


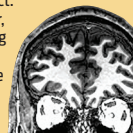















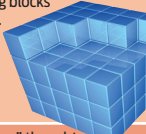





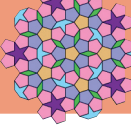




WHO	WHY WE SHOULD LISTEN	KIND OF SINGULARITY	BASIC ARGUMENT	WHAT THEY SAY	YOU MIGHT ALSO WANT TO KNOW	WHO	WHY WE SHOULD LISTEN	KIND OF SINGULARITY	BASIC ARGUMENT	WHAT THEY SAY	YOU MIGHT ALSO WANT TO KNOW	
<b>Raymond Kurzweil</b> Entrepreneur, inventor, author of popular books on the singularity 	Created several enormously successful inventions, including ones for pattern, character, and speech recognition; electronic musical instruments; and medical education.	Technotopia—accelerating change	The basic paradigm of Moore's Law—exponentially increasing improvement—will not only hold true indefinitely for logic circuits but will also apply to countless other technologies. It will lead to a singularity that will enable us to upload our consciousness into machines and, in effect, live indefinitely. This singularity will occur in about 15 years.	"I regard the freeing of the human mind from its severe physical limitations of scope and duration as the necessary next step in evolution." —"The Law of Accelerating Returns" ( <a href="http://www.kurzweilai.net/articles/art0134.html?printable=1">http://www.kurzweilai.net/articles/art0134.html?printable=1</a> ) 	Plans to release a movie later this year based on his book <i>The Singularity Is Near</i> . To maximize his chances of living until his consciousness can be downloaded, he reportedly takes upward of 200 vitamin pills a day.	<b>Marvin Minsky</b> Professor of media arts and sciences and of electrical engineering and computer science, MIT 	Built the first neural network simulator in 1951. Received the Turing Award in 1969. Developed a theory of mind involving small subunits of expertise.	Machine intelligence/uploading	Our increasing knowledge of the brain and increasing computing power will eventually intersect. Currently, however, the resources being devoted to such work are negligible compared with the amount of work to be done. 	"I asked the audience how many people wanted to live for 200 years, and no one raised their hand." —"It's 2001: Where Is HAL?" <i>Information Week</i> podcast, 11 March 2007 ( <a href="http://www.informationweek.com">http://www.informationweek.com</a> ; search on "minsky podcast")	Thinks that belief in a singularity could inspire people in the way that World War II inspired his generation.	
<b>Hans Moravec</b> Adjunct professor, Carnegie-Mellon Robotics Institute 	Wrote <i>Mind Children</i> and other books on artificial intelligence and robotics.	Machine intelligence 	The processes of the brain that give rise to consciousness arise inevitably and uniquely from chains of physical events and in accordance with physical principles. They are therefore reproducible with a sufficiently powerful computer.	"The 1500-cubic-centimeter human brain is about 100 000 times as large as the retina, suggesting that matching overall human behavior will take about 100 million MIPS of computer power." —"When Will Computer Hardware Match the Human Brain?" <i>Journal of Evolution and Technology</i> , 1998 ( <a href="http://www.jetpress.org/volume1/moravec.htm">http://www.jetpress.org/volume1/moravec.htm</a> )	Now works full-time at a small machine-vision start-up.	<b>Daniel Dennett</b> Codirector of the Center for Cognitive Studies and professor of philosophy, Tufts University 	Has written many books about the nature of consciousness and intelligence, including <i>Brainstorms: Philosophical Essays on Mind and Psychology</i> (1978), <i>The Intentional Stance</i> (1989), <i>Consciousness Explained</i> (1991), and <i>Brainchildren—Essays on Designing Minds</i> (1998).	Machine intelligence	Human-level AI may be inevitable, but don't expect it anytime soon. "I don't deny the possibility a priori; I just think it is vanishingly unlikely in the foreseeable future."	"The best reason for believing that robots might someday become conscious is that we human beings are conscious, and we are a sort of robot ourselves." (From Dennett's book <i>Cognition, Computation, and Consciousness</i> , 1994)	 Leading member of the Brights, a group dedicated to promoting "the civic understanding and acknowledgment of the naturalistic worldview, which is free of supernatural and mystical elements."	
<b>Nick Bostrom</b> Director, Future of Humanity Institute, Oxford University 	Cofounded World Transhumanist Association; writes extensively on machine intelligence, life extension, and philosophical issues surrounding the singularity.	Technotopia	Assuming that we don't render ourselves extinct, technological progress should lead to superintelligence and indefinitely extended life spans; once the singularity comes near, we will all be kicking ourselves for not having brought it about sooner.	"I would...assign less than a 50% probability to superintelligence being developed by 2033." —"How Long Before Superintelligence?" ( <a href="http://www.nickbostrom.com/superintelligence.html">http://www.nickbostrom.com/superintelligence.html</a> )	Had a short career as a stand-up comic. 	<b>Rodney Brooks</b> Professor of robotics, MIT 	Has built many robots that emulate the behavior of simple animals. One crucial aspect of his work is demonstrating how little cognition, understanding, or memory is required for apparently complex behaviors.	Technotopia	Any extrapolation of technological trends, particularly exponential ones, decades into the future isn't likely to be accurate.	"I am a machine. So are you." (p. 71 in this issue) 	Was featured in the 1997 movie <i>Fast, Cheap, and Out of Control</i> .	
<b>Vernor Vinge</b> Science-fiction author; professor of computer science at San Diego State University for 28 years (retired) 	Wrote the critically acclaimed novels <i>A Fire Upon the Deep</i> (Hugo Award, 1993) and <i>Deepness in the Sky</i> (1993); tackled issues that included human confrontation with superintelligent beings.	Event horizon—fundamentally unpredictable 	In a seminal 1993 essay, he wrote of the period following the development of machine intelligence: "From the human point of view this change will be a throwing away of all the previous rules, perhaps in the blink of an eye, an exponential runaway beyond any hope of control. Developments that before were thought might only happen in 'a million years' (if ever) will likely happen in the next century."	"In my 1993 essay, 'The Coming Technological Singularity,' I said I'd be surprised if the singularity had not happened by 2030. I'll stand by that claim." (p. 77 in this issue)	A chapter of his recent novel, <i>Rainbows End</i> , grew out of a story he published in the July 2004 issue of <i>IEEE Spectrum</i> .	<b>Jaron Lanier</b> Interdisciplinary scholar-in-residence at the Center for Entrepreneurship & Technology, University of California, Berkeley 	Coined the term <i>virtual reality</i> ; was CEO of VPL Research, a pioneering virtual-reality start-up.	Machine intelligence 	There's no real evidence that computer representation of a brain will actually produce a mind. Furthermore, true AI can't arise from current patterns of software development because those patterns are fatally flawed. Meanwhile, accelerating technological change could still make the world as we know it nearly unrecognizable.	"If computers are to become smart enough to design their own successors, initiating a process that will lead to God-like omniscience after a number of ever swifter passages from one generation of computers to the next, someone is going to have to write the software that gets the process going, and humans have given absolutely no evidence of being able to write such software." —"One-Half of a Manifesto," <i>Wired</i> , December 2000 ( <a href="http://www.wired.com/wired/archive/8.12/lanier.html">http://www.wired.com/wired/archive/8.12/lanier.html</a> )	Wrote an entertaining 1995 essay in the <i>Journal of Consciousness Studies</i> titled "You Can't Argue With a Zombie." It zinged Daniel Dennett, Daniel Dennett's critics, Dartmouth students, and philosopher David Chalmers, among others. Sample line: "Arrogance is always a bad strategy in science. In philosophy I suppose it's fine."	
<b>Eliezer Yudkowsky</b> Research fellow, Singularity Institute for Artificial Intelligence 	Has developed a theory for building "friendly" artificial intelligences whose goals will not change as they themselves evolve.	Intelligence explosion	Each generation of intelligent human-machine collaborations uses its increased intelligence to design the next generation. The more intelligence, the faster the cycle goes, until intelligences appear with capacities far beyond those of unaided humans. Eventually, essentially all the work will be done on the machine side.	"Deep Blue's engineers might say, 'We have no idea what chess moves the machine will make, but we know they'll be great moves.'" (Interview with author) 	Working on the development of machine metaethics so that superintelligences can develop moral reasoning beyond the current level while still retaining predictable characteristics.	<b>John Holland</b> Professor of computer science and engineering and professor of psychology, University of Michigan 	Invented genetic programming, one of the key technologies the singularity has claimed will lead to superintelligent machines.	Machine intelligence 	In order to combine building blocks and selectively "breed" them to form ever more intelligent programs, you have to have the building blocks in the first place.	Uncritical believers in the singularity "think evolution is like monkeys at the typewriter, and if you just type fast enough you'll get somewhere." (Interview with author)	Believes that the more people know about the technologies that will supposedly bring about the singularity, the more aware they will be of the limitations of those technologies.	
<b>Christof Koch</b> Professor of cognitive and behavioral biology, Caltech 	Studies the neural basis of consciousness.	Machine intelligence	To create thinking machines, we have to understand what it is—both biologically and philosophically—that makes humans conscious beings.	"Consciousness does not seem to require many of the things we associate most deeply with being human." (p. 56 in this issue) 	Runs up a 1700-meter-high mountain "every couple of weeks."	<b>John Searle</b> Professor of philosophy, University of California, Berkeley 	Has written many books on the brain and consciousness, including <i>Minds, Brains, and Science</i> (1985) and <i>The Mystery of Consciousness</i> (1997).	—	His "Chinese Room" thought experiment explains why a bare CPU cannot understand the intent of the program it executes. Only organisms embodied in the real world, with real experiences to draw on, can become truly conscious and intelligent. All else is merely symbol manipulation.	"I believe that there is no objection in principle to constructing an artificial hardware system that would duplicate the powers of the brain to cause consciousness using some chemistry different from neurons. But to produce consciousness any such system would have to duplicate the actual causal powers of the brain." —"I Married a Computer," review of Kurzweil's <i>The Age of Spiritual Machines</i> , <i>The New York Review of Books</i> , 8 April 1999	Was a member of the UC Berkeley Free Speech movement in the early 1960s.	
<b>Kevin Kelly</b> Senior maverick, <i>Wired</i> magazine 	Wrote <i>Out of Control: The New Biology of Machines, Social Systems, and the Economic World</i> (1995) and other works on technology and society.	Phase change	Singularities are pervasive changes in the state of the world that are often recognizable only in retrospect. As a result, the singularity is always near, but whether the current wave of technological progress constitutes a singularity is impossible to tell.	"The word <i>technology</i> was only coined in 1829, but they'd been doing it for centuries by then." (Interview with author)	Popularized the "Maes-Garreau point," which observes that most predictions of positive technological revolution fall just within the life span of the person doing the predicting.	<b>Roger Penrose</b> Physicist and professor of mathematics, Oxford University 	Expounded on Albert Einstein's general theory of relativity. Wrote <i>The Nature of Space and Time</i> with Stephen Hawking and <i>The Emperor's New Mind</i> , about computation and consciousness.	—	Consciousness cannot be duplicated in computational machines, because it depends on "noncomputational physical processes." Does not know what these might be but suggests it emerges from "large-scale" quantum-mechanical phenomena in microtubules in the brain's neurons.	"I'm not saying that consciousness is beyond physics...although I'm saying that it's beyond the physics we know now." (From Penrose's book <i>The Third Culture: Beyond the Scientific Revolution</i> , 1995)	Invented Penrose tiles, a shape capable of covering an infinite plane in a nonperiodic fashion. 	
<b>Bill Joy</b> Venture capitalist 	Cofounded Sun Microsystems. 	Event horizon	As computer science, biotech, and nanotechnology advance, it will become easier and easier for small groups or even individuals to create incredibly destructive things. Unless measures are taken, eventually this will happen.	"The future doesn't need us." —"Why the Future Doesn't Need Us," <i>Wired</i> , April 2000 ( <a href="http://www.wired.com/wired/archive/8.04/joy.html">http://www.wired.com/wired/archive/8.04/joy.html</a> )	Has called for voluntary renunciation or at least a delay of research into fields that offer easy extinction of humanity.							

Singularity color code: ■ True believer: thinks it will happen within 30 years ■ Yes, but... ■ Maybe someday ■ No way