hubble Accidentally Discovers a New Galaxy in Cosmic Neighborhood (NASA). The Fantastic Machine That Found the Higgs Boson (The Atlantic). AI Finds Disease-Related Genes (ScienceDaily.com). Machines Learn Chemistry (ScienceDaily.com)."

Many writers voiced the alternate view that computers were powerful tools that could amplify, augment, empower, and enhance humans. However, that view never became as popular as the seductive notion that computers were gaining capabilities to match or exceed humans [37]. Nevertheless, some researchers produced influential results, such as Engelbart [12], who gave an early vision of what it meant to augment human intellect and made a famed demonstration at the Fall Joint Computer Conference [13]. Markoff [35] carefully traces the history of AI versus intelligence augmentation (AI versus IA), describing controversies, personalities, and motivations. There is a growing belief that there are productive ways to pursue both the emulation and application goals.

Application goal developers were more likely to be influenced by designs of tool-like products and services. They were influenced by many guidelines documents, such as The Apple Human Interface Design Guidelines [1], which included two clear principles:

> "*user control:* ... people—not apps—are in control... it is usually a mistake for the app to take over the decision making" and "*flexibility:* ... (give) users complete, fine-grained control over their work."

The attraction of terminology and metaphors from the emulation goal remain prominent, as in numerous AI conferences. The conference papers may describe applications, but their approach is often to have a computer carry out the task automatically, as in reading mammograms or self-driving cars. However, there are strong application goal viewpoints, which describe designs in which humans operate tool-like devices and mechanoid appliances. These viewpoints are more likely to emerge at conferences, such as augmented humans (https://augmented-humans.org/) and many HCI conferences.

Application developers, who produce three million applications in the Apple and Android stores have largely built tool-like user interfaces, even when there is ample AI technology at work internally. These developers appreciate that users often expect a device that is comprehensible, predictable, and under their control.

The compromise design could be to build emulation goal technologies for internal operations, while the users see empowering interfaces that give them clear choices, as in GPS navigation systems, Web search, e-commerce, and recommender systems. Heer [18] showed three ways of using AI-guided methods in support of human control in data cleaning, exploratory data visualization, and natural language translation. A familiar example of well-designed integration of automated features and human control is the cell phone digital camera. These widely used devices employ AI-guided features, such as high dynamic range lighting control, jitter removal, and automatic focus, but give users control over composition, portrait modes, filters, and social media postings.

IV. Simulated Teammate or Teleoperated Device?

A common theme in design for robots and advanced technologies is that human–human interaction is a good model for human–robot interaction [24], [26] and that emotional attachment to embodied robots is an asset [69]. Many designers never consider alternatives, believing that the way people communicate with each other, coordinate activities, and form teams are the only model for design. The repeated missteps stemming from this assumption do not deter others who believe that this time will be different, that the technology is now more advanced, and that their approach is novel. Klein *et al.* [25] clarified the realistic challenges of making machines that behave as effectively as human teammates.

My objection is that human partners, teammates, and collaborators are very different from computers. I believe that it is helpful to remember that "computers are not people and people are not computers." Boden *et al.* [4] also made a simple clear statement: "robots are simply not people." The differences include the following.

Responsibility: Computers are not responsible participants, neither legally nor morally. They are never liable or accountable. They are a different category from humans. This continues to be true in our legal system and I think it will remain so. Boden *et al.* [4] offered this straightforward principle: "humans, not robots, are responsible agents." This principle is especially true in the military, where the chain of command and responsibility is taken seriously [57].

Pilots of advanced fighter jets with ample automation still think of themselves as in the control of the plane and responsible for their successful missions, even though they must adhere to their commander's orders and the rules of engagement. Astronauts rejected designs of the early Mercury capsules which had no window to eyeball the reentry if they had to do it manually—they wanted to be in control when necessary, yet responsive to mission control's rules. Neil Armstrong landed the Lunar Module on the moon—he was in charge, even though there was ample automation. The Lunar Module was not his partner. The Mars Rovers are not teammates; they are advanced automation with excellent integration of human teleoperation with high levels of automatic operation.

It is instructive that the U.S. Air Force shifted from using the term unmanned autonomous/aerial vehicles (UAVs) to remotely piloted vehicles (RPVs) to clarify responsibility. The Canadian Government [10] has a rich set of knowledge requirements that candidates must have to be granted a license to operate a remotely piloted aircraft system (RPAS).

Designers and marketers of commercial products and services take into account legal issues of accountability and liability, in which humans or organizations are the responsible parties [9]. Commercial activity is further shaped by independent oversight mechanisms, such as government regulation, industry voluntary standards, and insurance requirements.

Distinctive Capabilities: Computers have distinctive capabilities of algorithms, databases, sensors, effectors, etc. To buy into the metaphor of "teammate" has led to many design mistakes, which produce suboptimal performance. One robot rescue design team described their project to program natural language text messages that the robot would send to the operators. The messages described what the robot was "seeing," when a video or photograph could deliver much more detailed information more rapidly. Why settle for a human-like design when designs that make full use of distinctive computer capabilities would be more effective.

Designers who pursue advanced technologies will creatively find ways to empower people to be a 1000 times as effective as they have been—that is what microscopes, telescopes, bulldozers, ships, and planes have done and it is what digital cameras, Google Maps, Web search, etc., have done for people. Cameras, telescopes, cars, dishwashers, and pacemakers are not seen as teammates—they are tools that empower, enhance, amplify, and augment people.

Human Creativity: The human is always the creative force—for discovery, innovation, art, music, etc. Scientific papers are always authored by people, even when powerful computers are used. Artworks and music compositions are credited to humans, even if rich technologies with AI are heavily used.

Those who push the teammate metaphor to the limits are often led down the path of making humanoid designs, which have a long history of appealing robots, but limited commercial successes. I do not think this will change. I do not think rescue robots, bomb disposal, or eldercare robots will be human like—there are better design possibilities. The DaVinci surgical robot is nothing like a human in form or performance; it is not a teammate. It is a well-designed teleoperated machine that enables surgeons to perform precise actions in difficult to reach small body cavities. As Mumford reminds designers, successful technologies diverge from human forms.

In fact, many so-called robotic devices have a high degree of teleoperation, in which an operator controls many aspects. For example, drones, which are often described as an AI-guided technology, are generally teleoperated, even though they have the capacity to automatically hover or orbit at a fixed altitude, return to their take-off point, or follow a series of operator-chosen GPS waypoints. The NASA Mars Rover vehicles also have a rich mixture of teleoperated features and independent movement capabilities, guided by sensors to detect obstacles or precipices, with plans to circumvent them.

The language of "teleoperated instruments" or "telepresence" suggests alternative design possibilities that go beyond autonomy. These instruments enable remote operation, more careful control of devices such as when telepathologists control a remote microscope to study tissue samples. Other terms favored by the application goal community include "orthotics," such as eyeglasses or foot supports to improve performance and "prosthetics" such as replacements for missing limbs and exoskeletons that increase a human's capacity to lift heavy objects. While terms, such as "implement" or "utensil" convey modest capabilities, they constructively clarify that the user is in control.

The compromise design might be to take limited, yet mature and proven features of teammate models and embed them in designs that support human augmentation by direct or teleoperated devices. Emulation goal results can be put to work to handle sensor inputs, make well-understood decisions, and take actions whose results are predictable, leaving higher level choices to human operators.

V. Autonomous System or Supervisory Control?

Computer autonomy is an attractive emulation goal for many AI researchers, developers, journalists, and promoters. Computer autonomy has become a widely used term to describe an independently functioning machine, not directly under human control. The U.S. Defense Science Board [66] makes this definition:

> "Autonomy results from delegation of a decision to an authorized entity to take action within specific boundaries. An important distinction is that systems governed by prescriptive rules that permit no deviations are *automated*, but they are not *autonomous*. To be fully autonomous, a system must have the capability to independently compose and select among different courses of action to accomplish goals based on its knowledge and understanding of the world, itself, and the situation."

However, the U.S. Defense Science Board [65] cautioned that:

> "Unfortunately, the word "autonomy" often conjures images in the press and the minds of some military leaders of computers making independent decisions and taking uncontrolled action. It should be made clear that all autonomous systems are supervised by human operators at some level, and autonomous systems' software embodies the designed limits on the actions and decisions delegated to the computer. Autonomy is, by itself, not a solution to any problem."

This warning highlights the reality that humans and machines are embedded in complex organizational and social systems, making interdependence an important goal as well. Since humans remain as responsible actors (legally, morally, and ethically), should not computers be designed in ways that assure user control? The compromise design is that some features can be made autonomous if they are comprehensible, predictable, and controllable while giving the users the overall control that they expect.

While enthusiasm for fully autonomous systems remains high and may be valuable as a research goal, the realities of usage have been troubling. Autonomous high-speed financial trading systems have produced several billion-dollar financial crashes, but more troubling are deadly outcomes, such as the Patriot missile system shooting down two friendly aircraft during the Iraq War [3] or the 2016 crash of a Tesla while on Autopilot [67]. Maybe the most dramatic examples are the 2018 and 2019 crashes of the Boeing 737 MAX crashes, caused by the autonomous MCAS system, which took over some aircraft controls without even informing the pilots [45].

Some of the problems caused by autonomy are captured in Robin Murphy's Law of autonomous robots: "any deployment of robotic systems will fall short of the target level of autonomy, creating or exacerbating a shortfall in mechanisms for coordination with human problem holders" [68].

Those who faced the realities of dealing with application goals have repeatedly described the dangers of full computer autonomy. An early commentary in 1983 gently described the ironies of autonomy, which instead of lightening the operator's workload increased their workload because continuous monitoring of the autonomous computer was necessary [2]. These operators are unsure of what the computer will do, yet they are responsible for the outcome [3], [16], [41].

Other concerns were the difficulty of humans remaining vigilant when there was little for them to do, the challenge of rapidly taking over when problems arose, and the struggle to maintain skills when they need to take over operations. These ironies of vigilance, rapid transition, and deskilling of operators remain relevant because the operators are responsible for the outcomes [62].

Bradshaw *et al.* [5] made more forceful comments in a strongly worded paper on the "seven deadly myths of autonomous systems," which makes the bold statement that "there is nothing worse than a so-called smart machine that cannot tell you what it is doing, why it is doing something, or when it will finish. Even more frustrating—or dangerous—is a machine that is incapable of responding to human direction when something (inevitably) goes wrong." Bradshaw *et al.* also made the devastating remark that believers in full computer autonomy "have succumbed to myths of autonomy that is not only damaging in their own right but are also damaging by their continued propagation, because they engender a host of other serious misconceptions and consequences."

Even human factor specialists who support the autonomy goal describe conundrums: "as more autonomy is added to a system, and its reliability and robustness increase, the lower the situation awareness of human operators and the less likely that they will be able to take over manual control when needed" [11].

A consequential debate continues around the dangers of lethal autonomous weapons (LAWS), which could select targets and launch deadly missiles without human intervention. A vigorous effort to ban these weapons, much as land mines have been banned, has attracted almost 5000 signatures (https://autonomousweapons.org/). A regular United Nations Convention on Certain Conventional Weapons in Geneva attracts representatives of 125 countries who are drafting a treaty restricting the use of LAWS. Their case is bolstered by reports from cognitive science researchers [19], who document the failures, dangers, and costs of autonomous weapons. However, some military leaders do not wish to be limited, when they fear that adversaries will adopt autonomous weapons.

In contrast, supervisory control supports human operation and oversight by providing continuous situation awareness, clear mental model, rich control panel, and extensive feedback from actions. Supervisory control, telerobotics, and automation were extensively described by Sheridan [54], who sought to define the space between detailed manual and fully automatic control, to clarify human responsibility for the operation of industrial control rooms, robots, elevators, and washing machines.

Supervisory control suggests human decision making for setting goals with computers carrying out predictable tasks with low-level physical actions guided by sensors and carried out by effectors. Automobile automatic transmissions are a familiar example. In electronic systems, such as e-mail or e-commerce, users carry out their tasks of sending messages or ordering products, getting feedback on what has happened, with alerts if an e-mail bounces or a product shipment is delayed. In mature systems, users have a clear mental model of what the device or system is doing, with interlocks to prevent unintended actions, alerts when problems arise, and the capacity to intervene when undesired actions occur.

Contemporary versions of supervisory control have richer designs in which there may be several forms of human control. For example, aircraft pilots and co-pilots in airplanes work closely with air traffic controllers based in local centers (Terminal Radar Approach Control, TRACON) and 20 regional control rooms (Air Route Traffic Control Center, ARTCC) to coordinate flights in the national airspace. Similarly, hospital, transportation, power, stock market, military, and other complex systems have multiple layers of supervisory control, within which there may be many AI-guided components.

VI. Humanoid Robots or Mechanoid Appliances?

Visions of animated human-like robots go back at least to ancient Greek sources, but maybe one of the most startling successes was in the 1770s. Swiss Watchmaker Pierre Jaquet-Droz created elaborate programmable mechanical devices with human faces, limbs, and clothes. The writer used a quill pen on paper, the musician played a piano, and the draughtsman drew pictures, but these became only museum pieces for the Art and History Museum in Neufchatel.

The idea of human-created characters gained acceptance with classic stories such as the Golem created by the 16th-century rabbi of Prague and Mary Shelley's Frankenstein in 1818. Children stories tell of the puppet-maker Geppetto whose wooden Pinocchio comes to life and the anthropomorphic Tootle the Train character who refuses to follow the rules of staying on the track. In Goethe's *Sorcerer's Apprentice,* the protagonist conjures up a broomstick character to fetch pails of water, but when the water begins to flood the workshop, he cannot shut it down. Worse still, splitting it in half only generates twice as many broomsticks. In the 20th century, the metaphors and the language used to discuss animated human-like *robots* are usually traced back to Karel Capek's 1920 play *Rossum's Universal Robots*.

These examples illustrate the idea of humanoid robots, some are mechanical, biological, or made from materials, such as clay or wood, but they usually have human characteristics, such as two legs, a torso, arms, and a face with eyes, nose, mouth, and ears. They may make facial expressions and head gestures, while speaking in human-like voices, expressing emotion and showing personality [42].

These captivating humanoid robots have strong entertainment value that went beyond mere puppets because they seemed to operate autonomously. Children and many adults are enthusiastic about robots as film characters, engaged with robot toys, and eager to build their own robots [40]. But moving from entertainment to devices that serve application goals has proven to be difficult, except for medical mannequins and crash test dummies.

One example of how humanoid robot concepts were misleading is the design or early robot arms. The arms were typically 18 in long, had five fingers, had a wrist that rotated only 270°, and could lift about 20 pounds. Eventually, the demands of industrial automation led to flexible manufacturing systems and powerful dexterous robot arms, without humanoid forms, just as Lewis Mumford would predict.

Fig. 2. Mobile mechanoid robot for moving boxes in a warehouse from Boston Dynamics. (Handle robot image provided courtesy of Boston Dynamics, Inc.)

Serious researchers, companies, and even government agencies have created humanoid robots. The U.S. Postal Service created a life-sized human-like Postal Buddy in 1993 with plans to install 10 000 machines. However, they shut down the project after consumers rejected the 183 Postal Buddy kiosks [38]. Many designs for anthropomorphic bank tellers disappeared because of consumer disapproval. Contemporary banking systems, usually shun the name automatic teller machines, in favor of automatic transaction machines or cash machines, which support patrons getting their tasks done quickly without distracting conversation or human bank teller avatars.

Other related missteps were Microsoft's 1995 BOB, in which friendly characters would help users do their tasks and Microsoft's Office 1997 Clippy (Clippit), a too chatty character that popped up to offer help. Ananova, a Web-based news reading avatar, launched in 2000, was terminated, but the idea was revived by Chinese developers for the state news agency Xinhua in 2018 [27]. Even cheerful on-screen characters, such as Ken the Butler in Apple's famed 1987 Knowledge Navigator video and avatars in Intelligent Tutoring Systems have vanished. They distracted users from the tasks they were trying to accomplish. A careful review of 52 studies of robot failures provides guidance that could lead to greater success [21].

Manufacturer Honda created an almost life-sized humanoid robot named Asimo, which was featured at trade events and widely reported in the media, but no commercial products are planned [20]. A recent dramatic news event was when David Hanson's Social Robotics company, whose motto is "we bring robots to life," produced a talking robot named Sophia, which gained Saudi Arabian citizenship [17]. These publicity stunts draw wide attention from the media, but they have not led to commercial successes. Cynthia Breazeal's two decades of heavily promoted demonstrations of emotive robots, such as Kismet [6], culminated in a business startup, Jibo, which closed in 2019. Another social robot startup, Mayfield Robotics, produced Kuri, but it also closed in 2019.

Some companies are managing to turn impressive demonstrations into promising products. Boston Dynamics (https://www.bostondynamics.com), which began with two-legged two-armed humanoid robots, have shifted to wheel-driven mechanoid robots with vacuum suction for picking up packages in warehouses (Fig. 2).

Since its 2014 introduction, the Pepper robot, a four-foot high humanoid shape, with expressive head, arm, and hand movements, has a three-wheeled base for mobility. It is described as "pepper was optimized for human interaction and is able to engage with people through conversation and

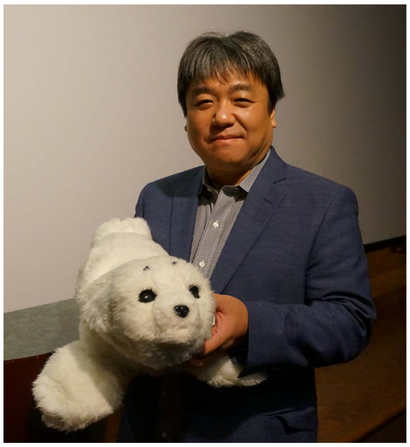

Fig. 3. Dr. Takanori Shibata holding his creation, PARO, a robot therapy device in June 2018.

Fig. 4. Roomba 700 series robotic vacuum cleaner sold by iRobot.

his touch screen" (https://www.softbankrobotics.com/emea/en/pepper). Its appealing design and conversational capacity generated strong interest, leading to sales of 15 000 units. Pepper is promoted for tasks, such as customer welcoming, product information delivery, exhibit or store guide, and satisfaction survey administration [47].

In Japan, which is often portrayed as eager for gadgets and robots, a robot-staffed hotel, closed in 2019 after a few months. The company president remarked "when you actually use robots you realize that there are places where they are not needed—or just annoy people" [15]. At the same time, traditional mechanoid automated soft drink, candy, and other dispensers are widely successful in Japan and elsewhere.

Controversy continues around the uses of humanoid robots for autism therapy. Some studies report benefits from using robots with children who may have difficulty with human relationships, but eagerly engage with robots, possibly paving the way for improved relationships with people [50]. Critics suggest that the focus on technology, rather than the child, leads to early successes, but less durable outcomes. McBride [39] worried that "if we view the human as something more than a machine, we cannot possibly devolve the responsibility for a therapeutic relationship to a mechanical toy."

Other controversies deal with eldercare social robots, especially for those with cognitive disorders and dementia. Small studies of newly introduced social robots (some humanoid and some mechanoid) have elicited sympathy from many users, especially in nursing homes, but long-term evidence is still needed to respond to the critics. The PARO therapeutic robot (www.parobots.com) is a synthetic fur-covered white seal-like robot (Fig. 3) that has touch, light sound, temperature, and posture sensors so that it "responds as if it is alive, moving its head and legs, making sounds, and imitates the voice of a real baby harp seal."

PARO has been approved by the U.S. Food and Drug Administration as a Class 2 medical device. Some studies conducted during 15 years report successes in producing positive responses from patients ("it is like a buddy," "it is a conversation piece," and "it makes me happy") and indications of potential therapeutic improvements [23]. However, these are typically short-term studies at the introduction of PARO, when patients are intrigued by the novelty, but the long-term use is still to be studied. SONY's dog robot, AIBO, remains a popular demonstration but has limited commercial success.

A notable consumer success comes from irobot (www.irobot.com), especially, Roomba, a floor-cleaning machine, and related products for mopping floors, mowing lawns, and cleaning swimming pools (Fig. 4). These mechanoid robots have mobility, sensors, and complex algorithms to map spaces while avoiding obstacles. Many kinds of mechanoid robots, often remotely controlled, have been used for bomb disposal, military applications in Afghanistan, and disaster response.

A notable success for the emulation goal is the speech-based virtual assistants, such as Apple's Siri, Amazon's Alexa, Google's Home, or Microsoft's Cortana. The designers have produced modestly reliable speech recognition and question answering systems, with high-quality speech generation that have gained consumer acceptance. This success may suggest other opportunities [32], [53], but these devices are typically simple cylinders, but no human forms. Humanoid virtual assistants have yet to prove successful. Even talking dolls have failed to draw consumer success, from Thomas Edison's efforts in the 1880s to the Mattell Talking Barbie in 1992 and a more ambitious Hello Barbie version in 2015 [31]. Mattell has no plans to pursue a talking Barbie.

Despite the modest commercial adoption of humanoid robots, many researchers and entrepreneurs still believe that they will eventually succeed. The academic research reports present a mixed view with studies from developers showing user satisfaction and sometimes delight, while others studies show the preference for more tool-like mechanoid designs, which adhere to the principles of giving users control of comprehensible, predictable, and controllable interfaces [7]. An academic survey of 1489 participants studied fear of autonomous robots and AI (FARAI). This fear, dubbed *robophobia* [59], ranged from slight afraid to afraid for 20.1% and afraid to very afraid for 18.5% [29]. Milder forms of concern over the uncanny valley, where near-human designs are distrusted, are a more common response [61].

Human willingness to engage with social robots was the focus of dozens of studies conducted by Clifford Nass, a Stanford psychologist, and his students [49]. They found that people were quick to respond to social robots, accepting them as valid partners in interactions. However, the central question remains: would people perform more effectively and prefer a more tool-like mechanoid design? Human control and operation of interfaces is a key concept in millions of applications that are based on direct manipulation [55], [56]. It is also the lesson from developers of banking machines and almost all other mobile devices, household appliances, office technologies, and e-commerce websites.

Early anthropomorphic designs and humanoid robots gave way to functional banking machines that support user control, without the deception of having a human-like bank teller machine or screen representation of a human bank teller.

This deception led bank users to wonder how else their bank was deceiving them, thereby undermining the trust needed in commercial transactions. Even leading AI researchers, such as Russell [51], clearly state that: "there is no good reason for robots to have humanoid form…they represent a form of dishonesty."

These historical precedents can provide useful guidance for contemporary designers pursuing the application goal since many still believe that improved designs of humanoid robots based on the emulation goal will eventually succeed. A commonly mentioned application is elder care in which users wishing to live independently at home will need a humanoid robot to use a kitchen with implements designed for humans, to navigate hallways and stairs, and to accommodate tasks, such as administering medications or offering a cup of tea. Inspired by the emulation goal, they see humanoid robots as the only way to provide these services, especially when the requirement is to work in environments built for human activity. However, I believe that if the imagination of these designers was more open they would see new possibilities. A pointed scenario is that if transported back to 1880, they would have proposed clothes washing robots that picked up a bar of soap and a washboard to scrub clothes, rinsed them in a sink, and hung them on a clothesline to dry. Designers of modern washer/dryers have gone well beyond humanoid robots to make mechanoid successes.

Similarly, Amazon fulfillment centers have many robots for moving products and packing boxes, but none of them are humanoid (https://www.amazonrobotics.com). Robin Murphy, a leader in developing, testing, and fielding rescue robots advocates agile mechanoid robots that can go under buildings or through ventilation ducts and teleoperated drones that can fly into dangerous places to provide video for human decision makers [44].

The compromise strategy could be to use limited humanoid services, which have proven acceptance, such as voice-operated virtual assistants embedded in mechanoid designs. Exploration of pet-like or human-like devices for therapeutic and commercial services could be further refined with long-term studies to understand what solutions remain appealing over time while being safe and effective.

VII. Conclusion

In summary, this article focused on a simplified version of just two prominent AI research goals. The first is the emulation goal for understanding human perceptual, cognitive, and motor skills to build computers that match or exceed human performance. The second is the application goal, which does HCAI research for developing successful and widely used commercial products and services. Both make valuable contributions, which researchers should pursue to bring societal benefits.

Problems arise when assumptions from one goal are used to drive work on the other goal. For example, humanoid robots remain a popular emulation goal, but humanoid robots have had far less commercial success than tool-like appliances or teleoperated devices. Understanding the mismatches could lead to designs for widely used products and services. Four such mismatches in conception and terminology were discussed: 1) intelligent agent or powerful tool; 2) simulated teammate or teleoperated device; 3) autonomous system or supervisory control; and 4) humanoid robot or mechanoid appliance.

The emulation goal inspires many researchers and creates widespread public interest. Powerful AI methods, such as machine learning, make possible recommender systems, speech recognition, image understanding, and natural language processing. When designer combine these AI methods with HCI-based user requirements gathering, design iteration, guidelines reviews, and usability testing, valuable products and services often emerge. Many other principles guide successful outcomes, such as supporting human self-efficacy, encouraging human creativity, and facilitating social participation.

Design compromises, which combine AI with HCI methods, need to be further shaped by the contextual needs of each application domain and thoroughly tested with real users. Then, the resulting products and services have a high chance of serving human needs in business, education, healthcare, environmental preservation, and community safety.

Acknowledgment

The author would like to thank Roba Abbas, Nancy Cooke, Mary Czerwinski, Hal Daume, Bonnie Dorr, Gerhard Fischer, Robert Fraser, Harry Hochheiser, Eric Hughes, Jonathan Lazar, Doug Oard, Catherine Plaisant, Jennifer Preece, Yvonne Rogers, Arnie Rosenthal, Astrid Rosenthal-von der Pütten, and Helen Sarid for their thoughtful comments on earlier drafts. The author also would like to thank the Editor-in-Chief, Katina Michael, for her support and thoughtful handling of this submission.

References

[1] Apple Computer, Inc., *Human Interface Guidelines*, Apple, Cupertino, CA, USA, 2020. [Online]. Available: https://developer.apple.com/design/human-interface-guidelines/ios/overview/themes/
[2] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.
[3] J. L. Blackhurst, J. S. Gresham, and M. O. Stone, "The autonomy paradox," *Armed Forces J.*, vol. 24, no. 5, pp. 20–40, 2011.
[4] M. Boden *et al.*, "Principles of robotics: Regulating robots in the real world," *Connection Sci.*, vol. 29, no. 2, pp. 124–129, 2017.
[5] J. M. Bradshaw, R. R. Hoffman, D. D. Woods, and M. Johnson, "The seven deadly myths of autonomous systems," *IEEE Intell. Syst.*, vol. 28, no. 3, pp. 54–61, May/Jun. 2013.
[6] C. L. Breazeal, *Designing Sociable Robots*. Cambridge, MA, USA: MIT Press, 2002.
[7] E. Broadbent, "Interactions with robots: The truths we reveal about ourselves," *Annu. Rev. Psychol.*, vol. 68, pp. 627–652, Jan. 2017.
[8] J. J. Bryson, M. E. Diamantis, and T. D. Grant, "Of, for, and by the people: The legal lacuna of synthetic persons," *Artif. Intell. Law*, vol. 25, pp. 273–291, Sep. 2017. [Online]. Available: https://doi.org/10.1007/s10506-017-9214-9
[9] R. Calo, *Robots in American Law*, Univ. Washington, Seattle, WA, USA, Mar. 2016. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2737598
[10] Canadian Government. (2019). *Knowledge Requirements for Pilots of Remotely Piloted Aircraft Systems 250G up to and Including 25 kg, Operating Within Visual Line-of-Sight (VLOS) (TP 15263)*. [Online]. Available: https://www.tc.gc.ca/en/services/aviation/publications/tp-15263.html
[11] M. R. Endsley, "From here to autonomy: Lessons learned from human–automation research," *Human Factors*, vol. 59, no. 1, pp. 5–27, 2017.
[12] D. C. Engelbart, "Augmenting human intellect: A conceptual framework," in *Multimedia From Wagner to Virtual Reality*, R. Packer and K. Jordan, Eds. New York, NY, USA: W. W. Norton and Company, 2001, pp. 64–90.
[13] D. C. Engelbart and W. K. English, "A research center for augmenting human intellect," in *Proc. Fall Joint Comput. Conf.*, 1968, pp. 395–410.
[14] B. Friedman and D. G. Hendry, *Value Sensitive Design: Shaping Technology With Moral Imagination*. Cambridge, MA, USA: MIT Press, 2019.
[15] A. Gale and T. Mochizuki, "Robot hotel loses love for robots," *Wall Street J.*, Jan. 2019. [Online]. Available: https://www.wsj.com/articles/robot-hotel-loses-love-for-robots-11547484628
[16] P. A. Hancock, "Imposing limits on autonomous systems," *Ergonomics*, vol. 60, no. 2, pp. 284–291, 2017.

[17] (2019). *Hanson Social Robotics*. [Online]. Available: https://www.hansonrobotics.com/news-meet-sophia-the-robot-that-looks-almost-human/

[18] J. Heer, "Agency plus automation: Designing artificial intelligence into interactive systems," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 6, pp. 1844–1850, 2019. [Online]. Available: https://www.pnas.org/content/116/6/1844

[19] R. R. Hoffman, T. M. Cullen, and J. K. Hawley, "The myths and costs of autonomous weapon systems," *Bull. Atomic Sci.*, vol. 72, no. 4, pp. 247–255, 2016. [Online]. Available: https://doi.org/10.1080/00963402.2016.1194619

[20] Honda. (2019). *ASIMO: The World's Most Advanced Humanoid Robot*. Accessed: May 15, 2020. [Online]. Available: https://asimo.honda.com/

[21] S. Honig and T. Oron-Gilad, "Understanding and resolving failures in human–robot interaction: Literature review and model development," *Front. Psychol.*, vol. 9, p. 861, Jun. 2018, doi: 10.3389/fpsyg.2018.00861

[22] F.-H. Hsu, *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion*. Princeton, NJ, USA: Princeton Univ. Press, 2002.

[23] L. Hung *et al.*, "Exploring the perceptions of people with dementia about the social robot PARO in a hospital setting," *Dementia*, to be published. [Online]. Available: https://doi.org/10.1177/1471301219894141

[24] P. H. Kahn *et al.*, "Design patterns for sociality in human–robot interaction," in *Proc. 3rd ACM/IEEE Int. Conf. Human Robot Interact.*, Mar. 2008, pp. 97–104.

[25] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich, "Ten challenges for making automation a 'team player' in joint human-agent activity,' *IEEE Intell. Syst.*, vol. 19, no. 6, pp. 91–95, Nov./Dec. 2004.

[26] N. C. Krämer, A. von der Pütten, and S. Eimler, "Human-agent and human–robot interaction theory: Similarities to and differences from human–human interaction," in *Human-Computer Interaction: The Agency Perspective*, Berlin, Germany: Springer, 2012, pp. 215–240, doi: 10.1007/978-3-642-25691-2_9.

[27] L. Kuo, *World's First AI News Anchor Unveiled in China*, Guardian, London, U.K., Nov. 2018. [Online]. Available: https://www.theguardian.com/world/2018/nov/09/worlds-first-ai-news-anchor-unveiled-in-china

[28] P. Laborie, J. Rogerie, P. Shaw, and P. Vilím, "IBM ILOG CP optimizer for scheduling," *Constraints*, vol. 23, no. 2, pp. 210–250, 2018.

[29] Y. Liang and S. A. Lee, "Fear of autonomous robots and artificial intelligence: Evidence from national representative data with probability sampling," *Int. J. Soc. Robot.*, vol. 9, no. 3, pp. 379–384, 2017.

[30] J. C. R. Licklider, "Man–computer symbiosis," *Trans. Human Factors Electron.*, vol. 1, pp. 4–11, Mar. 1960.

[31] K. Lobosco, *Talking Barbie Is Too Creepy for Some Parents*, CNN Bus., Atlanta, GA, USA, Mar. 12, 2015. [Online]. Available: https://money.cnn.com/2015/03/11/news/companies/creepy-hello-barbie/

[32] L. López, L. Quesada, and L. A. Guerrero, "Alexa vs. Siri vs. Cortana vs. Google assistant: A comparison of speech-based natural user interfaces," in *Advances in Human Factors and Systems Interaction. AHFE 2017* (Advances in Intelligent Systems and Computing), vol. 592, I. Nunes, Ed. Cham, Switzerland: Springer, 2018.

[33] G. Marcus and E. Davis, *Rebooting AI: Building Artificial Intelligence We Can Trust*, Penguin Random House, New York, NY, USA, 2019.

[34] G. Marcus, F. Rossi, and M. Veloso, Eds. "Beyond the Turing test," *AI Mag.*, vol. 37, no. 1, pp. 3–4, 2016. [Online]. Available: https://doi.org/10.1609/aimag.v37i1.2650

[35] J. Markoff, *Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots*. New York, NY, USA: HarperCollins, 2016.

[36] C. D. Martin, "The myth of the awesome thinking machine," *Commun. ACM*, vol. 36, no. 4, pp. 120–133, Apr. 1993, doi: 10.1145/255950.153587.

[37] C. D. Martin, "ENIAC: Press conference that shook the world," *IEEE Technol. Soc. Mag.*, vol. 14, no. 4, pp. 3–10, 1995.

[38] B. McAllister, *"Postal Buddy" Get an Unfriendly Push*, Washington Post, Washington, DC, USA, Sep. 1993. [Online]. Available: https://www.washingtonpost.com/archive/politics/1993/09/21/postal-buddy-gets-an-unfriendly-push/66554041-0218-44d0-b6d7-92a89b1da5e6/

[39] N. McBride, "Robot enhanced therapy for autistic children: An ethical analysis," *IEEE Technol. Soc. Mag.*, vol. 39, no. 1, pp. 51–60, Mar. 2020.

[40] K. Michael, D. Bowman, M. L. Jones, and R. Pringle, "Robots and socio-ethical implications," *IEEE Technol. Soc. Mag.*, vol. 37, no. 1, pp. 19–21, Mar. 2018.

[41] D. Mindell, *Our Robots, Ourselves: Robotics and the Myths of Autonomy*. New York, NY, USA: Viking Press, 2015.

[42] Y. Mou, C. Shi, T. Shen, and K. Xu, "A systematic review of the personality of robot: Mapping its conceptualization, operationalization, contextualization and effects," *Int. J. Human Comput. Interact.*, vol. 36, no. 6, pp. 591–605, 2020. [Online]. Available: https://doi.org/10.1080/10447318.2019.1663008

[43] L. Mumford, *Technics and Civilization*. Chicago, IL, USA: Univ. Chicago Press, 1934.

[44] R. R. Murphy, *Disaster Robotics*. Cambridge, MA, USA: MIT Press, 2014.

[45] J. Nicas, N. Kitroeff, D. Gelles, and J. Glanz, *Boeing Built Deadly Assumptions Into 737 Max, Blind to a Late Design Change*, New York Times, New York, NY, USA, Jun. 2019. [Online]. Available: https://www.nytimes.com/2019/06/01/business/boeing-737-max-crash.html

[46] C. O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY, USA: Crown, 2016.

[47] A. K. Pandey and R. Gelin, "A mass-produced sociable humanoid robot: Pepper: The first machine of its kind," *IEEE Robot. Autom. Mag.*, vol. 25, no. 3, pp. 40–48, Sep. 2018.

[48] D. Poole and A. Mackworth, *Artificial Intelligence: Foundations of Computational Agents*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2017.

[49] B. Reeves and C. Nass, *How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge, MA, USA: MIT Press, 1996.

[50] K. Richardson *et al.*, "Robot enhanced therapy for children with autism (DREAM): A social model of autism," *IEEE Technol. Soc. Mag.*, vol. 37, no. 1, pp. 30–39, Sep. 2018.

[51] S. Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*, Penguin Group, London, U.K., 2019.

[52] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ, USA: Prentice-Hall, 2009.

[53] A. Sciuto, A. Saini, J. Forlizzi, and J. I. Hong, "'Hey Alexa, What's Up?': A mixed-methods studies of in-home conversational agent usage," in *Proc. ACM Design. Interact. Syst. Conf.*, 2018, pp. 857–868.

[54] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA, USA: MIT Press, 1992.

[55] B. Shneiderman, "The future of interactive systems and the emergence of direct manipulation," *Behav. Inf. Technol.*, vol. 1, no. 3, pp. 237–256, 1982.

[56] B. Shneiderman, "Direct manipulation: A step beyond programming languages," *IEEE Comput.*, vol. 16, no. 8, pp. 57–69, Aug. 1983.

[57] B. Shneiderman, "Human responsibility for autonomous agents," *IEEE Intell. Syst.*, vol. 22, no. 2, pp. 60–61, Mar./Apr. 2007.

[58] B. Shneiderman, "Human-centered artificial intelligence: Reliable, safe & trustworthy," *Int. J. Human Comput. Interact.*, vol. 36, no. 6, pp. 495–504, Mar. 2020. [Online]. Available: https://doi.org/10.1080/10447318.2020.1741118

[59] T. Sprinkle. (2017). *Robophobia: Bridging the Uncanny Valley, American Society of Mechanical Engineers*. [Online]. Available: https://www.asme.org/topics-resources/content/robophobia-bridging-the-uncanny-valley

[60] Stanford University. (2016). *Artificial Intelligence and Life in 2030: One Hundred Year Study on Artificial Intelligence*. [Online]. Available: https://ai100.stanford.edu/2016-report

[61] M. K. Strait, C. Aguillon, V. Contreras, and N. Garcia, "The public's perception of humanlike robots: Online social commentary reflects an appearance-based uncanny valley, a general fear of a 'technology takeover,' and the unabashed sexualization of female-gendered robots," in *Proc. 26th IEEE Int. Symp. Robot Human Interact. Commun. (RO-MAN)*, 2017, pp. 1418–1423.

[62] B. Strauch, "Ironies of automation: Still unresolved after all these years," *IEEE Trans. Human–Mach. Syst.*, vol. 48, no. 5, pp. 419–433, Oct. 2018.

[63] S. S. Sundar, T. F. Waddell, and E. H. Jung, "The hollywood robot syndrome media effects on older adults' attitudes toward robots and adoption intentions," in *Proc. 11th ACM/IEEE Int. Conf. Human–Robot Interact. (HRI)*, Mar. 2016, pp. 343–350.

[64] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. 49, pp. 433–460, Oct. 1950.

[65] *The Role of Autonomy in DoD Systems*, U.S. Dept. Defense, Washington, DC, USA, Jul. 2012.

[66] *Summer Study on Autonomy*, U.S. Dept. Defense, Washington, DC, USA, Jun. 2016.

[67] "Collision between a car operating with automated vehicle control systems and a tractor-semitrailer truck near Williston, Florida, May 7, 2016," U.S. Nat. Transport. Safety Board, Washington, DC, USA, Rep. NTSB/HAR-17/02, Sep. 2017. [Online]. Available: https://www.ntsb.gov/investigations/accidentreports/pages/har1702.aspx

[68] D. D. Woods, J. Tittle, M. Feil, and A. Roesler, "Envisioning human–robot coordination in future operations," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 34, no. 2, pp. 210–218, May 2004. [Online]. Available: https://doi.org/10.1109/TSMCC.2004.826272

[69] S. You and L. P. Robert, "Emotional attachment, performance, and viability in teams collaborating with embodied physical action (EPA) robots," *J. Assoc. Inf. Syst.*, vol. 19, no. 5, pp. 377–407, 2018, doi: 10.17705/1jais.00496.