

The Economic Singularity Artificial Intelligence And The Death Of Capitalism

"What does AI mean for your business? Read this book to find out." -- Hal Varian, Chief Economist, Google Artificial intelligence does the seemingly impossible, magically bringing machines to life--driving cars, trading stocks, and teaching children. But facing the sea change that AI will bring can be paralyzing. How should companies set strategies, governments design policies, and people plan their lives for a world so different from what we know? In the face of such uncertainty, many analysts either cower in fear or predict an impossibly sunny future. But in *Prediction Machines*, three eminent economists recast the rise of AI as a drop in the cost of prediction. With this single, masterful stroke, they lift the curtain on the AI-is-magic hype and show how basic tools from economics provide clarity about the AI revolution and a basis for action by CEOs, managers, policy makers, investors, and entrepreneurs. When AI is framed as cheap prediction, its extraordinary potential becomes clear: Prediction is at the heart of making decisions under uncertainty. Our businesses and personal lives are riddled with such decisions. Prediction tools increase productivity--operating machines, handling documents, communicating with customers. Uncertainty constrains strategy. Better prediction creates opportunities for new business structures and strategies to compete. Penetrating, fun, and always insightful and practical, *Prediction Machines* follows its inescapable logic to explain how to navigate the changes on the horizon. The impact of AI will be profound, but the economic framework for understanding it is surprisingly simple.

In Ray Kurzweil's New York Times bestseller *The Singularity is Near*, the futurist and entrepreneur describes the Singularity, a likely future utterly different than anything we can imagine. The Singularity is triggered by the tremendous growth of human and computing intelligence that is an almost inevitable outcome of Moore's Law. Since the book's publication, the coming of the Singularity is now eagerly anticipated by many of the leading thinkers in Silicon Valley, from PayPal mastermind Peter Thiel to Google co-founder Larry Page. The formation of the Singularity University, and the huge popularity of the Singularity website kurzweilai.com, speak to the importance of this intellectual movement. But what about the average person? How will the Singularity affect our daily lives—our jobs, our families, and our wealth? *Singularity Rising: Surviving and Thriving in a Smarter, Richer, and More Dangerous World* focuses on the implications of a future society faced with an abundance of human and artificial intelligence. James D. Miller, an economics professor and popular speaker on the Singularity, reveals how natural selection has been increasing human intelligence over the past few thousand years and speculates on how intelligence enhancements will shape civilization over the next forty years. Miller considers several possible scenarios in this coming singularity: • A merger of man and machine making society fantastically wealthy and nearly immortal • Competition with billions of cheap AIs drive human wages to almost nothing while making investors rich • Businesses rethink investment decisions to take into account an expected future period of intense creative destruction • Inequality drops worldwide as technologies mitigate the cognitive cost of living in impoverished environments • Drugs designed to fight Alzheimer's disease and keep soldiers alert on battlefields have the fortunate side effect of increasing all of their users' IQs, which, in turn, adds a percentage points to worldwide economic growth *Singularity Rising* offers predictions about the economic implications for a future of widely expanding intelligence and practical career and investment advice on flourishing on the way to the Singularity.

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.

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.

The idea of technological singularity, and what it would mean if ordinary human intelligence were enhanced or overtaken by artificial intelligence. The idea that human history is approaching a "singularity"—that ordinary humans will someday be overtaken by artificially intelligent machines or cognitively enhanced biological intelligence, or both—has moved from the realm of science fiction to serious debate. Some singularity theorists predict that if the field of artificial intelligence (AI) continues to develop at its current dizzying rate, the singularity could come about in the middle of the present century. Murray Shanahan offers an introduction to the idea of the singularity and considers the ramifications of such a potentially seismic event. Shanahan's aim is not to make predictions but rather to investigate a range of scenarios. Whether we believe that singularity is near or far, likely or impossible, apocalypse or utopia, the very idea raises crucial philosophical and pragmatic questions, forcing us to think seriously about what we want as a species. Shanahan describes technological advances in AI, both biologically inspired and engineered from scratch. Once human-level AI—theoretically possible, but difficult to accomplish—has been achieved, he explains, the transition to superintelligent AI could be very rapid. Shanahan considers what the existence of superintelligent machines could mean for such matters as personhood, responsibility, rights, and identity. Some superhuman AI agents might be created to benefit humankind; some might go rogue. (Is Siri the template, or HAL?) The singularity presents both an existential threat to humanity and an existential opportunity for humanity to transcend its limitations. Shanahan makes it clear that we need to imagine both possibilities if we want to bring about the better outcome.

"The rise of AI must be better managed in the near term in order to mitigate longer term risks and to ensure that AI does not reinforce existing inequalities"--Publisher.

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.

This volume contains a selection of authoritative essays exploring the central questions raised by the conjectured technological singularity. In informed yet jargon-free contributions written by active research scientists, philosophers and sociologists, it goes beyond philosophical discussion to provide a detailed account of the risks that the singularity poses to human society and, perhaps most usefully, the possible actions that society and technologists can take to manage the journey to any singularity in a way that ensures a positive rather than a negative impact on society. The discussions provide perspectives that

cover technological, political and business issues. The aim is to bring clarity and rigor to the debate in a way that will inform and stimulate both experts and interested general readers.

The science of AI was born a little over 60 years ago, but for most of that time its achievements were modest. In 2012 it experienced a big bang, when a branch of statistics called Machine Learning (and a sub-branch called Deep Learning) was applied to it. Now machines have surpassed humans in image recognition, and they are catching up with us at speech recognition and natural language processing. Every day, the media reports the launch of a new service, a new product, and a new demonstration powered by AI. When will it end? The surprising truth is, the AI revolution has only just begun. Artificial Intelligence and the Two Singularities argues that in the course of this century, the exponential growth in the capability of AI is likely to bring about two "singularities" - points at which conditions are so extreme that the normal rules break down. The first is the economic singularity, when machine skill reaches a level that renders many of us unemployable and requires an overhaul of our current economic and social systems. The second is the technological singularity, when machine intelligence reaches and then surpasses the cognitive abilities of an adult human, relegating us to the second smartest species on the planet. These singularities will present huge challenges, but this book argues that we can meet these challenges and overcome them. If we do, the rewards could be almost unimaginable. This book covers:

- Recent developments in AI and its future potential
- The economic singularity and the technological singularity in depth
- The risks and opportunities presented by AI
- What actions we should take

Artificial intelligence can turn out to be the best thing ever to happen to humanity, making our future wonderful almost beyond imagination. But only if we address head-on the challenges that it will raise.

Calum Chace is a best-selling author of fiction and non-fiction books and articles, focusing on the subject of artificial intelligence. He is a regular speaker on artificial intelligence and related technologies, and runs a blog on the subject at www.pandoras-brain.com. Prior to becoming a full-time writer and speaker, he spent 30 years in business as a marketer, a strategy consultant, and a CEO. He studied philosophy at Oxford University, where he discovered that the science fiction he had been reading since boyhood was simply philosophy in fancy dress.

This interdisciplinary and international handbook captures and shapes much needed reflection on normative frameworks for the production, application, and use of artificial intelligence in all spheres of individual, commercial, social, and public life.

Artificial intelligence will not necessarily create a super-intelligent "human robot"; however, it is very probable that intelligent robots and intelligent informats will bring about a form of super-globalization, in which money and goods are prioritized over people and democracy and where the widespread use of casual labour – that is, short-term contracts – will become the most common form of employment relationship. It is also very likely that artificial intelligence will bring about what is known as singularity. This term is used to describe a situation where intelligent robots, from a rational and logical perspective, are smarter than humans, i.e. the development of AI. This book explores the impact that these intelligent robots and intelligent informats will have on social and societal development. The author tackles the question of singularity from three distinct standpoints: technological singularity – the intelligence of machines compared to that of humans – which he argues will bring about a qualitatively new labour market; economic singularity – the consequences for work relationships, value creation and employment – which he asserts will promote full automation, result in precarious contracts with low salaries, and, in some countries, possibly lead to the introduction of a universal basic income; and social singularity – the consequences of technological and economic singularity for democratic processes, bureaucratic procedures for exercising authority and control, and the direction in which society will develop, in addition to the emergence of new social institutions – which Johannessen says will promote a transition from representative democracy to genuine democracy. The book will appeal to academics, researchers and students of economic sociology and political economy, as well as those focusing upon the emerging innovation economy. It will also find an audience among professionals and policymakers keen to understand the impact the Fourth Industrial Revolution will have on organizations, individuals and society at large.

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.

Extraordinary innovations in technology promise to transform the world, but how realistic is the claim that AI will change our lives? In this much needed book the acclaimed economist Roger Bootle responds to the fascinating economic questions posed by the age of the robot, steering a path away from tech jargon and alarmism towards a rational explanation of the ways in which the AI revolution will affect us all. Tackling the implications of Artificial Intelligence on growth, productivity, inflation and the distribution of wealth and power, *THE AI ECONOMY* also examines coming changes to the the way we educate, work and spend our leisure time. A fundamentally optimistic view which will help you plan for changing times, this book explains AI and leads you towards a more certain future. Extraordinary innovations in technology promise to transform the world, but how realistic is the claim that AI will change our lives? In this much needed book the acclaimed economist Roger Bootle responds to the fascinating economic questions posed by the age of the robot, steering a path away from tech jargon and alarmism towards a rational explanation of the ways in which the AI revolution will affect us all. Tackling the implications of Artificial Intelligence on growth, productivity, inflation and the distribution of wealth and power, *THE AI ECONOMY* also examines coming changes to the the way we educate, work and spend our leisure time. A fundamentally optimistic view which will help you plan for changing times, this book explains AI and leads you towards a more certain future.

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.

“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.

The human brain has some capabilities that the brains of other animals lack. It is to these distinctive capabilities that our species owes its dominant position. Other animals have stronger muscles or sharper claws, but we have cleverer brains. If machine brains one day come to surpass human brains in general intelligence, then this new superintelligence could become very powerful. As the fate of the gorillas now depends more on us humans than on the gorillas themselves, so the fate of our species then would come to depend on the actions of the machine superintelligence. But we have one advantage: we get to make the first move. Will it be possible to construct a seed AI or otherwise to engineer initial conditions so as to make an intelligence explosion survivable? How could one achieve a controlled detonation? To get closer to an answer to this question, we must make our way through a fascinating landscape of topics and considerations. Read the book and learn about oracles, genies, singletons; about boxing methods, tripwires, and mind crime; about humanity's cosmic endowment and differential technological development; indirect normativity, instrumental convergence, whole brain emulation and technology couplings; Malthusian economics and dystopian evolution; artificial intelligence, and biological cognitive enhancement, and collective intelligence.

Pointing to the triumph of artificial intelligence over unaided humans in everything from games such as chess and Go to vital tasks such as protein folding and securities trading, many experts uphold the theory of a "singularity." This is the trigger point when human history ends and artificial intelligence prevails in an exponential cascade of self-replicating machines rocketing toward godlike supremacy in the universe. *Gaming AI* suggests that this belief is both dumb and self-defeating. Displaying a profound and crippling case of professional amnesia, the computer science establishment shows an ignorance of the most important findings of its own science, from Kurt Gödel's "incompleteness" to Alan Turing's "oracle" to Claude Shannon's "entropy." Dabbling in quantum machines, these believers in machine transcendence defy the deepest findings of quantum theory.

Claiming to create minds, they are clinically "out of their minds." Despite the quasi-religious pretensions of techno-elites nobly saving the planet from their own devices, their faith in a techno-utopian singularity is a serious threat to real progress. An industry utterly dependent on human minds will not prosper by obsoleting both their customers and their creators. *Gaming AI* calls for a remedial immersion in the industry's own heroic history and an understanding of the actual science of their own human minds.

In the year 2035, robots, artificial intelligences, and neural implants have become commonplace. The Institute for Ethics keeps the peace, using social reputation to ensure that robots and humans don't harm society or each other. But a powerful AI named Adam has found a way around the restrictions. Catherine Matthews, nineteen years old, has a unique gift: the ability to manipulate the net with her neural implant. Yanked out of her perfectly ordinary life, Catherine becomes the last firewall standing between Adam and his quest for world domination. **PRAISE FOR THE LAST FIREWALL** “Awesome near-term science fiction.” – Brad Feld, Foundry Group managing director “An insightful and adrenaline-inducing tale of what humanity could become and the machines we could spawn.” – Ben Huh, CEO of Cheezburger “A fun read and tantalizing study of the future of technology: both inviting and alarming.” – Harper Reed, former CTO of Obama for America, Threadless “A fascinating and prescient take on what the world will look like once computers become smarter than people. Highly recommended.” – Mat Ellis, Founder & CEO Cloudability “A phenomenal ride through a post-scarcity world where humans are caught between rogue AIs. If you like having your mind blown, read this book!” – Gene Kim, author of *The Phoenix Project: A Novel About IT, DevOps, and Helping Your Business Win*

Advances in Artificial Intelligence (AI) technology have opened up new markets and new opportunities for progress in critical areas such as health, education, energy, and the environment. In recent years, machines have surpassed humans in the performance of certain specific tasks, such as some aspects of image recognition. Experts forecast that rapid progress in the field of specialized artificial intelligence will continue. Although it is very unlikely that machines will exhibit broadly-applicable intelligence comparable to or exceeding that of humans in the next 20 years, it is to be expected that machines will reach and exceed human performance on more and more tasks. As a contribution toward preparing the United States for a future in which AI plays a growing role, this report surveys the current state of AI, its existing and potential applications, and the questions that are raised for society and public policy by progress in AI. The report also makes recommendations for specific further actions by Federal agencies and other actors.

A documentary filmmaker, bringing together Artificial Intelligence experts from around the world, explores the terrifying possibility of catastrophic outcomes once we share the planet with intelligent machines who are smarter and more powerful than we could ever have imagined. 25,000 first printing.

This authoritative reference work will provide readers with a complete overview of artificial intelligence (AI), including its historic development and current status; existing and projected AI applications; and present and potential future impact on the United States and the world. Some people believe that artificial intelligence (AI) will revolutionize modern life in ways that improve human existence. Others say that the promise of AI is overblown. Still others contend that AI applications could pose a grave threat to the economic security of millions of people by taking their jobs and otherwise rendering them "obsolete"—or, even worse, that AI could actually spell the end of the human race. This volume will help users understand the reasons AI development has both spirited defenders and alarmed critics; explain theories and innovations like Moore's Law, mindcloning, and Technological Singularity that drive AI research and debate; and give readers the information they need to make their own informed judgment about the promise and peril of this technology. All of this coverage is presented using language and terminology accessible to a lay audience. Introduction explaining the historical evolution of AI Chronology of important AI-related events Authoritative entries on leading pioneers, entrepreneurs, and thinkers; AI concepts and theories; AI's potential impact on different facets of society; and major movies and other cultural touchstones exploring AI technology

"Read *The Economic Singularity* if you want to think intelligently about the future." Aubrey de Grey Artificial intelligence (AI) is overtaking our human ability to absorb and process information. Robots are becoming increasingly dextrous, flexible, and safe to be around (except the military ones). It is our most powerful technology, and you need to understand it. This new book from best-selling AI writer Calum Chace argues that within a few decades, most humans will not be able to work for money. Self-driving cars will probably be the canary in the coal mine, providing a wake-up call for everyone who isn't yet paying attention. All jobs will be affected, from fast food McJobs to lawyers and journalists. This is the single most important development facing humanity in the first half of the 21st century. The fashionable belief that Universal Basic Income is the solution is only partly correct. We are probably going to need an entirely new economic system, and we better start planning soon - for the Economic Singularity! The outcome can be very good - a world in which machines do all the boring jobs and humans do pretty much what they please. But there are major risks, which we can only avoid by being alert to the possible futures and planning how to avoid the negative ones."

'A compelling invitation to imagine the future we want' —BRIAN CHRISTIAN, author of *The Most Human Human* By 2062 we will have built machines as intelligent as us – so the leading artificial intelligence and robotics experts predict. But what will this future look like? In 2062, world-leading researcher Toby Walsh considers the impact AI will have on work, war, economics, politics, everyday life and even death. Will automation take away most jobs? Will robots become conscious and take over? Will we become immortal machines ourselves, uploading our brains to the cloud? How will politics adjust to the post-truth, post-privacy digitised world? When we have succeeded in building intelligent machines, how will life on this planet unfold? Based on a deep understanding of technology, *2062* describes the choices we need to make today to ensure that the future remains bright. 'Clarity and sanity in a world full of fog and uncertainty – a timely book about the race to remain human.' —RICHARD WATSON, author of *Digital Vs. Human* and futurist-in-residence at Imperial College, London 'One of the deepest questions facing humanity, pondered by a mind well and truly up to the task.' —ADAM SPENCER, broadcaster

You are about to become obsolete. You think you are special, unique, and that whatever it is that you are doing is impossible to replace. You are wrong. As we speak, millions of algorithms created by computer scientists are frantically running on servers all over the world, with one sole purpose: do whatever humans can do, but better. That is the argument for a phenomenon called technological unemployment, one that is pervading modern society. But is that really the case? Or is it just a futuristic fantasy? What will become of us in the coming years, and what can we do to prevent a catastrophic collapse of society? *Robots Will Steal Your Job, But That's OK: how to survive the economic collapse and be happy* explores the impact of technological advances on our lives, what it means to be happy, and provides suggestions on how to avoid a systemic collapse.

Advances in artificial intelligence (AI) highlight the potential of this technology to affect productivity, growth, inequality, market power, innovation, and employment. This volume seeks to set the agenda for economic research on the impact of AI. It covers four broad themes: AI as a general purpose technology; the relationships between AI, growth, jobs, and inequality; regulatory responses to changes brought on by AI; and the effects of AI on the way economic research is conducted. It explores the economic influence of machine learning, the branch of computational statistics that has driven much of the recent excitement around AI, as well as the economic impact of robotics and automation and the potential economic consequences of a still-hypothetical artificial general intelligence. The volume provides frameworks for understanding the economic impact of AI and identifies a number of open research questions. Contributors: Daron Acemoglu, Massachusetts Institute of Technology Philippe Aghion, Collège de France Ajay Agrawal, University of Toronto Susan Athey, Stanford University James Bessen, Boston University School of Law Erik Brynjolfsson, MIT Sloan School of Management Colin F. Camerer, California Institute of Technology Judith Chevalier, Yale School of Management Iain M. Cockburn, Boston University Tyler Cowen, George Mason University Jason Furman, Harvard Kennedy School Patrick Francois, University of British Columbia Alberto Galasso, University of Toronto Joshua Gans, University of Toronto Avi Goldfarb, University of Toronto Austan Goolsbee, University of Chicago Booth School of Business Rebecca Henderson, Harvard Business School Ginger Zhe Jin, University of Maryland Benjamin F. Jones, Northwestern University Charles I. Jones, Stanford University Daniel Kahneman, Princeton University Anton Korinek, Johns Hopkins University Mara Lederman, University of Toronto Hong Luo, Harvard Business School John McHale, National University of Ireland Paul R. Milgrom, Stanford University Matthew Mitchell, University of Toronto Alexander Oettl, Georgia Institute of Technology Andrea Prat, Columbia Business School Manav Raj, New York University Pascual Restrepo, Boston University Daniel Rock, MIT Sloan School of Management Jeffrey D. Sachs, Columbia University Robert Seamans, New York University Scott Stern, MIT Sloan School of Management Betsey Stevenson, University of Michigan Joseph E. Stiglitz, Columbia University Chad Syverson, University of Chicago Booth School of Business Matt Taddy, University of Chicago Booth School of Business Steven Tadelis, University of California, Berkeley Manuel Trajtenberg, Tel Aviv University Daniel Trefler, University of Toronto Catherine Tucker, MIT Sloan School of Management Hal Varian, University of California, Berkeley

A day does not go by without a news article reporting some amazing breakthrough in artificial intelligence (AI). Many philosophers, futurists, and AI researchers have conjectured that human-

level AI will be developed in the next 20 to 200 years. If these predictions are correct, it raises new and sinister issues related to our future in the age of intelligent machines. Artificial Superintelligence: A Futuristic Approach directly addresses these issues and consolidates research aimed at making sure that emerging superintelligence is beneficial to humanity. While specific predictions regarding the consequences of superintelligent AI vary from potential economic hardship to the complete extinction of humankind, many researchers agree that the issue is of utmost importance and needs to be seriously addressed. Artificial Superintelligence: A Futuristic Approach discusses key topics such as: AI-Completeness theory and how it can be used to see if an artificial intelligent agent has attained human level intelligence Methods for safeguarding the invention of a superintelligent system that could theoretically be worth trillions of dollars Self-improving AI systems: definition, types, and limits The science of AI safety engineering, including machine ethics and robot rights Solutions for ensuring safe and secure confinement of superintelligent systems The future of superintelligence and why long-term prospects for humanity to remain as the dominant species on Earth are not great Artificial Superintelligence: A Futuristic Approach is designed to become a foundational text for the new science of AI safety engineering. AI researchers and students, computer security researchers, futurists, and philosophers should find this an invaluable resource.

Over the coming decades, Artificial Intelligence will profoundly impact the way we live, work, wage war, play, seek a mate, educate our young, and care for our elderly. It is likely to greatly increase our aggregate wealth, but it will also upend our labor markets, reshuffle our social order, and strain our private and public institutions. Eventually it may alter how we see our place in the universe, as machines pursue goals independent of their creators and outperform us in domains previously believed to be the sole dominion of humans. Whether we regard them as conscious or unwitting, revere them as a new form of life or dismiss them as mere clever appliances, is beside the point. They are likely to play an increasingly critical and intimate role in many aspects of our lives. The emergence of systems capable of independent reasoning and action raises serious questions about just whose interests they are permitted to serve, and what limits our society should place on their creation and use. Deep ethical questions that have bedeviled philosophers for ages will suddenly arrive on the steps of our courthouses. Can a machine be held accountable for its actions? Should intelligent systems enjoy independent rights and responsibilities, or are they simple property? Who should be held responsible when a self-driving car kills a pedestrian? Can your personal robot hold your place in line, or be compelled to testify against you? If it turns out to be possible to upload your mind into a machine, is that still you? The answers may surprise you.

David Ryan is the designer of ELOPe, an email language optimization program, that if successful, will make his career. But when the project is suddenly in danger of being canceled, David embeds a hidden directive in the software accidentally creating a runaway artificial intelligence. David and his team are initially thrilled when the project is allocated extra servers and programmers. But excitement turns to fear as the team realizes that they are being manipulated by an A.I. who is redirecting corporate funds, reassigning personnel and arming itself in pursuit of its own agenda. WINNER SCIENCE FICTION DIY BOOK FESTIVAL 2011-2012 "Avogadro Corp is a tremendous book that every single person needs to read. In the vein of Daniel Suarez's Daemon and Freedom(TM), William's book shows that science fiction is becoming science fact. Avogadro Corp describes issues, in solid technical detail, that we are dealing with today that will impact us by 2015, if not sooner. Not enough people have read these books. It's a problem for them, but not for the [emergent] machines." -- Brad Feld, managing director Foundry Group, co-founder Techstars "Highly entertaining, gripping, thought inspiring book. Don't start without the time to finish — it won't let you go." -- Gifford Pinchot III, founder Bainbridge Graduate Institute, author THE INTELLIGENT ORGANIZATION "An alarming and jaw-dropping tale about how something as innocuous as email can subvert an entire organization. I found myself reading with a sense of awe, and read it way too late into the night." -- Gene Kim, author of VISIBLE OPS "A fictional world where Portland is the hub for the most exciting advancements in technology... [J]am packed with great references to deep Portland culture...and Portlandia-type references" -- SILICON FLORIST

Pandora's Brain Description Around half the scientists researching artificial intelligence (AI) think that a conscious machine with cognitive abilities at or beyond human level will be created by 2050. If they are right, the consequence could be an intelligence explosion, in which the AI rapidly and enormously exceeds human competence. Pandora's Brain is a science thriller by best-selling writer Calum Chace. It uses the issues raised by the possible coming machine intelligence explosion as a platform for a fast-paced and thought-provoking adventure story. The story is set in the very near future, and features Matt, a shy but engaging and resourceful student who discovers that his recently-deceased father was involved in research that could enable the construction of the world's first conscious machine. Matt's enquiries lead to him being kidnapped, as he is caught in the crossfire between two groups pursuing that goal - one led by a Russian billionaire, and another backed by the US military. Matt has to do more than simply survive: he has to harness these powerful forces to his own ends. At stake is his own life and those of his family and friends. A dramatic seaborne rescue operation, a series of brutal murders and other filmic action scenes follow. In the course of his adventures, Matt discovers that the potential upside of creating machine intelligence includes immortality, and godlike powers of understanding and being - but the potential downside is immediate extinction, or worse. As he is drawn deeper into his adventure, he becomes both the symbol and the victim of a global struggle over the approach to be taken towards this powerful new technology. A landmark decision at a meeting of the UN General Assembly forces Matt to make a fateful decision which sparks the story's final twist. Selected reviews for Pandora's Brain: I love the concepts in this book! Peter James, author of the best-selling Roy Grace series Awesome! Count me as a fan. Brad Feld, co-founder of Techstars Pandora's Brain is a captivating tale of developments in artificial intelligence that could, conceivably, be just around the corner. Mainly set in the present day, the plot unfolds in an environment that seems reassuringly familiar, but which is overshadowed by a combination of both menace and promise. Carefully crafted, and absorbing from its very start, the book held my rapt attention throughout a series of surprise twists, as various personalities react in different ways to a growing awareness of that menace and promise. David Wood, Chairman of London Futurists I was eagerly anticipating a fiction adventure book on precisely

this topic! Very well done, Calum Chace. A timely, suspenseful, and balanced portrayal of AI and the most important decisions humanity will make in the near future. Hank Pellissier, producer of the Transhuman Visions conferences Pandora's Brain is a tour de force that neatly explains the key concepts behind the likely future of artificial intelligence in the context of a thriller novel. Ambitious and well executed, it will appeal to a broad range of readers. In the same way that Suarez's Daemon and Naam's Nexus leaped onto the scene, redefining what it meant to write about technology, Pandora's Brain will do the same for artificial intelligence. "Mind uploading? Check. Human equivalent AI? Check. Hard takeoff singularity? Check. Strap in, this is one heck of a ride. William Hertling, author of the Singularity series of novels "

The Economic Singularity Artificial Intelligence and the Death of Capitalism Three CS

Artificial intelligence is our most powerful technology, and in the coming decades it will change everything in our lives. If we get it right it will make humans almost godlike. If we get it wrong... well, extinction is not the worst possible outcome. "Surviving AI" is a concise, easy-to-read guide to what's coming, taking you through technological unemployment (the economic singularity) and the possible creation of a superintelligence (the technological singularity). Here's what some of the leading thinkers in the field have to say about it: A sober and easy-to-read review of the risks and opportunities that humanity will face from AI. Jaan Tallinn - co-founder of Skype Understanding AI - its promise and its dangers - is emerging as one of the great challenges of coming decades and this is an invaluable guide to anyone who's interested, confused, excited or scared. David Shukman - BBC Science Editor We have recently seen a surge in the volume of scholarly analysis of this topic; Chace impressively augments that with this high-quality, more general-audience discussion. Aubrey de Grey - CSO of SENS Research Foundation; former AI researcher It's rare to see a book about the potential End of the World that is fun to read without descending into sensationalism or crass oversimplification. Ben Goertzel - chairman of Novamente LLC Calum Chace is a prescient messenger of the risks and rewards of artificial intelligence. In "Surviving AI" he has identified the most essential issues and developed them with insight and wit - so that the very framing of the questions aids our search for answers. Chace's sensible balance between AI's promise and peril makes "Surviving AI" an excellent primer for anyone interested in what's happening, how we got here, and where we are headed. Kenneth Cukier - co-author of "Big Data" If you're not thinking about AI, you're not thinking. "Surviving AI" combines an essential grounding in the state of the art with a survey of scenarios that will be discussed with equal vigor at cocktail parties and academic colloquia. Chris Meyer - author of "Blur," "It's Alive," and "Standing on the Sun" The appearance of Calum Chace's book is of some considerable personal satisfaction to me, because it signifies the fact that the level of social awareness of the rise of massively intelligent machines has finally reached the mainstream. If you want to survive the next few decades, you cannot afford NOT to read Chace's book. Prof. Dr. Hugo de Garis - former director of the Artificial Brain Lab, Xiamen University, China "Surviving AI" is an exceptionally clear, well-researched and balanced introduction to a complex and controversial topic, and is a compelling read to boot. Sean O hEigeartaigh - executive director of Cambridge Centre for the Study of Existential Risk In "Surviving AI," Calum Chace provides a marvellously accessible guide to the swirls of controversy that surround discussion of what is likely to be the single most important event in human history - the emergence of artificial superintelligence. Throughout, "Surviving AI" remains clear and jargon-free. David Wood - chair of London Futurists Artificial intelligence is the most important technology of our era. Technological unemployment could force us to adopt an entirely new economic structure, and the creation of superintelligence would be the biggest event in human history. "Surviving AI" is a first-class introduction to all of this. Brad Feld - co-founder of Techstars"

What are the prospects for long-run economic growth? The present study looks at a recently launched hypothesis, which I label Singularity. The idea here is that rapid growth in computation and artificial intelligence will cross some boundary or Singularity after which economic growth will accelerate sharply as an ever-accelerating pace of improvements cascade through the economy. The paper develops a growth model that features Singularity and presents several tests of whether we are rapidly approaching Singularity. The key question for Singularity is the substitutability between information and conventional inputs. The tests suggest that the Singularity is not near.

Book Description How will AI evolve and what major innovations are on the horizon? What will its impact be on the job market, economy, and society? What is the path toward human-level machine intelligence? What should we be concerned about as artificial intelligence advances? Architects of Intelligence contains a series of in-depth, one-to-one interviews where New York Times bestselling author, Martin Ford, uncovers the truth behind these questions from some of the brightest minds in the Artificial Intelligence community. Martin has wide-ranging conversations with twenty-three of the world's foremost researchers and entrepreneurs working in AI and robotics: Demis Hassabis (DeepMind), Ray Kurzweil (Google), Geoffrey Hinton (Univ. of Toronto and Google), Rodney Brooks (Rethink Robotics), Yann LeCun (Facebook) , Fei-Fei Li (Stanford and Google), Yoshua Bengio (Univ. of Montreal), Andrew Ng (AI Fund), Daphne Koller (Stanford), Stuart Russell (UC Berkeley), Nick Bostrom (Univ. of Oxford), Barbara Grosz (Harvard), David Ferrucci (Elemental Cognition), James Manyika (McKinsey), Judea Pearl (UCLA), Josh Tenenbaum (MIT), Rana el Kaliouby (Affectiva), Daniela Rus (MIT), Jeff Dean (Google), Cynthia Breazeal (MIT), Oren Etzioni (Allen Institute for AI), Gary Marcus (NYU), and Bryan Johnson (Kernel). Martin Ford is a prominent futurist, and author of Financial Times Business Book of the Year, Rise of the Robots. He speaks at conferences and companies around the world on what AI and automation might mean for the future.

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.

Singularity Hypotheses: A Scientific and Philosophical Assessment offers authoritative, jargon-free essays and critical commentaries on accelerating technological progress and the notion of technological singularity. It focuses on conjectures about the intelligence explosion, transhumanism, and whole brain emulation. Recent years have seen a plethora of forecasts about the profound, disruptive impact that is likely to result from further progress in these areas. Many commentators however doubt the scientific rigor of these forecasts, rejecting them as speculative and unfounded. We therefore invited prominent computer scientists, physicists, philosophers, biologists, economists and other thinkers to assess the singularity hypotheses. Their contributions go beyond speculation, providing deep insights into the main issues and a balanced picture of the debate.

"AI and the Technological Singularity: A Fallacy or a Great Opportunity" is a collection of essays that addresses the question of whether the technological singularity--the notion that AI-based computers can program the next generation of AI-based computers until a singularity is achieved, where an AI-based computer can exceed human intelligence--is a fallacy or a great opportunity. The group of scholars that address this question have a variety of positions on the singularity, ranging from advocates to skeptics. No conclusion can be reached, as the development of artificial intelligence is still in its infancy, and there is much wishful thinking and imagination in this issue rather than trustworthy data. The reader will find a cogent summary of the issues faced by researchers who are working to develop the field of artificial intelligence and, in particular, artificial general intelligence. The only conclusion that can be reached is that there exists a variety of well-argued positions as to where AI research is headed.

Intelligent algorithms are already well on their way to making white collar jobs obsolete: travel agents, data-analysts, and paralegals are currently in the firing line. In the near future, doctors, taxi-drivers and ironically even computer programmers are poised to be replaced by 'robots'. Without a radical reassessment of our economic and political structures, we risk the very implosion of the capitalist economy itself. In *The Rise of the Robots*, technology expert Martin Ford systematically outlines the achievements of artificial intelligence and uses a wealth of economic data to illustrate the terrifying societal implications. From health and education to finance and technology, his warning is stark – all jobs that are on some level routine are likely to eventually be automated, resulting in the death of traditional careers and a hollowed-out middle class. The robots are coming and we have to decide – now – whether the future will bring prosperity or catastrophe.

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.

How to make liberal democracies more inclusive and the digital economy more equitable: a guide for the coming Fourth Industrial Revolution. Around the world, liberal democracies are in crisis. Citizens have lost faith in their government; right-wing nationalist movements frame the political debate. At the same time, economic inequality is increasing dramatically; digital technologies have created a new class of super-rich entrepreneurs. Automation threatens to transform the free economy into a zero-sum game in which capital wins and labor loses. But is this digital dystopia inevitable? In *Cyber Republic*, George Zarkadakis presents an alternative, outlining a plan for using technology to make liberal democracies more inclusive and the digital economy more equitable. *Cyber Republic* is no less than a guide for the coming Fourth Industrial Revolution. Zarkadakis, an expert on technology and management, explains how artificial intelligence, together with intelligent robotics, sophisticated sensors, communication networks, and big data, will fundamentally reshape the global economy; a new "intelligent machine age" will force us to adopt new forms of economic and political organization. He envisions a future liberal democracy in which intelligent machines facilitate citizen assemblies, helping to extend citizen rights, and blockchains and cryptoeconomics enable new forms of democratic governance and business collaboration. Moreover, the same technologies can be applied to scientific research and technological innovation. We need not fear automation, Zarkadakis argues; in a post-work future, intelligent machines can collaborate with humans to achieve the human goals of inclusivity and equality.

"The most important book on AI this year." --The Guardian "Mr. Russell's exciting book goes deep, while sparkling with dry witticisms." --The Wall Street Journal "The most important book I have read in quite some time" (Daniel Kahneman); "A must-read" (Max Tegmark); "The book we've all been waiting for" (Sam Harris) A leading artificial intelligence researcher lays out a new approach to AI that will enable us to coexist successfully with increasingly intelligent machines In the popular imagination, superhuman artificial intelligence is an approaching tidal wave that threatens not just jobs and human relationships, but civilization itself. Conflict between humans and machines is seen as inevitable and its outcome all too predictable. In this groundbreaking book, distinguished AI researcher Stuart Russell argues that this scenario can be avoided, but only if we rethink AI from the ground up. Russell begins by exploring the idea of intelligence in humans and in machines. He describes the near-term benefits we can expect, from intelligent personal assistants to vastly accelerated scientific research, and outlines the AI breakthroughs that still have to happen before we reach superhuman AI. He also spells out the ways humans are already finding to misuse AI, from lethal autonomous weapons to viral sabotage. If the predicted breakthroughs occur and superhuman AI emerges, we will have created entities far more powerful than ourselves. How can we ensure they never, ever, have power over us? Russell suggests that we can rebuild AI on a new foundation, according to which machines are designed to be inherently uncertain

about the human preferences they are required to satisfy. Such machines would be humble, altruistic, and committed to pursue our objectives, not theirs. This new foundation would allow us to create machines that are provably deferential and provably beneficial.

A look at how new technologies can be put to use in the creation of a more just society. Artificial Intelligence (AI) is not likely to make humans redundant. Nor will it create superintelligence anytime soon. But it will make huge advances in the next two decades, revolutionize medicine, entertainment, and transport, transform jobs and markets, and vastly increase the amount of information that governments and companies have about individuals. AI for Good leads off with economist and best-selling author Daron Acemoglu, who argues that there are reasons to be concerned about these developments. AI research today pays too much attention to the technological hurdles ahead without enough attention to its disruptive effects on the fabric of society: displacing workers while failing to create new opportunities for them and threatening to undermine democratic governance itself. But the direction of AI development is not preordained. Acemoglu argues for its potential to create shared prosperity and bolster democratic freedoms. But directing it to that task will take great effort: It will require new funding and regulation, new norms and priorities for developers themselves, and regulations over new technologies and their applications. At the intersection of technology and economic justice, this book will bring together experts--economists, legal scholars, policy makers, and developers--to debate these challenges and consider what steps tech companies can do take to ensure the advancement of AI does not further diminish economic prospects of the most vulnerable groups of population.

[Copyright: b20abd5221b85446f44cbfea682ddb06](#)