

From: John Brockman <[REDACTED]>
To: Anderson Chris <[REDACTED]>, Anderson Christopher <[REDACTED]>, Bezos Jeff <[REDACTED]>, Brin Sergey <[REDACTED]>, Deutsch David <david.deutsch@qubit.org>, Doctorow Cory <[REDACTED]>, Dyson Esther <[REDACTED]>, Dyson Freeman <dyson@ias.edu>, Epstein Jeffrey <jeevacation@gmail.com>, Gelernter David <[REDACTED]>, Greene Brian <greene@math.columbia.edu>, Jacquet Jennifer <[REDACTED]>, Lloyd Seth <[REDACTED]>, Kamen Dean <[REDACTED]>, Swisher Kara <kara@recode.net>, Heffernan Virginia <[REDACTED]>, Weinberger Russell <weinberger@edge.org>, Kreye Andrian <[REDACTED]>, Lanier Jaron <[REDACTED]>, Levin Janna <janna@astro.columbia.edu>, Marletto Chiara <[REDACTED]>, "Markoff, John" <[REDACTED]>, Mossberg Walter <walt@recode.net>, Tim O'Reilly <[REDACTED]>, Levy Steven <[REDACTED]>, Rushkoff Doug <[REDACTED]>, Steinhardt Paul <[REDACTED]>, Victoria Stodden <[REDACTED]>, Jimmy Wales <[REDACTED]>, Bowles Nellie <[REDACTED]>, West Geoffrey <[REDACTED]>, Brand Stewart <[REDACTED]>, Kelly Kevin <[REDACTED]>, Dyson George <[REDACTED]>, Minsky Marvin <[REDACTED]>, Brockman Max <[REDACTED]>, Gershenfeld Neil <[REDACTED]>, "Dennett Daniel C." <[REDACTED]>, Daniel Kahneman <[REDACTED]>, Thaler Richard <richard.thaler@chicagogobusiness.edu>, Mullainathan Sendhil <[REDACTED]>, Haidt Jonathan <[REDACTED]>, Pinker Steven <[REDACTED]>, Harris Sam <sam@samharris.org>, Coyne Jerry <j-coyne@nyu.edu>, Myhrvold Nathan <[REDACTED]>, Pescovitz David <[REDACTED]>, Williams Evan <[REDACTED]>, Thrun Sebastian <[REDACTED]>, Fadell Tony <[REDACTED]>, Norvig Peter <[REDACTED]>, Kamangar Salar <[REDACTED]>, Page Larry <[REDACTED]>, Diamandis Peter <[REDACTED]>, Yuri Milner <[REDACTED]>, Khosla Vinod <[REDACTED]>, Strogatz Steven <[REDACTED]>, Shirky Clay <[REDACTED]>, Dawkins Richard <[REDACTED]>, "Hillis W. Daniel" <[REDACTED]>, Eno Brian <[REDACTED]>, Morozov Evgeny <[REDACTED]>, Rees Martin <[REDACTED]>, Parker Bruce <[REDACTED]>, Shermer Michael <[REDACTED]>, Smoot George <[REDACTED]>, Church George <[REDACTED]>, Venter Craig <[REDACTED]>, Musk Elon <[REDACTED]>, "Grayling A.C." <[REDACTED]>, Krauss Lawrence <[REDACTED]>, Smolin Lee <[REDACTED]>, Leonard Susskind <[REDACTED]>, Turok Neil <[REDACTED]>, Zeilinger Anton <[REDACTED]>, Sejnowski Terrence <[REDACTED]>, Gigerenzer Gerd <[REDACTED]>, Diamond Jared <[REDACTED]>, Davies Paul <[REDACTED]>, Guth Alan <[REDACTED]>, Linde Andrei <alinde@stanford.edu>, Rodney Brooks <[REDACTED]>

Subject: The Reality Club Discourse (a digest) on Jaron Lanier's "The Myth of AI"

Date: Tue, 18 Nov 2014 23:20:09 +0000

To: Chris Anderson, Chris Anderson, Ross Anderson, Jeff Bezos, Stewart Brand, Sergey Brin, Rodney Brooks, Jerry Coyne, Paul Davies, Richard Dawkins, Daniel Dennett, David Deutsch, Peter Diamandis, Jared Diamond, Cory Doctorow, Esther Dyson, Freeman Dyson, George Dyson, Brian Eno, Jeffrey Epstein, Tony Fadell, Bill GateDavid Gelernter, Niel Gershenfeld, Gerd Gigerenzer, Anthony Grayling, Brian Greene, Alan Guth, Jonathan Haidt, Sam Harris, Virginia Heffernan, Danny Hillis, Jennifer Jacquet, Seth Lloyd, Daniel Kahneman, Salar Kamangar, Dean Kamen, Kelly Kevin, Vinod Khosla, Lawrence Krauss, Andrian Kreye, Janna Levin, Steven Levy, Andrei Linde, Chiara Marletto, John Markoff, Yuri Milner, Marvin Minsky, Evgeny Morozov, Walt Mossberg, Sendhil Mullainathan, Nathan Myhrvold, Peter Norvig, Hans Ulrich Obrist, Tim O'Reilly, Larry Page,

Bruce Parker, David Pescovitz, Steven Pinker, Martin Rees, Douglas Rushkoff, Terrence Sejnowski, Michael Shermer, Clay Shirky, Lee Smolin, George Smoot, Paul Steinhardt, Bruce Sterling, Victoria Stodden, Steven Strogatz, Leonard Susskind, Kara Swisher, Richard Thaler, Sebastian Thrun, Neil Turok, Craig Venter, Jimmy Wales, Geoffrey West, Evan Williams, Anton Zeilinger

Neil Gershenfeld, one of the participants in this discussion, followed the rules and sent his comments to me for posting on the EDGE website, sparing all 77 of you from reading yet another email. He made a point when he said that "I doubt anyone will be following the link" and he asked if I would do some kind of collective update. So, here it is below, a digest with individual links to each comment. Of course, there is a link, and you may want to follow it: <http://edge.org/conversation/the-myth-of-ai#rc>

JB

John Brockman
917-744-8920 Mobile

THE MYTH OF AI: JARON LANIER

<http://edge.org/conversation/the-myth-of-ai>

The idea that computers are people has a long and storied history. It goes back to the very origins of computers, and even from before. There's always been a question about whether a program is something alive or not since it intrinsically has some kind of autonomy at the very least, or it wouldn't be a program. There has been a domineering subculture—that's been the most wealthy, prolific, and influential subculture in the technical world—that for a long time has not only promoted the idea that there's an equivalence between algorithms and life, and certain algorithms and people, but a historical determinism that we're inevitably making computers that will be smarter and better than us and will take over from us. ... That mythology, in turn, has spurred a reactionary, perpetual spasm from people who are horrified by what they hear. You'll have a figure say, "The computers will take over the Earth, but that's a good thing, because people had their chance and now we should give it to the machines." Then you'll have other people say, "Oh, that's horrible, we must stop these computers." Most recently, some of the most beloved and respected figures in the tech and science world, including Stephen Hawking and **Elon Musk**, have taken that position of: "Oh my God, these things are an existential threat. They must be stopped."

THE REALITY CLUB

<http://edge.org/conversation/the-myth-of-ai#rc>

THE REALITY CLUB: George Church, Peter Diamandis, Lee Smolin, Rodney Brooks, Nathan Myhrvold, George Dyson, Pamela McCorduck, Sendhil Mullainathan, Steven Pinker, Neal Gershenfeld, D.A. Wallach, Michael Shermer, Stuart Kauffman, Kevin Kelly, Lawrence Krauss

George Church

Professor, Harvard University, Director, Personal Genome Project.

Thanks Jaron and John,

We are now growing at a pace that's fully exponential—with a doubling time of 1.5 years. If we are concerned with exponentials, then we must also consider biotech—improving with an even faster rate of change. Synthetic neurobiology (BRAIN initiative) and AI are now competing and synergizing. Moving beyond mere warnings of existential risks to strategies for risk reduction and scenario testing—join us at: <http://cser.org>, <http://thefutureoflife.org>. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25979>

Peter Diamandis

Chairman/CEO, X PRIZE Foundation

Three thoughts:

- (1) I'm not concerned about the long-term, "adult" General A.I.... It's the 3-5 year old child version that concerns me most as the A.I. grows up. ...
- (2) The government's first reaction is always to regulate...
- (3) Best analogy I know is what happened back in 1975 with the Asilomar Conference on Recombinant DNA. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25980>

Lee Smolin

Physicist, Perimeter Institute; Author, TIME REBORN

I am puzzled by the arguments put forward by those who say we should worry about a coming AI, singularity, because all they seem to offer is a prediction based on Moore's law. But, an exponential increase is not enough to demonstrate that a qualitative change in behavior will take place. Besides which, the zeroth law of economics is that exponential change never goes on forever. What specific capacities do they fear computers may acquire before Moore's law runs out, and why do they think these could get "out of control"? Is there any concrete evidence for a programmable digital computer evolving the ability of taking initiatives or making choices which are not on a list of options programmed in by a human programmer? Finally, is there any detailed reason to think that a programmable digital computer is a good model for what goes on in the brain? ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25981>

Rodney A. Brooks

Robotist; Panasonic Professor of Robotics (emeritus), MIT; Founder, Chairman & CTO, Rethink Robotics; Author, FLESH AND MACHINES

...Recently there has been a spate of articles in the mainstream press, and a spate of high profile people who are in tech but not AI, speculating about the dangers of malevolent AI being developed, and how we should be worried about that possibility. I say relax. Chill. This all comes from some fundamental misunderstandings of the nature of the undeniable progress that is being made in AI, and from a misunderstanding of how far we really are from having volitional or intentional artificially intelligent beings, whether they be deeply benevolent or malevolent. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25982>

Nathan Myhrvold

CEO and Managing Director, Intellectual Ventures; Coauthor (with Bill Gates), THE ROAD AHEAD; Author, MODERNIST CUISINE

Somebody has to play skeptic or naysayer. It is truly bizarre that that role seems to have fallen on me, but here goes.

I would really love if AI was working so damn well that it was about to get scary. I think that day may well come—I have no objection to the idea that a machine can ultimately think as well or better than humans. Computers have been on track for exponential increase in at least some measures of computing power (basic operations, accessing memory, floating point multiplication...). Algorithms have built on that to let s do some amazing things, and conceptual progress on algorithms has arguably been as fast or faster than the raw hardware power. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25983>

George Dyson

Science Historian; Author, TURING'S CATHEDRAL: THE ORIGINS OF THE DIGITAL UNIVERSE; DARWIN AMONG THE MACHINES

Jaron, as always, is articulate, and I agree with most of what he says. He (and others) however, seem to be ignoring the real elephant in the room: analog computing.

The brain (of a human or of a fruit fly) is not a digital computer, and intelligence is not an algorithm. The difficulty of turning this around, despite some initial optimism, and achieving even fruit fly level intelligence with algorithms running on digital computers should have put this fear to rest by now. Listen to Jaron, and relax. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25984>

Pamela Mccorduck

Author, MACHINES WHO THINK, THE UNIVERSAL MACHINE, BOUNDED RATIONALITY, THIS COULD BE IMPORTANT; Coauthor (with Edward Feigenbaum), THE FIFTH GENERATION

Corporations aren't people and machines aren't people either. In the more than half century that I've been watching AI, I've never heard a researcher say they were equivalent. Sadly, I've heard outsiders attribute such beliefs to AI researchers, even to me, but it wasn't and isn't so. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25985>

Sendhil Mullainathan

Professor of Economics, Harvard; Assistant Director for Research, The Consumer Financial Protection Bureau (CFPB), U.S. Treasury Department (2011-2013); Coauthor, SCARCITY: WHY HAVING TOO LITTLE MEANS SO MUCH

I would add a +1 to everything Nathan said.

In addition I would make a distinction between machine intelligence and machine decision-making.

We should be afraid. Not of intelligent machines. But of machines making decisions that they do not have the intelligence to make. I am far more afraid of machine stupidity than of machine intelligence. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25986>

Steven Pinker

Johnstone Family Professor, Department of Psychology; Harvard University; Author, THE BETTER ANGELS OF OUR NATURE

Jaron Lanier has pointed out one reason that paranoid worries about artificial intelligence are a waste of time: Human-level AI is still the proverbial 15-to-25 years away, just as it always has been, and many of its recently touted advances have shallow (and human-nourished) roots. But there are other reasons not to worry about killer bots and other machines running amok. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25987>

Neil Gershenfeld

Physicist, Director, MIT's Center for Bits and Atoms; Author, FAB

I find the discussion of killer AI to be a bit (so to speak) silly.

The history of technology advancing has been one of sigmoids that begin as exponentials. The exponential phase comes with utopian dreams paired with fears of existential threats, followed by the sigmoidal crossover which is identified by both arguments fading into irrelevance. Both the fears and dreams have value in inspiring and moderating progress, but both are best viewed as markers of a transitional evolutionary stage. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25988>

■. Wallach

Recording Artist; Songwriter; Artist in Residence, Spotify

George Church makes a very important point in his comment on your Edge discussion about the Lanier piece: that synthetic neurobiology and computing are going to be increasingly merging. While the human brain and body might not do many things as well as digital supercomputers, they are pretty good substrates for lots of complex activity, very little of which we understand in any detail today. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25989>

Michael Shermer

Publisher, Skeptic magazine; monthly columnist, Scientific American; Author, THE MORAL ARC, and Presidential Fellow at Chapman University.

The latest round of handwringing over the potential for computers, machines, or robots to turn evil overlooks the fundamental difference between artificial intelligence (AI) and natural intelligence (NI). AI is intelligently designed whereas NI is the product of natural selection that produced emotions—both good and evil—to direct behavior. Machines with AI are not subject to the pressures of natural selection and therefore will not evolve emotions of any kind, good or evil. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25990>

Stuart A. Kauffman

Professor of Biological Sciences, Physics, Astronomy, University of Calgary; Author, REINVENTING THE SACRED

Since Turing and the explosive growth of algorithmic artificial intelligence, many of us think we are machines. I will argue we are surely not machines at all, but rather, trapped in an inadequate theory. Turing machines are discrete state, (0,1), discrete time (T, T + 1) subsets of continuous state continuous time classical physics. We have made amazing advances with universal computers, and with continuous models of neural systems as nonlinear dynamical systems. In all these cases the present state of the system entirely determines the next state of the system, so that next state is "entailed" by the laws of motion of the computer or classical dynamical system. Many hope that consciousness might emerge in such a system. Were that to happen, which is possible, the causal closure of classical physics demands that there is nothing for such a conscious mind to do, for the current state of the system suffices entirely for the next state. Worse, there is no way such a mind could alter the behavior of the classical physical system. At best such a mind could only be epiphenomenal. Why then did mind evolve to use so much real estate in us? ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25991>

Kevin Kelly

Senior Maverick, Wired, Author, COOL TOOLS; WHAT TECHNOLOGY WANTS; "The Three Breakthroughs That Have Finally Unleashed AI on the World" (Wired)

It is wise to think through the implications of new technology. I understand the good intentions of Jaron Lanier and others who have raised an alarm about AI. But I think their method of considering the challenges of AI relies too much on fear, and is not based on the evidence we have so far. I propose a counterview with four parts:

1. AI is not improving exponentially.
2. We'll reprogram the AIs if we are not satisfied with their performance.
3. Reprogramming themselves, on their own, is the least likely of many scenarios.
4. Rather than hype fear, this is a great opportunity.

I expand each point below. ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25992>

Lawrence M. Krauss

Physicist; Cosmologist, ASU; Author, A UNIVERSE FROM NOTHING

Motivated by this discussion, my institute at ASU, The Origins Project, will run a high level workshop and associated public event on "The Dangers of AI?"—exact title to be determined—during the 2015-2016 academic year, when our Origins theme will be "Life and Death in the 21st Century", (and during which we will host other workshops on subjects that will likely include The Origin of Life, and the Origin of Disease). ...

Permalink: <http://edge.org/conversation/the-myth-of-ai#25993>

