THE HUMACHINE

Humankind, Machines, and the Future of Enterprise

Nada R. Sanders, PhD
John D. Wood, Esq.

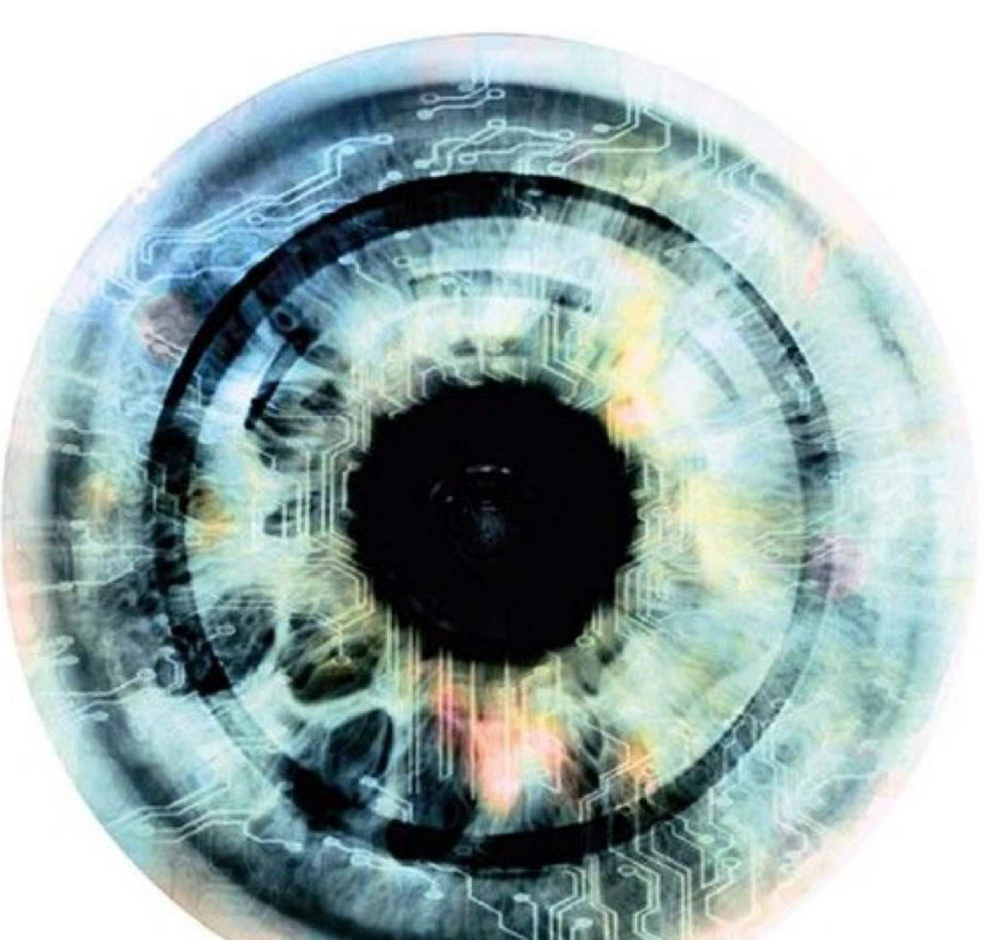

ROUTLEDGE

First published 2020
by Routledge
52 Vanderbilt Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

Library of Congress Cataloging-in-Publication Data
A catalog record for this title has been requested

ISBN: 978-1-138-57135-8 (hbk)
ISBN: 978-1-138-57134-1 (pbk)
ISBN: 978-0-429-00117-8 (ebk)

Typeset in Sabon
by Lumina Datamatics Limited

CONTENTS

LIST OF FIGURES

LIST OF TABLES

PREFACE

The competitive playing field is shifting, not slightly but seismically. A gaping fault line has opened in the ground, creating a branch in the evolutionary trajectory of human enterprise, dividing the past from the future.

On The Past side of the fault line are those companies that mistakenly treat AI as another piece of technology to tack on to your company (like Eeyore's tail). But AI does not fit the plug-and-play model of technology adaptation. It is not like updating your workforce's laptops or installing a new CRM platform. Indeed, it is unlike any technological change that has ever occurred in recorded human history.

On The Future side of the fault line are those companies implementing AI at the enterprise level, mutating into a new form of enterprise entirely—one that combines the highest capabilities of humankind with the newfound and continually emerging powers of AI.

You might think of AI as a major environmental factor being introduced to the Earth's ecosphere, which will create a sort of evolutionary pressure that will force enterprise to evolve or perish. And for those companies on The Future side of the fault line, the issue becomes how to keep their humanity amidst the tumult.

What started as an inquiry for executives into how companies could adopt AI to harness the best of human and machine capabilities turned into a much more profound rumination on the future of humanity.

This is a wake-up call for executives and business leaders. Not only should you consider implementation of AI regardless of your industry, but once you do, you should consider how to stay true to your purpose, your ethical convictions, indeed your humanity, even as our organizations continue to evolve.

INTRODUCTION

Defining the Humachine

> 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.
>
> Nick Bostrom, author of *Superintelligence: Paths, Dangers, Strategies*[1]

"Corporations are People, My Friend"

When Mitt Romney was running for President of the United States, he used a phrase during a campaign speech that got a lot of airtime, for better or worse. On August 11, 2011, at the Iowa State Fair, responding to a protestor claiming that the government should tax corporations instead of people, Romney candidly remarked, "Corporations are people, my friend."

To be fair, a corporation very much just *is* the directors, the officers, the managers, the laborers, the shareholders, and the individual purchasers—all of these *people*. When we think of "corporate activity," it really is all of these people engaging in a bundle of interconnected commercial activities that we have in mind.

When Romney says, "Corporations are people," he may have been using a linguistic convention called ***synecdoche***, "A figure of speech in which a part is made to represent the whole or vice versa."[2] Further, in a legal sense, his comment is accurate. About two hundred years ago, the Supreme Court of the United States held that corporations are legal persons:

> An aggregate corporation… is, in short, an artificial person, existing in contemplation of law and endowed with certain powers and franchises which, though they must be exercised through the medium of its natural members, are yet considered as subsisting in the corporation itself as distinctly as if it were a real personage.[3]

According to Black's Law Dictionary, an "artificial person" has been understood since the 1600s to mean an entity created by law and given certain legal rights and duties of a human being.

What makes a person, a person? The least controversial answers that are more intuitive might be that personhood means having a sense of having a mental life, selfhood or having an "I," making meaningful choices, some notion of interiority, an internal mental theater. For the sake of clarity, we have to distinguish our project from those common sense answers. We just have to set to the side the whole debate about phenomenal consciousness. We are *not* talking about the first-person subjective experience of consciousness, like what it means to see a color, taste a flavor, or hear music—all of the notions of "personhood" and "interiority" and the first-person subjective feelings of being conscious with which we are familiar. Our focus is on *rational agency*—the mental activity of thinking, of processing information, of making logical inferences, of making decisions—that are relevant to legal and business contexts. What if we took the phrase, "Corporations are people," in a different way, combining the *legal* sense with the *rational* attributes associated with human personhood—the ability to learn, reason, plan, and process complex facts?

To hold that "Corporations are people," allows us to justify attributing civil liberties to a commercial enterprise. Generally, most people don't think an organization has a mind. Relatedly, it is not common to suggest that a corporation possesses qualities of personhood. Corporations are in a certain sense merely legal fictions formed by filing some paperwork with a secretary of state in a given jurisdiction. No one thinks that a corporation has a soul, or religious beliefs, or could fall in love, or could befriend a dog, or could have a favorite flavor of ice cream, or do other things we think of when we think of a unique "person."

That said, maybe we should start to take this mere figure of speech as something literally true: a corporation can marry (merger) or divorce (divestiture), or adopt (acquisition) or procreate (subsidiaries), or die (dissolution). A corporation can speak (campaign finance contributions). A corporation can break a law or serve its nation. A corporation can have an idea. A corporation can be moral. If corporations are people, and people can attain enlightenment, then would it follow that a corporation can attain enlightenment? Can we create enterprises that operate free from bias, ignorance, greed, and fear?

We believe an organization could have a mental life in the sense that an enterprise could possess intelligence and make decisions that are not simply reducible to the decisions of specific people within the company.

What happens when humans institute procedures and practices that are functionally equivalent to human mental states but at the enterprise level? Have we created something new, a collective consciousness? What new creature is this that combines human qualities of content-rich thoughts, wisdom, and creativity with mechanical efficiency in a way that is scalable? That, we contend, would be a Humachine.

What is a Humachine?

To go back to Philosophy 101, the proper way to go about rational discourse is to first define the key terms so that one is clear about what we *mean*. Then, based on those definitions, to distinguish those terms from other terms, so that one is clear about what we do *not* mean. Then, upon those definitions and distinctions, we may make deductions. Definition, distinction, deduction—these are the "three d's" of proper philosophy. We hope by following proper philosophical method we can avoid junk science, hype, and hyperbole, which are all too common when discussing the frontiers of technology and AI in particular.

To our knowledge, the word "Humachine" first appeared on the cover of a 1999 *MIT Technology Review Special Edition*.[4] The Editor in Chief understood this was a unique phrasing: "Don't be surprised if you didn't recognize the title of this column. The word isn't in dictionaries yet. But it may be soon. Or some other word like it, coined to describe the symbiosis that is currently developing between human beings and machines. Humachines."

There are meaningful distinctions between that use of the term versus how it is deployed in this book. All of the examples given in the *MIT Technology Review* article were mere brute physical combinations of humans with machinery, such as visual implants and wearable technology—in other words, cyborgs and androids, which we are careful to distinguish.

We are talking about the emergence of a new form of intelligence in the history of life on Earth. We are *not* talking about cyborgs, androids, or AI that can parrot a human personality. We are not talking about any specific physical example that combines some human attribute with a machine, or combines some mechanical attribute with that of a human.

We begin with the Oxford English Dictionary.[5]

Human $=_{df}$ "Relating to or characteristic of humankind. … Of or characteristic of people as opposed to God or animals or machines, especially in being susceptible to weaknesses. … Showing the better qualities of humankind, such as kindness."

Machine $=_{df}$ "An apparatus using mechanical power and having several parts, each with a definite function and together performing a particular task. … Any device that transmits a force or directs its application. … An efficient and well-organized group of powerful people. … A person who acts with the mechanical efficiency of a machine."

Now combine those terms into "Humachine."

Humachine $=_{df}$ the combination of the better qualities of humankind—creativity, intuition, compassion, judgment—with the mechanical efficiency of a machine—economies of scale, big data processing capabilities, augmented by artificial intelligence, in such a way as to shed the limitations and vices of both humans and machines while maintaining the virtues of both.

In this book, we are interested in exploring combining human and machine virtues at the enterprise level—an organization, company, corporation, or other kind of organized economic undertaking. In other words, when we say "Humachines," we are talking about harnessing the power of machines to amplify human capabilities to create benign superintelligence at an enterprise level.

Humachine is, of course, a ***portmanteau***, "A word blending the sounds and combining the meanings of two others."[6]

Moravec's Paradox holds that what machines are good at, humans are not, and vice versa.[7] ***Kasparov's Law*** holds that "*weak human + machine + better process* [is] superior to a strong computer alone and, more remarkably, superior to a *strong human + machine + inferior process*.''[8] Bostrom's definition of "collective'' or ***organizational network superintelligence*** holds that superintelligence could emerge "through the gradual enhancement of networks and organizations that link individual human minds with one another and with various artifacts and bots."[9]

The Humachine is created by implementing an organizational management framework that applies Kasparov's Law to solve the problem posed by Moravec's Paradox in a way that satisfies the conditions of Bostrom's collective superintelligence.

Our book means to lay out a roadmap for the future to create some day that entity which combines the highest capabilities of humanity with the highest capabilities of machinery. By machinery we mean the nonhuman mechanisms powering the underlying processes responsible for work.

Machinery can be computational—for instance, information processing tools that extend the computational capabilities of human minds—or they can be physical—such as an industrial-sized metal cutting tool that extends the physical capabilities of human hands.

With the Industrial Revolution, physical machinery in certain ways supported and in other ways supplanted human labor. The physical capabilities of a steam-powered engine outstripped the capabilities of human muscle and sweat, not to mention animal labor. Machinery was used to operate at scales unachievable by human hands and to perform tasks involving high temperatures or health and safety issues that presented unacceptable physical risk.

Physical machinery supplanted human labor again in part, when national economies evolved from "industrial" to "service" as the primary output. Certain jobs in the service industry became co-opted by robots in disguise. The bank teller became the ATM, which is really just a very rudimentary wall-mounted robot that dispenses money for us. With the evolution from service to information-based economies, physical machinery is once again supplanting human work. Airborne drones with special sensors piloted from thousands of miles away are replacing soldiers' boots on the ground. As the economy evolved from agricultural to industrial, to service, to information-based, physical machinery continued to evolve too, shaping the economy and being shaped by it.

As the marketplace has evolved, ***computational machinery*** has also supported and supplanted human mental activity. Perhaps this displacement began with the abacus or the beloved TI-83 calculator, but certainly with the quantum computers in the nuclear laboratories: computational power is no longer bounded by the physical limits of our grey matter. Instead of a room full of junior attorneys conducting document review for a month, we can now feed documents into a natural language processing software that finds key words in a matter of seconds. This liberates well-educated human resources to focus on higher value work.

The ***humachinist*** finds ways to harness the powers of physical and computational machinery to provide optimum enhancements to human action. Our emphasis in this book is on computational machinery, such as sensing technology (e.g., RFID chips) as well as processing technology (e.g., big data analytics), rather than physical machinery. We contend that humans can harness computational machinery to create mentality at the enterprise level that transcends present-day human mental limitations.

Of course, the coming self-awareness of an enterprise portends great changes, not all of them positive. We will get into the risks posed by artificial intelligence instantiated in large-scale networks more in Chapter 6, wherein we discuss ways in which legal frameworks are challenged to help mitigate these risks and keep AI on the right track.

An incremental step in the right direction is to understand what a Humachine is and how to build one. The omega point is, in our view, a sustainable equilibrium that enables the human species to thrive on this planet and beyond.

As exciting as the endeavor to create Humachines may be, it is also dangerous. There is a very real existential concern about the "***control problem***," or building something that is so powerful it cannot be contained.

"If some day we build machine brains that surpass human brains in general intelligence, then this new superintelligence could become very powerful. And, as the fate of the gorillas now depends more on humans than on the gorillas themselves, so the fate of our species would depend on the actions of the machine superintelligence."[10] By imbuing superintelligence with humane qualities, we hope to mitigate the dangers of the control problem by creating superintelligence that regulates itself, vetoes dangerous impulses, and even protects humanity from its own vices.

We think the risk factor of expanding AI at the enterprise level is, at present, like the risk of hitting an iceberg in a vast sea of opportunity. That is to say, AI is a known risk that we should see coming with ample opportunity to avoid the danger. We know about the risk of rogue actors weaponizing AI. Lawmakers are already having hearings on this subject. We hope tech giants are taking that risk seriously because they can actually do something about it. But, because the gravity of the harm and its irreversible nature, we wish to give priority to those AI researchers like Professor Nick Bostrom who forecast an inevitable lurch toward a singleton AI that would hold the human condition in its virtual hands:

> "In principle, we could build a kind of superintelligence that would protect human values. [But] in practice, the control problem—the problem of how to control what the superintelligence would do—looks quite difficult. It also looks like we will only get one chance. Once unfriendly superintelligence exists, it would prevent us from replacing it or changing its preferences. … 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."[11]

While Bostrom's grave concerns may be, in his own modest words, "seriously wrong and misleading," nonetheless we agree with Bostrom that “the alternative views that have been presented in the literature are substantially worse—including the default view, or ‘null hypothesis,’ according to which we can for the time being safely or reasonably ignore the prospect of superintelligence.”[12] Bostrom is concerned that if we do not give the control problem sufficient attention we will make no provision for the dangers posed thereby, and we agree it is much more prudent to be overly cautious than to proceed with no safeguards in this context.

It is with respect to the control problem—howsoever remote it may be at this early stage of AI research—that we include a fairly generous helping of normative, rather than simply descriptive and tactical, ideas in this book. In our healthy and good-spirited quest for strategic competitive advantages, let’s not lose sight of what happens if we were to slip on this tightrope walk toward the future.

Summary of each Chapter

Chapter 1: The Fourth Industrial Revolution. We explain what the Fourth Industrial Revolution means for human work. We introduce Kasparov’s Law, which tells us that superior performance does not necessarily require human genius or superior machines; rather, it can be achieved through better processes using ordinary people and machines. Next, we explain Moravec’s Paradox to illustrate the symbiotic, mutually dependent relationship between human and machine. In contrast to the promise of combining the best of human and machine, we look at the peril of creating superintelligent machines. We discuss the control problem to underscore the gravity of what’s at stake in managing AI. We conclude by introducing the form of enterprise that leverages the collaboration between human and machine to achieve “superintelligence”—the Humachine.

Chapter 2: Pathways to Superintelligence. We introduce the concept of superintelligence as any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest. We lay out the various well-defined pathways to attaining superintelligence: biological cognitive enhancement, neural lace, whole brain emulation, and collective superintelligence. We argue that the latter is the most promising pathway. This appreciation for the potential of collective superintelligence motivates the rest of this book. We refute the prejudice of “species chauvinism” that suggests only human beings can have mental states. We introduce the notion of collective intentionality to bolster our contention that an enterprise could conceivably have a mind of its own. This chapter concludes with an excerpt from Teilhard de Chardin’s *Phenomenon of Man*, a beautiful depiction of his vision for the noosphere, that layer of consciousness enveloping Earth that provides the next pathway for the evolution of life itself.

Chapter 3: The Limits of Machine Capabilities. We push back against some of the hyperbole surrounding AI by illustrating the extent of its powers and its limits. We explain the role of big data, algorithms, cloud computing, and dark data as a predicate for understanding AI, machine learning, deep learning, and neural networks. While AI excels in pattern recognition and the ability to process data with super-human velocity, accuracy, volume, and consistency, AI is significantly hampered when compared to human general intelligence. AI systems lack common sense, fail to appreciate context, are data hungry and brittle, lack intuition, limit suggestions to local optima, make decisions in a black box, and are limited by the quality of the data upon which they feed.

Chapter 4: The Limits of Human Capabilities. We look at the limits of human capabilities at work. First, we discuss how demographic changes are impacting the labor force in ways that are at least as profound as botsourcing—or the displacement of human laborers with robotic labor or automated processes. Then we look at both our strengths and our weaknesses, with an eye toward implications for the work environment. We will think about thinking with guidance from Daniel

Kahneman's take on System 1 and System 2 type thinking. We try to understand in more depth what human general intelligence brings to the table when compared to AI. Our originality, intuition, emotions, care, playfulness, ethical conviction, and aesthetic taste, for example, are irreplaceable by machines. That said, human rationality is prone to specific kinds of breakdowns. We look at some well-documented human biases that affect how we work with others. We conclude by identifying uniquely human qualities that we need to cultivate, including creativity, emotional intelligence, intuition, care, ethical convictions, aesthetics, and playfulness. As machines become increasingly intelligent and enterprises shift responsibilities away from humans, we must cultivate certain human traits to remain relevant.

Chapter 5: Integration of People and Technology. We look at the evolution of work as people and artificially intelligent machines deepen collaborative relationships. We consider AI not a replacement to human workers but an enabling layer that is complementary to uniquely human skills. If we can successfully manage the human-machine interface, human learning and performance will be significantly amplified. Leveraging our relationship with AI requires us to place trust in technology. We consider the potential risk of atrophy resulting from an ongoing reliance on machines to perform work previously requiring certain human skills.

Chapter 6: Legal Issues in the Humachine Era. We lay out the control problem posed by artificial intelligence and urge readers to take this problem seriously, even as we pursue headlong the creation of a Humachine. Let's consider what laws might do to manage the risk of these AI-augmented powers. As law is a reflection of our appetite for risk, we contrast the precautionary principle with the more traditional cost-benefit analysis risk management approach. We recommend the former for dealing with the risks posed by AI. We take a deep dive into the problem of deepfakes to illustrate the challenge AI can pose to lawmakers. Issues of data privacy and security become increasingly urgent in light of the control problem because AI feeds upon, and could potentially weaponize, consumer data. The black box nature of AI decision-making and issues posed by ingraining bias within AI systems also throw curve balls at traditional legal notions of liability and fault.

Chapter 7: Breaking the Paradigm. We describe four shifts needed to detach from the traditional plug-and-play model of technology adoption. The Humachine is driven by technology but nonetheless human centric. Following the Four I's of Intentionality, Integration, Implementation, and Indication, we can create an organization that is greater than the sum of its parts. From profits to purpose, siloes to integration, and rigid performance quotas to aspirational metrics, a paradigm shift about the role of technology in the workforce lays the foundation for the Humachine.

Chapter 8: Mutations. We see the Humachine emerging in front of our eyes. Human-centric orientation, flat and fluid organizational structures, entrepreneurial and innovative cultures, and enterprise-level self-awareness define the Humachine and distinguish it from traditional business structures. We explain the necessity of purposeful, mission-driven disruption to organizational structure in order to create a cultural and technological ecosystem that supports mutation. We introduce the concepts of rendanheyi and holacracy as examples of flat and fluid organizational structures. The three variables of Kasparov's Law—people, machine, and process—need to be unified around the intentionality of the enterprise. The Humachine is capable of delivering products and services that excel in the experiential economy because they are co-created by humans and machines, workers, and consumers. Drawing intelligence through its porous organizational boundaries, the Humachine delivers innovation rapidly to market. Ultimately possessing self-awareness at the enterprise level, the Humachine is not business transformed; it is business mutated, dynamically adapting to evolutionary pressures.

Chapter 9: Reflections on the Humachine. We prognosticate on where all this is headed. We sum up some of our findings about the future of enterprise, offer some normative reflections on what executives can and should be doing to navigate the transition to create humane superintelligence, review changes to educational curriculum organized around the new discipline of "humanics" to prepare students across educational departments, and conclude with a look at what the future of work may hold in the Humachine era.

Conclusion

To distinguish this project from Bostrom, we propose that *collective* superintelligence is not only more likely, it is more desirable than *individual* superintelligence. We set forth an organizational network theory of superintelligence motivated by Kasparov's Law. Ordinary humans combined with ordinary computers with the right process can outperform human genius and supercomputing combined. Along the way, we survey technology, philosophy, psychology, economics, and other disciplines to help illuminate this multifaceted area of inquiry.

We are not fixated on some silver bullet technological breakthrough like the pill from the movie *Limitless*, or some genius programmer who writes the script for general artificial intelligence. We do not even suggest that the pot of gold at the end of this rainbow requires hiring the very best human resources. We offer a roadmap that involves using the human and technological resources available to us *now*, implementing certain processes to manage those resources effectively, and by application of Kasparov's Law, turning what is ordinary into something extraordinary. To take humans and machines, combine them in certain ways, and yield a *Humachine* capable of carrying humanity into a future that is both healthy and humane.

We share the core conviction that an enterprise can have a mind, and that, as a consequence, an enterprise can and should pursue superintelligence for the betterment of our species. That is a lofty goal. Being preoccupied with lofty goals for AI is better than to have idle hands where AI is used merely for entertainment. We need no more of the devil's playthings than our species already has.

Notes

1 Bostrom, preface, p. v.
2 Oxford English Dictionary. Available at https://en.oxforddictionaries.com/definition/synecdoche
3 *Trustees of Dartmouth Coll. v. Woodward*, 17 U.S. 518, 667 (1819) (Story, J.)
4 Intelligent Machines—Humachines, From the Editor in Chief John Benditt. *MIT Technology Review*, May 1, 1999. Available at www.technologyreview.com/s/400387/humachines/
5 Oxford English Dictionary. "human" is defined at https://en.oxforddictionaries.com/definition/human and "machine" is defined at https://en.oxforddictionaries.com/definition/machine
6 Oxford English Dictionary. https://en.oxforddictionaries.com/definition/portmanteau
7 Moravec, Hans P. 1988. *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press, Cambridge, MA.
8 Kasparov, Garry. *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins*, p. 246. Hatchet Books, New York, 2017.
9 Bostrom, pp. 58–59.
10 Bostrom, preface, page v.
11 Ibid.
12 Ibid., p. vii.

1

THE FOURTH INDUSTRIAL REVOLUTION

Previous industrial revolutions liberated humankind from animal power, made mass production possible and brought digital capabilities to billions of people. This Fourth Industrial Revolution is, however, fundamentally different. It is characterized by a range of new technologies that are fusing the physical, digital and biological worlds, impacting all disciplines, economies and industries, and even challenging ideas about what it means to be human. The resulting shifts and disruptions mean that we live in a time of great promise and great peril.

Professor Klaus Schwab, Founder and Executive Chairman of the World Economic Forum, author of *The Fourth Industrial Revolution*[1]

The business plans of the next 10,000 start-ups are easy to forecast: Take X and add AI.

Kevin Kelly, founder of Wired and former publisher of *Whole Earth Review*[2]

"Deep Blue" Blues

It was 1997 and Garry Kasparov was the greatest chess player in the world. At the age of twenty-two, he was the youngest ever undisputed World Chess Champion. He had been beating chess-playing computers since the 1980s and had just prevailed over an early version of IBM's supercomputer ***Deep Blue*** a year earlier. Now he was going for a rematch.

Garry went into the match confident. He was considered unbeatable. Now in front of a global audience, playing the fateful second match with Deep Blue, Garry was becoming visibly frustrated. He fidgeted and shook his head, waiting for his opponent's next move. After only nineteen moves, the audience saw Garry jump up and race away from the board. He had just been beaten by a machine.[3]

The match against Deep Blue put Kasparov in a philosophical mood[4]:

> I got my first glimpse of artificial intelligence on Feb. 10, 1996, at 4:45 p.m. EST, when in the first game of my match with Deep Blue, the computer nudged a pawn forward to a square where it could easily be captured. It was a wonderful and extremely human move … Humans do this sort of thing all the time. But computers generally calculate each line of play so far as possible within the time allotted … So I was stunned by this pawn sacrifice. What could it mean? I had played a lot of computers but had never experienced anything like this. I could feel—I could smell—a new kind of intelligence across the table. While I played through the rest of the game as best I could, I was lost; it played beautiful, flawless chess the rest of the way and won easily. Later I discovered the truth. Deep Blue's computational powers were so great that it did in fact calculate every possible move all the way to the actual recovery of the pawn six moves later. The computer didn't view the pawn sacrifice as a sacrifice at all. So the question is, if the computer makes the same move that I would make for completely different reasons, has it made an "intelligent" move? Is the intelligence of an action dependent on who (or what) takes it?

Did the triumph of Deep Blue over the GOAT human chess player signal a "superhuman" level of intelligence in game-playing artificial intelligence (AI)?

A special-purpose chess-playing algorithm is extremely limited: "It plays chess; it can do no other."[5] We are well-advised to bear in mind the extraordinary *narrowness* of artificial intelligence—even when rising to a superhuman level in one area of mental endeavor, that excellence does not necessarily translate into any other area of mental activity.

Yet, let's not forget the unique and singular focus and ruthless resilience of artificial intelligence when it is playing to its strength: Deep Blue does not get overcome with emotion when encountering an outrageous move by its opponent. In Kasparov's words, "Had I not melted down during game two and resigned prematurely, none of this would have mattered. Not only was the early resignation my fault, but allowing it to ruin my composure was the real fatal mistake."[6] Absent extreme heat, computers do not "melt down" under mental strain or emotional stress.

What if we could combine the labile generality of human intellects with the power and rigor of narrow AI?

Like an adamantine needle, artificial intelligence is currently as narrow as it is rigid. AI, like Deep Blue, suffers from a "lack of a purpose beyond its narrow goal," which is set out by its programmers.[7] Another way to frame this is, "Expertise does not necessarily translate into applicable understanding, let alone wisdom."[8]

Nonetheless, this was a revolutionary moment—the beginning of the world seeing the capability of thinking machines. The IBM supercomputer Deep Blue was a machine capable of processing over 100 million positions per second, nothing any human could do.

Kasparov's Law: Triumph of Process

In the fallout of the Deep Blue match, Kasparov was motivated to analyze the potential interactions, and collaborations, of human and machine thinking: "What if, instead of human versus machine, we played as partners? My brainchild saw the light of day in a match in 1998 in Leon, Spain, and we called it Advanced Chess. Each player had a PC at hand running the chess software of his choice during the game. The idea was to create the highest level of chess ever played, a synthesis of the best of man and machine."[9]

The results of this experiment were both predictable and surprising. Predictably, humans with access to machine support were less likely to make tactical blunders, because the computer would analyze potential moves and countermoves with speed and accuracy surpassing human ability. This, in turn, freed up the human player to deploy mental bandwidth on strategic analysis and creative ideation, instead of using precious (that is, far more limited) brainpower doing labor-intensive computations of the various permutations on the board. Even Kasparov, who was perhaps unparalleled among chess players for his powerful and accurate calculations (relative to other humans), lost a competitive edge under these conditions.[10]

Computers compute; when human competitors both have access to computers, the human with greater computational power loses his edge, as the computer has an equalizing effect along that dimension. As a consequence, when aided with machine computation, the human with greater creative and strategic skills will tend to prevail, all other things being equal.

Another foreseeable result was that the teams that combined amateur chess players with ordinary computers prevailed over a superhuman chess-specialist computer with no human teammate. "Human strategic guidance combined with the tactical acuity of a computer was overwhelming."[11]

However, as the "Advanced Chess" tournament came to a climax, something remarkable occurred. Based on the foregoing, you would be forgiven for betting that the champion of the Advanced Chess tournament would be a grand master (among the greatest human chess players alive) partnered up with a high-powered chess-specialist computer. You would be wrong.

The winning team was actually composed of *two amateur chess players* who deployed *three ordinary computers* simultaneously. The champions were amateurs Steven Crampton and Zackary Stephen, "chess buddies who met at a local club in New Hampshire in the US," who "had spent a few years honing their skills at the game," but still "had day jobs and were effectively unknown in the world of competitive chess."[12] Steven and Zackary entered the freestyle tournament up against no less than "several teams of grandmasters aided by computers."[13] Based on historical precedent and the merits of the competitive playing field, they should have lost. But they didn't, thanks to a unique method.

Steven and Zackary happened to have developed a database fed with over four years of data of their own personal strategies. This database showed "which of the two players typically had greater success when faced with similar situations," so they knew when to defer to their teammate and when to take the initiative.[14] They also fully utilized the rules allowing for the use of personal computers running optimization algorithms for any given arrangement of the board they happened to find themselves in.

The secret ingredient to the success was, according to one of the players, "We had really good methodology for when to use the computer and when to use our human judgement, that elevated our advantage."[15]

The shocking outcome demonstrated that "certain human skills were still unmatched by machines when it came to chess and using those skills cleverly and co-operatively could make a team unbeatable."[16]

We contend that the results of the Advanced Chess tournament are instructive well beyond the realm of chess. Kasparov's own takeaway from the surprising results actually provides a major theme for our book: "A clever process beat superior knowledge and superior technology. It didn't render knowledge and technology obsolete, of course, but it illustrated the power of efficiency and coordination to dramatically improve results."[17]

In what has become known as ***Kasparov's Law***, we can formulate the insight as: "*weak human + machine + better process* was superior to a strong computer alone and, more remarkably, superior to a *strong human + machine + inferior process*."[18]

One of the objectives for this book is to elucidate what this "better process" looks like at the enterprise level. Even chess enthusiasts can appreciate that this type of human + machine cooperation is not limited to chess and extends from diagnostic medicine to manufacturing. How can ordinary ("weak") humans combine ordinary ("weak") machines to achieve extraordinary results? By using a better process. That process is laid out in Chapters 7 and 8. We do not need to be geniuses or have access to quantum computer power in order to achieve extraordinary results. We simply need to follow Kasparov's Law.

Using ordinary people and ordinary machinery, combined with the right process, we can create the extraordinary—the Humachine. That is the goal of this book.

The New Kid in Town

While the forecasts of the extent of blue-collar and white-collar professional displacement by machines vary in magnitude, they all share the same insight: the machines are coming for human jobs of all kinds. We may call it the ***silicon wave***, as computerization, bots, autonomous machines, and so forth gobble up human jobs. At each successive interval of machine innovation, more and more jobs appear to be "low-hanging fruit," easily susceptible to displacement by machines. The silicon wave promises to cause a flood of biblical proportions, leaving behind a world transformed, no less profound than the changes wrought by the onset of electricity.

Research shows that while automation will eliminate very few occupations entirely in the next decade, it will affect portions of almost all jobs to a greater or lesser degree, depending on the type of work they entail.[19] Unless forbidden by law, no job is sacrosanct—a machine could theoretically displace *any job*.

Could a robot be President of the United States? These days, it would appear nothing is too controversial for that office. The US Constitution, Article II, Section 1 provides that, "No person except a natural born citizen, or a citizen of the United States, at the time of the adoption of this Constitution, shall be eligible to the office of President; neither shall any person be eligible to that office who shall not have attained to the age of thirty-five years, and been fourteen years a resident within the United States." At first blush, we cannot elect a robot to be the President because a robot is not a "natural-born citizen."

However, we need just a little bit of interpretive wriggle room to get there. Does the word "born" include "*assembled*"? If yes, then a computer that is "Made in America" is "born" in the USA.

Can we pass a law that grants citizenship to a robot? The humanoid AI bot christened Sophia has already been granted citizenship by Saudi Arabia, so we know that robot citizenship is possible.

Therefore, a robot assembled in the USA and granted citizenship by law, that has been in existence for no less than 35 years and which has been within the territorial limits of the USA for no less than 14 years, could theoretically run for the office of the President.

Not that President Tron would be guaranteed a strong chance of successfully navigating the political landscape, but in theory it is possible. Thirty-five years provides a lot of time for upgrades and feeding algorithms with political data.

Imagine a robot mind running the calculations of IBM's Watson to determine the absolutely pitch-perfect talking points at any given political moment, optimized to persuade a critical mass of voters. This hypothetical is just to illustrate the point that any job, even President of the United States, could theoretically be displaced by a machine. So we better make preparations for that silicon wave to wash over every sector and every level of the corporate hierarchy.

According to Ed Hess, Professor of Business Administration and Batten Executive-in-Residence at the Darden Graduate School of Business, "Because AI will be a far more formidable competitor than any human, we will be in a frantic race to stay relevant. That will require us to take our cognitive and emotional skills to a much higher level."[20]

Unfortunately, the very traits that make humans succeed where robots struggle—that is, innovative thinking and emotional intelligence—are stymied by "our natural cognitive and emotional proclivities: We are confirmation-seeking thinkers and ego-affirmation-seeking defensive reasoners." According to Hess[21]:

> We will spend more time training to be open-minded and learning to update our beliefs in response to new data. We will practice adjusting after our mistakes, and we will invest more in the skills traditionally associated with emotional intelligence. The new smart will be about trying to overcome the two big inhibitors of critical thinking and team collaboration: our ego and our fears. Doing so will make it easier to perceive reality as it is, rather than as we wish it to be. In short, we will embrace humility. That is how we humans will add value in a world of smart technology.

We simply cannot compete with machines in terms of processing speed, calculation accuracy, pattern recognition across big data sets, or the quantity of computations per second. We need to recognize the playing field is forever unleveled along those dimensions.

Let's shift our focus to where we can compete. Hess suggests we redefine what it means to be "smart," determined "not by what or how you know but by the quality of your thinking, listening, relating, collaborating, and learning." In a phrase, Hess recommends *humility* as the life raft to avoid drowning in the silicon wave.

Artificial intelligence has arrived in multifarious forms. We now live in a world of smartphones with predictive text and voice recognition, digital assistants in our home stereos, and self-steering vacuum cleaners. A recent study estimates that analytics, AI, and automation will wipe out half of today's workforce by 2055, or possibly even twenty years sooner.[22]

Those familiar with AI research may find the timing of these predictions amusing in one sense. Since the invention of the computer in the 1940s, the "expected arrival date" of human-like machine intelligence (what we can call strong general artificial intelligence) has generally been about twenty years in the future. This forecast "has been receding at a rate of one year per year."[23] "Two decades is a sweet spot for prognosticators of radical change: near enough to be attention-grabbing and relevant, yet far enough to make it possible to suppose that a string of breakthroughs, currently only vaguely imaginable, might by then have occurred."[24]

No matter what year it is, we cannot ignore the very real fact that displacement of human labor by automation has already begun. Another study by Oxford University estimates that one in every two jobs will be automated.[25] We discuss the issues around labor displacement more in subsequent chapters. Suffice it so say, robotics will revolutionize the workplace and the workforce, in every kind of service industry, in manufacturing, in policing and military force, in transportation and logistics, and even in the lesser-discussed (but no less significant) areas of the economy such as the black market.

Vernor Vinge, a math and computer scientist and science fiction author, is perhaps the original prophet of disastrous artificial intelligence predictions. He coined the term *singularity*, by which he means the tipping point occurring after we create intelligence greater than our own. We have avoided use of the term "singularity" to describe the integration of human and machine virtues. We are not throwing shade at Kurzweil or Vinge. We use "combine" and "integrate" as opposed to "singularity" because singularity has an almost rapture-esque, apocalyptical sound to it.

According to Vinge, when the singularity happens, "Human history will have reached a kind of singularity, an intellectual transition as impenetrable as the knotted space-time at the center of a black hole, and the world will pass far beyond our understanding."[26] Once the "technological means to create superhuman intelligence" emerge, "shortly thereafter, the human era will be ended."[27]

While this sounds ominous, it need not be. The end of the anthropocene era could also mean the end of living in the fear of mutually assured destruction from nuclear war, the end of famine, and the end of environmental degradation because superintelligence would have given us the security, prosperity, and abundance, respectively, required to avoid these all-too-human problems and begin to live in dignified peace.

We are optimistic about the era that happens after the human era has ended, insofar as it will be the Humachine era.

That said, these forecasts are startling and augur socioeconomic disasters of various forms.

Despite the fun of fearmongering, it is not the primary goal of this book to identify the various kinds of risks or benefits that automation or superintelligence brings, nor to point out potential safety nets (such as universal basic income or a return to more Luddite modes of being). We will have a better time managing these challenges if we follow Kasparov's Law and create a Humachine. We simply take for granted that these changes are coming and operate under the assumption that competing in this new world will require a different approach to educating future workers, and require reimagining the structure of enterprise. Keep in mind an implication of Kasparov's Law: *we need not attain genius to improve our performance; we simply need better processes.*

Details on how firms can improve process are available in the closing chapters of this book, in teach humans to become "Robot Proof," by deploying what he coins as a new educational discipline: humanics.[36]

Educational institutions should give in to the power of Moravec's Paradox instead of trying to fight against it. Aoun's theory of humanics would be a step in the right direction. ***Humanics***:

> prepares students to perform the future jobs that only human beings can do. It achieves this by fostering purposeful integration of technical literacies, such as coding and data literacy, with human literacies, such as creativity, ethics, cultural agility, and entrepreneurship. When students combine these literacies with experiential components, they integrate their knowledge with real life settings, leading to deep learning. Experiential learning is a powerful delivery system for the humanics curriculum.[37]

An education in humanics teaches us how to work alongside high-performing technologies while accentuating our uniquely human strengths. This blended approach to educational policy would empower us to do what neither the smartest person or the most advanced AI system could do alone. Our effort in this book is in furtherance of Aoun's objective of promoting humanics, not only in the classroom but in the office bullpen as well.

To be fair to humans, even those on the front lines of AI research and implementation still believe humans are indispensable from some of the most important kind of work to be done. For instance, John E. Kelly III, director of IBM Research, contends that jobs requiring higher-order critical thinking, creative thinking, innovation, and high levels of emotional engagement will need to be filled by humans for some time to come. Of course, human performance in these kinds of jobs would be greatly enhanced by harnessing the power of "cognitive systems" that enable decision-makers to penetrate complexity and crunch through big data sets to make optimal decisions.[38]

The urgency of significant educational policy reform called for by Aoun is echoed by the sentiment expressed by Kasparov: "That our classrooms still mostly look like they did a hundred years ago isn't quaint; it's absurd… Wealthy nations approach education in the same way a wealthy aristocratic family approaches investing. The status quo has been good for a long time; why rock the boat? … The prevailing attitude is that education is too important to take risks. My response is that education is too important *not* to take risks."[39]

Machines Can't Do Everything

We are in an age of infatuation with evolving technological capability. However, it is still humans—executives, managers, and other decision-makers—who use the output of algorithms to make decisions within an organizational context. These decision-makers bring their human judgment, individual personalities, opinions, and biases to the decision-making process, deciding how to use the analytically generated output. For example, UPS drivers are authorized to override the route optimization algorithm.

Humans and technology are often seen as competitors in the new world. However, the reality of the situation is more nuanced and should be approached with cautious optimism rather than fear. The strengths of one are the weaknesses of the other. Automation technologies such as machine learning and robotics play an increasingly greater role in everyday life and have a huge potential effect on the workplace. Today automation has gone beyond repetitive manufacturing activities. Robots run factories, they work side by side with physicians in operating rooms, they read X-rays, and render medical diagnoses. Analytics is used in everything from fraud detection to driving autonomous vehicles.

Yes, machines are far superior at both repetitive and nonrepetitive tasks. They have precision, strength, and do not fatigue. Humans are imprecise, overconfident, and prone to place too much faith in their intuitions. They are highly biased.

However, all machine intelligence is based on data and is *only as good as the data upon which it is based*. Machines are not good at "out-of-the-box" thinking. They are not creative or develop innovative solutions. How can an algorithm develop innovative strategy or a unique marketing campaign? Yes, they can identify photos of dogs, but they can also confuse a dachshund dog with a hot dog, and they cannot extrapolate that one is a pet while the other is a ballpark treat.

Context matters and machines don't have context. Consider the "broken-leg problem."[40] An actuarial formula based on historical data might be highly accurate in predicting the odds of someone going to the movies in a given week. But that model should be abandoned if it turns out the person in question has a broken femur. The broken leg problem demonstrates the importance of paying attention to context. We might be able to use historical, actuarial data to predict odds, but the forecast will fail if we don't appreciate the context. Obviously, without that information, the algorithm would be way off. All the data provided to the algorithm suddenly becomes irrelevant because of this change in context.

The fact is that in today's rapidly changing economy, the "broken-leg problem" occurs often. Disruptions that require understanding context and providing interpretation are simply a part of today's business environment. This could be a storm delaying a shipment, a political event, a competitor launching a new product, or a union strike. Today companies operate in volatile markets and environments, and "broken-leg problem" is a part of corporate life. Yes, algorithms can offer a far better answer provided the future facts look exactly like the data upon which forecasts are based. However, they cannot help when the environment and context change.

Machines are stronger, better, faster, more precise. Humans are intuitive, creative, and understand context. Harnessing the strengths of both—and finding the right way for them to work together—is the recipe for success. That is what this book is about. By all measures machines are getting smarter, better, faster, and stronger than humans at a variety of "mental" processes. Machines assess customer credit card risk. They detect fraud, fly airplanes, and drive cars. They can grade school essays, sort mail by reading handwritten characteristics, and diagnose illness by comparing images.

A number of inventions that enjoy widespread commercial applications were invented by sheer humanity. Accidental good fortune, fun, and humor are foreign notions to computers, but they sure drive a lot of inventions.

Consider the invention of the Post-it Note (the yellow sticky note). Post-it Notes are ubiquitous in the office environment, yet they were invented through creativity, innovation, and out-of-the-box thinking. In 1968, Dr. Spencer Silver, a scientist at 3M, was working on developing a super-strong adhesive but accidentally created a "low-tack" reusable adhesive. It was a "solution without a problem." Although he promoted it within 3M, there was no apparent use for this reusable adhesive. Then in 1974, Arthur Fry, a colleague at 3M, came up with the idea of using it as a bookmark for his hymnbook while singing in the church choir. The low-tack adhesive would not tear the fine Bible paper used in the hymnal. Thus the Post-it Note was born. If it were not for the real-world lived experience of the 3M employee, this novel application of a low-tack adhesive would not have been conceived. We cannot rely on machines to invent like this.

How about the invention of saccharin, the oldest artificial sweetener? It was discovered by accident when a researcher at Johns Hopkins forgot to wash his hands before lunch. He accidently spilled a chemical on his hands and discovered that it made the bread he was eating taste sweet. It was an "aha moment" and saccharin was born. Initial adoption was slow but then received widespread use when sugar was rationed during the First World War. Its popularity skyrocketed with the increase of diet-conscious consumers and the manufacture of Sweet'N Low and diet soft drinks. Without the happy accident of discovering the sweet taste of saccharine, Sweet'N Low would not exist. Machines do not take lunch breaks, and as a result, are unlikely to discover a new flavor by accident.

What about the beloved Slinky toy? The idea came in 1943 from a naval engineer, Richard James. He superintelligence is how we will build Humachines. Humachines are the pathway for the evolution of life on Earth.

This book is a roadmap to that end.

Notes

1 Schitzgebel, Eric. If Materialism Is True, the United States Is Probably Conscious. *Philosophical Studies*, 172:1697–1721, 1706, 2015.

2 Bostrom, *Superintelligence*, p. 26 (italics omitted).

3 Bostrom, *Superintelligence*, p. 27.

4 Bostrom, p. 63.

5 Bostrom, pp. 64, 65, and 68, respectively.

6 Bostrom, *Superintelligence*, p. 31.

7 Bostrom, *Superintelligence*, p. 31.

8 Newman, Daniel. In a World Where Data Rules, All Companies Are Tech Companies. *Forbes*, June 6, 2018. Available at www.forbes.com/sites/danielnewman/2018/06/06/in-a-world-where-data-rules-all-companies-are-tech-companies/#7124025140bf

9 www.researchgate.net/publication/320225452_Total_Consumer_Power_Consumption_Forecast

10 www.forbes.com/sites/forbestechcouncil/2017/12/15/why-energy-is-a-big-and-rapidly-growing-problem-for-data-centers/#32bac8205a30

11 Climate Home News, Part of the Guardian Environment Network. "Tsunami of Data" Could Consume One Fifth of Global Electricity by 2025. *The Guardian*, December 11, 2017. Available at www.theguardian.com/environment/2017/dec/11/tsunami-of-data-could-consume-fifth-global-electricity-by-2025

12 Andrae, Anders S.G. (Huawei). Total Consumer Power Consumption Forecast. *Presented at the Nordic Digital Business Summit*, Helsinki, Finland, October 5, 2017. Available at www.researchgate.net/publication/320225452_Total_Consumer_Power_Consumption_Forecast [accessed November 9, 2018]. Climate Home News, Part of the Guardian Environment Network. "Tsunami of Data" Could Consume One Fifth of Global Electricity by 2025. *The Guardian*, December 11, 2017. Available at www.theguardian.com/environment/2017/dec/11/tsunami-of-data-could-consume-fifth-global-electricity-by-2025

13 Danilak, Radoslav. Why Energy Is a Big and Rapidly Growing Problem for Data Centers. *FORBES*, December 15, 2017. Available at www.forbes.com/sites/forbestechcouncil/2017/12/15/why-energy-is-a-big-and-rapidly-growing-problem-for-data-centers/#e13e2ba5a307

14 Lo, Andrew W. The Origin of Bounded Rationality and Intelligence. *Proceedings of the American Philosophical Society*, 157(3):269–280, 2013 (*emphasis added*).

15 Bostrom, p. 81.

16 Bostrom, *Superintelligence*, p. 54.

17 Bostrom, *Superintelligence*, p. 53.

18 Winkler, Rolfe. Elon Musk Launches Neuralink to Connect Brains with Computers. *The Wall Street Journal*, March 27, 2017.

19 Bostrom, *Superintelligence*, p. 54.

20 Bostrom, pp. 54–55.

21 Bostrom, p. 55.

22 Nisselson, Evan. The War Over Artificial Intelligence Will Be Won with Visual Data. *TechCrunch*, May 17, 2017. Available at https://techcrunch.com/2017/05/17/the-war-over-artificial-intelligence-will-be-won-with-visual-data/

23 Ibid.

24 Bostrom, p. 55.

25 Ibid., p. 56.

26 Ibid.

27 Ibid., p. 36.

28 Ibid.

29 This may be more of a philosophical question than a practical one, but let's take whole brain emulation seriously for a second and assume it was actually done at some point in the near future. Although the hypothetical involves a genius donating her brain to science (and thus, being biologically deceased), we may be forgiven for asking if she would, in some meaningful way, still be alive? If her mind is in all respects still "intact," then would she have a first-person

perspective on the world—a personal "I," a sense of self, emotions, desires, and so forth, even after being uploaded?

30 Bostrom, p. 36.

31 Ibid.

32 Ibid., p. 40.

33 Ibid., p. 39.

34 Ibid., Sandberg, Anders. Feasibility of Whole Brain Emulation. In *Philosophy and Theory of Artificial Intelligence*, edited by Vincent C. Müller, Vol. 5, pp. 251–264. Studies in Applied Philosophy, Epistemology and Rational Ethics. Springer, New York, 2013.

35 Bostrom, pp. 58–59.

36 Ibid., p. 59.

37 Ibid.

38 Ibid., p. 60.

39 Nisselson, Evan. The War Over Artificial Intelligence Will Be Won with Visual Data. *TechCrunch*, May 17, 2017. Available at https://techcrunch.com/2017/05/17/the-war-over-artificial-intelligence-will-be-won-with-visual-data/

40 Bostrom, p. 60.

41 Ibid.

42 Ibid., 62.

43 Ibid., 68.

44 Ibid., 66.

45 Ibid.

46 Ibid.

47 Ibid., 67.

48 Ibid., 62.

49 Hofstadter, Douglass. *Gödel, Escher, Bach: An Eternal Golden Braid* (Twentieth-Anniversary Preface), p. 3. Basic Books, New York, 1999.

50 Ibid., p. 4.

51 Bostrom, p. 336.

52 Ravenscroft, Ian. *Philosophy of Mind: A Beginner's Guide*, p. 53. Oxford University Press, New York, 2005.

53 Ibid., p. 82.

54 Hofstadter, Douglass. *Gödel, Escher, Bach: An Eternal Golden Braid* (Twentieth-Anniversary Preface), p. 4. Basic Books, New York, 1999.

55 According to Jonathan Simon, postdoctoral associate in the NYU Department of Philosophy and research fellow of the NYU Center for Mind, Brain and Consciousness, "The skeptics proceed by highlighting the differences between individual minds and group mindedness, or by arguing that one can explain the relevant data (what groups do) without positing that they have cognitive states. It's generally controversial territory."

56 Rupert, Robert. Against Group Cognitive States. In *From Individual to Collective Intentionality: New Essays*, edited by Sara Rachel Chant et al., p. 98. Oxford University Press.

57 Ibid.

58 Ibid.

59 Rupert, Robert. Against Group Cognitive States. In *From Individual to Collective Intentionality: New Essays*, edited by Sara Rachel Chant et al., pp. 98–99. Oxford University Press.

60 Another way to frame Rupert's argument is that we can take some collection of things—A, B and C—then define D as the combination of A, B, and C (so, D = A + B + C). We can then talk about the causal impact of what D does in the world, but that's not a reason to take D to be a genuine causal actor in the world—whether D has a causal role depends entirely on whether A, B, and C form a suitably integrated system, or whatever criterion we might have for explaining the role of constituent parts in forming a whole. According to Simon, "It's debatable what the criterion might be, but Rupert's point would be, it's not enough to just give more examples of cases where we like to say 'D does such and so' because we could say that even if D didn't fit the criterion." That is, as long as A, B, and C are present, the world would look the same whether or not we attribute causation to D.

> Even if we concede that conscious states are reducible to brain states and are nothing over and above brain states, Professor Schitzgebel would say that it does not follow from the reducibility of group minds to individual minds (if they are reducible) and the ability to explain group mental states in terms of individual mental states (if they are explainable) that groups would not be conscious or fail to have mental states. In other words, our pains might be reducible to the activity of neurons, but that does not make them any less painful.

61 Schitzgebel, Eric. If Materialism Is True, the United States Is Probably Conscious. *Philosophical Studies*, 172:1697–1721, 2015. Abstract.

62 Ibid., p. 1706.

63 Ibid., p. 1707.

64 Reiser, Oliver. *Cosmic Humanism*, p. 557.

65 www.noosphere.princeton.edu.com.

66 Pierre Teilhard de Chardin. *The Phenomenon of Man*. Harper Perennial, New York, 1959.

> Between the last strata of the Pliocene period in which man is absent, and the next, in which the geologist is dumbfounded to find the first chipped flints, what has happened? And what is the true measure of this leap? It is our task to divine and to measure the answers to these questions before we follow step by step the march of mankind right down to the decisive stage in which it is involved today. (164–165)
>
> We have been following the successive stages of the same grand progression from the fluid contours of the early earth. Beneath the pulsations of geo-chemistry, of geo-tectonics and of geo-biology, we have detected one and the same fundamental process, always recognizable—the one which was given material form in the first cells and was continued in the construction of nervous systems. We saw geogenesis promoted to biogenesis, which turned out in the end to be nothing else than psychogenesis.
>
> With and within the crisis of reflection, the next term in the series manifests itself. Psychogenesis has led to man. Now it effaces itself, relieved or absorbed by another and a higher function—the engendering and subsequent development of the mind, in one word *noogenesis*. When for the first time in a living creature instinct perceived itself in its own mirror, the whole world took a pace forward. (181)
>
> Geologists have for long agreed in admitting the zonal composition of our planet. We have already spoken of the barysphere, central and metallic, surrounded by the rocky lithosphere that in turn is surrounded by the fluid layers of the hydrosphere and the atmosphere. Since Suess, science has rightly become accustomed to add another to these four concentric layers, the living membrane composed of the fauna and flora of the globe, the biosphere, so often mentioned in these pages, an envelope as definitely universal as the other "spheres" and even more definitely individualized than them. For, instead of representing a more or less vague grouping, it forms a single piece, of the very tissue of the genetic relations which delineate the tree of life. (182)
>
> The paradox of man resolves itself by passing beyond measure. Despite the relief and harmony it brings to things, this perspective is at first sight disconcerting, running counter as it does to the illusion and habits which incline us to measure events by their material face. It also seems to us extravagant because, steeped as we are in what is human like a fish in the sea, we have difficulty in emerging from it in our minds so as to appreciate its specificness and breadth. But let us look round us a little more carefully. This sudden deluge of cerebralisation, this biological invasion of a new animal type which gradually eliminates or subjects all forms of life that are not human, this irresistible tide of fields and factories, this immense and growing edifice of matter and ideas—all these signs that we look at, for days on end—to proclaim that there has been a change on the earth and a change of planetary magnitude.
>
> There can indeed be no doubt that, to an imaginary geologist coming one day far in the future to inspect our fossilised globe, the most astounding of the revolutions undergone by the earth would be that which took place at the beginning of what has so rightly been called the psychozoic era. And even today, to a Martian capable of analysing sidereal radiations psychically no less than physically, the first characteristic of our planet would be, not the blue of the seas or the green of the forests, but the phosphorescence of thought. (183)

67 Bostrom, preface, page v.

INDEX

Note: Page numbers in italic and bold refer to figures and tables, respectively.