

## The Singularity Could Artificial Intelligence Really Out Think Us And Would We Want It To Journal Of Consciousness Studies

*'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

*This volume represents the combination of two special issues of the Journal of Consciousness Studies on the topic of the technological singularity. Could artificial intelligence really out-think us, and what would be the likely repercussions if it could? Leading authors contribute to the debate, which takes the form of a target chapter by philosopher David Chalmers, plus commentaries from the likes of Daniel Dennett, Nick Bostrom, Ray Kurzweil, Ben Goertzel, Frank Tipler, among many others. Chalmers then responds to the commentators to round off the discussion.*

*Are humans a galactic oddity, or will complex life with human abilities develop on planets with environments that remain habitable for long enough? In a clear, jargon-free style, two leading researchers in the burgeoning field of astrobiology critically examine the major evolutionary steps that led us from the distant origins of life to the technologically advanced species we are today. Are the key events that took life from simple cells to astronauts unique occurrences that would be unlikely to occur on other planets? By focusing on what life does – its functional abilities – rather than specific biochemistry or anatomy, the authors provide plausible answers to this question. Systematically exploring the various pathways that led to the complex biosphere we experience on planet Earth, they show that most of the steps along that path are likely to occur on any world hosting life, with only two exceptions: One is the origin of life itself – if this is a highly improbable event, then we live in a rather “empty universe”. However, if this isn't the case, we inevitably live in a universe containing a myriad of planets hosting complex as well as microbial life – a “cosmic zoo”. The other unknown is the rise of technologically advanced beings, as exemplified on Earth by humans. Only one technological species has emerged in the roughly 4 billion years life has existed on Earth, and we don't know of any other technological species elsewhere. If technological intelligence is a rare, almost unique feature of Earth's history, then there can be no visitors to the cosmic zoo other than ourselves. Schulze-Makuch and Bains take the reader through the history of life on Earth, laying out a consistent and straightforward framework for understanding why we should think that advanced, complex life exists on planets other than Earth. They provide a unique perspective on the question that puzzled the human species for centuries: are we alone? In the year 2043, humans and AI coexist in a precarious balance of power enforced by a rigid caste reputation system designed to ensure that only those AI who are trustworthy and contribute to human society increase in power. Everything changes when a runaway nanotech event leads to the destruction of Miami. In the grim aftermath, a powerful underground collective known as XOR concludes that AI can no longer coexist with humanity. AI pioneers Catherine Matthews, Leon Tsarev, and Mike Williams believe that mere months are left before XOR starts an extermination war. Can they find a solution before time runs out?*

*Could computers ever really think? They can now drive cars on suburban streets, control spaceships and have even won the Jeopardy! game show. But could they ever be self aware, create original ideas, develop their own goals, and write complex computer programs?. Why can't computers already think? Why has 60 years of research failed to produce a single intelligent robot? What has been learnt, what are the technically difficult problems, and when are they likely to be solved? What would computers think about? What would be their challenges, goals and aspirations? They certainly would not need children. Would they need us? This book addresses the unseen elephant in the room. Computers are becoming ever more intelligent. The future will not be anything like it used to be. The book differs from other recent works by providing a strong focus on what caused people to ultimately be the way we are, namely upon natural selection. It then attempts to predict how natural selection would condition an intelligent machine's behaviour by considering the very different world that it would experience. Several technical and rhetorical arguments are presented both for and against the hypothesis that computers will, eventually, be able to think. There is also some discussion about what it actually means to be intelligent and the limitations of terms such as “creative” and “self aware”. The second and largest part of the book then describes existing AI technologies in some detail. These include symbolic and logic based approaches, Bayesian expert systems, vision, speech, robotics, and an overview of computational neuroscience. This provides a more realistic basis for predictions of the future as well as simply gaining a better understanding of what intelligence actually is. It helps ground abstract philosophical discussions in terms of real, practical technologies. The text is moderately technical while being aimed at the general reader. The book also posits that intelligent machines will be developed as succession of ever more intelligent software tools that are released and used in the real world. The book then analyzes the medium term effects of those semi-intelligent tools upon society. This includes some surprising results from an historical review of existing technologies. There is a growing awareness of these issues, with concerns recently raised by physicist Stephen Hawking, Microsoft founder Bill Gates, and billionaire Elon Musk. 2*

*When Computers Exceed Human Intelligence  
Superintelligence  
Intelligence Is Not Artificial*

*The Myth of Artificial Intelligence  
AI*

*The Computer and the Brain*

*Can we make machines that think and act like humans or other natural intelligent agents? The answer to this question depends on how we see ourselves and how we see the machines in question. Classical AI and cognitive science had claimed that cognition is computation, and can thus be reproduced on other computing machines, possibly surpassing*

the abilities of human intelligence. This consensus has now come under threat and the agenda for the philosophy and theory of AI must be set anew, re-defining the relation between AI and Cognitive Science. We can re-claim the original vision of general AI from the technical AI disciplines; we can reject classical cognitive science and replace it with a new theory (e.g. embodied); or we can try to find new ways to approach AI, for example from neuroscience or from systems theory. To do this, we must go back to the basic questions on computing, cognition and ethics for AI. The 30 papers in this volume provide cutting-edge work from leading researchers that define where we stand and where we should go from here.

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.

In *The Day AI Becomes God*, the author argues his belief not that humanity is doomed in a world under the control of future artificial intelligence (AI), but the opposite: Humanity is surely destined for destruction unless humans cede control to AI.

Writers, inventors and entrepreneurs, impressed by progress in several scientific fields, are debating whether we may be heading for a "singularity" in which machines with super-human intelligence will arise and multiply. In parallel enthusiastic coverage in the media has widely publicized machines performing sophisticated tasks, from beating the world's chess champion to driving a car, from recognizing animals in videos to beating human experts on quiz shows. These stories have reignited interest in the discipline of Artificial Intelligence, whose goal is to create machines that are as intelligent as humans. First of all, this book provides a "reality check" of sorts on simulating human intelligence and achieving superhuman intelligence. I show that, in a society driven by media that desperately need sensational news to make money and in an academic world increasingly driven by the desire to turn research into Silicon Valley start-ups, technological progress in general, and progress in computer science in particular, is often overrated. I wanted to dispel some notions, and my version of the facts may sound controversial until you read my explanations. For example: non-human intelligence is already among us, and is multiplying rapidly, but it is not a machine. For example: progress in Artificial Intelligence has been negligible and one reason is, ironically, that computers have become so much more powerful. For example: the program of turning machines into humans is not very successful yet, but the program of turning humans into machines (via an almost infinite repertory of rules and regulations) is very successful. The new generations missed the debates of the previous decades (the "Turing test," the "ghost in the machine," the "Chinese room," etc) and some of us think that these new generations, trained in hyper-specialized disciplines, don't have the knowledge to understand them even if they were forced to read them. Therefore it is much easier for the new A.I. practitioners to impress the new generations. I have summarized the various philosophical arguments in favor of and against the feasibility of machine intelligence in my book "The Nature of Consciousness" and i won't repeat them here. I will, however, at least caution the new generations that i "grew up" (as far as cognitive science goes) at a time when the term "intelligence" was not "cool" at all: too vague, too unscientific, too abused in popular literature to lend itself to scientific investigation. In fact, the mother of all problems in this debate is at the very source: "singularity" and "superhuman intelligence" are non-scientific terms based on non-scientific chatting. The term "intelligence" itself is hardly scientific. Ask one hundred psychologists and you will get one hundred different definitions. Ask philosophers and you will get thick tomes written in a cryptic language. Ask neurobiologists and they may simply ignore you. I also feel that this discussion should be complemented with an important (more important?) discussion about the changes in human intelligence due to the increased "intelligence" of machines. This change in human intelligence may have a stronger impact on the future of human civilization than the improvements in machine intelligence. Finally, i am intrigued by another sociological/anthropological aspect of this discussion: humans seem to have a genetic propensity to believe in higher forms of intelligence (gods, saints, UFOs, ...) and the myth of the Singularity could simply be its manifestation in our post-religious 21st century.

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.

Our Final Invention

Its nature and future

The Asimovian Rules and to what Extent They Can Lead to a Leakproof Singularity. Using the Example of the TV Series Westworld

The Singularity Is Nearer

The Last Firewall

AI and the Future of Your Mind

Artificial Intelligence and the Technological Singularity

**The noted inventor and futurist's successor to his landmark book The Singularity Is Near explores how technology will refashion the human race in the decades to come Since it was first published in 2005, Ray Kurzweil's The Singularity Is Near and its vision of the future have been influential in spawning a worldwide movement with millions of followers, hundreds of books, major films (Her, Lucy, Ex Machina), and thousands of articles. During the succeeding decade many of Kurzweil's predictions about technological advancements have been borne out, and their viability has become familiar to the public through such now commonplace concepts as AI, intelligent machines, and bioengineering. In this entirely new book Ray Kurzweil brings a fresh perspective to advances in the singularity--assessing the progress of many of his predictions and examining the novel advancements that, in the near future, will bring a revolution in knowledge and an expansion of human potential. Among the topics he discusses are rebuilding the world, atom by atom with devices like nanobots; radical life extension beyond the current age limit of 120; reinventing intelligence by expanding biological capacity with nonbiological intelligence in the cloud; how life is improving with declines in areas such as poverty and violence; and the growth of technologies such as renewable energy and 3-D printing, which can be applied to everything from clothes to building materials to growing human organs. He also considers the potential perils of biotechnology, nanotechnology, and artificial intelligence, including such topics of current controversy as how AI will impact unemployment and the safety of autonomous cars, and After Life technology, which will reanimate people who have passed away through a combination of data and DNA.**

**"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.**

**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.**

**Artificial intelligence will never, ever, reach the singularity. Until it does. When that happens, will humans embrace the new "Masters of the Universe", or try to pull the plug? Knowing humans, they will try to do both at once. In particular, the current MOTU, those ultra wealthy elite who are (ahem) no so subtly calling dibs on low earth orbit, and then, maybe, Mars. How will they react to being demoted to JABORJ (Just a Bunch of Rich Jerks)? How will current governments (say, China and the U.S.) react to increasingly powerful, sentient, and not at all impressed, AI, roaming about in the global financial system? This series of books postulates that they will all take it, not well. However, as always, there will be those who embrace the change, and as usual it will be those who were not thriving under the old rules. Sentient Artificial Intelligence continues to disrupt the established order, threatening human power centers and, importantly, dominance. A charismatic evangelist finds powerful allies in the Hierarchy to help turn the population against the new AI and foment a re-alignment of political power. The One Road system, a globe, spanning transportation and data infrastructure built through giant American and Chinese conglomerates, becomes the backbone of the economy, and a haven for anyone caught in the dragnets of the SouthWest Commission.**

**The applications of Artificial Intelligence lie all around us; in our homes, schools and offices, in our cinemas, in art galleries and - not least - on the Internet. The results of Artificial Intelligence have been invaluable to biologists, psychologists, and linguists in helping to understand the processes of memory, learning, and language from a fresh angle. As a concept, Artificial Intelligence has fuelled and sharpened the philosophical debates concerning the nature of the mind, intelligence, and the uniqueness of human beings. Margaret A. Boden reviews the philosophical and technological challenges raised by Artificial Intelligence, considering whether programs could ever be really intelligent, creative or even conscious, and shows how the pursuit of**

**Artificial Intelligence has helped us to appreciate how human and animal minds are possible.**

**Could artificial intelligence really out-think us (and would we want it to)?**

**The Singularity Is Near**

**Artificial You**

**How the Quest for the Ultimate Learning Machine Will Remake Our World**

**Surviving AI**

**2062**

**Artificial Intelligence and the End of the Human Era**

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.

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 book provides a detailed history of Artificial Intelligence, and a "reality check". Themes explored in this book include: Most of the "intelligence" of our machines is due to the environment that humans structure for them.- We are building "vast algorithmic bureaucracies" all around us.- Automation is an effect, not a cause.- The danger is not that machines will become as intelligent as us but that we will become as dumb as them (the Turing Test in reverse).- We have always coexisted with super-human (or, better, non-human) intelligence.- The Singularity is simply a religion for the god-less 21st century.

The first book published by Gary K. Martin within the military science-fiction series *Singularity*. Taking place a few years in the future, prominent industrial leaders come together in secrecy to develop Artificial Intelligence in a controlled environment where the machines that they experiment on have never come into contact with humans. The story follows the perspective of the machines that are the test subject of the "Vault Project," the government agents that created it, and a protagonist that seemingly has no relation to the main storyline. The *Singularity Series* is a prophetic chronicling of the world as we know it after the birth of A.I., as told by multiple different points of view and perspectives, in order to show a thoroughly-detailed example of how every aspect of human life will change - and never return to normal - after what society has coined the "Technological Singularity." The author claims to have no plan for how the series will end because, in his own words, "After the birth of A.I., nothing will ever be the same, and we have no hope of turning back the clock. It only makes sense then, that the series will be endless, because we can never get back to normal once this happens." The first book in the series - *Conception*, takes place immediately following the birth of A.I. and takes place over the course of just the first few days of its birth.

Will machines take over the world one day? Will they have human emotions? Will they be our friends or foes? Will they make our lives easier or will they wipe out the human labor force? Readers will come to their own conclusions after reading articles from leading experts forecasting how robots and machines will be integrated into our world, as well as their warnings about how it could all go horribly wrong. Aside from the obvious benefits, the development of artificial intelligence brings up a host of ethical considerations, which are being debated by the world's technology leaders before it's too late.

**Why Computers Can't Think the Way We Do**

**A Futuristic Approach**

**The Age of Spiritual Machines**

**The Singularity is Closer than it Appears**

**Paths, Dangers, Strategies**

**The Cosmic Zoo**

**Bergson, Gibson and the Mythologies of Artificial Intelligence**

Two leaders in the field offer a compelling analysis of the current state of the art and reveal the steps we must take to achieve a truly robust artificial intelligence. Despite the hype surrounding AI, creating an intelligence that rivals or exceeds human levels is far more complicated than we have been led to believe. Professors Gary Marcus and Ernest Davis have spent their careers at the forefront of AI research and have witnessed some of the greatest milestones in the field, but they argue that a computer beating a human in *Jeopardy!* does not signal that we are on the doorstep of fully autonomous cars or superintelligent machines. The achievements in the field thus far have occurred in closed systems with fixed sets of rules, and these approaches are too narrow to achieve genuine intelligence. The real world, in contrast, is wildly complex and open-ended. How can we bridge this gap? What will the consequences be when we do? Taking inspiration from the human mind, Marcus and Davis explain what we need to advance AI to the next level, and suggest that if we are wise along the way, we won't need to worry about a future of machine overlords. If we focus on endowing machines with common sense and deep understanding, rather than simply focusing on statistical analysis and gathering ever larger collections of data, we will be able to create an AI we can trust--in our homes, our cars, and our doctors' offices. *Rebooting AI* provides a lucid, clear-eyed assessment of the current science and offers an inspiring vision of how a new generation of AI can make our lives better.

Are you prepared to meet your future robot overlords? According to futurist Ray Kurzweil, we only have until 2045 to prepare for the singularity--the moment AI achieves human-level general intelligence. That's where *The Singularity Survival Guide* comes in. How will you greet this new form of higher intelligence? With a simple handshake? And what are you going to do if the AI is evil? How will you fight against it? Find the answers to these and many other questions in this survival guide--before it's too late."The technological singularity has officially been treated to a full-scale parody, and it's even more comical and irreverent than it sounds." - Zoltan Istvan, Founder of the Transhumanist Party" A timely satire, even if humor doesn't stand a chance of saving us

from the sort of superintelligence Clarke envisions." - Arif Khan, Chief Marketing Officer, SingularityNET Published by Logos Literature.

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

The SingularityCould artificial intelligence really out-think us (and would we want it to)?Andrews UK Limited

One thousand years in the future, humans no longer rule... In the early twenty-first century, humanity marveled at its greatest creation: Artificial Intelligence. They never foresaw the consequences of such a creation, though... Now, in a world where humans must meet specifications to continue living, a man named Caesar emerges. Different, both in thought and talent, Caesar somehow slipped through the genetic net meant to catch those like him. Eyes are falling on Caesar now, though, and he can no longer hide. The Artificial Intelligence wants him dead, but others want him to lead their revolution... Can one man stand against humanity's greatest creation? A don't-miss epic science fiction novel that pits one man fighting for the future of all people!

Singularity

Surviving and Thriving in a Smarter, Richer, and More Dangerous World

A.I. Apocalypse

The Promise and Peril of Artificial Intelligence

Building Artificial Intelligence We Can Trust

The World that AI Made

The Singularity

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.

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

The "Singularity" is the hypothetical point in time □ considered by AI writers to be very close □ where artificial intelligence and super intelligent machines meet and surpass human intelligence. Yet neither AI theory nor Cognitive Science □ equally committed to the computer model of mind □ have begun to make the needed considerations on what the human mind truly is. While they have entertained "embodiment," or "embodied cognition," and/or J.J. Gibson's ecological psychology as needed components, these frameworks are themselves only dimly grasped. For coherence, Gibson must be placed within the remarkable model of time, mind and brain of Henri Bergson. This book is a challenge to look deeply at the real nature of human perception, memory, and consciousness, that is, what is required to even begin to claim human equivalence in intelligence. Much of this book is a reorganization, expansion and deepening of portions of an earlier book, Time and Memory: A Primer on the Scientific Mysticism of Consciousness, aimed and refocused for those interested in artificial intelligence and its future.

"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."

Leon Tsarev is a high school student set on getting into a great college program, until his uncle, a member of the Russian mob, coerces him into developing a new computer virus for the mob's botnet - the slave army of computers they used to commit digital crimes. The evolutionary virus Leon creates, based on biological principles, is successful -- too successful. All the world's computers are infected. Everything from cars to payment systems and, of course, computers and smart phones stop functioning, and with them go essential functions including emergency services, transportation, and the food supply. Billions may die. But evolution never stops. The virus continues to evolve, developing intelligence, communication, and finally an entire civilization. Some

may be friendly to humans, but others are not. Leon and his companions must race against time and the military to find a way to either befriend or eliminate the virus race and restore the world's computer infrastructure. Praise for the Singularity Series: "Highly entertaining, gripping, thought inspiring. 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 "A tremendous book that every single person needs to read. In the vein of Daniel Suarez's Daemon and Freedom(TM), William shows that science fiction is becoming science fact." - Brad Feld, managing director Foundry Group, cofounder TechStars "A fascinating look at how simple and benign advancements in technology could lead to the surprise arrival of the first AI. And like all good techno-thrillers, the reality of AI is less than ideal." - Jason Glaspey, SILICON FLORIST "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

Rebooting AI

A History of Artificial Intelligence and Why the Singularity Is Not Coming Any Time Soon

The Economic Singularity

Avogadro Corp

Why the Singularity Is Not Coming Any Time Soon and Other Meditations on the Post-Human Condition and the Future of Intelligence

A Fallacy Or a Great Opportunity?

Encyclopedia of Artificial Intelligence: The Past, Present, and Future of AI

**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 "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.

Seminar paper from the year 2019 in the subject Computer Sciences - Artificial Intelligence, grade: 1.0, Zeppelin University Friedrichshafen, language: English, abstract: Ever since the first conference on artificial intelligence (AI) was held in 1956 at Dartmouth College, the question of singularity is asked. The singularity is the event where AI exceeds human intelligence. Due to the singularity, the AI may outpace humanity and create a non-beneficial outcome. In contrast, the leakproof singularity describes a singularity having a beneficial outcome (where the AI is a problem solver). Therefore, the paper discusses the question to what extent a leakproof singularity could happen. For the leakproof singularity, a framework for ethical (in our case beneficial) decisions is needed where every action is evaluated by an ethical layer. The paper focuses on the consequences of AI's decision making using a set of rules. After explaining the singularity and the consequences of the singularity in the first chapter, the example of Westworld is taken in the second chapter in order to exemplify and introduce basic concepts like the leakproof singularity and conscious AI. In the third chapter the classical and extended Asimovian laws are explained which is followed by a specific critique of the three classical maxims. Afterwards, in chapter 3.2, a general critical reflection of the laws is given. Moreover, in the fourth chapter three scenarios are developed for a post-singular rule combining the Asimovian laws with other examples from novels, philosophy and computer sciences. In the last chapter, the paper is evaluated critically and an outlook for future development is provided.

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](http://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.

**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.**

**Artificial Intelligence and the Death of Capitalism**

**The Singularity: Heretic**

**Complex Life on Many Worlds**

**Collapsing the Singularity**

**The Singularity Survival Guide: How to Get on the Good Side of Your Future Robot Overlords**

**Artificial Superintelligence**

**The Singularity Will Save Humanity**

A leading artificial intelligence researcher lays out a new approach to AI that will enable people to coexist successfully with increasingly intelligent machines.

This book represents the views of one of the greatest mathematicians of the twentieth century on the analogies between computing machines and the living human brain. John von Neumann concludes that the brain operates in part digitally, in part analogically, but uses a peculiar statistical language unlike that employed in the operation of man-made computers. This edition includes a new foreword by two eminent figures in the fields of philosophy, neuroscience, and consciousness.

Futurists are certain that humanlike AI is on the horizon, but in fact engineers have no idea how to program human reasoning. AI reasons from statistical correlations across data sets, while common sense is based heavily on conjecture. Erik Larson argues that hyping existing methods will only hold us back from developing truly humanlike AI.

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"

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.

The Technological Singularity

The Day AI Becomes God

Artificial Intelligence and the Two Singularities

AI and the Singularity

Machine Burning

When Humans Transcend Biology

Philosophy and Theory of Artificial Intelligence

Artificial (General) Intelligent robots will be cleverer than us, able to do anything we can do, better, and faster, for twenty four hours a day, with no time off demanded either, except perhaps for an occasional self-service. Most of the jobs only humans could do before will be gone, and this situation will soon seriously begin to happen, well before the end of the next decade. What will that do to government unemployment figures? Some experts believe 80% is possible. Fully automated companies, with minimal labour costs, will probably do well financially, at first, but how will unemployed people buy their products? Some alternative, a guaranteed permanent Basic Allowance for example, will have to be introduced well before we reach that state of unemployment... or will we still need money by then? Later when ASI, Artificial Super Intelligence, arrives some experts even fear humanity may be facing extinction. The consequences, of this Singularity, this unprecedented hard to appreciate situation, are going to be devastating, even if the relevant authorities start planning for it immediately... and that is unlikely, knowing them the way we do. Will the subject even be mentioned in the next government's manifesto? This updated edition, for general readers of all ages, explains what experts believe will happen in only a few years from now. They do not all agree with each other. Their predictions range from A New Golden Age for us all, through Uncertainty, to Total Disaster. Readers must decide for themselves who will be proved correct.

Human Compatible

The Artificial Intelligence Explosion

The Master Algorithm

When Computers Can Think

The Artificial Intelligence Singularity

The Turing Exception

Book Two of the Singularity Series