


The book was found

# Life 3.0: Being Human In The Age Of Artificial Intelligence



## Synopsis

How will Artificial Intelligence affect crime, war, justice, jobs, society and our very sense of being human? The rise of AI has the potential to transform our future more than any other technology—and there's nobody better qualified or situated to explore that future than Max Tegmark, an MIT professor who's helped mainstream research on how to keep AI beneficial. How can we grow our prosperity through automation without leaving people lacking income or purpose? What career advice should we give today's kids? How can we make future AI systems more robust, so that they do what we want without crashing, malfunctioning or getting hacked? Should we fear an arms race in lethal autonomous weapons? Will machines eventually outsmart us at all tasks, replacing humans on the job market and perhaps altogether? Will AI help life flourish like never before or give us more power than we can handle? What sort of future do you want? This book empowers you to join what may be the most important conversation of our time. It doesn't shy away from the full range of viewpoints or from the most controversial issues—from superintelligence to meaning, consciousness and the ultimate physical limits on life in the cosmos.

## Book Information

Hardcover: 384 pages

Publisher: Knopf (August 29, 2017)

Language: English

ISBN-10: 1101946598

ISBN-13: 978-1101946596

Product Dimensions: 6.5 x 1.3 x 9.7 inches

Shipping Weight: 1.6 pounds (View shipping rates and policies)

Average Customer Review: 4.4 out of 5 stars 17 customer reviews

Best Sellers Rank: #84 in Books (See Top 100 in Books) #1 in Books > Computers & Technology > Computer Science > AI & Machine Learning #1 in Books > Textbooks > Business & Finance #2 in Books > Computers & Technology > Computer Science > Robotics

## Customer Reviews

"Original, accessible, and provocative. Tegmark successfully gives clarity to the many faces of AI, creating a highly readable book that complements *The Second Machine Age's* economic perspective on the near-term implications of recent accomplishments in AI and the more detailed analysis of how we might get from where we are today to AGI and even the

superhuman AI in *Life 3.0: Superintelligence*. At one point, Tegmark quotes Emerson: "Life is a journey, not a destination." The same may be said of the book itself. Enjoy the ride, and you will come out the other end with a greater appreciation of where people might take technology and themselves in the years ahead. "Science" In [Tegmark's] magnificent brain, each fact or idea appears to slip neatly into its appointed place like another little silver globe in an orrery the size of the universe. There are spaces for Kant, Cold War history and Dostoyevsky, for the behaviour of subatomic particles and the neuroscience of consciousness.... Tegmark describes the present, near-future and distant possibilities of AI through a series of highly original thought experiments.... Tegmark is not personally wedded to any of these ideas. He asks only that his readers make up their own minds. In the meantime, he has forged a remarkable consensus on the need for AI researchers to work on the mind-bogglingly complex task of building digital chains that are strong and durable enough to hold a superintelligent machine to our bidding.... This is a rich and visionary book and everyone should read it. "The Times (UK)" "This is a compelling guide to the challenges and choices in our quest for a great future of life, intelligence and consciousness on Earth and beyond." "Elon Musk, Founder, CEO and CTO of SpaceX and co-founder and CEO of Tesla Motors" "All of us not only scientists, industrialists and generals should ask ourselves what can we do now to improve the chances of reaping the benefits of future AI and avoiding the risks. This is the most important conversation of our time, and Tegmark's thought-provoking book will help you join it." "Professor Stephen Hawking, Director of Research, Cambridge Centre for Theoretical Cosmology" "Tegmark's new book is a deeply thoughtful guide to the most important conversation of our time, about how to create a benevolent future civilization as we merge our biological thinking with an even greater intelligence of our own creation." "Ray Kurzweil, Inventor, Author and Futurist, author of *The Singularity is Near* and *How to Create a Mind*" "Being an eminent physicist and the leader of the Future of Life Institute has given Max Tegmark a unique vantage point from which to give the reader an inside scoop on the most important issue of our time, in a way that is approachable without being dumbed down." "Jaan Tallinn, co-founder of Skype" "This is an exhilarating book that will change the way we think about AI, intelligence, and the future of humanity." "Bart Selman, Professor of Computer Science, Cornell University" "The unprecedented power unleashed by artificial intelligence means the next decade could be humanity's best or worst."

Max Tegmark has written the most insightful and just plain fun exploration of AI's implications that I've ever read. If you haven't been exposed to Tegmark's joyful mind yet, you're in for a huge treat.

Professor Erik Brynjolfsson, Director of the MIT Initiative on the Digital Economy and co-author of *The Second Machine Age* "Tegmark seeks to facilitate a much wider conversation about what kind of future we, as a species, would want to create. Though the topics he covers—AI, cosmology, values, even the nature of conscious experience—can be fairly challenging, he presents them in an unintimidating manner that invites the reader to form her own opinions."

Nick Bostrom, Founder of Oxford's Future of Humanity Institute, author of *Superintelligence* "I was riveted by this book. The transformational consequences of AI may soon be upon us—but will they be utopian or catastrophic? The jury is out, but this enlightening, lively and accessible book by a distinguished scientist helps us to assess the odds."

Professor Martin Rees, Astronomer Royal, cosmology pioneer, author of *Our Final Hour* "Exhilarating! MIT physicist Tegmark surveys advances in artificial intelligence such as self-driving cars and Jeopardy-winning software, but focuses on the looming prospect of "recursive self-improvement"—AI systems that build smarter versions of themselves at an accelerating pace until their intellects surpass ours. Tegmark's smart, freewheeling discussion leads to fascinating speculations on AI-based civilizations spanning galaxies and eons. Engrossing." Publishers Weekly

MAX TEGMARK is a professor of physics at MIT and the co-founder of the Future of Life Institute. Tegmark has been featured in dozens of science documentaries. His passion for ideas, adventure, and entrepreneurship is infectious.

The first chapter of Tegmark's new book is called "Welcome to the most important conversation of our time, and that's exactly what this book is. Before diving into the book, a few words about why this conversation is so important and why Tegmark is a central agent helping make it happen and, through the book, the perfect guide. Have you notice how you don't solve CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) anymore? That's because computers now can. Artificial Intelligence, from being a fairly niche area of mostly academic study a decade ago has exploded in the last five years. Much more quickly than many anticipated, machine learning (a subset of AI) systems have defeated the

best human Go players, are piloting self-driving cars, usefully if imperfectly translating documents, labeling your photos, understanding your speech, and so on. This has led to huge investment in AI by companies and governments, with every sign that progress will continue. This book is about what happens if and when it does. But why hear about it from Tegmark, an accomplished MIT physicist and cosmologist, rather than (say) an AI researcher? First, Tegmark has over the past few years \*become\* an AI researcher, with 5 published technical papers in the past two years. But he's also got a lifetime of experience thinking carefully, rigorously, generally (and entertainingly to boot) about the "big picture" of what is possible, and what is not, over long timescales and cosmic distances (see his last book!) which most AI researchers do not. Finally, he's played an active and very key role (as you can read about in the book's epilogue) in actually creating conversation and research about the impacts and safety of AI in the long-term. I don't think anyone is more comprehensively aware of the full spectrum of important aspects of the issue. So now the book.

Chapter 1 lays out why AI is suddenly on everyone's radar, and very likely to be extremely important over the coming decades, situating present-day as a crucial point within the wider sweep of human and evolutionary history on Earth. Chapter 2 takes the question of "what is intelligence?" and abstracts it from its customary human application, to "what is intelligence \*in general\*?" How can we define it in a useful way to cover both biological and artificial forms, and how do these tie to a basic understanding of the physical world? This lays the groundwork for the question of what happens as artificial intelligences grow ever more powerful. Chapter 3 addresses this question in the near future: what happens as more and more human jobs can be done by AIs? What about AI weapons replacing human-directed ones? How will we cope when more and more decisions are made by AIs which may be flawed or biased? This is about a lot of important changes occurring \*right now\* to which society is, for the most part, asleep at the wheel. Chapter 4 gets into what is exciting and terrifying about AI: as a designed intelligence, it can in principle \*re\*design itself to get better and better, potentially on a relatively short timescale. This raises a lot of rich, important, and extremely difficult questions that not that many people have thought through carefully (another in-print example is the excellent book by Bostrom). Chapter 5 discusses where what happens to humans as a species after an "intelligence explosion" takes place. Here Tegmark is making a call to start thinking about where we want to be, as we may end up somewhere sooner than we think, and some of the possibilities are pretty awful. Chapter 6 exhibits Tegmark's unique talent for tackling the big

questions, looking at the \*ultimate\* limits and promise of intelligent life in the universe, and how stupefyingly high the stakes might be for getting the next few decades right. It's both a sobering and an exhilarating prospect. Chapters 7 and 8 then dig into some of the deep and interesting questions about AI: what does it mean for a machine to have "goals"? What are our goals as individuals and a society, and how can we best aim toward them in the long term? Can a machine we design have consciousness? What is the long-term future of consciousness? Is there a danger of relapsing into a universe \*without\* consciousness if we aren't careful? Finally, an epilogue describes Tegmark's own experience which I've had the privilege to personally witness as a key player in an effort to focus thought and effort on AI and its long-term implications, of which writing this book is a part. (And I should also mention the prologue, which gives a fictional but less \*science\*fictional depiction of an artificial superintelligence being used by a small group to seize control of human society. The book is written in a very lively and engaging style. The explanations are clear, and Tegmark develops a lot of material at a level that is understandable to a general audience, but rigorous enough to give readers a real understanding of the issues relevant to thinking about the future impact of AI. There are a lot of new ideas in the book, and although it is sometimes written in a breezy and engaging style, that belies a lot of careful thinking about the issues. It's possible that real, general artificial intelligence (AGI) is 100 or more years away, a problem for the next generation, with large but manageable effects of "narrow" AI to deal with over a span of decades. But it's also quite possible that it's going to happen 10, 15, 20, or 30 years from now, in which case society is going to have to make a lot of very wise and very important (literally of cosmic import) decisions very quickly. It's important to start the conversation now, and there's no better way.

Tegmark's new book is a really excellent introduction for anyone who wants to understand why AI researchers are both excited and worried about the rise of artificial intelligence. Most of the other books I've read about AI are either more technical, talk about the impact of AI over the short-term, or are already outdated. AI is developing at a rapid pace right now, and it's hard for most people to wrap their heads around how dramatically this could impact society. In fact, many people I talk to don't even know what AI stands for. Tegmark explains AI in easily accessible language with fun personal stories. He talks about what breakthroughs have occurred recently that are leading to such rapid development, as well as

how we can prepare for jobs in the coming decades. But then he goes on to look at how AI could impact humanity in 1,000 years, in 10,000 years, and even further into the future. And then there are the chapters on space exploration and consciousness, which Tegmark obviously had a lot of fun writing. Throughout the book, he asks two important questions: How do you want AI to impact your life? How do you think AI should impact future society? These are questions we all need to consider if we want to make sure AI helps the many and not the few, and Tegmark's book is a great jumping off point.

Tegmark brings a refreshing perspective to what likely is one of, if not the, most important conversation of our time. After setting the stage and clearing the field of common myths and misconceptions regarding AI, Tegmark methodically moves through the emergence of intelligence in our cosmos some 4 billion years ago to the implications of what he calls "Life 3.0," entities which can both redesign their hardware and software. Couching the emergence and development of both consciousness and intelligence within the more cosmological world view of a physicist offers a truly inspiring narrative. Moving through hundreds, thousands, and billions of years with Tegmark solidifies the development of intelligence as a cosmological phenomenon which you come to realize that you yourself are a part of. Learning how intelligence is an emergent expression of more basic physical laws feels, at least to me, thoroughly grounding and deeply reestablishing of a real connection between me and the world. Chapters 6, 7, and 8 are my favorite and cover territory seldom explored in similar literature. In 6, he explores the capacity of artificial superintelligence to colonize the universe and the implications for the about 10 billion galaxies he estimates might be able to be colonized by it. Chapter 7 explores the emergence and evolution of goals at different levels of reality, ranging from thermodynamics to wet and squishy intelligences like you and me. In chapter 8, Tegmark explores consciousness and the mysteries and questions surrounding it, a topic I believe deserves far more attention than it is currently getting. This book is aimed at a wider audience than Bostrom's Superintelligence, but even if you are an avid reader of all that concerns technology, the deep future, and AI, I'm certain you'll find novel content and an enjoyable recontextualization of AI from the perspective a physicist. On a less serious note, I really love that the author summarizes the most important points at the end of each chapter. It helps me remember everything. :p

A book to read and reread, to discuss, to consider. The author invites his reader to develop an informed opinion about the role of AI in the future. Well presented by an author who is not just

concerned about the consciousness of artificial intelligence, but humanity's consciousness of how its world is changing. Challenging, informative and hopeful, the book enables the reader to boldly go where he or she may never have gone before.

[Download to continue reading...](#)

Readings in Medical Artificial Intelligence. The First Decade (Addison-Wesley Series in Artificial Intelligence) Life 3.0: Being Human in the Age of Artificial Intelligence The Most Human Human: What Artificial Intelligence Teaches Us About Being Alive Emotional Intelligence: Why You're Smarter But They Are More Successful (Emotional intelligence leadership, Emotional Quotient, emotional intelligence depression, emotional intelligence workbook) Emotional Intelligence: 3 Manuscripts - Emotional Intelligence Definitive Guide, Mastery, Complete Step by Step Guide (Social Engineering, Leadership, ... (Emotional Intelligence Series Book 4) Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence Forbidden Gates: How Genetics, Robotics, Artificial Intelligence, Synthetic Biology, Nanotechnology, & Human Enhancement Herald The Dawn Of Techno-Dimensional Spiritual Warfare Human Computation (Synthesis Lectures on Artificial Intelligence and Machine Learning) Our Final Invention: Artificial Intelligence and the End of the Human Era Artificial Intelligence for Games Robots and Artificial Intelligence (Technology Behind) The Fourth Transformation: How Augmented Reality & Artificial Intelligence Will Change Everything Artificial Intelligence: What Everyone Needs to Know Economy Monitor Guide to Smart Contracts: Blockchain Examples (Artificial Intelligence, Law and Finance) Applications of Artificial Intelligence for Decision-Making: Multi-Strategy Reasoning Under Uncertainty Essentials of Game Theory: A Concise, Multidisciplinary Introduction (Synthesis Lectures on Artificial Intelligence and Machine Learning) Artificial Intelligence in Label-free Microscopy: Biological Cell Classification by Time Stretch Artificial Intelligence In Medicine (Aaas Selected Symposium) The Technological Elixir: Invoking Artificial Intelligence & Understanding the Arcana of Mysticism Singularity Now! The Artificial Intelligence Timeline: Automation, robotics, tech titans, and their brave new world (September 2017)

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)