


(Read free) Artificial Superintelligence: A Futuristic Approach

# Artificial Superintelligence: A Futuristic Approach

*Roman V. Yampolskiy*

*DOC | \*audiobook | ebooks | Download PDF | ePub*



[Download](#)

[Read Online](#)

#1107770 in Books Roman V Yampolskiy 2015-06-21 Original language: English PDF # 1 10.00 x .60 x 7.10l, .0 #File Name: 1482234432227 pages Artificial Superintelligence A Futuristic Approach | File size: 65.Mb

**Roman V. Yampolskiy : Artificial Superintelligence: A Futuristic Approach** before purchasing it in order to gage whether or not it would be worth my time, and all praised Artificial Superintelligence: A Futuristic Approach:

2 of 2 people found the following review helpful. I think the greatest service it provides is to show future generations where ...By Customer This book reviews the difficulties of securing safe AGI. Overall, I think the greatest service it provides is to show future generations where others have already looked, and failed. And in this regard this book was tremendous. I also appreciated the comprehensive references on every page, which really lets someone new to the field get caught up rather quickly. I also loved how this books tries to bring more computer scientists into the

conversation. My only critique is how the author's personal fears sometimes leak into the main text. While editorial views are common place in most books, I thought it felt out of place in this otherwise scholarly work. 0 of 0 people found the following review helpful. Disappointing By A Reader I bought this expecting more - I was deeply disappointed. It's a hodge podge of semi-baked ideas and papers cobbled together in the hopes of cashing in on the AI hype at the moment. There are much better sources of information on this topic out there. Avoid. 2 of 2 people found the following review helpful. Important By Calum Dr Roman Yampolskiy is a tenured computer scientist at the University of Louisville. He has published over 100 papers and books on artificial intelligence, genetic algorithms and behavioural biometrics. This is obviously a strong pedigree for a book about the subject of how to make sure that the arrival of superintelligence on the planet is an event that works out well for humans which is probably the single most important challenge facing humanity this century. Yampolskiy's preference for a safety engineering approach over an ethics approach to the Friendly AI problem is refreshing. The book faces up squarely to the immense difficulty of controlling an entity which is many times smarter than its would-be controllers, and is an important contribution to a vital field.

A day does not go by without a news article reporting some amazing breakthrough in artificial intelligence (AI). Many philosophers, futurists, and AI researchers have conjectured that human-level AI will be developed in the next 20 to 200 years. If these predictions are correct, it raises new and sinister issues related to our future in the age of intelligent machines. *Artificial Superintelligence: A Futuristic Approach* directly addresses these issues and consolidates research aimed at making sure that emerging superintelligence is beneficial to humanity. While specific predictions regarding the consequences of superintelligent AI vary from potential economic hardship to the complete extinction of humankind, many researchers agree that the issue is of utmost importance and needs to be seriously addressed. *Artificial Superintelligence: A Futuristic Approach* discusses key topics such as: AI-Completeness theory and how it can be used to see if an artificial intelligent agent has attained human level intelligence Methods for safeguarding the invention of a superintelligent system that could theoretically be worth trillions of dollars Self-improving AI systems: definition, types, and limits The science of AI safety engineering, including machine ethics and robot rights Solutions for ensuring safe and secure confinement of superintelligent systems The future of superintelligence and why long-term prospects for humanity to remain as the dominant species on Earth are not great *Artificial Superintelligence: A Futuristic Approach* is designed to become a foundational text for the new science of AI safety engineering. AI researchers and students, computer security researchers, futurists, and philosophers should find this an invaluable resource.

"Dr. Roman Yampolskiy is a tenured computer scientist at the University of Louisville. He has published over 100 papers and books on artificial intelligence, genetic algorithms, and behavioral biometrics. This is obviously a strong pedigree for a book about the subject of how to make sure that the arrival of superintelligence on the planet is an event that works out well for humans, which is probably the single most important challenge facing humanity this century. Yampolskiy's preference for a safety engineering approach over an ethics approach to the Friendly AI problem is refreshing. The book faces up squarely to the immense difficulty of controlling an entity that is many times smarter than its would-be controllers, and is an important contribution to a vital field." Calum Chace, Author of *Surviving AI* "... a very interesting book. Crammed into some 200 pages, index included, the book tries to establish a method of measuring progress in artificial intelligence (AI) by creating an AI analogy to the work of Stephen Cook and others in computational complexity. Specifically, the book introduces the author's concepts of AI-complete and AI-hard as analogies to the computational complexity categories of NP-complete and NP-hard. Yampolskiy (Univ. of Louisville) makes his case in just ten chapters. Chapter 1 introduces the topic of AI-Completeness. Chapters 2 through 8 elaborate the details of the author's vision of superintelligences. Chapter 9, 'Efficiency Theory: A Unifying Theory for Information, Computation, and Intelligence,' brings together the diversity of issues presented in the earlier chapters and does a good job of unifying the book. Yampolskiy presents his thoughts on AI's future in the final chapter. Each chapter includes an impressive collection of references, and the text has a healthy index. In general, this work should interest researchers in both AI and computational complexity. Readers may also wish to consult Nick Bostrom's *Superintelligence* (CH, Mar'15, 52-3620). Summing up: Highly recommended. Upper-division undergraduates through professionals/practitioners." J. Beidler, University of Scranton, CHOICE, March 2016 "Concerns over the existential risks of artificial superintelligence have spawned multiple vectors of research and development into specification, validation, security, and control. Roman Yampolskiy's *Artificial Superintelligence: A Futuristic Approach* reviews the relevant literature and stakes out the territory of AI safety engineering. Specifically, Yampolskiy advocates formal approaches to characterizing AIs and systematic confinement of superintelligent AIs. Serious students of AI and artificial general intelligence should study this work, and consider its recommendations." Neil Jacobstein, Chair, AI and Robotics, Singularity University at NASA Research Park, and Distinguished Visiting Scholar, MediaX Program at Stanford University "There are those of us who philosophize and debate the finer points surrounding the dangers of artificial intelligence (AI). And then there are those who dare go in the trenches and get their hands dirty by doing the

actual work that may just end up making the difference. So if AI turns out to be like the terminator then Prof. Roman Yampolskiy may turn out to be like John Connor but better. Because instead of fighting by using guns and brawn, he is utilizing computer science, human intelligence, and code." Nikola Danaylov, SingularityWeblog.com, September 7, 2015 "In his new book Artificial Superintelligence, Yampolsky argues for addressing AI potential dangers with a safety engineering approach, rather than with loosely defined ethics, since human values are inconsistent and dynamic. Yampolsky acknowledges the concern of AI escaping confines and takes the reader on a tour of AI taxonomies with a general overview of the field of Intelligence Yampolsky proposes initiation of an AI hazard symbol, which could prove useful for constraining AI to designated containment areas For readers intrigued in what safe variety of AI might be possible, the section of Artificial Superintelligence early in the book will be of great interest." Cynthia Sue Larson, RealityShifters Blog, September 1, 2015 " the hot topic that seems to have come straight from science fiction ... vigorous academic analysis pursued by the author produced an awesome textbook that should attract everyone's attention: from high school to graduate school students to professionals." Leon Reznik, Professor of Computer Science, Rochester Institute of Technology "This new book by Roman Yampolskiy is truly futuristic. I have had the chance to see some of his previous works, and this one is his best so far. Not to be missed by anyone really interested in artificial intelligence and the future of humanity. This book is a tour-de-force with deep insights into artificial intelligence and the future by one of the young experts in this fascinating field." Jose Cordeiro, Director, The Millennium Project, Venezuela Node About the Author Roman V. Yampolskiy holds a PhD from the Department of Computer Science and Engineering at the University at Buffalo (Buffalo, NY). There, he was a recipient of a four-year National Science Foundation (NSF) Integrative Graduate Education and Research Traineeship (IGERT) fellowship. Before beginning his doctoral studies, Dr. Yampolskiy received a BS/MS (High Honors) combined degree in computer science from the Rochester Institute of Technology in New York State. Dr. Yampolskiy's main areas of interest are behavioral biometrics, digital forensics, pattern recognition, genetic algorithms, neural networks, artificial intelligence, and games. Dr. Yampolskiy is an author of over 100 publications, including multiple journal articles and books. His research has been cited by numerous scientists and profiled in popular magazines, both American and foreign (New Scientist, Poker Magazine, Science World Magazine), dozens of websites (BBC, MSNBC, Yahoo! News), and on radio (German National Radio, Alex Jones Show).