


The book was found

Our Final Invention: Artificial Intelligence And The End Of The Human Era



Synopsis

A Huffington Post Definitive Tech Book of 2013 In as little as a decade, artificial intelligence could match and then surpass human intelligence. Corporations and government agencies around the world are pouring billions into achieving AI's Holy Grail—human-level intelligence. Once AI has attained it, scientists argue, it will have survival drives much like our own. We may be forced to compete with a rival more cunning, more powerful, and more alien than we can imagine. Through profiles of tech visionaries, industry watchdogs, and groundbreaking AI systems, James Barrat's *Our Final Invention* explores the perils of the heedless pursuit of advanced AI. Until now, human intelligence has had no rival. Can we coexist with beings whose intelligence dwarfs our own? And will they allow us to?

Book Information

Paperback: 336 pages

Publisher: St. Martin's Griffin; Reprint edition (February 17, 2015)

Language: English

ISBN-10: 1250058783

ISBN-13: 978-1250058782

Product Dimensions: 5.5 x 0.9 x 8.2 inches

Shipping Weight: 12 ounces (View shipping rates and policies)

Average Customer Review: 4.1 out of 5 stars— See all reviews— (267 customer reviews)

Best Sellers Rank: #36,227 in Books (See Top 100 in Books) #12 in Books > Computers & Technology > Computer Science > Human-Computer Interaction #31 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics

Customer Reviews

"Our Final Invention" is a fascinating and well-written look at the risks posed by artificial super-intelligence. As other reviewers have pointed out, this book offers a relatively pessimistic take on the subject, but there is a lot of value in that perspective. There are plenty of other books, by Ray Kurzweil and others, that offer the optimistic viewpoint. The danger highlighted by the book is that an intelligent machine would turn its energies toward building even better versions of itself—creating an accelerating feedback loop that could culminate in a machine THOUSANDS of times more intelligent than any human. Once such an intelligence "escaped from its box" there would be no way to protect ourselves. This book focuses entirely on the long term risk of super-intelligence and does not touch at all on the near term consequences of less advanced and more specialized AI. For

example, millions of routine jobs will be lost and the economy will be transformed, and this could happen quite soon. In the longer run, the points raised in *Our Final Invention* are well worth thinking about. Some experts feel that an advanced AI would be controlled by programming in "friendliness" right from the start. Just as humans have basic drives (food, shelter, sex, etc.) a machine might be programmed to have an essential need to help humanity. As the author points out, however, in humans these basic drives often produce unpleasant and unexpected consequences -- like for example suicide bombers. A truly advanced, alien intelligence might exhibit some qualities that are not unlike mental illness in humans. A machine might by nature be a sociopath.

If you don't know much about real-world AI research and/or you're totally unfamiliar with the nonfiction concerns about the risks it poses, then this book is a quick and easy read that will make you aware of the basics. However, the author is himself clearly non-technical and has a sensationalist style that feels too much like tabloid writing. When I started reading it, I began bookmarking pages with passages that struck me as problematic. I thought I might write a short review on my wife's tech blog, or perhaps for LessWrong. But as I read further, I realized there were so many problem areas that I'd never bother to sit down and address them individually. Again, these problems would only matter to a technical audience -- experienced programmers, people with a more-than-passing-interest in AI, and so on. This is my big problem with the book: It's a critically important subject which deserves better treatment than this. Barrat seems to understand the basic problem well enough, but much of the time I had the feeling his primary goal was hitting a page-count target. For example, most of the section about malware is largely irrelevant to the real problem, but it felt like one of the longer chapters in the book (I didn't bother to confirm this, that's just my impression). His TV documentary background shows at the start of each paragraph: each time I felt like I was coming back from a commercial break. He'll shoot somebody down in one chapter, then use that same person to support his argument in the next. He tosses around concepts like cognitive bias and logical fallacies apparently without realizing the book is mostly one big appeal to authority. There is a very good, very important story here waiting to be told. This book only scratches the surface.

Just done the new-ish book *Our Final Invention: Artificial Intelligence and the End of the Human Era* by James Barrat. It explains the inevitability of super-intelligent machines evolving to the point of wiping out all biological life in the galaxy - with opening day coming soon to a species near you (yours). First off I have to say this is a very enjoyable read. This guy has the kind of snappy, crisp,

slightly sarcastic, slightly smartass style that I enjoy. He has some sense of humor. (That's a human trait right there which I bet our smarty-pants AI Overlords won't be able to replicate convincingly.) So it's fun. And though as somebody with a doctorate from MIT earned through cross-disciplinary work in Theoretical Linguistics, Computational Linguistics at the MIT AI Lab, and speech modeling at the MIT Research Laboratory of Electronics, not to mention my 25 years as a Senior Researcher in high tech for companies including IBM, Apple, and Microsoft I can claim to know some few things about this subject, yet still I learned a lot about the current state of the art from this guy. He particularly emphasizes the small attempted counterweight efforts to offset Kurzweil's manic robotic boosterism for his utopian Singularity, which boils down basically to a few guys chatting over the internet about how to create "Friendly AI". Well ... good luck suckers! ... seems to be the author's final conclusion on the dim hope that super intelligent systems could be constrained to maintain a commitment to honor any kind of human moral values over many iterations of recursive upgrading and exponentially awesome self-aggrandizement. Basically these machines will end up as gods.

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

Our Final Invention: Artificial Intelligence and the End of the Human Era
Java: Artificial Intelligence; Made Easy, w/ Java Programming; Learn to Create your * Problem Solving * Algorithms! TODAY!
w/ Machine Learning & Data Structures (Artificial Intelligence Series)
Javascript Artificial Intelligence: Made Easy, w/ Essential Programming; Create your * Problem Solving * Algorithms! TODAY!
w/ Machine Learning & Data Structures (Artificial Intelligence Series)
The Official Guide To Selling Final Expense Insurance: The Proven Final Expense Insurance Sales And Lead Generation System Used By Top Final Expense Agents Across The Country
The Most Human Human: What Artificial Intelligence Teaches Us About Being Alive
Gene Expression Programming: Mathematical Modeling by an Artificial Intelligence (Studies in Computational Intelligence)
Social Intelligence: A Practical Guide to Social Intelligence: Communication Skills - Social Skills - Communication Theory - Emotional Intelligence - 3D Parametric Intensity Models for the Localization of 3D Anatomical Point Landmarks and 3D Segmentation of Human Vessels (Dissertations in Artificial Intelligence: Infix, Vol. 299)
Nursing: Human Science And Human Care (Watson, Nursing: Human Science and Human Care)
Final Fantasy Type-0 Side Story, Vol. 5: The Ice Reaper (Final Fantasy 0-Type)
Final Fantasy Type-0 Side Story, Vol. 4: The Ice Reaper (Final Fantasy 0-Type)
Prolog and Databases: Implementations and Applications (Ellis Horwood Series in Artificial Intelligence)
Neural Network Training Using Genetic Algorithms (Series in Machine Perception and Artificial Intelligence)
Artificial Intelligence with Common Lisp: Fundamentals of Symbolic and Numeric Processing
Common Lisp and Artificial Intelligence
Prolog Programming for Students: With Expert Systems and Artificial

Intelligence Topics CommonKADS Library for Expertise Modelling, Reusable Problem Solving
Components (Frontiers in Artificial Intelligence and Applications) Artificial Intelligence to Bus
Systems and Computer Interfacing, Volume 2, Encyclopedia of Applied Physics Artificial Intelligence
(Cutting-Edge Science and Technology) Humans Need Not Apply: A Guide to Wealth and Work in
the Age of Artificial Intelligence

[Dmca](#)