

Human And Machine Consciousness

The book is interdisciplinary and focuses on the topic of artificial consciousness: from neuroscience to artificial intelligence, from bioengineering to robotics. It provides an overview on the current state of the art of research in the field of artificial consciousness and includes extended and revised versions of the papers presented at the International Workshop on ‘Artificial Consciousness’, held in November 2005 at Agrigento (Italy). A terrific book - essential reading for everyone seeking to make sense of Artificial Intelligence! Professor Sir Adrian Smith, Director and Chief Executive of the Alan Turing Institute In this myth-busting guide to AI past and present, one of the world’s leading researchers shows you our fears for the future are misplaced. The ultimate dream of AI is to build machines that are like us: conscious and self-aware. While this remains a remote possibility, rapid progress in AI is already transforming our world. Yet the public debate is still largely centred on unlikely prospects. From sentient machines to dystopian robot takeovers. In this lively and clear-headed guide, Michael Woodruff challenges the prevailing narrative, revealing how AI truly distracts us from both the more immediate risks that this technology poses - from algorithmic bias to fake news - and the more long-term implications of the field. The Road to Conscious Machines elucidates the discoveries of AI’s greatest pioneers from Alan Turing to Demis Hassabis, and what today’s researchers actually think and do. ‘Nobody understands the past, the present, the promise and the peril of this new technology better than Michael Woodruff. The definitive account’ Matt Ridley, author of The Rational Optimist ‘Effortlessly readable. The perfect guide to the history and future of AI’ Tom Chivers, author of The AI Does Not Hate You

It seems obvious that phenomenally conscious experience is something of great value, and that this value maps onto a range of important ethical issues. For example, claims about the value of life for those in Permanent Vegetative State (PVS); debates about treatment and study of disorders of consciousness; controversies about end-of-life care for those with advanced dementia; and arguments about the moral status of embryos, fetuses, and non-human animals arguably turn on the moral significance of various facts about consciousness. However, though work has been done on the moral significance of elements of consciousness, such as pain and pleasure, little explicit attention has been devoted to the ethical significance of consciousness. In this book Joshua Shepherd presents a systematic account of the value present within conscious experience. This account emphasizes not only the nature of consciousness, but also the importance of items within experience such as affect, valence, and the complex overall shape of particular valuable experiences. Shepherd also relates this account to difficult cases involving non-humans and humans with disorders of consciousness, arguing that the value of consciousness influences and partially explains the degree of moral status a being possesses, without fully determining it. The upshot is a deeper understanding of both the moral importance of phenomenal consciousness and its relations to moral status. This book will be of great interest to philosophers and students of ethics, bioethics, philosophy of psychology, philosophy of mind, and cognitive science. THIS BOOK is the fully revised and updated second edition of ‘Consciousness and Robot Sentience’. With lots of new material, it will provide new insights into artificial intelligence (AI) and machine consciousness, beyond materials published in the first edition. The organization of the book has been streamlined for better clarity and continuity of the lines of arguments. The perspective of AI has been added to this edition. It is shown that contemporary AI has a hidden problem, which prevents it from becoming a true intelligent agent. A self-evident solution to this problem is given in this book. This solution is surprisingly connected with the concepts of qualia, the mind-body problem and consciousness. These are the hard problems of consciousness that so far have been without viable solution. Unfortunately, the solution to the hidden problem of AI cannot be satisfactorily implemented, unless the phenomena of qualia and consciousness are first understood. In this book an explanation of consciousness is presented, one that rejects material and immaterial substances, dualism, pansychism, emergence and metaphysics. What remains is obvious. This explanation excludes consciousness in digital computers, but allows the artificial creation of consciousness in one natural-like way, by associative non-computational neural networks.The proof of a theory calls for empirical verification. In this case, the proof could be in the form of a sentient robot. This book describes a step towards this in the form of the author’s small experimental robot XCR-1. This robot has evolved through the years, and has now new cognitive abilities, which are described.

As we approach a great turning point in history when technology is poised to redefine what it means to be human, The Fourth Age offers fascinating insight into AI, robotics, and their extraordinary implications for our species. “If you only read just one book about the AI revolution, make it this one!” (John Mackey, cofounder and CEO, Whole Foods Market). In The Fourth Age, Byron Reese makes the case that technology has not only changed the way we live, it has changed the way we think. The Fourth Age is a bold, eye-opening look at the future of humanity. Reese traces the history of AI from its humble beginnings to the present, and he shows us how we can prepare for the future. The Fourth Age provides an essential background on how we got to this point, and how—rather than what—we should think about the topics we’ll soon all be facing: machine consciousness, automation, changes in employment, creative computers, radical life extension, artificial life, AI ethics, the future of warfare, superintelligence, and the implications of extreme prosperity. By asking questions like “Are you a machine?” and “Could a computer feel anything?”, Reese leads you through a discussion along the cutting edge in robotics and AI, and provides a framework by which we can all understand, discuss, and act on the issues of the Fourth Age and how they’ll transform humanity.

Consciousness is widely perceived as one of the most fundamental, interesting and difficult problems of our time. However, we still know next to nothing about the relationship between consciousness and the brain and we can only speculate about the consciousness of animals and machines.Human and Machine Consciousness presents a new foundation for the scientific study of consciousness. It sets out a bold interpretation of consciousness that neutralizes the philosophical problems and explains how we can make scientific predictions about the consciousness of animals, brain-damaged patients and machines.Gamez interprets the scientific study of consciousness as a search for mathematical theories that map between measurements of consciousness and measurements of the physical world. We can use artificial intelligence to discover these theories and they could make accurate predictions about the consciousness of humans, animals and artificial systems. Human and Machine Consciousness also provides original insights into unusual conscious experiences, such as hallucinations, religious experiences and out-of-body states, and demonstrates how ‘designer’ states of consciousness could be created in the future.Gamez explains difficult concepts in a clear way that closely engages with scientific research. His punchy, concise prose is packed with vivid examples, making it suitable for the educated general reader as well as philosophers and scientists. Problems are brought to life in colourful illustrations and a helpful summary is given at the end of each chapter. The endnotes provide detailed discussions of individual points and full references to the scientific and philosophical literature. This work was published by Saint Philip Street Press pursuant to a Creative Commons license permitting commercial use. All rights not granted by the work’s license are retained by the author or authors.

Machine Dreaming and Consciousness is the first book to discuss the questions raised by the advent of machine dreaming. Artificial intelligence (AI) systems meeting criteria of primary and self-reflexive consciousness are often utilized to extend the human internet, creating waking experiences that resemble the human dream. Surprisingly, AI systems also easily meet all human-based operational criteria for dreaming. These ‘dreams are far different from anthropomorphic dreaming, including such processes as fuzzy logic, liquid illogic, and integration instability, all processes that may be necessary in both biologic and artificial systems to extend creative capacity. Today, multi-linear AI systems are being built to resemble the structural framework of the human central nervous system. The creation of the biologic framework of dreaming (emotions, associative memories, and visual imagery) is well within our technical capacity. AI dreams potentially portend the further development of consciousness in these systems. This focus on AI dreaming raises even larger questions. In many ways, dreaming defines our humanity. What is humanly special about the states of dreaming? And what are we losing when we limit our focus to its technical and biologic structure, and extend the capacity for dreaming into our artificial creations? Machine Dreaming and Consciousness provides thorough discussion of these issues for neuroscientists and other researchers investigating consciousness and cognition. Addresses the function and role of dream-like processing in AI systems Describes the functions of dreaming in the creative process of both humans and machines Presents an alternative approach to the philosophy of machine consciousness Provides thorough discussion of machine dreaming and consciousness for neuroscientists and other researchers investigating consciousness and cognition

An argument that consciousness, more widespread than previously assumed, is the feeling of being alive, not a type of computation or a clever hack. In The Feeling of Life Itself, Christof Koch offers a straightforward definition of consciousness as any subjective experience, from the most mundane to the most exalted—the feeling of being alive. Psychologists study which cognitive operations underpin a given conscious perception. Neuroscientists track the neural correlates of consciousness in the brain, the organ of the mind. But why the brain and not, say, the liver? How can the brain, three pounds of highly excitable matter, a piece of furniture in the universe, subject to the same laws of physics as any other piece, give rise to subjective experience? Koch argues that what is needed to answer these questions is a quantitative theory that starts with experience and proceeds to the brain. In The Feeling of Life Itself, Koch outlines such a theory, based on integrated information. Koch describes how the theory explains many facts about the neurology of consciousness and how it has been used to build a clinically useful consciousness meter. The theory predicts that many, and perhaps all, animals experience the sights and sounds of life; consciousness is much more widespread than conventionally assumed. Contrary to received wisdom, however, Koch argues that programmable computers will not have consciousness. Even a perfect software model of the brain is not conscious. Its simulation is fake consciousness. Consciousness is not a special type of computation—it is not a clever hack. Consciousness is about being.

Your Brain Is a Time Machine: The Neuroscience and Physics of Time

Consciousness in Humanoid Robots

The Fourth Age

The Hidden Spring

From the Dawn of Computing to Digital Consciousness

Mind and Mechanism

The Feeling of Life Itself

Consciousness and Moral Status

The Story of AI

The Cognitive Approach to Conscious Machines

How to Grow a Robot

What Philosophy Is

An exploration of the mind-body problem from the perspective of artificialintelligence.

A bestselling author, neuroscientist, and computer engineer unveils a theory of intelligence that will revolutionize our understanding of the brain and the future of AI For all of neuroscience’s advances, we’ve made little progress on its biggest question: How do simple cells in the brain create intelligence? Jeff Hawkins and his team discovered that the brain uses maplike structures to build a model of the world—not just one model, but hundreds of thousands of models of everything we know. This discovery allows Hawkins to answer important questions about how we perceive the world, why we have a sense of self, and the origin of high-level thought. A Thousand Brains heralds a revolution in the understanding of intelligence. It is a big-think book, in every sense of the word.

A lively, engaging and provocative examination of the limits of philosophy and science through studies of perception, time, madness and knowledge. >

An authoritative, up-to-date survey of the state of the art in artificial intelligence, written for non-specialists.

“Beautifully written, eloquently reasoned...Mr. Buonomano takes us off and running on an edifying scientific journey.” —Carol Tavris, Wall Street Journal In Your Brain Is a Time Machine, leading neuroscientist Dean Buonomano embarks on an “immensely engaging” exploration of how time works inside the brain (Barbara Kiser, Nature). The human brain, he argues, is a complex system that not only tells time, but creates it; it constructs our sense of chronological movement and enables “mental time travel”—simulations of future and past events. These functions are essential not only to our daily lives but to the evolution of the human race: without the ability to anticipate the future, mankind would never have crafted tools or invented agriculture. This virtuous work of popular science will lead you to a revelation as strange as it is true: your brain is, at its core, a time machine.

“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.

How to develop robots that will be more like humans and less like computers, more social than machine-like, and more playful and less programmed. Most robots are not very friendly. They vacuum the rug, mow the lawn, dispose of bombs, even perform surgery—but they aren’t good conversationalists. It’s difficult to make eye contact. If the future promises more human-robot collaboration in both work and play, wouldn’t it be better if the robots were less mechanical and more social? In How to Grow a Robot, Mark Lee explores how robots can be more human-like, friendly, and engaging. Developments in artificial intelligence—notably Deep Learning—are widely seen as the foundation on which our robot future will be built. These advances have already brought us self-driving cars and chess match-winning algorithms. But, Lee writes, we need robots that are perceptive, animated, and responsive—more like humans and less like computers, more social than machine-like, and more playful and less programmed. The way to achieve this, he argues, is to “grow” a robot so that it learns from experience—just as infants do. After describing “what’s wrong with artificial intelligence” (one key shortcoming: it’s not embodied), Lee presents a different approach to building human-like robots: developmental robotics, inspired by developmental psychology and its accounts of early infant behavior. He describes his own experiments with the iCub humanoid robot and its development from newborn helplessness to ability levels equal to a nine-month-old, explaining how the iCub learns from its own experiences. AI robots are designed to know humans as objects; developmental robots will learn empathy. Developmental robots, with an internal model of “self,” will be better interactive partners with humans. That is the kind of future technology we should work toward.

Consciousness is widely perceived as one of the most fundamental, interesting and difficult problems of our time. However, we still know next to nothing about the relationship between consciousness and the brain and we can only speculate about the consciousness of animals and machines. Human and Machine Consciousness presents a new foundation for the scientific study of consciousness. It sets out a bold interpretation of consciousness that neutralizes the philosophical problems and explains how we can make scientific predictions about the consciousness of animals, brain-damaged patients and machines. Gamez interprets the scientific study of consciousness as a search for mathematical theories that map between measurements of consciousness and measurements of the physical world. We can use artificial intelligence to discover these theories and they could make accurate predictions about the consciousness of humans, animals and artificial systems. Human and Machine Consciousness also provides original insights into unusual conscious experiences, such as hallucinations, religious experiences and out-of-body states, and demonstrates how ‘designer’ states of consciousness could be created in the future. Gamez explains difficult concepts in a clear way that closely engages with scientific research. His punchy, concise prose is packed with vivid examples, making it suitable for the educated general reader as well as philosophers and scientists. Problems are brought to life in colourful illustrations and a helpful summary is given at the end of each chapter. The endnotes provide detailed discussions of individual points and full references to the scientific and philosophical literature.

Life 3.0

The truth about AI from the people building it

Consciousness Beyond the Brain

AI and the Future of Your Mind

Why Consciousness Is Widespread But Can’t Be Computed

The Cambridge Handbook of Artificial Intelligence

The Emperor’s New Mind

Architects of Intelligence

Blindspots in Philosophy and Science

Machine Consciousness

AI and the Problem of Control

Consciousness and Robot Sentience

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, singlets, 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.

A survey of the exciting technologies that will bring us closer to dialogue with neuroscience in order to challenge the whiz-bang computational views of the human and machine sentience that dominate the headlines. Giving robots LSD sounds like a joke, but Smart is dead serious in his critique of the hidden and sometimes dangerous biases that underlie both popular and scientific fantasies of digital minds. “ —Erik Davis, host of “Expanding Mind” and author, Technosis: Myth, Magic, and Mysticism in the Age of Information “ Philosophy, psychedelics, robots, and the future, consciousness and intelligence, what else do you desire? Here you will see why those machines that reach singularity will be smarter than us and take over the world—and shall need to be conscious... and maybe they can only be conscious if they are human enough. The thesis of the book, and the path shown us by Smart, leads to a great trip, of imagination and philosophy, of maths and neuroscience. ” —Dr. Tristan Bekinschtein, Lecturer, Department of Psychology of Cambridge Can we build a robot that trips on acid? This is not a frivolous question, according to neuroscientist Andrew Smart. If we can ’ t, he argues, we haven ’ t really created artificial intelligence. In an exposition reminiscent of crossover works such as Gödel, Escher, Bach and Fermat “ s Last Theorem, Andrew Smart weaves together Mangaravan binary numbers, the discovery of LSD, Leibniz, computer programming, and much more to connect the vast but largely forgotten world of psychedelic research with the resurgent field of AI and the attempt to build conscious robots. A book that draws on the history of mathematics, philosophy, and digital technology, Beyond Zero and One challenges fundamental assumptions underlying artificial intelligence. Is the human brain based on computation? Can information alone explain human consciousness and intelligence? Smart convincingly makes the case that true intelligence, and artificial intelligence, requires an appreciation of what is beyond the computational.

Originating from a Dagstuhl seminar, the collection of papers presented in this book constitutes on the one hand a representative state-of-the-art survey of embodied artificial intelligence, and on the other hand the papers identify the important research trends and directions in the field. Following an introductory overview, the 23 papers are organized into topical sections on - philosophical and conceptual issues - information dynamics, and morphology - principles of real-world applications - developmental approaches - artificial evolution and self-reconfiguration

This text addresses the impact of new technology on our ideas about art, science, philosophy and what it is to be human. It argues that many of our beliefs are no longer useful or relevant and we must develop new ways of thinking about and understanding the complexity of contemporary existence.

When Computers Exceed Human Intelligence

What We Can Never Know

In the Shadow of the Machine

A Journey to the Source of Consciousness

A Thousand Brains

Embodied Artificial Intelligence

A Search for the Missing Science of Consciousness

Being Human in the Age of Artificial Intelligence

The Universal Machine

Superintelligence

Consciousness And Robot Sentience (Second Edition)

The Road to Conscious Machines

In this age of DNA computers and artificial intelligence, information is becoming disembodied even as the “bodies” that once carried it vanish into virtuality. While some marvel at these changes, envisioning consciousness downloaded into a computer or humans “beamed” Star Trek-style, others view them with horror, seeing monsters brooding in the machines. In How We Became Posthuman, N. Katherine Hayles separates hype from fact, investigating the fate of embodiment in an information age. Hayles relates three interwoven stories: how information lost its body, that is, how it came to be conceptualized as an entity separate from the material forms that carry it; the cultural and technological construction of the cyborg; and the dismantling of the liberal humanist “subject” in cybernetic discourse, along with the emergence of the “posthuman.” Ranging widely across the history of technology, cultural studies, and literary criticism, Hayles shows what had to be erased, forgotten, and elided to conceive of information as a disembodied entity. Thus she moves from the post-World War II Macy Conferences on cybernetics to the 1952 novel Limbo by cybernetics aficionado Bernard Wolfe; from the concept of self-making to Philip K. Dick’s literary explorations of hallucination and reality; and from artificial life to postmodern novels exploring the implications of seeing humans as cybernetic systems. Although becoming posthuman can be nightmarish, Hayles shows how it can also be liberating. From the birth of cybernetics to artificial life, How We Became Posthuman provides an indispensable account of how we arrived in our virtual age, and of where we might go from here. A survey of AI’s theories of the most innovative and compelling technology of our era, an international authority on artificial intelligence, and 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-enhanced technology, bringing about the demise or suffering of conscious beings. Examining the philosophical questions lying beneath the algorithms, Schneider takes on AI’s thorniest implications.

A FINANCIAL TIMES BEST BOOK OF THE YEAR: “The most important book I have read in quite some time” Daniel Kahneman; “A must-read” Max Tegmark; “The book we’ve all been waiting for” Sam Harris Humans dream of super-intelligent machines. But what happens if we actually succeed? Creating superior intelligence would be the biggest event in human history. Unfortunately, according to the world’s pre-eminent AI expert, it could also be the last. In this groundbreaking book, Stuart Russell sets out why he has come to consider his own discipline an existential threat to humanity, and how we can change course before it’s too late. In brilliant and lucid prose, he explains how AI actually works and its enormous capacity to improve our lives - and why we must never lose control of machines more powerful than we are. Russell contends that we can avert the worst threats by reshaping the foundations of AI to guarantee that machines pursue our objectives, not theirs. Profound, urgent and visionary, Human Compatible is the one book everyone needs to read to understand a future that is coming sooner than we think. LONGLISTED FOR THE FINANCIAL TIMES & MCKINSEY BUSINESS BOOK OF THE YEAR Thought-provoking Financial Times ‘Fascinating and significant’ Sunday Times ‘The most important book on AI this year’ Guardian Winner of the Wolf Prize for his contribution to our understanding of the universe, Penrose takes on the question of whether artificial intelligence will ever approach the intricacy of the human mind. 144 illustrations.

Contemporary life is so deeply reliant upon digital technology that the computer has come to dominate almost every aspect of our culture. What is the philosophical and spiritual significance of this dependence on electronic technology, both for our relationship to nature and for the future of humanity? And, what processes in human perception and awareness have produced the situation we find ourselves in? As Jeremy Naydler elucidates in this penetrating study, we cannot understand the emergence of the computer without seeing within the wider context of the evolution of human consciousness, which has taken place over millennia. Modern consciousness, he shows, has evolved in conjunction with the development of machines and under their intensifying shadow. The computer was the product of a long historical development, culminating in the scientific revolution of the 17th century. It was during this period that the first mechanical calculators were invented and the project to create more complex ‘thinking machines’ began in earnest. But the seeds were sown many hundreds of years earlier, deep in antiquity. Naydler paints a vast panorama depicting human development and the emergence of electronic technology. His painstaking research illuminates an urgent question that concerns every living person today: What does it mean to be human and what, if anything, distinguishes us from machines?

This new book by the award-winning scientist and author of The Emperor’s New Mind is a profound exploration of what modern physics has to tell us about the mind. It also gives a visionary description of what a new physics might look like.

This book addresses the question What is Philosophy? by gathering together responses from philosophers working in a variety of areas. The resulting collection provides focused discussions of the character and methods of philosophy and its relationship with other disciplines.

Humans are extraordinary creatures, with the unique ability among animals to imitate and so copy from one another ideas, habits, skills, behaviours, inventions, songs, and stories. These are all memes, a term first coined by Richard Dawkins in 1976 in his book The Selfish Gene. Memes, like genes, are replicators, and this enthralling book is an investigation of whether this link between genes and memes can lead to important discoveries about the nature of the inner self. Confronting the deepest questions about our inner selves, with all our emotions, memories, beliefs, and decisions, Susan Blackmore makes a compelling case for the theory that the inner self is merely an illusion created by the memes for the sake of replication.

The Posthuman Condition

Beyond Zero and One

When Humans Transcend Biology

Conscious Robots

How We Became Posthuman

The Meme Machine

Human Compatible

The Age of Spiritual Machines

Developing Human-Friendly, Social AI

Human and Machine Consciousness

A New Theory of Intelligence

Concerning Computers, Minds, and the Laws of Physics

The author argues that true conscious machines can be built, but rejects artificial intelligence and classical neural networks in favour of the emulation of the cognitive processes of the brain. Novel views on consciousness and the mind-body problem are presented. This books is a must for anyone interested in consciousness research and the latest ideas in the forthcoming technology of mind.

Robots are becoming more human, but could they also become sentient and have human-like consciousness? What is consciousness, exactly? It is a fact that our thoughts and consciousness are based on the neural activity of the brain. It is also a fact that we do not perceive our brain activity as it really is — patterns of neural firings. Instead, we perceive our sensations and thoughts apparently as they are. What kind of condition would transform the neural activity into this kind of internal appearance? This is the basic problem of consciousness. The author proposes an explanation that also provides preconditions for true conscious cognition — the requirement of a direct perceptive system with inherent sub-symbolic and symbolic information processing. Associative neural information processing with distributed signal representations is introduced as a method that satisfies these requirements. Conscious robot cognition also calls for information integration and sensorimotor integration. This requirement is satisfied by the Haikonen Cognitive Architecture (HCA). This book demystifies both the enigmatic philosophical issues of consciousness and the practical engineering issues of conscious robots by presenting them in an easy-to-understand manner for the benefit of students, researchers, philosophers and engineers in the field. Contents: IntroductionThe Problem of ConsciousnessConsciousness and Subjective ExperiencePerception and QualiaFrom Perception to ConsciousnessEmotions and ConsciousnessInner Speech and ConsciousnessQualia and Machine ConsciousnessTesting ConsciousnessArtificial Conscious CognitionAssociative Information ProcessingNeural Realization of Associative ProcessingDesigning a Cognitive Perception SystemExamples of Perception/Response Feedback LoopsThe Transition to Symbolic ProcessingInformation Integration with Multiple ModulesEmotional Significance of PerceptsThe Outline of the Haikonen Cognitive Architecture (HCA)Mind Reading ApplicationsThe Comparison of Some Cognitive ArchitecturesExample: An Experimental Robot with the HCAConcluding Notes Readership: Enthusiasts in cognitive robot research (including not only experts but also hobbyists), as well as university students, researchers and engineers on robots and/or cognitive machines. Keywords:(Artificial) Intelligence; Consciousness; Machine Consciousness; RobotsKey Features:Explains consciousness and delves into the treatment of qualia in philosophy and practical conscious robotsPresents an architecture for conscious robotsPresents a simple cognitive robot as an illustrative example

An authority on artificial intelligence introduces a theory that explores the workings of the human mind and the mysteries of thought

Building a conscious robot is a scientific and technological challenge. Debates about the possibility of conscious robots and the related positive outcomes and hazards for human beings are today no longer confined to philosophical circles. Robot consciousness is a research field aimed at a two-part goal: on the one hand, scholars working in robot consciousness take inspiration from biological consciousness to build robots that present forms of experiential and functional consciousness. On the other hand, scholars employ robots as tools to better understand biological consciousness. Thus, part one of the goal concerns the replication of aspects of biological consciousness in robots, by uniting a variety of approaches from AI and robotics, cognitive robotics, epigenetic and affective robotics, situated and embodied robotics, developmental robotics, anticipatory systems, and biomimetic robotics. Part two of the goal is pursued by employing robots to advance and mark progress in the study of consciousness in humans and animals. Notably, neuroscientists involved in the study of consciousness do not exclude the possibility that robots may be conscious. This eBook comprises a collection of thirteen manuscripts and an Editorial published by Frontiers in Robotics and Artificial Intelligence, under the section Humanoid Robotics, and Frontiers in Neurorobotics, on the topic “Consciousness in Humanoid Robots.” This compendium aims at collating the most recent theoretical studies, models, and case studies of machine consciousness that take the humanoid robot as a frame of reference. The content in the articles may be applied to many different kinds of robots, and to software agents as well.

Consciousness is widely perceived as one of the most fundamental, interesting and difficult problems of our time. However, we still know next to nothing about the relationship between consciousness and the brain and we can only speculate about the consciousness of animals and machines. Human and Machine Consciousness presents a new foundation for the scientific study of consciousness. It sets out a bold interpretation of consciousness that neutralizes the philosophical problems and explains how we can make scientific predictions about the consciousness of animals, brain-damaged patients and machines. Gamez interprets the scientific study of consciousness as a search for mathematical theories that map between measurements of consciousness and measurements of the physical world. We can use artificial intelligence to discover these theories and they could make accurate predictions about the consciousness of humans, animals and artificial systems. Human and Machine Consciousness also provides original insights into unusual conscious experiences, such as hallucinations, religious experiences and out-of-body states, and demonstrates how ‘designer’ states of consciousness could be created in the future. Gamez explains difficult concepts in a clear way that closely engages with scientific research. His punchy, concise prose is packed with vivid examples, making it suitable for the educated general reader as well as philosophers and scientists. Problems are brought to life in colourful illustrations and a helpful summary is given at the end of each chapter. The endnotes provide detailed discussions of individual points and full references to the scientific and philosophical literature. This is the most important conversation of our time, and Tegmark’s thought-provoking book will help you join it! Stephen Hawking THE INTERNATIONAL BESTSELLER. DAILY TELEGRAPH AND THE TIMES BOOKS OF THE YEAR SELECTED AS ONE OF BARACK OBAMA’S FAVOURITE BOOKS OF 2018 AI is the future - but what will that future look like? Will superhuman intelligence be our slave, or become our god? Taking us to the heart of the latest thinking about AI, Max Tegmark, the MIT professor whose work has helped mainstream research on how to keep AI beneficial, separates myths from reality, utopias from dystopias, to explore the next phase of our existence. How can we grow our prosperity through automation, without leaving people lacking income or purpose? How can we ensure that future AI systems do what we want without crashing, malfunctioning or getting hacked? Should we fear an arms race in lethal autonomous weapons? Will AI help life flourish as never before, or will machines eventually outsmart us at all tasks, and even, perhaps, replace us altogether? ‘This is a rich and visionary book and everyone should read it’ The Times

The era-defining book that will forever change the way you understand your mind. ‘Required reading’ - Susie Orbach ‘Truly pioneering’ - Eric Kandel ‘It changes everything’ - Brian Eno How does the mind connect to the body? Why does it feel like something to be us? For one of the boldest thinkers in neuroscience, solving this puzzle has been a lifetime’s quest. Now at last, the man who discovered the brain mechanism for dreaming appears to have made a breakthrough. The very idea that a solution is at hand may seem outrageous. Isn’t consciousness intangible, beyond the reach of science? Yet Mark Solms shows how misguided fears and suppositions have concealed its true nature. Stick to the medical facts, pay close attention to the eerie testimony of hundreds of neurosurgery patients, and a way past our obstacles reveals itself. Join Solms on a voyage into the extraordinary realms beyond. More than just a philosophical argument, The Hidden Spring will forever alter how you understand your own experience. There is a secret buried in the brain’s

“Lives up to all the hype!” “An absolutely necessary book.” “Should be taught in schools.” “Dynamite, this is a brilliant book.” - see the reviews on Amazon.com. The Future’s Most Important Book -- Why we’re so convinced that we’re in charge when we’re really just carrying out evolution’s instructions -- Why our lives, as Buddha suggested, are inherently unsatisfactory, despite our luxurious homes, successful careers and loving families -- How humans will one day take control of their conscious minds, get happy and stay happy. And the real reason Aliens haven’t visited the Earth yet... 107 minutes (average read time) to change the way you think about everything. “Easy to understand and persuasive.” “Fun, short, insightful!” “Bad Ass ”

International Seminar, Dagstuhl Castle, Germany, July 7-11, 2003, Revised Selected Papers

Artificial You

Artificial Consciousness

Virtual Bodies in Cybernetics, Literature, and Informatics

Smart Robots, Conscious Computers, and the Future of Humanity

The Singularity Is Near

Machine Dreaming and Consciousness

Paths, Dangers, Strategies

Machines, Psychedelics, and Consciousness

Shadows of the Mind

Society Of Mind