What To Think About Machines That Think Todays Leading Thinkers On The Age Of Machine Intelligence

A guide to AI’s thorniest implications that asks: How shall we navigate our brave new world? We are at a monumental turning point in human history. AI is taking intelligence in new directions. The strongest human competitors in chess, go, and Jeopardy! have been beaten by AIs, and AI is getting more sophisticated by the day. Further, AI research is going inside the human brain itself, attempting to augment human minds. It may even create greater-than-human-level intelligence, leading to a new generation of artificial minds—Minds 2.0. Susan Schneider, a philosopher, argues that these undertakings must not be attempted without a richer understanding of the nature of the mind. An insufficient grasp of the underlying philosophical issues could undermine the use of AI and brain enhancement technology, bringing about the demise or suffering of conscious beings. Examining the philosophical questions lying beneath the algorithms, Schneider takes on AI’s thorniest implications.

For Readers of Ray Kurzweil and Michio Kaku, a New Look at the Cutting Edge of Artificial Intelligence Imagine a robotic stuffed animal that can read and respond to a child’s emotional state, a commercial that can recognize and change based on a customer’s facial expression, or a company that can actually create feelings as though a person were experiencing them naturally. Heart of the Machine explores the next giant step in the relationship between humans and technology: the ability of computers to recognize, respond to, and even replicate emotions. Computers have long been integral to our lives, and their advances continue at an exponential rate. Many believe that artificial intelligence equal or superior to human intelligence will happen in the not-too-distance future; some even think machine consciousness will follow. Futurist Richard Yonck argues that emotion, the first, most basic, and most natural form of communication, is at the heart of how we will soon work with and use computers. Instilling emotions into computers is the next leap in our centuries-old obsession with creating machines that replicate humans. But for every benefit this progress may bring to our lives, there is a possible pitfall. Emotion recognition could lead to advanced surveillance, and the same technology that can manipulate our feelings could become a method of mass control. And, as shown in movies like Her and Ex Machina, our society already holds a deep-seated anxiety about what might happen if machines could actually feel and break free from our control. Heart of the Machine is an exploration of the new and inevitable ways in which mankind and technology will interact. The paperback edition has a new foreword by Rana el Kaliouby, PhD, a pioneer in artificial emotional intelligence, as well as the cofounder and CEO of Affectiva, the acclaimed AI startup spun off from the MIT Media Lab.

The director of the MIT Artificial Intelligence Lab speculates about the future of humankind as it explores the relationship between humans and technologically engineered robots and examines the vast capabilities of such machines.

"A globe-spanning investigation into the Transhumanist movement, considering the tech billionaires, scientific luminaries, and DIY body-hackers attempting to prolong, improve, and ultimately transcend the limits of human life"--

AI is radically transforming business. Are you ready? Look around you. Artificial intelligence is no longer just a futuristic notion. It's here right now--in software that senses what we need, supply chains that "think" in real time, and robots that respond to changes in their environment. Twenty-first-century pioneer companies are already using AI to innovate and grow fast. The bottom line is this: Businesses that understand how to harness AI can surge ahead. Those that neglect it will fall behind. Which side are you on? In Human + Machine, Accenture leaders Paul R. Daugherty and H. James (Jim) Wilson show that the essence of the AI paradigm shift is the transformation of all business processes within an organization--whether related to breakthrough innovation, everyday customer service, or personal productivity habits. As humans and smart machines collaborate ever more closely, work processes become more fluid and adaptive, enabling companies to change them on the fly--or to completely reimagine them. AI is changing all the rules of how companies operate. Based on the authors' experience and research with 1,500 organizations, the book reveals how companies are using the new rules of AI to leap ahead on innovation and profitability, as well as what you can do to achieve similar results. It describes six entirely new types of hybrid human + machine roles that every company must develop, and it includes a "leader’s guide" with the five crucial principles required to become an AI-fueled business. Human + Machine provides the missing and much-needed management playbook for success in our new age of AI. BOOK PROCEEDS FOR THE AI GENERATION The authors' goal in publishing Human + Machine is to help executives, workers, students and others navigate the changes that AI is making to business and the economy. They believe AI will bring innovations that truly improve the way the world works and lives. However, AI will cause disruption, and many people will need education, training and support to prepare for the newly created jobs. To support this need, the authors are donating the royalties received from the sale of this book to fund education and retraining programs focused on developing fusion skills for the age of artificial intelligence.

A Wharton professor and tech entrepreneur examines how algorithms and artificial intelligence are starting to run every aspect of our lives, and how we can shape the way they impact us Through the technology embedded in almost every major tech platform and every web-enabled device, algorithms and the artificial intelligence that underlies them make a staggering number of everyday decisions for us, from what products we buy, to where we decide to eat, to how we consume our news, to whom we date, and how we find a job. We've even delegated life-and-death decisions to algorithms--decisions once made by doctors, pilots, and judges. In his new book, Kartik Hosanagar surveys the brave new world of algorithmic decision-making and reveals the potentially dangerous biases they can give rise to as they increasingly run our lives. He makes the compelling case that we need to arm ourselves with a better, deeper, more nuanced understanding of the phenomenon of algorithmic thinking. And he gives us a route in, pointing out that algorithms

often think a lot like their creators--that is, like you and me. Hosanagar draws on his experiences designing algorithms professionally--as well as on history, computer science, and psychology--to explore how algorithms work and why they occasionally go rogue, what drives our trust in them, and the many ramifications of algorithmic decision-making. He examines episodes like Microsoft's chatbot Tay, which was designed to converse on social media like a teenage girl, but instead turned sexist and racist; the fatal accidents of self-driving cars; and even our own common, and often frustrating, experiences on services like Netflix and Amazon. A Human's Guide to Machine Intelligence is an entertaining and provocative look at one of the most important developments of our time and a practical user's guide to this first wave of practical artificial intelligence.

Explores universal questions about humanity's capacity for living and thriving in the coming age of sentient machines and AI, examining debates from opposing perspectives while discussing emerging intellectual diversity and its potential role in enabling a positive life.

Sometime in the future the intelligence of machines will exceed that of human brain power. So are we on the edge of an AI-pocalypse, with superintelligent devices superseding humanity, as predicted by Stephen Hawking? Or will this herald a kind of Utopia, with machines doing a far better job at complex tasks than us? You might not realise it, but you interact with AIs every day. They route your phone calls, approve your credit card transactions and help your doctor interpret results. Driverless cars will soon be on the roads with a decision-making computer in charge. But how do machines actually think and learn? In Machines That Think, AI experts and New Scientist explore how artificial intelligence helps us understand human intelligence, machines that compose music and write stories - and ask if AI is really a threat. ABOUT THE SERIES New Scientist Instant Expert books are definitive and accessible entry points to the most important subjects in science; subjects that challenge, attract debate, invite controversy and engage the most enquiring minds. Designed for curious readers who want to know how things work and why, the Instant Expert series explores the topics that really matter and their impact on individuals, society, and the planet, translating the scientific complexities around us into language that's open to everyone, and putting new ideas and discoveries into perspective and context.

Award-winning author Don Brown explores computers and technology in book two of the Big Ideas series Machines That Think! explores machines from ancient history to today that perform a multitude of tasks, from making mind-numbing calculations to working on assembly lines. Included are fascinating looks at the world’s earliest calculators, the birth of computer programming, and the arrival of smartphones. Contributors discussed include Muhammad ibn Musa al-Khwarizmi, Ada Lovelace, and Bill Gates. From the abacus to artificial intelligence, machines through the ages have pushed the boundaries of human capability and creativity. Back matter includes a timeline, endnotes, a bibliography, an author’s note, and an index.

Algorithms will soon know more about us than we know ourselves Where should machine automation end? Is it acceptable to have a digital assistant arrange your calendar, but not to have a robot spouse? Are companion robots acceptable for seniors in need of comfort, but not okay for toddlers exposed to emotional software that could influence their behavior? Is it desirable to live a life within the virtual reality of Facebook's Oculus Rift, but not if your thoughts are sold to advertisers who manipulate your purchases? We've entered an era where a myriad of personalization algorithms influence our every decision, and the lines between human assistance, automation, and extinction have blurred. We need to create ethical standards for the Artificial Intelligence usurping our lives, and allow individuals to control their identity based on their values. Otherwise, we sacrifice our humanity for productivity versus purpose and for profits versus people. Featuring pragmatic solutions drawing on economics, emerging technologies, and positive psychology, Heartificial Intelligence provides the first values-driven approach to algorithmic living--a definitive roadmap to help humanity embrace the present and positively define their future. Each chapter opens with a fictional vignette, helping readers imagine how they would respond to various Artificial Intelligence scenarios while demonstrating the need to codify their values, as the algorithms dominating society today are already doing. Funny, poignant, and accessible, this book paints a vivid portrait of how our lives might look in either a dystopia of robotic and corporate dominance, or a utopia where humans use technology to enhance our natural abilities to evolve into a long-lived, super-intelligent, and altruistic species.

An update edition of Solomon’s Code—now The A.I. Generation—the thought-provoking examination of artificial intelligence and how it reshapes human values, trust, and power around the world. Whether in medicine, money, or love, technologies powered by forms of artificial intelligence are playing an increasingly prominent role in our lives. As we cede more decisions to thinking machines, we face new questions about staying safe, keeping a job and having a say over the direction of our lives. The answers to those questions might depend on your race, gender, age, behavior, or nationality. New AI technologies can drive cars, treat damaged brains and nudge workers to be more productive, but they also can threaten, manipulate, and alienate us from others. They can pit nation against nation, but they also can help the global community tackle some of its greatest challenges—from food crises to global climate change. In clear and accessible prose, global trends and strategy adviser Olaf Groth, AI scientist and social entrepreneur Mark Nitzberg, along with seasoned economics reporter Dan Zehr, provide a unique human-focused, global view of humanity in a world of thinking machines.

A thought-provoking and wide-ranging exploration of machine learning and the race to build computer intelligences as flexible as our own In the world's top research labs and universities, the race is on to invent the ultimate learning algorithm: one capable of discovering any knowledge from data, and doing anything we want, before we even ask. In The Master Algorithm, Pedro Domingos lifts the veil to give us a peek inside the learning machines that power Google, Amazon, and your smartphone. He assembles a blueprint for the future universal learner--the Master Algorithm--and discusses what it will mean for business, science, and society. If data-ism is today's philosophy, this book is its bible.

A public policy leader addresses how artificial intelligence is transforming the future of labor—and what we can do to protect the role of workers. As computer technology advances with dizzying speed, human workers face an ever-increasing threat of obsolescence. In Human Work In the Age of Smart Machines, Jamie Merisotis argues that we can—and must—rise to this challenge by preparing to work alongside smart machines doing that which only humans can: thinking critically, reasoning ethically, interacting interpersonally, and serving others with empathy. The president and CEO of Lumina Foundation, Merisotis offers a roadmap for the large-scale, radical changes we must make in order to find abundant and meaningful work for ourselves in the 21st century. His vision centers on developing our unique capabilities as humans through learning opportunities that deliver fair results and offer a broad range of credentials. By challenging long-held assumptions and expanding our concept of work, Merisotis argues that we can harness the population’s potential, encourage a deeper sense of community, and erase a centuries-long system of inequality.

New from Ian McEwan, Booker Prize winner and international bestselling author of Atonement and The Children Act Machines Like Me takes place in an alternative 1980s London. Charlie, drifting through life and dodging full-time employment, is in love with Miranda, a bright student who lives with a terrible secret. When Charlie comes into money, he buys Adam, one of the first synthetic humans and—with Miranda's help—he designs Adam's personality. The near-perfect human that emerges is beautiful, strong, and clever. It isn't long before a love triangle soon forms, and these three beings confront a profound moral dilemma. In his subversive new novel, Ian McEwan asks whether a machine can understand the human heart—or whether we are the ones who lack understanding.

Weighing in from the cutting-edge frontiers of science, today’s most forward-thinking minds explore the rise of “machines that think.” Stephen Hawking recently made headlines by noting, “The development of full artificial intelligence could spell the end of the human race.” Others, conversely, have trumpeted a new age of “superintelligence” in which smart devices will exponentially extend human capacities. No longer just a matter of science-fiction fantasy (2001, Blade Runner, The Terminator, Her, etc.), it is time to seriously consider the reality of intelligent technology, many forms of which are already being integrated into our daily lives. In that spirit, John Brockman, publisher of Edge. org (“the world’s smartest website” – The Guardian), asked the world’s most influential scientists, philosophers, and artists one of today’s most consequential questions: What do you think about machines that think?

This book is about the story of modern machines, and how they do the world’s work. We shall see how raw materials, such as metals, wool and cotton fibers, and wood, are turned into automobiles, ships, typewriters, fabrics, and other useful products. Our machine age has made this possible, for only by machines can these many products be made cheaply and quickly. We shall also see how machines are able to do such varied jobs as threshing grain, machining engine blocks, or multiplying long rows of numbers. And later on, we shall learn how automation, the newest marvel of the machine age, teaches machines almost to “think” for themselves.

This book explores justice in the age of artificial intelligence. It argues that current AI tools used in connection with liberty decisions are based on utilitarian frameworks of justice and inconsistent with individual fairness reflected in the US Constitution and Declaration of Independence. It uses AI risk assessment tools and lethal autonomous weapons as examples of how AI influences liberty decisions. The algorithmic design of AI risk assessment tools can and does embed human biases. Designers and users of these AI tools have allowed some degree of compromise to exist between accuracy and individual fairness.Written by a former federal judge who lectures widely and frequently on AI and the justice system, this book is the first comprehensive presentation of the theoretical framework of AI tools in the criminal justice system and lethal autonomous weapons utilized in decision-making. The book then provides a comprehensive explanation as to why, tracing the evolution of the debate regarding racial and other biases embedded in such tools. No other book delves as comprehensively into the theory and practice of AI risk assessment tools.

Thinking Machines: Machine Learning and Its Hardware Implementation covers the theory and application of machine learning, neuromorphic computing and neural networks. This is the first book that focuses on machine learning accelerators and hardware development for machine learning. It presents not only a summary of the latest trends and examples of machine learning hardware and basic knowledge of machine learning in general, but also the main issues involved in its implementation. Readers will learn what is required for the design of machine learning hardware for neuromorphic computing and/or neural networks. This is a recommended book for those who have basic knowledge of machine learning or those who want to learn more about the current trends of machine learning. Presents a clear understanding of various available machine learning hardware accelerator solutions that can be applied to selected machine learning algorithms Offers key insights into the development of hardware, from algorithms, software, logic circuits, to hardware accelerators Introduces the baseline characteristics of deep neural network models that should be treated by hardware as well Presents readers with a thorough review of past research and products, explaining how to design through ASIC and FPGA approaches for target machine learning models Surveys current trends and models in neuromorphic computing and neural network hardware architectures Outlines the strategy for advanced hardware development through the example of deep learning accelerators

Garry Kasparov's 1997 chess match against the IBM supercomputer Deep Blue was a watershed moment in the history of technology. It was the dawn of a new era in artificial intelligence: a machine capable of beating the reigning human champion at this most cerebral game. That moment was more than a century in the making, and in this breakthrough book, Kasparov reveals his astonishing side of the story for the first time. He describes how it felt to strategize against an implacable, untiring opponent with the whole world watching, and recounts the history of machine intelligence through the microcosm of chess, considered by generations of scientific pioneers to be a key to unlocking the secrets of human and machine cognition. Kasparov uses his unrivaled experience to look into the future of intelligent machines

and sees it bright with possibility. As many critics decry artificial intelligence as a menace, particularly to human jobs, Kasparov shows how humanity can rise to new heights with the help of our most extraordinary creations, rather than fear them. Deep Thinking is a tightly argued case for technological progress, from the man who stood at its precipice with his own career at stake.

Today, a scientific explanation is not meant to ascribe agency to natural phenomena: we would not say a rock falls because it seeks the center of the earth. Even for living things, in the natural sciences and often in the social sciences, the same is true. A modern botanist would not say that plants pursue sunlight. This has not always been the case, nor, perhaps, was it inevitable. Since the seventeenth century, many thinkers have made agency, in various forms, central to science. The Restless Clock examines the history of this principle, banning agency, in the life sciences. It also tells the story of dissenters embracing the opposite idea: that agency is essential to nature. The story begins with the automata of early modern Europe, as models for the new science of living things, and traces questions of science and agency through Descartes, Leibniz, Lamarck, and Darwin, among many others. Mechanist science, Jessica Riskin shows, had an associated theology: the argument from design, which found evidence for a designer in the mechanisms of nature. Rejecting such appeals to a supernatural God, the dissenters sought to naturalize agency rather than outsourcing it to a "divine engineer." Their model cast living things not as passive but as active, self-making machines. The conflict between passive- and active-mechanist approaches maintains a subterranean life in current science, shaping debates in fields such as evolutionary biology, cognitive science, and artificial intelligence. This history promises not only to inform such debates, but also our sense of the possibilities for what it means to engage in science—and even what it means to be alive.

Everything you've always wanted to know about self-driving cars, Netflix recommendations, IBM's Watson, and video game-playing computer programs. The future is here: Self-driving cars are on the streets, an algorithm gives you movie and TV recommendations, IBM's Watson triumphed on Jeopardy over puny human brains, computer programs can be trained to play Atari games. But how do all these things work? In this book, Sean Gerrish offers an engaging and accessible overview of the breakthroughs in artificial intelligence and machine learning that have made today's machines so smart. Gerrish outlines some of the key ideas that enable intelligent machines to perceive and interact with the world. He describes the software architecture that allows self-driving cars to stay on the road and to navigate crowded urban environments; the million-dollar Netflix competition for a better recommendation engine (which had an unexpected ending); and how programmers trained computers to perform certain behaviors by offering them treats, as if they were training a dog. He explains how artificial neural networks enable computers to perceive the world—and to play Atari video games better than humans. He explains Watson's famous victory on Jeopardy, and he looks at how computers play games, describing AlphaGo and Deep Blue, which beat reigning world champions at the strategy games of Go and chess. Computers have not yet mastered everything, however; Gerrish outlines the difficulties in creating intelligent agents that can successfully play video games like StarCraft that have evaded solution—at least for now. Gerrish weaves the stories behind these breakthroughs into the narrative, introducing readers to many of the researchers involved, and keeping technical details to a minimum. Science and technology buffs will find this book an essential guide to a future in which machines can outsmart people.

"A concise, insightful and sophisticated guide to maintaining humane values in an age of new machines."—The New York Times Book Review "While we need to rewrite the rules of the twenty-first-century economy, Kevin's book is a great look at how people can do this on a personal level to always put humanity first."—Andrew Yang You are being automated. After decades of hype and sci-fi fantasies, artificial intelligence is leaping out of research labs and into the center of our lives. Automation doesn't just threaten our jobs. It shapes our entire human experience, with AI and algorithms influencing the TV shows we watch, the music we listen to, the beliefs we hold, and the relationships we form. And while the age-old debate over whether automation will destroy jobs rages on, an even more important question is being ignored: How can we be happy, successful humans in a world that is increasingly built by and for machines? In Futureproof: 9 Rules for Humans in the Age of Automation, New York Times technology columnist Kevin Roose lays out a hopeful, pragmatic vision for how we can thrive in the age of AI and automation. He shares the secrets of people and organizations that have survived previous waves of technological change, and explains what skills are necessary to stay ahead of today's intelligent machines, with lessons like • Be surprising, social, and scarce. • Resist machine drift. • Leave handprints. • Demote your devices. • Treat AI like a chimp army. Roose rejects the conventional wisdom that in order to succeed in the AI age, we have to become more like machines ourselves—hyper-efficient, data-driven workhorses. Instead, he says, we should focus on being more human, and doing the kinds of creative, inspiring, and meaningful things even the most advanced robots can't do.

Offers an illustrated telling of the story of Ada Byron Lovelace, from her early creative fascination with mathematics and science and her devastating bout with measles, to the ground-breaking algorithm she wrote for Charles Babbage's analytical engine.

"The Knowledge Machine is the most stunningly illuminating book of the last several decades regarding the all-important scientific enterprise." —Rebecca Newberger Goldstein, author of Plato at the Googleplex A paradigm-shifting work, The Knowledge Machine revolutionizes our understanding of the origins and structure of science. • Why is science so powerful? • Why did it take so long—two thousand years after the invention of philosophy and mathematics—for the human race to start using science to learn the secrets of the universe? In a groundbreaking work that blends science, philosophy, and history, leading philosopher of science Michael Strevens answers these challenging questions, showing how science came about only once thinkers stumbled upon the astonishing idea that scientific breakthroughs could be accomplished by breaking the rules of logical argument. Like such classic works as Karl Popper's The Logic of Scientific Discovery and Thomas Kuhn's The Structure of Scientific Revolutions, The Knowledge Machine grapples with the meaning and origins of science, using a plethora of vivid historical examples to demonstrate that scientists willfully ignore religion, theoretical beauty, and even philosophy to embrace a constricted code of argument whose very narrowness channels unprecedented energy into empirical observation and experimentation. Strevens calls this scientific code the iron rule of explanation, and reveals the way in which the rule, precisely because it is unreasonably close-minded, overcomes individual prejudices to lead humanity inexorably toward the secrets of nature. "With a mixture of philosophical and historical argument, and written in an engrossing style" (Alan Ryan), The Knowledge Machine provides captivating portraits of some of the greatest luminaries in science's history, including Isaac Newton, the chief architect of modern science and its foundational theories of motion and gravitation; William Whewell, perhaps the greatest philosopher-scientist of the early nineteenth century; and Murray Gell-Mann, discoverer of the quark. Today, Strevens argues, in the face of threats from a changing climate and global pandemics, the idiosyncratic but highly effective scientific knowledge machine must be protected from politicians, commercial interests, and even scientists themselves who seek to open it up, to make it less narrow and more rational—and thus to undermine its devotedly empirical search for truth. Rich with illuminating and often delightfully quirky illustrations, The Knowledge Machine, written in a winningly accessible style that belies the import of its revisionist and groundbreaking concepts, radically reframes much of what we thought we knew about the origins of the modern world.

A thought-provoking examination of artificial intelligence and how it reshapes human values, trust, and power around the world. Whether in medicine, money, or love, technologies powered by forms of artificial intelligence are playing an increasingly prominent role in our lives. As we cede more decisions to thinking machines, we face new questions about staying safe, keeping a job and having a say over the direction of our lives. The answers to those questions might depend on your race, gender, age, behavior, or nationality. New AI technologies can drive cars, treat damaged brains and nudge workers to be more productive, but they also can threaten, manipulate, and alienate us from others. They can pit nation against nation, but they also can help the global community tackle some of its greatest

challenges—from food crises to global climate change. In clear and accessible prose, global trends and strategy adviser Olaf Groth, AI scientist and social entrepreneur Mark Nitzberg, along with seasoned economics reporter Dan Zehr, provide a unique human-focused, global view of humanity in a world of thinking machines.

From the former president of MIT, the story of the next technology revolution, and how it will change our lives. A century ago, discoveries in physics came together with engineering to produce an array of astonishing new technologies: radios, telephones, televisions, aircraft, radar, nuclear power, computers, the Internet, and a host of still-evolving digital tools. These technologies so radically reshaped our world that we can no longer conceive of life without them. Today, the world's population is projected to rise to well over 9.5 billion by 2050, and we are currently faced with the consequences of producing the energy that fuels, heats, and cools us. With temperatures and sea levels rising, and large portions of the globe plagued with drought, famine, and drug-resistant diseases, we need new technologies to tackle these problems. But we are on the cusp of a new convergence, argues world-renowned neuroscientist Susan Hockfield, with discoveries in biology coming together with engineering to produce another array of almost inconceivable technologies—next-generation products that have the potential to be every bit as paradigm shifting as the twentieth century's digital wonders. The Age of Living Machines describes some of the most exciting new developments and the scientists and engineers who helped create them. Virus-built batteries. Protein-based water filters. Cancer-detecting nanoparticles. Mind-reading bionic limbs. Computer-engineered crops. Together they highlight the promise of the technology revolution of the twenty-first century to overcome some of the greatest humanitarian, medical, and environmental challenges of our time.

Visionary designer and technologist John Maeda defines the fundamental laws of how computers think, and why you should care even if you aren't a programmer. "Maeda is to design what Warren Buffett is to finance." --Wired John Maeda is one of the world's preeminent interdisciplinary thinkers on technology and design. In How to Speak Machine, he offers a set of simple laws that govern not only the computers of today, but the unimaginable machines of the future. Technology is already more powerful than we can comprehend, and getting more powerful at an exponential pace. Once set in motion, algorithms never tire. And when a program's size, speed, and tirelessness combine with its ability to learn and transform itself, the outcome can be unpredictable and dangerous. Take the seemingly instant transformation of Microsoft's chatbot Tay into a hate-spewing racist, or how crime-predicting algorithms reinforce racial bias. How to Speak Machine provides a coherent framework for today's product designers, business leaders, and policymakers to grasp this brave new world. Drawing on his wide-ranging experience from engineering to computer science to design, Maeda shows how businesses and individuals can identify opportunities afforded by technology to make world-changing and inclusive products--while avoiding the pitfalls inherent to the medium.

Artificial Knowing challenges the masculine slant in the Artificial Intelligence (AI) view of the world. Alison Adam admirably fills the large gap in science and technology studies by showing us that gender bias is inscribed in AI-based computer systems. Her treatment of feminist epistemology, focusing on the ideas of the knowing subject, the nature of knowledge, rationality and language, are bound to make a significant and powerful contribution to AI studies. Drawing from theories by Donna Haraway and Sherry Turkle, and using tools of feminist epistemology, Adam provides a sustained critique of AI which interestingly re-enforces many of the traditional criticisms of the AI project. Artificial Knowing is an esential read for those interested in gender studies, science and technology studies, and philosophical debates in AI.

A scientist who has spent a career developing Artificial Intelligence takes a realistic look at the technological challenges and assesses the likely effect of AI on the future. How will Artificial Intelligence (AI) impact our lives? Toby Walsh, one of the leading AI researchers in the world, takes a critical look at the many ways in which "thinking machines" will change our world. Based on a deep understanding of the technology, Walsh describes where Artificial Intelligence is today, and where it will take us. * Will automation take away most of our jobs? * Is a "technological singularity" near? * What is the chance that robots will take over? * How do we best prepare for this future? The author concludes that, if we plan well, AI could be our greatest legacy, the last invention human beings will ever need to make.

A call-to-arms about the broken nature of artificial intelligence, and the powerful corporations that are turning the human-machine relationship on its head. We like to think that we are in control of the future of "artificial" intelligence. The reality, though, is that we--the everyday people whose data powers AI--aren't actually in control of anything. When, for example, we speak with Alexa, we contribute that data to a system we can't see and have no input into--one largely free from regulation or oversight. The big nine corporations--Amazon, Google, Facebook, Tencent, Baidu, Alibaba, Microsoft, IBM and Apple--are the new gods of AI and are short-changing our futures to reap immediate financial gain. In this book, Amy Webb reveals the pervasive, invisible ways in which the foundations of AI--the people working on the system, their motivations, the technology itself--is broken. Within our lifetimes, AI will, by design, begin to behave unpredictably, thinking and acting in ways which defy human logic. The big nine corporations may be inadvertently building and enabling vast arrays of intelligent systems that don't share our motivations, desires, or hopes for the future of humanity. Much more than a passionate, human-centered call-to-arms, this book delivers a strategy for changing course, and provides a path for liberating us from algorithmic decision-makers and powerful corporations.

An investigation into the assignment of moral responsibilities and rights to intelligent and autonomous machines of our own making. One of the enduring concerns of moral philosophy is deciding who or what is deserving of ethical consideration. Much recent attention has been devoted to the "animal question"—consideration of the moral status of nonhuman animals. In this book, David Gunkel takes up the "machine question": whether and to what extent intelligent and autonomous machines of our own making can be considered to have legitimate moral responsibilities and any legitimate claim to moral consideration. The machine question poses a fundamental challenge to moral thinking, questioning the traditional philosophical conceptualization of technology as a tool or instrument to be used by human agents. Gunkel begins by addressing the question of machine moral agency: whether a machine might be considered a legitimate moral agent that could be held responsible for decisions and actions. He then approaches the machine question from the other side, considering whether a machine might be a moral patient due legitimate moral consideration. Finally, Gunkel considers some recent innovations in moral philosophy and critical theory that complicate the machine question, deconstructing the binary agent–patient opposition itself. Technological advances may prompt us to wonder if the science fiction of computers and robots whose actions affect their human companions (think of HAL in 2001: A Space Odyssey) could become science fact. Gunkel's argument promises to influence future considerations of ethics, ourselves, and the other entities who inhabit this world.

As robots are increasingly integrated into modern society—on the battlefield and the road, in business, education, and health—Pulitzer-Prize-winning New York Times science writer John

Markoff searches for an answer to one of the most important questions of our age: will these machines help us, or will they replace us? In the past decade alone, Google introduced us to driverless cars, Apple debuted a personal assistant that we keep in our pockets, and an Internet of Things connected the smaller tasks of everyday life to the farthest reaches of the internet. There is little doubt that robots are now an integral part of society, and cheap sensors and powerful computers will ensure that, in the coming years, these robots will soon act on their own. This new era offers the promise of immense computing power, but it also reframes a question first raised more than half a century ago, at the birth of the intelligent machine: Will we control these systems, or will they control us? In Machines of Loving Grace, New York Times reporter John Markoff, the first reporter to cover the World Wide Web, offers a sweeping history of the complicated and evolving relationship between humans and computers. Over the recent years, the pace of technological change has accelerated dramatically, reintroducing this difficult ethical quandary with newer and far weightier consequences. As Markoff chronicles the history of automation, from the birth of the artificial intelligence and intelligence augmentation communities in the 1950s, to the modern day brain trusts at Google and Apple in Silicon Valley, and on to the expanding tech corridor between Boston and New York, he traces the different ways developers have addressed this fundamental problem and urges them to carefully consider the consequences of their work. We are on the verge of a technological revolution, Markoff argues, and robots will profoundly transform the way our lives are organized. Developers must now draw a bright line between what is human and what is machine, or risk upsetting the delicate balance between them.

A revelatory and timely look at how technology boosts our cognitive abilities—making us smarter, more productive, and more creative than ever It’s undeniable—technology is changing the way we think. But is it for the better? Amid a chorus of doomsayers, Clive Thompson delivers a resounding “yes.” In Smarter Than You Think, Thompson shows that every technological innovation—from the written word to the printing press to the telegraph—has provoked the very same anxieties that plague us today. We panic that life will never be the same, that our attentions are eroding, that culture is being trivialized. But, as in the past, we adapt—learning to use the new and retaining what is good of the old. Smarter Than You Think embraces and extols this transformation, presenting an exciting vision of the present and the future.

Stories suggest some unexpected results of using computers and robots in insurance, transportation, and sales

A pair of technology experts describe how humans will have to keep pace with machines in order to become prosperous in the future and identify strategies and policies for business and individuals to use to combine digital processing power with human ingenuity.

This book is a history of artificial intelligence, that audacious effort to duplicate in an artifact what we consider to be our most important property—our intelligence. It is an invitation for anybody with an interest in the future of the human race to participate in the inquiry.

“Startling in scope and bravado.” —Janet Maslin, The New York Times “Artfully envisions a breathtakingly better world.” —Los Angeles Times “Elaborate, smart and persuasive.” —The Boston Globe “A pleasure to read.” —The Wall Street Journal One of CBS News’s Best Fall Books of 2005 • Among St Louis Post-Dispatch’s Best Nonfiction Books of 2005 • One of Amazon.com’s Best Science Books of 2005 A radical and optimistic view of the future course of human development from the bestselling author of How to Create a Mind and The Singularity is Nearer who Bill Gates calls “the best person I know at predicting the future of artificial intelligence” For over three decades, Ray Kurzweil has been one of the most respected and provocative advocates of the role of technology in our future. In his classic The Age of Spiritual Machines, he argued that computers would soon rival the full range of human intelligence at its best. Now he examines the next step in this inexorable evolutionary process: the union of human and machine, in which the knowledge and skills embedded in our brains will be combined with the vastly greater capacity, speed, and knowledge-sharing ability of our creations.

“Refreshingly thought-provoking...” – The Financial Times The essential playbook for the future of your business What To Do When Machines Do Everything is a guidebook to succeeding in the next generation of the digital economy. When systems running on Artificial Intelligence can drive our cars, diagnose medical patients, and manage our finances more effectively than humans it raises profound questions on the future of work and how companies compete. Illustrated with real-world cases, data, and insight, the authors provide clear strategic guidance and actionable steps to help you and your organization move ahead in a world where exponentially developing new technologies are changing how value is created. Written by a team of business and technology expert practitioners—who also authored Code Halos: How the Digital Lives of People, Things, and Organizations are Changing the Rules of Business—this book provides a clear path to the future of your work. The first part of the book examines the once in a generation upheaval most every organization will soon face as systems of intelligence go mainstream. The authors argue that contrary to the doom and gloom that surrounds much of IT and business at the moment, we are in fact on the cusp of the biggest wave of opportunity creation since the Industrial Revolution. Next, the authors detail a clear-cut business model to help leaders take part in this coming boom; the AHEAD model outlines five strategic initiatives—Automate, Halos, Enhance, Abundance, and Discovery—that are central to competing in the next phase of global business by driving new levels of efficiency, customer intimacy and innovation. Business leaders today have two options: be swallowed up by the ongoing technological evolution, or ride the crest of the wave to new profits and better business. This book shows you how to avoid your own extinction event, and will help you; Understand the untold full extent of technology's impact on the way we work and live. Find out where we're headed, and how soon the future will arrive Leverage the new emerging paradigm into a sustainable business advantage Adopt a strategic model for winning in the new economy The digital world is already transforming how we work, live, and shop, how we are governed and entertained, and how we manage our money, health, security, and relationships. Don't let your business—or your career—get left behind. What To Do When Machines Do Everything is your strategic roadmap to a future full of possibility and success. Or peril.

Ray Kurzweil is the inventor of the most innovative and compelling technology of our era, an international authority on artificial intelligence, and one of our greatest living visionaries. Now he offers a framework for envisioning the twenty-first century--an age in which the marriage of human sensitivity and artificial intelligence fundamentally alters and improves the way we live. Kurzweil's prophetic blueprint for the future takes us through the advances that inexorably result in computers exceeding the memory capacity and computational ability of the human brain by the year 2020 (with human-level capabilities not far behind); in relationships with automated personalities who will be our teachers, companions, and lovers; and in information fed straight into

our brains along direct neural pathways. Optimistic and challenging, thought-provoking and engaging, The Age of Spiritual Machines is the ultimate guide on our road into the next century.
What to Think About Machines That ThinkToday's Leading Thinkers on the Age of Machine IntelligenceHarperCollins
A fascinating look at Artificial Intelligence, from its humble Cold War beginnings to the dazzling future that is just around the corner. When most of us think about Artificial Intelligence, our minds go straight to cyborgs, robots, and sci-fi thrillers where machines take over the world. But the truth is that Artificial Intelligence is already among us. It exists in our smartphones, fitness trackers, and refrigerators that tell us when the milk will expire. In some ways, the future people dreamed of at the World's Fair in the 1960s is already here. We're teaching our machines how to think like humans, and they're learning at an incredible rate. In Thinking Machines, technology journalist Luke Dormehl takes you through the history of AI and how it makes up the foundations of the machines that think for us today. Furthermore, Dormehl speculates on the incredible--and possibly terrifying--future that's much closer than many would imagine. This remarkable book will invite you to marvel at what now seems commonplace and to dream about a future in which the scope of humanity may need to broaden itself to include intelligent machines.