

Beyond AI Creating The Conscience Of The Machine

'If AI is outside your field, or you know something of the subject and would like to know more then Artificial Intelligence: The Basics is a brilliant primer.' - Nick Smith, Engineering and Technology Magazine November 2011
Artificial Intelligence: The Basics is a concise and cutting-edge introduction to the fast moving world of AI. The author Kevin Warwick, a pioneer in the field, examines issues of what it means to be man or machine and looks at advances in robotics which have blurred the boundaries. Topics covered include: how intelligence can be defined whether machines can 'think' sensory input in machine systems, the nature of consciousness, the controversial culturing of human neurons. Exploring issues at the heart of the subject, this book is suitable for anyone interested in AI, and provides an illuminating and accessible introduction to this fascinating subject.

Focused on mapping out the frontiers of technology, this volume serves as an excellent, forward-looking resource in the field and in cognate areas of study. The 32 chapters, all of them appearing in print here for the first time, were written by both established scholars and fresh voices. They cover topics ranging from data discrimination and engineering design, to art and technology, space junk, and beyond. Spaces for the Future: A Companion to the Philosophy of Technology is structured in six parts: (1) Ethical Space and Experience; (2) Political Space and Agency; (3) Virtual Space and Property; (4) Personal Space and Design; (5) Inner Space and Environment; and (6) Outer Space and Imagination. The organization maps out current and emerging spaces of activity in the field and anticipates the big issues that we soon will face.

Can you imagine swapping your job for a virtual reality? This technology-based look at the afterlife chronicles America's fascination with death and reveals how digital immortality may become a reality. ¶ Reveals the period in American history that established cultural views about the afterlife. ¶ Discusses how technology aids in achieving and designing perspectives on heaven and immortality. ¶ Reviews technologies surrounding digital mortality, including Dmitry Itskov's 2045 Initiative and Martine Rothblatt's Terasen Foundation. ¶ Illustrates how the Internet helps those who have lost loved ones to grieve and cope in new ways. ¶ Links perceptions about death and the hereafter to typical American attitudes, including optimism, confidence, self-reliance, and innovation. ¶ Examines the use of virtual memorials, online obituaries, Facebook pages of deceased users, and avatars

'God from the machine' (deus ex machina) refers to an ancient dramatic device where a god was mechanically brought onto the stage to save the hero from a difficult situation. But here, William Sims Bainbridge uses the term in a strikingly different way. Instead of looking to a machine to deliver an already known god, he asks what a computing machine and its simulations might teach us how religion and religious beliefs come to being. Bainbridge posits the virtual town of Cyburg, population 44,100. Then, using rules for individual and social behavior taken from the social sciences, he models a complex community where residents form groups, learn to trust or distrust each other, and develop religious faith. Bainbridge's straightforward arguments point to many more applications of computer simulation in the study of religion. God from the Machine will show how religion can serve as an important text in any class with a social scientific approach to religion.

Crimes, Contracts, and Torts

Robotethics

Social Robots

How to Understand Ourselves and Other Beings, from Animals to AI to Aliens

Impact of Industry 4.0 on Architecture and Cultural Heritage

Proceedings of the First AGI Conference

Machine Ethics and Robot Ethics

Popular science writer Philip Ball explores a range of sciences to help our answers to a huge, philosophical rich question: How do we even begin to think about minds that are not human? Sciences from zoology to astrobiology, computer science to neuroscience, are seeking to understand minds in their own distinct disciplinary realms. Taking a uniquely broad view of minds and where to find them—including in plants, aliens, and God—Philip Ball pulls the pieces together to explore what sorts of minds we might expect to find in the universe. In so doing, he offers for the first time a unified way of thinking about what minds are and what they can do, by locating them in what he calls the “space of possible minds.” By identifying and mapping out properties of mind without prioritizing the human, Ball sheds new light on a host of fascinating questions: What moral rights should we afford animals, and can we understand their thoughts? Should we worry that AI is going to take over society? If there are intelligent aliens out there, how could we communicate with them? Should we? Understanding the space of possible minds is a key to understanding some of the most challenging questions in contemporary science: What is thought? What is consciousness? And what (if anything) is free will? Informed by conversations with leading researchers, Ball’s brilliant survey of current views about the nature and existence of minds is more mind-expanding than we could imagine. In this fascinating panorama of other minds, we come to better know our own.

The history of robotics and artificial intelligence in many ways is also the history of humanity's attempts to control such technologies. From the Golem of Prague to the military robots of modernity, the debate continues as to what degree of independence such entities should have and how to make sure that they do not turn on us, its inventors. Numerous recent advancements in all aspects of research, development and deployment of intelligent systems are well publicized but safety and security issues related to AI are rarely addressed. This book is proposed to mitigate this fundamental problem. It is comprised of chapters from leading AI Safety researchers addressing different aspects of the AI control problem as it relates to the development of safe and secure artificial intelligence. The book is the first edited volume dedicated to addressing challenges of constructing safe and secure advanced machine intelligence. The chapters vary in length and technical content from broad interest opinion essays to highly formalized algorithmic approaches to specific problems. All chapters are self-contained and could be read in any order or skipped without a loss of comprehension.

Artificial Probability and Friends: Proceedings of the Ray Solomonoff 85th memorial conference is a collection of original work and surveys. The Solomonoff 85th memorial conference was held at Monash University's Clayton campus in Melbourne, Australia as a tribute to pioneer, Ray Solomonoff (1926-2009), honouring his various pioneering works - most particularly, his revolutionary insight in the early 1960s that the universality of Universal Turing Machines (UTMs) could be used for universal Bayesian prediction and artificial intelligence (machine learning). This work continues to increasingly influence and under-pin statistics, econometrics, machine learning, data mining, inductive inference, search algorithms, data compression, theories of (general) intelligence and philosophy of science - and applications of these areas. Ray not only envisioned this as the path to genuine artificial intelligence, but also, still in the 1960s, anticipated stages of progress in machine intelligence which would ultimately lead to machines surpassing human intelligence. Ray warned of the need to anticipate and discuss the potential consequences - and dangers - sooner rather than later. Possibly foremostly, Ray Solomonoff was a fine, happy, frugal and adventurous human being of gentle resolve who managed to fund himself while electing to conduct so much of his paradigm-changing research outside of the university system. The volume contains 35 papers pertaining to the abovementioned topics in tribute to Ray Solomonoff and his legacy.

A provocative attempt to think about what was previously considered unthinkable: a serious philosophical case for the rights of robots. We are in the midst of a robot invasion, as devices of different configurations and capabilities slowly but surely come to take up increasingly important positions in everyday social reality—self-driving vehicles, recommendation algorithms, machine learning decision-making systems, and social robots of various forms and functions. Although considerable attention has already been devoted to the subject of robots and responsibility, the question concerning the social status of these artifacts has been largely overlooked. In this book, David Gunkel offers a provocative attempt to think about what has been previously regarded as unthinkable: whether and to what extent robots and other technological artifacts of our own making can and should have any claim to moral and legal standing. In his analysis, Gunkel invokes the philosophical distinction (developed by David Hume) between “is” and “ought” in order to evaluate and analyze the different arguments regarding the question of robot rights. In the course of his examination, Gunkel finds that none of the existing positions or proposals hold up under scrutiny. In response to this, he then offers an innovative alternative proposal that effectively flips the script on the is/ought problem by introducing another, altogether different way to conceptualize the social situation of robots and the opportunities and challenges they present to existing moral and legal systems.

Engaging With Stakeholders

Robot Theology and Prophetic: Humanity in Cyberculture, Cyberspace, and Science Fiction

Final Jeopardy

The Machine Question

Managing the Journey

The Future of Artificial Intelligence

Outsmarting AI

Most humanity will have some 5,000 years of recorded history has been experiencing growing pains, with no end in sight. It is high time for humanity to grow up and to transcend itself by embracing transhumanism. Transhumanism offers the most inclusive ideology for all ethnicities and races, the religious and the atheists, conservatives and liberals, the young and the old regardless of socioeconomic status, gender identity, or any other individual qualities. This book expands on contemporary views and practical advice from more than 70 transhumanists. Astronaut Neil Armstrong said on the Apollo 11 moon landing in 1969, “One small step for a man, one giant leap for mankind.” Transhumanism is the next logical step in the evolution of humankind, and it is the existential solution to the long-term survival of the human race.

Science fiction explores the wonderful, baffling and wildly entertaining aspects of a universe unimaginably old and vast, and with a future even more immense. It reaches into that endless cosmos with the tools of rational investigation and storytelling. At the core of both science and science fiction is the engaged human mind—a consciousness that sees and feels and thinks and loves. But what is this mind, this aware and self-aware consciousness that seems unlike anything else we experience? What makes consciousness the Hard Problem of philosophy, still unsolved after millennia of probing? This book looks into the heart of this mystery - at the science and philosophy of consciousness and at many inspiring fictional examples - and finds strange, challenging answers. The book's content and entertaining style will appeal equally to science fiction enthusiasts and scholars, including cognitive and neuroscientists, as well as philosophers of mind. It is a refreshing romp through the science and science fiction of consciousness.

Social robotics is a cutting edge research area gathering researchers and stakeholders from various disciplines and organizations. The transformational potential that these machines, in the form of, for example, caregiving, entertainment or partner robots, pose to our societies and to us as individuals seems to be limited by our technical limitations and phantasy alone. This collection contributes to the field of social robotics by exploring its boundaries from a philosophically informed standpoint. It constructively outlines central potentials and challenges and thereby also provides a stable foundation for further research of empirical, qualitative or methodological nature.

Today's wars have no definitive end in sight, are conducted among civilian populations, and are fought not only by soldiers but also by unmanned aerial vehicles. According to M. Shane Riza, this persistent conflict among the people and the trend toward robotic warfare has outpaced deliberate thought and debate about the deep moral issues affecting the military mission and the warrior spirit. The pace of change, Riza explains, is revolutionizing warfare in ways seldom discussed but vitally important. A key development is risk inversion, which occurs when all noncombatants are at greater risk th.

Boundaries, Potential, Challenges

Ethics in Research Practice and Innovation

Artificial Intelligence Models of Religious Cognition

Where is My Flying Car?

Papers from the Ray Solomonoff 85th Memorial Conference, Melbourne, VIC, Australia, November 30 – December 2, 2011

The Laws of Robots

This book is designed to offer a comprehensive high-level introduction to transhumanism, an international political and cultural movement that aims to produce a “paradigm shift” in our ethical and political understanding of human evolution. Transhumanist thinkers want the human species to take the course of evolution into its own hands, using advanced technologies currently under development—such as robotics, artificial intelligence, biotechnology, cognitive neurosciences, and nanotechnology—to overcome our present physical and mental limitations, improve our intelligence beyond the current maximum achievable level, acquire skills that are currently the preserve of other species, abolish involuntary aging and death, and ultimately achieve a post-human level of existence. The book covers transhumanism from a historical, philosophical, and scientific viewpoint, tracing its cultural roots, discussing the main philosophical, epistemological, and ethical issues, and reviewing the state of the art in scientific research on the topics of most interest to transhumanists. The writing style is clear and accessible for the general reader, but the book will also appeal to graduate and undergraduate students.

"Moral machines is a first introduction to the emerging field of robot ethics. There is much here that will interest ethicists, philosophers, cognitive scientists, and robotists."—Peter Danielson, Notre Dame Philosophical Reviews—

Artificial intelligence (AI) permeates Google searches, the personal assistants in our smartphones, and is all over our newsfeeds. Watson's machine learning has already started to revolutionize many important industries including oncology, law, finance, and entertainment. The idea that man is about to increase his immediate surroundings with exponential gains in the level of intelligence over the coming generations is based upon a technological revolution and the potential for artificial superintelligence (ASI). It is within this context that there is a prevailing need for a discussion of its ethical implications. As a Christian ethicist, Paul Golata believes that the need for this conversation to be informed by Christian principles is imperative. ASI as a move toward the proper handling of information. However, how a society interprets and applies this information is actually more pertinent than the raw amount of information it possesses. This important ethical conversation is being led by humanistic thinkers who assume that all of reality is just matter in motion and that mind is nothing more than electrochemical activity in the “wetware” of human brains. The Ethics of Superintelligent Design critically examines and challenges some of the most important trajectories of ASI while upholding the authority and inerrancy of the Bible, the supernatural creation account, a realistic view of the state of humanity, and biblical ethics.

Engaging with Stakeholders: A Relational Perspective on Responsible Business contends that meaningful and constructive stakeholder engagement efforts should be rooted in a deep relational process of shared understanding, expectations, and viewpoints, through honest, continued dialogue between stakeholders and company management. This anthology follows and reaffirms this view, which also establishes the increasing need to explore the subtleties of how companies can respectfully engage their stakeholders in ways that reflect the corporate strategy and contribute to the ongoing development of business activities and creation of value, for themselves and stakeholders, from social, environmental, and economic perspectives. Stakeholder engagement practices, however, remain highly complex and difficult to manage; their ability to generate value in an inclusive way requires critical consideration. Sound stakeholder engagement efforts also constitute a keystone for successful business activities. Drawing on a wide range of literature and studies, this book addresses key dimensions of stakeholder engagement, through a responsible business lens, and thereby contributes to identifying the opportunities, challenges, and key organizational implications associated with their unfolding. The four main topics covered are: • Delineating the nature and multiple reasons d'être of stakeholder engagement • Dialogical and communicational foundations of stakeholder engagement • Engaging with diverse stakeholders throughout the value chain • Reaping organizational returns and relational rewards of stakeholder engagement efforts

Power, Profit, and Leadership in the Age of Machines

A Navigating Overview

Consciousness and Science Fiction

Our Future with Social Robots

Creating the Conscience of the Machine

Cyber-Humans

The two-volume set LNAI 7267 and 7268 (together with LNCS 7269) constitutes the refereed proceedings of the 11th International Conference on Artificial Intelligence and Soft Computing, ICAISC 2012, held in Zakopane, Poland in April/ May 2012. The 212 revised full papers presented were carefully reviewed and selected from 483 submissions. The papers are organized in topical sections on neural networks and their applications, computer vision, image and speech analysis, data mining, hardware implementation, bioinformatics, biometrics and medical applications, concurrent parallel processing, agent systems, robotics and control, artificial intelligence in modeling and simulation, various problems of artificial intelligence.

This book explores how the design, construction, and use of robotics technology may affect today's legal systems and, more particularly, matters of responsibility and agency in criminal law, contractual obligations, and torts. By distinguishing between the behaviour of robots as tools of human interaction, and robots as proper agents in the legal arena, jurists will have to address a new generation of "hard cases." General disagreement may concern immunity in criminal law (e.g., the employment of robot soldiers in battle), personal accountability for certain robots in contracts (e.g., robo-traders), much as clauses of strict liability and negligence-based responsibility in extra-contractual obligations (e.g., service robots in tort law). Since robots are here to stay, the aim of the law should be to wisely govern our mutual relationships.

"An extraordinarily good synthesis from an amazing range of philosophical, legal, and technological sources. . . . the book will appeal to legal academics and students, lawyers involved in e-commerce and cyberspace legal issues, technologists, moral philosophers, and intelligent lay readers interested in high tech issues, privacy, [and] robotics."—Kevin Ashley, University of Pittsburgh School of Law
As corporations and government agencies replace human employees with online customer service and automated phone systems, we become accustomed to doing business with nonhuman agents. If artificial intelligence (AI) technology advances as today's leading researchers predict, these agents may soon function with such limited human input that they appear to act independently. When they achieve that level of autonomy, what legal status should they have? Samir Chopra and Laurence F. White present a carefully reasoned discussion of how existing philosophy and legal theory can accommodate increasingly sophisticated AI technology. Arguing for the legal personhood of an artificial agent, the authors discuss what it means to say it has "knowledge" and the ability to make a decision. They consider key questions such as who must take responsibility for an agent's actions, whom the agent serves, and whether it could face a conflict of interest.

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

Our Future with Machines

The Transhumanism Handbook

Killing Without Heart

The Story of Watson, the Computer That Will Transform Our World

Teaching Robots Right from Wrong

Advanced Introduction to Law and Artificial Intelligence

Artificial Superintelligence

From a pessimistic account of technological stagnation since the 1970s and an imaginative blueprint for a richer, more abundant future. The science fiction of the 1960s promised us a future remade by technological innovation. We'd vacation in godesmic domes on Mars, have meaningful conversations with computers, and drop our children off at school in flying cars. Fast-forward 60 years, and we're still stuck in traffic in gas-guzzling sedans and boarding the same types of planes we flew in over half a century ago. What happened to the future we were promised? In Where Is My Flying Car?, J. Storrs Hall sets out to answer this deceptively simple question. What starts as an examination of the technical limitations of building flying cars evolves into an investigation of the scientific, technological, and social roots of the economic stagnation that started in the 1970s. From the failure to adopt nuclear energy and the suppression of cold fusion technology to the rise of a counterculture hostile to progress, Hall recounts how our collective ambitions for the future were derailed, with devastating consequences for global wealth creation and distribution. He then outlines a framework for a future powered by exponential progress—one in which we build as much in the world of atoms as we do in the world of bits, one rich in abundance and wonder. Drawing on years of original research and personal engineering experience, Where Is My Flying Car?, originally published in 2018, is an urgent, timely analysis of technological progress over the last 50 years and a bold vision for a better future.

"The co-author of Moral Machines explores accountability challenges related to a world shaped by such technological innovations as combat drones, 3-D printers and synthetic organisms to consider how people of the near future can be protected."—Novelist.

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 predilections 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 Agreements 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.

Woodrow Barfield and Ugo Pagallo present a succinct introduction to the legal issues related to the design and use of artificial intelligence (AI). Exploring human rights, constitutional law, data protection, criminal law, tort law, and intellectual property law, they consider the laws of a number of jurisdictions including the US, the European Union, Japan, and China, making reference to case law and statutes.

Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence

History, Philosophy and Current Status

Spaces for the Future

Robot Ethics

Transhumanism - Engineering the Human Condition

Robotic Persons

The Book of Minds

The papers collected in this volume document the exchange and development of ideas that comprised the 5th Global Conference on Visions of Humanity in Cyberculture, Cyberspace, and Science Fiction, hosted at Mansfield College, Oxford, United Kingdom, in July 2010.

The modern age of the 4th Industrial Revolution, advancements in communication and connectivity are transforming the professional world as new technologies are being embedded into society. These innovations have triggered the development of a digitally driven world where adaptation is necessary. This is no different in the architectural field, where the changing paradigm has opened new methods and advancements that have yet to be researched. Impact of Industry 4.0 on Architecture and Cultural Heritage is a pivotal reference source that provides vital research on the application of new technological tools, such as digital modeling, within architectural design, and improves the understanding of the strategic role of Industry 4.0 as a tool to empower the role of architecture and cultural heritage in society. Moreover, the book provides insights and support concerned with advances in communication and connectivity among digital environments in different types of research and industry communities. While highlighting topics such as semantic processing, crowdsourcing, and interactive environments, this publication is ideally designed for architects, engineers, construction professionals, cultural researchers, academicians, and students.

It is predicted that robots will surpass human intelligence within the next fifty years. The ever increasing speed of advances in technology and neuroscience, coupled with the creation of super computers and enhanced body parts and artificial limbs, is paving the way for a merger of both human and machine. Devices which were once worn on the body are now being implanted into the body, and as a result, a class of true cyborgs, who are displaying a range of skills beyond those of normal humans-beings, are being created. There are cyborgs which can see colour by hearing sound, others have the ability to detect magnetic fields, others are equipped with telescopic lenses to aid their vision or implanted computers to monitor their heart, and some use thought to communicate with a computer or to manipulate a robotic arm. This is not science-fiction, these are developments that are really happening now, and will continue to develop in the future. However, a range of legal and policy questions has arisen alongside this rise of artificial intelligence. Cyber-Humans provides a deep and unique perspective on the technological future of humanity, and describes how law and policy will be particularly relevant in creating a fair and equal society and protecting the liberties of different life forms which will emerge in the 21st century. Dr Woodrow (Woody) Barfield previously headed up the Sensory Engineering Laboratory, holding the position of Industrial and Systems Engineering Professor at the University of Washington. His research revolves around the design and use of wearable computers and augmented reality systems and holds both JD and LL.M degrees in intellectual property law and policy. He has published over 350 articles and major presentations in the areas of computer science, engineering and law. He currently lives in Chapel Hill, NC, USA.

Prominent experts from science and the humanities explore issues in robot ethics that range from sex to war. Robots today serve in many roles, from entertainer to educator to executioner. As robotics technology advances, ethical concerns become more pressing: Should robots be programmed to follow a code of ethics, if this is even possible? Are there risks in forming emotional bonds with robots? How might society—and ethics—change with robotics? This volume is the first book to bring together prominent scholars and experts from both science and the humanities to explore these and other questions in this emerging field. Singularity nearby? 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 Ethical and Social Implications of Robotics

A Futuristic Approach

A Companion to Philosophy of Technology

Moral Machines

Internet Afterlife: Virtual Salvation in the 21st Century

God from the Machine

Old Questions through New Media

Once the stuff of science fiction, recent progress in artificial intelligence, robotics, and machine learning means that these rapidly advancing technologies are finally coming into widespread use within everyday life. Such rapid development in these areas also brings with it a host of social, political and legal issues, as well as a rise in public concern and academic interest in the ethical challenges these new technologies pose. This volume is a collection of scholarly work from leading figures in the development of both robot ethics and machine ethics; it includes essays of historical significance which have become foundational for research in these two new areas of study, as well as important recent articles. The research articles selected focus on the control and governance of computational systems; the exploration of ethical and moral theories using software and robots as experimental simulations; inquiry into the necessary requirements for moral agency and the basis and boundaries of rights; and questions of how best to design systems that are both useful and morally sound. Collectively the articles ask what the practical ethical and legal issues, arising from the development of robots, will be over the next twenty years and how best to address these future considerations.

This volume explores the ethical questions that arise in the development, creation and use of robots that are capable of semiautonomous or autonomous decision making and human-like action. It examines how ethical and moral theories can and must be applied to address the complex and critical issues of the application of these intelligent robots in society. Coverage first presents fundamental concepts and provides a general overview of ethics, artificial intelligence and robotics. Next, the book studies all principal ethical applications of robots, namely medical, assistive, socialized and war robotethics. It looks at such issues as robotic surgery, children-robot and elderly-robot therapeutic/social interactions and the use of robots, especially autonomous lethal ones, in warfare. In addition, a chapter also considers Japanese robotethics as well as key intercultrual and robot legislation issues. Overall, readers are provided with a thorough investigation into the moral responsibility (if any) of autonomous robots when doing harm. This volume will serve as an ideal educational source in engineering and robotics courses as well as an introductory reference for researchers in the field.

Beyond AICreating the Conscience of the MachinePrometheus Books

*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.*

Virtual Salvation in the 21st Century

A Relational Perspective on Responsible Business

Artificial Intelligence Safety and Security

Artificial General Intelligence 2008

Machines that Think

Robot Rights

Limits on Robotic Warfare in an Age of Persistent Conflict

Robotic Persons will introduce the evangelical community to the journey of Robotic Futurism and how current and forthcoming AI-driven robots will impact human value and dignity. This book will consider three key areas of robotic development and the existential risks on the horizon for humans in the fields of work, war, and sex. There are risks in the fields of work, because there is a temptation to replace human workers with automation. Current arguments for the benefit of war fighting robots posit that these robots will eliminate war and the risk of war, but there is much more to the story. Arguments for sex and companion robots proffer that they will benefit the fringe community or help those who do not have a relative to care for them, but again there are many ethical and logistical problems with these arguments. Robotic Persons will not only introduce the reader to these issues, but also gives an evangelical response to each. There is presently no evangelical work addressing these critical issues. Robotic Persons will argue that granting legal personhood to qualified robots will further prevent dehumanizing use of robots and protect human dignity and value.

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.

From factories to smartphones, Artificial Intelligence is already taking over. Outsmarting AI is not a how-to guide on making AI work, but making it work for YOU to boost profits and productivity. Each development in Artificial Intelligence (AI) technology brings about apprehension and panic for the future of society and for business. We´re bombarded with stories about the impending human-less workplace: it is no longer a question if man can be replaced by machine in certain tasks, but when. However, AI was not manufactured to destroy life as we know it. These emerging technologies were developed and are constantly updating with a particular goal in mind: optimization. AI feeds on data and information to improve outputs and increase potential. With this enhanced productivity, profit and productivity will be sure to follow. Written by Brennan Pursell, a business consultant and professor who hates jargon, and Joshua Walker, an AI pioneer with 18 years of experience in solutions and applications, Outsmarting AI is the first plain-English how-to guide on adapting AI for the non-coding proficient business leader. This book will help readers to Cut through the fog of AI hype See exactly what AI can actually do for people in business Identify the areas of their organization in most need of AI tools Prepare and control their data – AI is useless without it! Adopt AI and develop the right culture to support it! Track the productivity boost, cost savings, and increased profits! Manage and minimize the threat of crippling lawsuits

The field of Artificial Intelligence (AI) was initially directly aimed at the construction of "thinking machines" – that is, computers with human-like general intelligence. But this task proved more difficult than expected. As the years passed, AI researchers gradually shifted focus to producing AI systems that intelligently approached specific tasks in relatively narrow domains. In recent years, however, AI researchers have recognized the necessity – and the feasibility – of returning to the original goal of the field. Increasingly, there is a call to focus less on highly specialized "narrow AI" problem-solving systems, and more on confronting the difficult issues involved in creating "human-level intelligence", and ultimately general intelligence that goes beyond the human level in various ways. Artificial General Intelligence (AGI), as this renewed focus has come to be called, attempts to study and reproduce intelligence as a whole in a domain independent way. Encouraged by the recent success of several smaller-scale AGI-related meetings and special tracks at conferences, the initiative to organize the very first international conference on AGI was taken, with the goal to give researchers in the field an opportunity to present relevant research results and to exchange ideas on topics of common interest. In this collection you will find the conference papers, full-length papers, short position statements and also the papers presented in the post conference workshop on the sociocultural, ethical and futurological implications of AGI.

How to Keep Technology from Slipping Beyond Our Control

Artificial Intelligence: The Basics

Critical Perspectives on AI, Robots, and Ethics

The Technological Singularity

Beyond AI

A Dangerous Master

A Legal Theory for Autonomous Artificial Agents

Artificial intelligence (AI) is now advancing at such a rapid clip that it has the potential to transform our world in ways both exciting and disturbing. Computers have already been designed that are capable of driving cars, playing soccer, and finding and organizing information on the Web in ways that no human could. With each new gain in processing power, will scientists soon be able to create supercomputers that can read a newspaper with understanding, or write a news story, or create novels, or even formulate laws? And if machine intelligence advances beyond human intelligence, will we need to start talking about a computer's intentions?These are some of the questions discussed by computer scientistJ. Storrs Hall in this fascinating layperson's guide to the latest developments in artificial intelligence. Drawing on a thirty-year career in artificial intelligence and computer science, Hall reviews the history of AI, discussing some of the major roadblocks that the field has recently overcome, and predicting the probable achievements in the near future. There is new excitement in the field over the amazing capabilities of the latest robots and renewed optimism that achieving human-level intelligence is a reachable goal.But what will this mean for society and the relations between technology and human beings? Soon ethical concerns will arise and programmers will need to begin thinking about the computer counterparts of moral codes and how ethical interactions between humans and their machines will eventually affect society as a whole. Weaving disparate threads together in an enlightening manner from cybernetics, computer science, psychology, philosophy of mind, neuropsychology, game theory, and economics, Hall provides an intriguing glimpse into the astonishing possibilities and dilemmas on the horizon.J. Storrs Hall, Ph.D. (Laporte, PA), the founding chief scientist of Nanorex, Inc., is a research fellow for the Institute for Molecular Manufacturing and the author of Nanotechnology section for The Macmillan Encyclopedia of Energy, and a leading scientific expert on computer science.

The "charming and terrifying" story of IBM's breakthrough in artificial intelligence, from the *Business Week* technology writer and author of *The Numerati* (Publishers Weekly, starred review). For centuries, people have dreamed of creating a machine that thinks like a human. Scientists have made progress: computers can now beat chess grandmasters and help prevent terrorist attacks. Yet we still await a machine that exhibits the rich complexity of human thought—one that doesn't just crunch numbers, or take us to a relevant web page, but understands and communicates with us. With the creation of Watson, IBM's Jeopardy!-playing computer, we are one step closer to that goal. In *Final Jeopardy*, Stephen Baker traces the arc of Watson's "life," from its birth in the IBM labs to its big night on the podium. We meet Hollywood moguls and Jeopardy! masters, genius computer programmers and ambitious scientists, including Watson's eccentric creator, David Ferrucci. We see how Watson's breakthroughs and the future of artificial intelligence could transform medicine, law, marketing, and even science itself, as machines process huge amounts of data at lightning speed, answer our questions, and possibly come up with new hypotheses. As fast and fun as the game itself, *Final Jeopardy* shows how smart machines will fit into our world—and how they'll disrupt it. "The place to go if you're really interested in this version of the quest for creating Artificial Intelligence." —The Seattle Times "Like Tracy Kidder's *Soul of a New Machine*, Baker's book finds us at the dawn of a singularity. It's an excellent case study, and does good double duty as a Philip K. Dick scenario.

—Kirkus Reviews "Like a cross between *Born Yesterday* and *2001: A Space Odyssey*, Baker's narrative is both . . . an entertaining romp through the field of artificial intelligence—and a sobering glimpse of things to come." —Publishers Weekly, starred review

Artificially intelligent machines can refer to varied categories of practice - from the protection of human subjects involved in medical and social research to the publication of results research. More recently, AI researchers gradually shifted focus to producing AI systems that intelligently approached specific tasks in relatively narrow domains. In recent years, however, AI researchers have recognized the necessity – and the feasibility – of returning to the original goal of the field. Increasingly, there is a call to focus less on highly specialized "narrow AI" problem-solving systems, and more on confronting the difficult issues involved in creating "human-level intelligence", and ultimately general intelligence that goes beyond the human level in various ways. Artificial General Intelligence (AGI), as this renewed focus has come to be called, attempts to study and reproduce intelligence as a whole in a domain independent way. Encouraged by the recent success of several smaller-scale AGI-related meetings and special tracks at conferences, the initiative to organize the very first international conference on AGI was taken, with the goal to give researchers in the field an opportunity to present relevant research results and to exchange ideas on topics of common interest. In this collection you will find the conference papers, full-length papers, short position statements and also the papers presented in the post conference workshop on the sociocultural, ethical and futurological implications of AGI.

How to Keep Technology from Slipping Beyond Our Control

Artificial Intelligence: The Basics

Critical Perspectives on AI, Robots, and Ethics

The Technological Singularity

Beyond AI

A Dangerous Master

