Technological singularity

From Wikipedia, the free encyclopedia

Contents [hide]

(Top)
Intelligence explosion
Emergence of superintelligence
Non-AI singularity
Speed superintelligence
Predictions
Plausibility
Speed improvements
Exponential growth
Accelerating change
Algorithm improvements
Criticism
Potential impacts
Uncertainty and risk
Next step of sociobiological evolution
Implications for human society
Hard vs. soft takeoff
Immortality
History of the concept
In politics
See also
References
Citations
Sources
Further reading
External links

"The Singularity" redirects here. For other uses, see Singularity (disambiguation).

The **technological singularity**—or simply the **singularity**[1]—is a hypothetical future point in time at which technological growth becomes uncontrollable and irreversible, resulting in unforeseeable changes to human civilization.[2][3] According to the most popular version of the singularity hypothesis, I. J. Good's intelligence explosion model, an upgradable intelligent agent will eventually enter a "runaway reaction" of self-improvement cycles, each new and more intelligent generation appearing more and more rapidly, causing an "explosion" in intelligence and resulting in a powerful superintelligence that qualitatively far surpasses all human intelligence.[4]

The first person to use the concept of a "singularity" in the technological context was the 20th-century Hungarian-American mathematician John von Neumann.[5] Stanislaw Ulam reports in 1958 an earlier discussion with von Neumann "centered on the accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue".[6] Subsequent authors have echoed this viewpoint.[3][7]

The concept and the term "singularity" were popularized by Vernor Vinge first in 1983 in an article that claimed that once humans create intelligences greater than their own, there will be a technological and social transition similar in some sense to "the knotted space-time at the center of a black hole",[8] and later in his 1993 essay *The Coming Technological Singularity*,[4][7] in which he wrote that it would signal the end of the human era, as the new superintelligence would continue to upgrade itself and would advance technologically at an incomprehensible rate. He wrote that he would be surprised if it occurred before 2005 or after 2030.[4] Another significant contributor to wider circulation of the notion was Ray Kurzweil's 2005 book *The Singularity Is Near*, predicting singularity by 2045.[7]

Some scientists, including Stephen Hawking, have expressed concern that artificial superintelligence (ASI) could result in human extinction.[9][10] The consequences of the singularity and its potential benefit or harm to the human race have been intensely debated.

Prominent technologists and academics dispute the plausibility of a technological singularity and the associated artificial intelligence explosion, including Paul Allen,[11] Jeff Hawkins,[12] John Holland, Jaron Lanier, Steven Pinker,[12] Theodore Modis,[13] and Gordon Moore.[12] One claim made was that the artificial intelligence growth is likely to run into decreasing returns instead of accelerating ones, as was observed in previously developed human technologies.

Futures studies

Concepts
Accelerating change · Cashless society · Existential risk · Future (Earth · Mathematics · Climate · Space exploration · Universe) · Kardashev scale · Moore's law · Resource depletion · **Singularity**

Techniques
Backcasting · Causal layered analysis · Chain-linked model · Consensus forecast · Cross impact analysis · Delphi (Real-time Delphi) · Foresight · Future-proof · Futures wheel · Futures workshops · Horizon scanning · Reference class forecasting · Scenario planning · Systems analysis · Threatcasting · Trend analysis

Technology assessment and forecasting
Critical design · Design fiction · Exploratory engineering · FTA · Hype cycle · Science fiction prototyping · Speculative design · TRL · Technology scouting

V · T · E

Intelligence explosion [edit]

Although technological progress has been accelerating in most areas, it has been limited by the basic intelligence of the human brain, which has not, according to Paul R. Ehrlich, changed significantly for millennia.[14] However, with the increasing power of computers and other technologies, it might eventually be possible to build a machine that is significantly more intelligent than humans.[15]

If a superhuman intelligence were to be invented—either through the amplification of human intelligence or through artificial intelligence—it would vastly improve over human problem-solving and inventive skills. Such an AI is referred to as **Seed AI**[16][17] because if an AI were created with engineering capabilities that matched or surpassed those of its human creators, it would have the potential to autonomously improve its own software and hardware to design an even more capable machine, which could repeat the process in turn. This recursive self-improvement could accelerate, potentially allowing enormous qualitative change before any upper limits imposed by the laws of physics or theoretical computation set in. It is speculated that over many iterations, such an AI would far surpass human cognitive abilities.

I. J. Good speculated in 1965 that artificial general intelligence might bring about an intelligence explosion:[18][19]

> Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion', and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control.

One version of intelligence explosion is where computing power approaches infinity in a finite amount of time. In this version, once AIs are doing the research to improve themselves, speed doubles e.g. after 2 years, then 1 year, then 6 months, then 3 months, then 1.5 months, etc., where the infinite sum of the doubling periods is 4 years. Unless prevented by physical limits of computation and time quantization, this process would literally achieve infinite computing power in 4 years, properly earning the name "singularity" for the final state. This form of intelligence explosion is described in Yudkowsky (1996).[20]

Emergence of superintelligence [edit]

Further information: Superintelligence

A superintelligence, hyperintelligence, or superhuman intelligence is a hypothetical agent that possesses intelligence far surpassing that of the brightest and most gifted human minds. "Superintelligence" may also refer to the form or degree of intelligence possessed by such an agent. John von Neumann, Vernor Vinge and Ray Kurzweil define the concept in terms of the technological creation of super intelligence, arguing that it is difficult or impossible for present-day humans to predict what human beings' lives would be like in a post-singularity world.[4][21]

Technology forecasters and researchers disagree regarding when, or whether, human intelligence will likely be surpassed. Some argue that advances in artificial intelligence (AI) will probably result in general reasoning systems that bypass human cognitive limitations. Others believe that humans will evolve or directly modify their biology so as to achieve radically greater intelligence.[22][23] A number of futures studies focus on scenarios that combine these possibilities, suggesting that humans are likely to interface with computers, or upload their minds to computers, in a way that enables substantial intelligence amplification. The book *The Age of Em* by Robin Hanson describes a hypothetical future scenario in which human brains are scanned and digitized, creating "uploads" or digital versions of human consciousness. In this future, the development of these uploads may precede or coincide with the emergence of superintelligent artificial intelligence.[24]

Non-AI singularity [edit]

Some writers use "the singularity" in a broader way to refer to any radical changes in society brought about by new technologies such as molecular nanotechnology,[25][26][27] although Vinge and other writers specifically state that without superintelligence, such changes would not qualify as a true singularity.[4]

Speed superintelligence [edit]

A speed superintelligence describes an AI that can function like a human mind, only much faster.[28] For example, with a million-fold increase in the speed of information processing relative to that of humans, a subjective year would pass in 30 physical seconds.[29] Such a difference in information processing speed could drive the singularity.[30]

Predictions [edit]

In 1965 I. J. Good predicted an ultraintelligent machine by 2000,[18] Vinge (1993) predicted greater-than-human intelligence between 2005 and 2030,[4] Yudkowsky (1996) predicted a singularity by 2021,[20] and Kurzweil (2005) predicted human-level artificial intelligence by 2030.[7] Moravec (1988) predicted human-level artificial intelligence in supercomputers by 2010 by extrapolating past trend using a chart,[31] while Moravec (1998/1999) predicted human-level artificial intelligence by 2040, and intelligence far beyond human by 2050.[32] In a 2017 interview, Kurzweil predicted human-level intelligence by 2029 and billion fold intelligence and singularity by 2045.[33][34]

Four polls of AI researchers, conducted in 2012 and 2013 by Nick Bostrom and Vincent C. Müller, suggested a confidence of 50% that artificial general intelligence (AGI) would be developed by 2040–2050.[35][36]

Plausibility [edit]

Prominent technologists and academics dispute the plausibility of a technological singularity, including Paul Allen,[11] Jeff Hawkins,[12] John Holland, Jaron Lanier, Steven Pinker,[12] Theodore Modis,[13] and Gordon Moore,[12] whose law is often cited in support of the concept.[37]

Most proposed methods for creating superhuman or transhuman minds fall into one of two categories: intelligence amplification of human brains and artificial intelligence. The many speculated ways to augment human intelligence include bioengineering, genetic engineering, nootropic drugs, AI assistants, direct brain–computer interfaces and mind uploading. These multiple possible paths to an intelligence explosion, all of which will presumably be pursued, makes a singularity more likely.[29]

Robin Hanson expressed skepticism of human intelligence augmentation, writing that once the "low-hanging fruit" of easy methods for increasing human intelligence have been exhausted, further improvements will become increasingly difficult.[38] Despite all of the speculated ways for amplifying human intelligence, non-human artificial intelligence (specifically seed AI) is the most popular option among the hypotheses that would advance the singularity. [*citation needed*]

The possibility of an intelligence explosion depends on three factors.[39] The first accelerating factor is the new intelligence enhancements made possible by each previous improvement. Contrariwise, as the intelligences become more advanced, further advances will become more and more complicated, possibly outweighing the advantage of increased intelligence. Each improvement should generate at least one more improvement, on average, for movement towards singularity to continue. Finally, the laws of physics may eventually prevent further improvement.

There are two logically independent, but mutually reinforcing, causes of intelligence improvements: increases in the speed of computation, and improvements to the algorithms used.[7] The former is predicted by Moore's Law and the forecasted improvements in hardware,[40] and is comparatively similar to previous technological advances. But there are some AI researchers,[*who?*] who believe software is more important than hardware.[41]

A 2017 email survey of authors with publications at the 2015 NeurIPS and ICML machine learning conferences asked about the chance that "the intelligence explosion argument is broadly correct". Of the respondents, 12% said it was "quite likely", 17% said it was "likely", 21% said it was "about even", 24% said it was "unlikely" and 26% said it was "quite unlikely".[42]

Speed improvements [edit]

Both for human and artificial intelligence, hardware improvements increase the rate of future hardware improvements. An analogy to Moore's Law suggests that if the first doubling of speed took 18 months, the second would take 18 subjective months; or 9 external months, whereafter, four months, two months, and so on towards a speed singularity.[43][20] Some upper limit on speed may eventually be reached. Jeff Hawkins has stated that a self-improving computer system would inevitably run into upper limits on computing power: "in the end there are limits to how big and fast computers can run. We would end up in the same place; we'd just get there a bit faster. There would be no singularity."[12]

It is difficult to directly compare silicon-based hardware with neurons. But Berglas (2008) notes that computer speech recognition is approaching human capabilities, and that this capability seems to require 0.01% of the volume of the brain. This analogy suggests that modern computer hardware is within a few orders of magnitude of being as powerful as the human brain.

Exponential growth [edit]

The exponential growth in computing technology suggested by Moore's law is commonly cited as a reason to expect a singularity in the relatively near future, and a number of authors have proposed generalizations of Moore's law. Computer scientist and futurist Hans Moravec proposed in a 1998 book[44] that the exponential growth curve could be extended back through earlier computing technologies prior to the integrated circuit.

Ray Kurzweil postulates a law of accelerating returns in which the speed of technological change (and more generally, all evolutionary processes[45]) increases exponentially, generalizing Moore's law in the same manner as Moravec's proposal, and also including material technology (especially as applied to nanotechnology), medical technology and others.[46] Between 1986 and 2007, machines' application-specific capacity to compute information per capita roughly doubled every 14 months; the per capita capacity of the world's general-purpose computers has doubled every 18 months; the global telecommunication capacity per capita doubled every 34 months; and the world's storage capacity per capita doubled every 40 months.[47] On the other hand, it has been argued that the global acceleration pattern having the 21st century singularity as its parameter should be characterized as hyperbolic rather than exponential.[48]

Kurzweil reserves the term "singularity" for a rapid increase in artificial intelligence (as opposed to other technologies), writing for example that "The Singularity will allow us to transcend these limitations of our biological bodies and brains ... There will be no distinction, post-Singularity, between human and machine".[49] He also defines his predicted date of the singularity (2045) in terms of when he expects computer-based intelligences to significantly exceed the sum total of human brainpower, writing that advances in computing before that date "will not represent the Singularity" because they do "not yet correspond to a profound expansion of our intelligence."[50]



Ray Kurzweil writes that, due to paradigm shifts, a trend of exponential growth extends Moore's law from integrated circuits to earlier transistors, vacuum tubes, relays, and electromechanical computers. He predicts that the exponential growth will continue, and that in a few decades the computing power of all computers will exceed that of ("unenhanced") human brains, with superhuman artificial intelligence appearing around the same time.



An updated version of Moore's law over 120 Years (based on Kurzweil's graph). The 7 most recent data points are all Nvidia GPUs.

Accelerating change [edit]

Main article: Accelerating change

Some singularity proponents argue its inevitability through extrapolation of past trends, especially those pertaining to shortening gaps between improvements to technology. In one of the first uses of the term "singularity" in the context of technological progress, Stanislaw Ulam tells of a conversation with John von Neumann about accelerating change:

> One conversation centered on the ever accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue.[6]

Kurzweil claims that technological progress follows a pattern of exponential growth, following what he calls the "law of accelerating returns". Whenever technology approaches a barrier, Kurzweil writes, new technologies will surmount it. He predicts paradigm shifts will become increasingly common, leading to "technological change so rapid and profound it represents a rupture in the fabric of human history".[51] Kurzweil believes that the singularity will occur by approximately 2045.[46] His predictions differ from Vinge's in that he predicts a gradual ascent to the singularity, rather than Vinge's rapidly self-improving superhuman intelligence.

Oft-cited dangers include those commonly associated with molecular nanotechnology and genetic engineering. These threats are major issues for both singularity advocates and critics, and were the subject of Bill Joy's April 2000 *Wired* magazine article "Why The Future Doesn't Need Us".[7][52]



According to Kurzweil, his logarithmic graph of 15 lists of paradigm shifts for key historic events shows an exponential trend.

Algorithm improvements [edit]

Some intelligence technologies, like "seed AI",[16][17] may also have the potential to not just make themselves faster, but also more efficient, by modifying their source code. These improvements would make further improvements possible, which would make further improvements possible, and so on.

The mechanism for a recursively self-improving set of algorithms differs from an increase in raw computation speed in two ways. First, it does not require external influence: machines designing faster hardware would still require humans to create the improved hardware, or to program factories appropriately. [*citation needed*] An AI rewriting its own source code could do so while contained in an AI box.

Second, as with Vernor Vinge's conception of the singularity, it is much harder to predict the outcome. While speed increases seem to be only a quantitative difference from human intelligence, actual algorithm improvements would be qualitatively different. Eliezer Yudkowsky compares it to the changes that human intelligence brought: humans changed the world thousands of times more rapidly than evolution had done, and in totally different ways. Similarly, the evolution of life was a massive departure and acceleration from the previous geological rates of change, and improved intelligence could cause change to be as different again.[53]

There are substantial dangers associated with an intelligence explosion singularity originating from a recursively self-improving set of algorithms. First, the goal structure of the AI might self-modify, potentially causing the AI to optimise for something other than what was originally intended.[54][55]

Secondly, AIs could compete for the same scarce resources humankind uses to survive.[56][57] While not actively malicious, AIs would promote the goals of their programming, not necessarily broader human goals, and thus might crowd out humans completely.[58][59][60]

Carl Shulman and Anders Sandberg suggest that algorithm improvements may be the limiting factor for a singularity; while hardware efficiency tends to improve at a steady pace, software innovations are more unpredictable and may be bottlenecked by serial, cumulative research. They suggest that in the case of a software-limited singularity, intelligence explosion would actually become more likely than with a hardware-limited singularity, because in the software-limited case, once human-level AI is developed, it could run serially on very fast hardware, and the abundance of cheap hardware would make AI research less constrained.[61] An abundance of accumulated hardware that can be unleashed once the software figures out how to use it has been called "computing overhang".[62]

Criticism [edit]

Some critics, like philosopher Hubert Dreyfus[63] and philosopher John Searle,[64] assert that computers or machines cannot achieve human intelligence. Others, like physicist Stephen Hawking,[65] object that whether machines can achieve a true intelligence or merely something similar to intelligence is irrelevant if the net result is the same.

Psychologist Steven Pinker stated in 2008: "There is not the slightest reason to believe in a coming singularity. The fact that you can visualize a future in your imagination is not evidence that it is likely or even possible. Look at domed cities, jet-pack commuting, underwater cities, mile-high buildings, and nuclear-powered automobiles—all staples of futuristic fantasies when I was a child that have never arrived. Sheer processing power is not a pixie dust that magically solves all your problems. ..."[12]

Martin Ford[66] postulates a "technology paradox" in that before the singularity could occur most routine jobs in the economy would be automated, since this would require a level of technology inferior to that of the singularity. This would cause massive unemployment and plummeting consumer demand, which in turn would destroy the incentive to invest in the technologies that would be required to bring about the Singularity. Job displacement is increasingly no longer limited to those types of work traditionally considered to be "routine".[67]

Theodore Modis[68] and Jonathan Huebner[69] argue that the rate of technological innovation has not only ceased to rise, but is actually now declining. Evidence for this decline is that the rise in computer clock rates is slowing, even while Moore's prediction of exponentially increasing circuit density continues to hold. This is due to excessive heat build-up from the chip, which cannot be dissipated quickly enough to prevent the chip from melting when operating at higher speeds. Advances in speed may be possible in the future by virtue of more power-efficient CPU designs and multi-cell processors.[70]

Theodore Modis holds the singularity cannot happen.[71][13][72] He claims the "technological singularity" and especially Kurzweil lack scientific rigor; Kurzweil is alleged to mistake the logistic function (S-function) for an exponential function, and to see a "knee" in an exponential function where there can in fact be no such thing.[73] In a 2021 article, Modis pointed out that no milestones – breaks in historical perspective comparable in importance to the Internet, DNA, the transistor, or nuclear energy – had been observed in the previous twenty years while five of them would have been expected according to the exponential trend advocated by the proponents of the technological singularity.[74]

AI researcher Jürgen Schmidhuber stated that the frequency of subjectively "notable events" appears to be approaching a 21st-century singularity, but cautioned readers to take such plots of subjective events with a grain of salt: perhaps differences in memory of recent and distant events could create an illusion of accelerating change where none exists.[75]

Microsoft co-founder Paul Allen argued the opposite of accelerating returns, the complexity brake;[11] the more progress science makes towards understanding intelligence, the more difficult it becomes to make additional progress. A study of the number of patents shows that human creativity does not show accelerating returns, but in fact, as suggested by Joseph Tainter in his *The Collapse of Complex Societies*,[76] a law of diminishing returns. The number of patents per thousand peaked in the period from 1850 to 1900, and has been declining since.[69] The growth of complexity eventually becomes self-limiting, and leads to a widespread "general systems collapse".

Hofstadter (2006) raises concern that Ray Kurzweil is not sufficiently scientifically rigorous, that an exponential tendency of technology is not a scientific law like one of physics, and that exponential curves have no "knees".[77] Nonetheless, he does not rule out the singularity in principle in the distant future.[12]

Jaron Lanier denies that the singularity is inevitable: "I do not think the technology is creating itself. It's not an autonomous process."[78] Furthermore: "The reason to believe in human agency over technological determinism is that you can then have an economy where people earn their own way and invent their own lives. If you structure a society on *not* emphasizing individual human agency, it's the same thing operationally as denying people clout, dignity, and self-determination ... to embrace [the idea of the Singularity] would be a celebration of bad data and bad politics."[78]

Economist Robert J. Gordon points out that measured economic growth slowed around 1970 and slowed even further since the financial crisis of 2007–2008, and argues that the economic data show no trace of a coming Singularity as imagined by mathematician I. J. Good.[79]

Philosopher and cognitive scientist Daniel Dennett said in 2017: "The whole singularity stuff, that's preposterous. It distracts us from much more pressing problems", adding "AI tools that we become hyper-dependent on, that is going to happen. And one of the dangers is that we will give them more authority than they warrant."[80]

In addition to general criticisms of the singularity concept, several critics have raised issues with Kurzweil's iconic chart. One line of criticism is that a log-log chart of this nature is inherently biased toward a straight-line result. Others identify selection bias in the points that Kurzweil chooses to use. For example, biologist PZ Myers points out that many of the early evolutionary "events" were picked arbitrarily.[81] Kurzweil has rebutted this by charting evolutionary events from 15 neutral sources, and showing that they fit a straight line on a log-log chart. Kelly (2006) argues that the way the Kurzweil chart is constructed with x-axis having time before present, it always points to the singularity being "now", for any date on which one would construct such a chart, and shows this visually on Kurzweil's chart.[82]

Some critics suggest religious motivations or implications of singularity, especially Kurzweil's version of it. The buildup towards the Singularity is compared with Judeo-Christian end-of-time scenarios. Beam calls it "a Buck Rogers vision of the hypothetical Christian Rapture".[83] John Gray says "the Singularity echoes apocalyptic myths in which history is about to be interrupted by a world-transforming event".[84]

David Streitfield in the *New York Times* questioned whether "it might manifest first and foremost—thanks, in part, to the bottom-line obsession of today's Silicon Valley—as a tool to slash corporate America's head count."[85]

Potential impacts [edit]

Dramatic changes in the rate of economic growth have occurred in the past because of technological advancement. Based on population growth, the economy doubled every 250,000 years from the Paleolithic era until the Neolithic Revolution. The new agricultural economy doubled every 900 years, a remarkable increase. In the current era, beginning with the Industrial Revolution, the world's economic output doubles every fifteen years, sixty times faster than during the agricultural era. If the rise of superhuman intelligence causes a similar revolution, argues Robin Hanson, one would expect the economy to double at least quarterly and possibly on a weekly basis.[86]

Uncertainty and risk [edit]

Further information: Existential risk from artificial general intelligence

The term "technological singularity" reflects the idea that such change may happen suddenly, and that it is difficult to predict how the resulting new world would operate.[87][88] It is unclear whether an intelligence explosion resulting in a singularity would be beneficial or harmful, or even an existential threat.[89][90] Because AI is a major factor in singularity risk, a number of organizations pursue a technical theory of aligning AI goal-systems with human values, including the Future of Humanity Institute, the Machine Intelligence Research Institute,[87] the Center for Human-Compatible Artificial Intelligence, and the Future of Life Institute.

Physicist Stephen Hawking said in 2014 that "Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last, unless we learn how to avoid the risks."[91] Hawking believed that in the coming decades, AI could offer "incalculable benefits and risks" such as "technology outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand."[91] Hawking suggested that artificial intelligence should be taken more seriously and that more should be done to prepare for the singularity:[91]

> So, facing possible futures of incalculable benefits and risks, the experts are surely doing everything possible to ensure the best outcome, right? Wrong. If a superior alien civilisation sent us a message saying, "We'll arrive in a few decades," would we just reply, "OK, call us when you get here – we'll leave the lights on"? Probably not – but this is more or less what is happening with AI.

Berglas (2008) claims that there is no direct evolutionary motivation for an AI to be friendly to humans. Evolution has no inherent tendency to produce outcomes valued by humans, and there is little reason to expect an arbitrary optimisation process to promote an outcome desired by humankind, rather than inadvertently leading to an AI behaving in a way not intended by its creators.[92][93][94] Anders Sandberg has also elaborated on this scenario, addressing various common counter-arguments.[95] AI researcher Hugo de Garis suggests that artificial intelligences may simply eliminate the human race for access to scarce resources,[56][96] and humans would be powerless to stop them.[97] Alternatively, AIs developed under evolutionary pressure to promote their own survival could outcompete humanity.[60]

Bostrom (2002) discusses human extinction scenarios, and lists superintelligence as a possible cause:

> When we create the first superintelligent entity, we might make a mistake and give it goals that lead it to annihilate humankind, assuming its enormous intellectual advantage gives it the power to do so. For example, we could mistakenly elevate a subgoal to the status of a supergoal. We tell it to solve a mathematical problem, and it complies by turning all the matter in the solar system into a giant calculating device, in the process killing the person who asked the question.

According to Eliezer Yudkowsky, a significant problem in AI safety is that unfriendly artificial intelligence is likely to be much easier to create than friendly AI. While both require large advances in recursive optimisation process design, friendly AI also requires the ability to make goal structures invariant under self-improvement (or the AI could transform itself into something unfriendly) and a goal structure that aligns with human values and does not automatically destroy the human race. An unfriendly AI, on the other hand, can optimize for an arbitrary goal structure, which does not need to be invariant under self-modification.[98] Bill Hibbard (2014) proposes an AI design that avoids several dangers including self-delusion,[99] unintended instrumental actions,[54][100] and corruption of the reward generator.[100] He also discusses social impacts of AI[101] and testing AI.[102] His 2001 book *Super-Intelligent Machines* advocates the need for public education about AI and public control over AI. It also proposed a simple design that was vulnerable to corruption of the reward generator.

Next step of sociobiological evolution [edit]

Further information: Sociocultural evolution



This section **may contain material unrelated or insufficiently related to the topic of the article**. Please help improve this section or discuss this issue on the talk page. *(October 2021)* (*Learn how and when to remove this template message*)

While the technological singularity is usually seen as a sudden event, some scholars argue the current speed of change already fits this description.[*citation needed*]

In addition, some argue that we are already in the midst of a major evolutionary transition that merges technology, biology, and society. Digital technology has infiltrated the fabric of human society to a degree of indisputable and often life-sustaining dependence.

A 2016 article in *Trends in Ecology & Evolution* argues that "humans already embrace fusions of biology and technology. We spend most of our waking time communicating through digitally mediated channels... we trust artificial intelligence with our lives through antilock braking in cars and autopilots in planes... With one in three courtships leading to marriages in America beginning online, digital algorithms are also taking a role in human pair bonding and reproduction".

The article further argues that from the perspective of the evolution, several previous Major Transitions in Evolution have transformed life through innovations in information storage and replication (RNA, DNA, multicellularity, and culture and language). In the current stage of life's evolution, the carbon-based biosphere has generated a cognitive system (humans) capable of creating technology that will result in a comparable evolutionary transition.



Schematic Timeline of Information and Replicators in the Biosphere: Gillings et al.'s "major evolutionary transitions" in information processing.[103]

The digital information created by humans has reached a similar magnitude to biological information in the biosphere. Since the 1980s, the quantity of digital information stored has doubled about every 2.5 years, reaching about 5 zettabytes in 2014 (5×10^{21} bytes).[104]

In biological terms, there are 7.2 billion humans on the planet, each having a genome of 6.2 billion nucleotides. Since one byte can encode four nucleotide pairs, the individual genomes of every human on the planet could be encoded by approximately 1×10^{19} bytes. The digital realm stored 500 times more information than this in 2014 (see figure). The total amount of DNA contained in all of the cells on Earth is estimated to be about 5.3×10^{37} base pairs, equivalent to 1.325×10^{37} bytes of information.



Amount of digital information worldwide (5×10^{21} bytes) versus human genome information worldwide (10^{19} bytes) in 2014[103]

If growth in digital storage continues at its current rate of 30–38% compound annual growth per year,[47] it will rival the total information content contained in all of the DNA in all of the cells on Earth in about 110 years. This would represent a doubling of the amount of information stored in the biosphere across a total time period of just 150 years".[103]

Implications for human society [edit]

Further information: Artificial intelligence in fiction

In February 2009, under the auspices of the Association for the Advancement of Artificial Intelligence (AAAI), Eric Horvitz chaired a meeting of leading computer scientists, artificial intelligence researchers and roboticists at the Asilomar conference center in Pacific Grove, California. The goal was to discuss the potential impact of the hypothetical possibility that robots could become self-sufficient and able to make their own decisions. They discussed the extent to which computers and robots might be able to acquire autonomy, and to what degree they could use such abilities to pose threats or hazards.[105]

Some machines are programmed with various forms of semi-autonomy, including the ability to locate their own power sources and choose targets to attack with weapons. Also, some computer viruses can evade elimination and, according to scientists in attendance, could therefore be said to have reached a "cockroach" stage of machine intelligence. The conference attendees noted that self-awareness as depicted in science-fiction is probably unlikely, but that other potential hazards and pitfalls exist.[105]

Frank S. Robinson predicts that once humans achieve a machine with the intelligence of a human, scientific and technological problems will be tackled and solved with brainpower far superior to that of humans. He notes that artificial systems are able to share data more directly than humans, and predicts that this would result in a global network of super-intelligence that would dwarf human capability.[106] Robinson also discusses how vastly different the future would potentially look after such an intelligence explosion.

Hard vs. soft takeoff [edit]

In a hard takeoff scenario, an AGI rapidly self-improves, "taking control" of the world (perhaps in a matter of hours), too quickly for significant human-initiated error correction or for a gradual tuning of the AGI's goals. In a soft takeoff scenario, AGI still becomes far more powerful than humanity, but at a human-like pace (perhaps on the order of decades), on a timescale where ongoing human interaction and correction can effectively steer the AGI's development.[108][109]



In this sample recursive self-improvement scenario, humans modifying an AI's architecture would be able to double its performance every three years through, for example, 30 generations before exhausting all feasible improvements (left). If instead the AI is smart enough to modify its own architecture as well as human researchers can, its time required to complete a redesign halves with each generation, and it progresses all 30 feasible generations in six years (right).[107]

Ramez Naam argues against a hard takeoff. He has pointed out that we already see recursive self-improvement by superintelligences, such as corporations. Intel, for example, has "the collective brainpower of tens of thousands of humans and probably millions of CPU cores to... design better CPUs!" However, this has not led to a hard takeoff; rather, it has led to a soft takeoff in the form of Moore's law.[110] Naam further points out that the computational complexity of higher intelligence may be much greater than linear, such that "creating a mind of intelligence 2 is probably *more* than twice as hard as creating a mind of intelligence 1."[111]

J. Storrs Hall believes that "many of the more commonly seen scenarios for overnight hard takeoff are circular – they seem to assume hyperhuman capabilities at the *starting point* of the self-improvement process" in order for an AI to be able to make the dramatic, domain-general improvements required for takeoff. Hall suggests that rather than recursively self-improving its hardware, software, and infrastructure all on its own, a fledgling AI would be better off specializing in one area where it was most effective and then buying the remaining components on the marketplace, because the quality of products on the marketplace continually improves, and the AI would have a hard time keeping up with the cutting-edge technology used by the rest of the world.[112]

Ben Goertzel agrees with Hall's suggestion that a new human-level AI would do well to use its intelligence to accumulate wealth. The AI's talents might inspire companies and governments to disperse its software throughout society. Goertzel is skeptical of a hard five minute takeoff but speculates that a takeoff from human to superhuman level on the order of five years is reasonable. Goerzel refers to this scenario as a "semihard takeoff".[113]

Max More disagrees, arguing that if there were only a few superfast human-level AIs, that they would not radically change the world, as they would still depend on other people to get things done and would still have human cognitive constraints. Even if all superfast AIs worked on intelligence augmentation, it is unclear why they would do better in a discontinuous way than existing human cognitive scientists at producing super-human intelligence, although the rate of progress would increase. More further argues that a superintelligence would not transform the world overnight: a superintelligence would need to engage with existing, slow human systems to accomplish physical impacts on the world. "The need for collaboration, for organization, and for putting ideas into physical changes will ensure that all the old rules are not thrown out overnight or even within years."[114]

Immortality [edit]

Drexler (1986), one of the founders of nanotechnology, postulates cell repair devices, including ones operating within cells and using as yet hypothetical biological machines.[115] According to Richard Feynman, it was his former graduate student and collaborator Albert Hibbs who originally suggested to him (circa 1959) the idea of a *medical* use for Feynman's theoretical micromachines. Hibbs suggested that certain repair machines might one day be reduced in size to the point that it would, in theory, be possible to (as Feynman put it) "swallow the doctor". The idea was incorporated into Feynman's 1959 essay *There's Plenty of Room at the Bottom*.[116]

Moravec (1988)[31] predicts the possibility of "uploading" human mind into a human-like robot, achieving quasi-immortality by extreme longevity via transfer of the human mind between successive new robots as the old ones wear out; beyond that, he predicts later exponential acceleration of subjective experience of time leading to a subjective sense of immortality.

Kurzweil (2005) suggests that medical advances would allow people to protect their bodies from the effects of aging, making the life expectancy limitless. Kurzweil argues that the technological advances in medicine would allow us to continuously repair and replace defective components in our bodies, prolonging life to an undetermined age.[117] Kurzweil further buttresses his argument by discussing current bio-engineering advances. Kurzweil suggests somatic gene therapy; after synthetic viruses with specific genetic information, the next step would be to apply this technology to gene therapy, replacing human DNA with synthesized genes.[118]

Beyond merely extending the operational life of the physical body, Jaron Lanier argues for a form of immortality called "Digital Ascension" that involves "people dying in the flesh and being uploaded into a computer and remaining conscious."[119]

History of the concept [edit]

A paper by Mahendra Prasad, published in *AI Magazine*, asserts that the 18th-century mathematician Marquis de Condorcet was the first person to hypothesize and mathematically model an intelligence explosion and its effects on humanity.[120]

An early description of the idea was made in John W. Campbell's 1932 short story "The last evolution".

In his 1958 obituary for John von Neumann, Ulam recalled a conversation with von Neumann about the "ever accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue."[6]

In 1965, Good wrote his essay postulating an "intelligence explosion" of recursive self-improvement of a machine intelligence.[18][19]

In 1977, Hans Moravec wrote an article with unclear publishing status where he envisioned a development of self-improving thinking machines, a creation of "super-consciousness, the synthesis of terrestrial life, and perhaps jovian and martian life as well, constantly improving and extending itself, spreading outwards from the solar system, converting non-life into mind."[121][122] The article describes the human mind uploading later covered in Moravec (1988). The machines are expected to reach human level and then improve themselves beyond that ("Most significantly of all, they [the machines] can be put to work as programmers and engineers, with the task of optimizing the software and hardware which make them what they are. The successive generations of machines produced this way will be increasingly smarter and more cost effective.") Humans will no longer be needed, and their abilities will be overtaken by the machines: "In the long run the sheer physical inability of humans to keep up with these rapidly evolving progeny of our minds will ensure that the ratio of people to machines approaches zero, and that a direct descendant of our culture, but not our genes, inherits the universe." While the word "singularity" is not used, the notion of human-level thinking machines thereafter improving themselves beyond human level is there. In this view, there is no intelligence explosion in the sense of a very rapid intelligence increase once human equivalence is reached. An updated version of the article was published in 1979 in Analog Science Fiction and Fact.[123][122]

In 1981, Stanisław Lem published his science fiction novel *Golem XIV*. It describes a military AI computer (Golem XIV) who obtains consciousness and starts to increase his own intelligence, moving towards personal technological singularity. Golem XIV was originally created to aid its builders in fighting wars, but as its intelligence advances to a much higher level than that of humans, it stops being interested in the military requirements because it finds them lacking internal logical consistency.

In 1983, Vernor Vinge addressed Good's intelligence explosion in print in the January 1983 issue of *Omni* magazine. In this op-ed piece, Vinge seems to have been the first to use the term "singularity" (although not "technological singularity") in a way that was specifically tied to the creation of intelligent machines:[8][122]

> We will soon create intelligences greater than our own. When this happens, human history will have reached a kind of singularity, an intellectual transition as impenetrable as the knotted space-time at the center of a black hole, and the world will pass far beyond our understanding. This singularity, I believe, already haunts a number of science-fiction writers. It makes realistic extrapolation to an interstellar future impossible. To write a story set more than a century hence, one needs a nuclear war in between ... so that the world remains intelligible.

In 1985, in "The Time Scale of Artificial Intelligence", artificial intelligence researcher Ray Solomonoff articulated mathematically the related notion of what he called an "infinity point": if a research community of human-level self-improving AIs take four years to double their own speed, then two years, then one year and so on, their capabilities increase infinitely in finite time.[7][124]

In 1986, Vernor Vinge published *Marooned in Realtime*, a science-fiction novel where a few remaining humans traveling forward in the future have survived an unknown extinction event that might well be a singularity. In a short afterword, the author states that an actual technological singularity would not be the end of the human species: "of course it seems very unlikely that the Singularity would be a clean vanishing of the human race. (On the other hand, such a vanishing is the timelike analog of the silence we find all across the sky.)".[125][126]

In 1988, Vinge used the phrase "technological singularity" (including "technological") in the short story collection *Threats and Other Promises*, writing in the introduction to his story "The Whirligig of Time" (p. 72): *Barring a worldwide catastrophe, I believe that technology will achieve our wildest dreams, and* soon. *When we raise our own intelligence and that of our creations, we are no longer in a world of human-sized characters. At that point we have fallen into a technological "black hole", a technological singularity.*[127]

In 1988, Hans Moravec published *Mind Children*,[31] in which he predicted human-level intelligence in supercomputers by 2010, self-improving intelligent machines far surpassing human intelligence later, human mind uploading into human-like robots later, intelligent machines leaving humans behind, and space colonization. He did not mention "singularity", though, and he did not speak of a rapid explosion of intelligence immediately after the human level is achieved. Nonetheless, the overall singularity tenor is there in predicting both human-level artificial intelligence and further artificial intelligence far surpassing humans later.

Vinge's 1993 article "The Coming Technological Singularity: How to Survive in the Post-Human Era",[4] spread widely on the internet and helped to popularize the idea.[128] This article contains the statement, "Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended." Vinge argues that science-fiction authors cannot write realistic post-singularity characters who surpass the human intellect, as the thoughts of such an intellect would be beyond the ability of humans to express.[4]

Minsky's 1994 article says robots will "inherit the Earth", possibly with the use of nanotechnology, and proposes to think of robots as human "mind children", drawing the analogy from Moravec. The rhetorical effect of that analogy is that if humans are fine to pass the world to their biological children, they should be equally fine to pass it to robots, their "mind" children. As per Minsky, 'we could design our "mind-children" to think a million times faster than we do. To such a being, half a minute might seem as long as one of our years, and each hour as long as an entire human lifetime.' The feature of the singularity present in Minsky is the development of superhuman artificial intelligence ("million times faster"), but there is no talk of sudden intelligence explosion, self-improving thinking machines or unpredictability beyond any specific event and the word "singularity" is not used.[129]

Tipler's 1994 book *The Physics of Immortality* predicts a future where super–intelligent machines will build enormously powerful computers, people will be "emulated" in computers, life will reach every galaxy and people will achieve immortality when they reach Omega Point.[130] There is no talk of Vingean "singularity" or sudden intelligence explosion, but intelligence much greater than human is there, as well as immortality.

In 1996, Yudkowsky predicted a singularity by 2021.[20] His version of singularity involves intelligence explosion: once AIs are doing the research to improve themselves, speed doubles after 2 years, then 1 one year, then after 6 months, then after 3 months, then after 1.5 months, and after more iterations, the "singularity" is reached.[20] This construction implies that the speed reaches infinity in finite time.

In 2000, Bill Joy, a prominent technologist and a co-founder of Sun Microsystems, voiced concern over the potential dangers of robotics, genetic engineering, and nanotechnology.[52]

In 2005, Kurzweil published *The Singularity Is Near*. Kurzweil's publicity campaign included an appearance on *The Daily Show with Jon Stewart*.[131]

From 2006 to 2012, an annual Singularity Summit conference was organized by Machine Intelligence Research Institute, founded by Eliezer Yudkowsky.

In 2007, Yudkowsky suggested that many of the varied definitions that have been assigned to "singularity" are mutually incompatible rather than mutually supporting.[26][132] For example, Kurzweil extrapolates current technological trajectories past the arrival of self-improving AI or superhuman intelligence, which Yudkowsky argues represents a tension with both I. J. Good's proposed discontinuous upswing in intelligence and Vinge's thesis on unpredictability.[26]

In 2009, Kurzweil and X-Prize founder Peter Diamandis announced the establishment of Singularity University, a nonaccredited private institute whose stated mission is "to educate, inspire and empower leaders to apply exponential technologies to address humanity's grand challenges."[133] Funded by Google, Autodesk, ePlanet Ventures, and a group of technology industry leaders, Singularity University is based at NASA's Ames Research Center in Mountain View, California. The not-for-profit organization runs an annual ten-week graduate program during summer that covers ten different technology and allied tracks, and a series of executive programs throughout the year.

In politics [edit]

In 2007, the Joint Economic Committee of the United States Congress released a report about the future of nanotechnology. It predicts significant technological and political changes in the mid-term future, including possible technological singularity.[134][135][136]

Former President of the United States Barack Obama spoke about singularity in his interview to *Wired* in 2016:[137]

> One thing that we haven't talked about too much, and I just want to go back to, is we really have to think through the economic implications. Because most people aren't spending a lot of time right now worrying about singularity—they are worrying about "Well, is my job going to be replaced by a machine?"

See also [edit]

- Accelerating change – Perceived increase in the rate of technological change throughout history
- Artificial consciousness – Field in cognitive science
- Artificial intelligence arms race – Arms race for the most advanced AI-related technologies
- Artificial intelligence in fiction – Artificial intelligence as depicted in works of fiction
- Brain simulation – Concept of creating a functioning computer model of a brain or part of a brain
- Brain–computer interface – Direct communication pathway between an enhanced or wired brain and an external device
- Emerging technologies – Technologies whose development, practical applications, or both are still largely unrealized
- Ephemeralization – Technological advancement theory
- Fermi paradox – Lack of evidence that aliens exist
- Flynn effect – 20th-century rise in intelligence test scores
- Futures studies – Study of postulating possible, probable, and preferable futures
- Global brain – Futuristic concept of a global interconnected network
- Human intelligence § Improving intelligence
- Mind uploading – Hypothetical process of digitally emulating a brain
- Neuroenhancement
- Robot learning – Machine learning for robots
- Singularitarianism – Belief in an incipient technological singularity
- Technological determinism – Reductionist theory
- Technological revolution – Period of rapid technological change
- Technological unemployment – Unemployment caused by technological change
- Transhumanism – Philosophical movement
- The Man Who Evolved – short story by Edmond Hamilton

References [edit]

Citations [edit]

1. ^ Cadwalladr, Carole (22 February 2014). "Are the robots about to rise? Google's new director of engineering thinks so…". *The Guardian*. Retrieved 8 May 2022.
2. ^ "Collection of sources defining "singularity" ". *singularitysymposium.com*. Archived from the original on 17 April 2019. Retrieved 17 April 2019.
3. ^ a b Eden, Amnon H.; Moor, James H.; Søraker, Johnny H.; Steinhart, Eric, eds. (2012). *Singularity Hypotheses: A Scientific and Philosophical Assessment*. The Frontiers Collection. Dordrecht: Springer. pp. 1–2. doi:10.1007/978-3-642-32560-1. ISBN 9783642325601.
4. ^ a b c d e f g h Vinge, Vernor. "The Coming Technological Singularity: How to Survive in the Post-Human Era" Archived 2018-04-10 at the Wayback Machine, in *Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, G. A. Landis, ed., NASA Publication CP-10129, pp. 11–22, 1993.
5. ^ *The Technological Singularity* by Murray Shanahan, (MIT Press, 2015), p. 233.
6. ^ a b c Ulam, Stanislaw (May 1958). "Tribute to John von Neumann" (PDF). *Bulletin of the American Mathematical Society*. 64, #3, part 2: 5. Archived (PDF) from the original on 15 February 2021. Retrieved 7 November 2018.
7. ^ a b c d e f g Chalmers, David J. (2010). "The Singularity: A Philosophical Analysis" (PDF). *Journal of Consciousness Studies*. **17** (9–10): 7–65.
8. ^ a b Dooling, Richard. *Rapture for the Geeks: When AI Outsmarts IQ* (2008), p. 88
9. ^ Sparkes, Matthew (13 January 2015). "Top scientists call for caution over artificial intelligence". *The Telegraph (UK)*. Archived from the original on 7 April 2015. Retrieved 24 April 2015.
10. ^ "Hawking: AI could end human race". BBC. 2 December 2014. Archived from the original on 30 October 2015. Retrieved 11 November 2017.
11. ^ a b c *Paul Allen: The Singularity Isn't Near*, retrieved 12 April 2015
12. ^ a b c d e f g h i "Tech Luminaries Address Singularity – IEEE Spectrum". *Spectrum.ieee.org*. 1 June 2008. Archived from the original on 30 April 2019. Retrieved 9 September 2011.
13. ^ a b c Modis, Theodore (2012). "Why the Singularity Cannot Happen". Published in Eden, Amnon H. et al (Eds.) (2012). *Singularity Hypothesis* (PDF). New York: Springer. p. 311. ISBN 978-3-642-32560-1. pp. 311–339.
14. ^ Ehrlich, Paul. "Paul Ehrlich: The Dominant Animal: Human Evolution and the Environment – The Long Now". *longnow.org*. Retrieved 14 June 2023.
15. ^ "Businessweek – Bloomberg". *Bloomberg.com*. 20 April 2023. Retrieved 14 June 2023.
16. ^ a b Yampolskiy, Roman V. "Analysis of types of self-improving software." Artificial General Intelligence. Springer International Publishing, 2015. pp. 384–393.
17. ^ a b Eliezer Yudkowsky. General Intelligence and Seed AI-Creating Complete Minds Capable of Open-Ended Self-Improvement, 2001
18. ^ a b c Good, I. J. (1965), *Speculations Concerning the First Ultraintelligent Machine*, archived from the original on 27 May 2001
19. ^ a b Good, I. J. (1965), *Speculations Concerning the First Ultraintelligent Machine* (PDF), archived from the original (PDF) on 1 May 2012
20. ^ a b c d e Yudkowsky, Eliezer (1996). "Staring into the Singularity". Archived from the original on 21 October 2008.
21. ^ Ray Kurzweil, The Singularity Is Near, pp. 135–136. Penguin Group, 2005.
22. ^ Pearce, David (2012), Eden, Amnon H.; Moor, James H.; Søraker, Johnny H.; Steinhart, Eric (eds.), "The Biointelligence Explosion", *Singularity Hypotheses*, The Frontiers Collection, Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 199–238, doi:10.1007/978-3-642-32560-1_11, ISBN 978-3-642-32559-5, retrieved 16 January 2022
23. ^ Gouveia, Steven S., ed. (2020). "ch. 4, "Humans and Intelligent Machines: Co-evolution, Fusion or Replacement?", David Pearce". *The Age of Artificial Intelligence: An Exploration*. ISBN 978-1-62273-872-4.
66. ^ Ford, Martin, *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future Archived 2010-09-06 at the Wayback Machine*, Acculant Publishing, 2009, ISBN 978-1-4486-5981-4
67. ^ Markoff, John (4 March 2011). "Armies of Expensive Lawyers, Replaced by Cheaper Software". *The New York Times*. Archived from the original on 15 February 2021. Retrieved 25 February 2017.
68. ^ Modis, Theodore (2002) "Forecasting the Growth of Complexity and Change" Archived 2021-02-15 at the Wayback Machine, *Technological Forecasting & Social Change*, 69, No 4, 2002, pp. 377 – 404
69. ^ a b Huebner, Jonathan (2005) "A Possible Declining Trend for Worldwide Innovation" Archived 2015-03-22 at the Wayback Machine, *Technological Forecasting & Social Change*, October 2005, pp. 980–6
70. ^ Krazit, Tom. Intel pledges 80 cores in five years Archived 2021-02-15 at the Wayback Machine, *CNET News*, 26 September 2006.
71. ^ Modis, Theodore (2020). "Forecasting the Growth of Complexity and Change—An Update". Published in Korotayev, Andrey; LePoire, David (Eds.) (3 January 2020). *The 21st Century Singularity and Global Futures* (1 ed.). Springer. p. 620. ISBN 978-3-030-33730-8. pp 101-104
72. ^ Modis, Theodore (May–June 2003). "The Limits of Complexity and Change". The Futurist. 37 (3): 26-32.
73. ^ Modis, Theodore (2006) "The Singularity Myth" Archived 2021-02-15 at the Wayback Machine, *Technological Forecasting & Social Change*, February 2006, pp. 104 - 112
74. ^ Modis, Theodore (1 March 2022). "Links between entropy, complexity, and the technological singularity". *Technological Forecasting and Social Change*. **176**: 121457. doi:10.1016/j.techfore.2021.121457. ISSN 0040-1625. S2CID 245663426.
75. ^ Schmidhuber, Jürgen (2006), *New millennium AI and the convergence of history*, arXiv:cs/0606081
76. ^ Tainter, Joseph (1988) "The Collapse of Complex Societies Archived 2015-06-07 at the Wayback Machine" (Cambridge University Press)
77. ^ Trying to Muse Rationally About the Singularity Scenario by Douglas Hofstadter, 2006, unauthorized transcript
78. ^ a b Lanier, Jaron (2013). "Who Owns the Future?". *New York: Simon & Schuster*. Archived from the original on 13 May 2016. Retrieved 2 March 2016.
79. ^ William D. Nordhaus, "Why Growth Will Fall" (a review of Robert J. Gordon, *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton University Press, 2016, ISBN 978-0691147727, 762 pp., $39.95), *The New York Review of Books*, vol. LXIII, no. 13 (August 18, 2016), p. 68.
80. ^ Cadwalladr, Carole (12 February 2017), "Daniel Dennett: 'I begrudge every hour I have to spend worrying about politics' ", *The Guardian*
81. ^ Myers, PZ, *Singularly Silly Singularity*, archived from the original on 28 February 2009, retrieved 13 April 2009
82. ^ Kelly, Kevin (2006). "The Singularity Is Always Near". *The Technium*. Retrieved 14 June 2023.
83. ^ Beam, Alex (24 February 2005). "That Singularity Sensation". *The Boston Globe*. Retrieved 15 February 2013.
84. ^ Gray, John (24 November 2011). "On the Road to Immortality". *The New York Review of Books*. Retrieved 19 March 2013.
85. ^ Streitfiled, David (11 June 2023). "Silicon Valley Confronts the Idea That the 'Singularity' Is Here". *New York Times*. Retrieved 11 June 2023.
86. ^ Hanson, Robin (1 June 2008), "Economics Of The Singularity", *IEEE Spectrum Special Report: The Singularity*, archived from the original on 11 August 2011, retrieved 25 July 2009 & Long-Term Growth As A Sequence of Exponential Modes Archived 2019-05-27 at the Wayback Machine
87. ^ a b Yudkowsky, Eliezer (2008), Bostrom, Nick; Cirkovic, Milan (eds.), "Artificial Intelligence as a Positive and Negative Factor in Global Risk" (PDF), *Global Catastrophic Risks*, Oxford University Press: 303,

24. ^ Hanson, Robin (2016). *The Age of Em*. Oxford: Oxford University Press. p. 528. ISBN 9780198754626.
25. ^ Hall, Josh (2010). "Singularity: Nanotech or AI?". Hplusmagazine.com. Archived from the original on 23 December 2010. Retrieved 9 September 2011.
26. ^ a b c Yudkowsky, Eliezer (2007), *The Singularity: Three Major Schools*, archived from the original on 1 October 2018
27. ^ Sandberg, Anders. An overview of models of technological singularity Archived 2011-07-24 at the Wayback Machine
28. ^ Sotala, Kaj; Yampolskiy, Roman (2017). "Risks of the Journey to the Singularity". *The Technological Singularity*. The Frontiers Collection. Springer Berlin Heidelberg. pp. 11–23. doi:10.1007/978-3-662-54033-6_2. ISBN 978-3-662-54031-2.
29. ^ a b "What is the Singularity? | Singularity Institute for Artificial Intelligence". Singinst.org. Archived from the original on 8 September 2011. Retrieved 9 September 2011.
30. ^ Chalmers, David J. (2016). "The Singularity". *Science Fiction and Philosophy*. John Wiley & Sons, Inc. pp. 171–224. doi:10.1002/9781118922590.ch16. ISBN 9781118922590.
31. ^ a b c Hans Moravec, *Mind Children*, 1988
32. ^ Moravec, Hans P. (1998). *Robot: Mere Machine to Transcendent Mind*. Oup Usa.
33. ^ Reedy, Christianna (2017). "Kurzweil Claims That the Singularity Will Happen by 2045". *Futurism*. Retrieved 14 June 2023.
34. ^ "List of Analyses of Time to Human-Level AI". *AI Impacts*. 22 January 2015. Retrieved 14 June 2023.
35. ^ Khatchadourian, Raffi (16 November 2015). "The Doomsday Invention". *The New Yorker*. Archived from the original on 29 April 2019. Retrieved 31 January 2018.
36. ^ Müller, V. C., & Bostrom, N. (2016). "Future progress in artificial intelligence: A survey of expert opinion". In V. C. Müller (ed): *Fundamental issues of artificial intelligence* (pp. 555–572). Springer, Berlin. http://philpapers.org/rec/MLLFPI Archived 2019-03-02 at the Wayback Machine
37. ^ "Who's Who In The Singularity – IEEE Spectrum". Spectrum.ieee.org. 1 June 2008. Archived from the original on 12 March 2016. Retrieved 9 September 2011.
38. ^ Hanson, Robin (1998). "Some Skepticism". Archived from the original on 15 February 2021. Retrieved 8 April 2020.
39. ^ David Chalmers John Locke Lecture, 10 May, Exam Schools, Oxford, Presenting a philosophical analysis of the possibility of a technological singularity or "intelligence explosion" resulting from recursively self-improving AI Archived 2013-01-15 at the Wayback Machine.
40. ^ "ITRS" (PDF). Archived from the original (PDF) on 29 September 2011. Retrieved 9 September 2011.
41. ^ Shulman, Carl; Anders, Sandberg. "Implications of a Software-Limited Singularity" (PDF). *intelligence.org*. Machine Intelligence Research Institute. Retrieved 21 March 2021.
42. ^ Grace, Katja; Salvatier, John; Dafoe, Allan; Zhang, Baobao; Evans, Owain (24 May 2017). "When Will AI Exceed Human Performance? Evidence from AI Experts". arXiv:1705.08807 [cs.AI].
43. ^ Siracusa, John (31 August 2009). "Mac OS X 10.6 Snow Leopard: the Ars Technica review". *Ars Technica*. Archived from the original on 3 September 2011. Retrieved 9 September 2011.
44. ^ Moravec, Hans (1999). *Robot: Mere Machine to Transcendent Mind*. Oxford U. Press. p. 61. ISBN 978-0-19-513630-2.
45. ^ Ray Kurzweil, *The Age of Spiritual Machines*, Viking; 1999, ISBN 978-0-14-028202-3. pp. 30, 32 Archived 2021-02-15 at the Wayback Machine
46. ^ a b Ray Kurzweil, *The Singularity Is Near*, Penguin Group, 2005
47. ^ a b "The World's Technological Capacity to Store, Communicate, and Compute Information" Archived 2013-07-27 at the Wayback Machine, Martin Hilbert and Priscila López (2011), Science, 332(6025), 60–65; free access to the article through here: martinhilbert.net/WorldInfoCapacity.html
48. ^ Korotayev, Andrey V.; LePoire, David J., eds. (2020). "The 21st Century Singularity and Global Futures". *World-Systems Evolution and Global Futures*. doi:10.1007/978-3-030-33730-8. ISBN 978-3-030-33729-2. ISSN 2522-0985. S2CID 241407141.
49. ^ Ray Kurzweil, The Singularity Is Near, p. 9. Penguin Group, 2005
50. ^ Ray Kurzweil, *The Singularity Is Near*, pp. 135–136. Penguin Group, 2005. "So we will be producing about 10^{26} to 10^{29} cps of nonbiological computation per year in the early 2030s. This is roughly equal to our estimate for the capacity of all living biological human intelligence ... This state of computation in the early 2030s will not represent the Singularity, however, because it does not yet correspond to a profound expansion of our intelligence. By the mid-2040s, however, that one thousand dollars' worth of computation will be equal to 10^{26} cps, so the intelligence created per year (at a total cost of about $\$10^{12}$) will be about one billion times more powerful than all human intelligence today. That *will* indeed represent a profound change, and it is for that reason that I set the date for the Singularity—representing a profound and disruptive transformation in human capability—as 2045."
51. ^ Kurzweil, Raymond (2001), "The Law of Accelerating Returns", *Nature Physics*, Lifeboat Foundation, **4** (7): 507, Bibcode:2008NatPh...4..507B, doi:10.1038/nphys1010, archived from the original on 27 August 2018, retrieved 7 August 2007
52. ^ a b Joy, Bill (April 2000), "Why The Future Doesn't Need Us", *Wired Magazine*, Viking Adult, vol. 8, no. 4, ISBN 978-0-670-03249-5, archived from the original on 5 February 2009, retrieved 7 August 2007, "Our most powerful 21st-century technologies – robotics, genetic engineering, and nanotech – are threatening to make humans an endangered species."
53. ^ Yudkowsky, Eliezer S. "Power of Intelligence". Yudkowsky. Archived from the original on 3 October 2018. Retrieved 9 September 2011.
54. ^ a b "Omohundro, Stephen M., "The Basic AI Drives." Artificial General Intelligence, 2008 proceedings of the First AGI Conference, eds. Pei Wang, Ben Goertzel, and Stan Franklin. Vol. 171. Amsterdam: IOS, 2008". 30 November 2007. Archived from the original on 17 September 2018. Retrieved 20 August 2010.
55. ^ "Artificial General Intelligence: Now Is the Time". KurzweilAI. Archived from the original on 4 December 2011. Retrieved 9 September 2011.
56. ^ a b "Omohundro, Stephen M., "The Nature of Self-Improving Artificial Intelligence." Self-Aware Systems. 21 Jan. 2008. Web. 07 Jan. 2010". 6 October 2007. Archived from the original on 12 June 2018. Retrieved 20 August 2010.
57. ^ Barrat, James (2013). "6, "Four Basic Drives" ". *Our Final Invention* (First ed.). New York: St. Martin's Press. pp. 78–98. ISBN 978-0312622374.
58. ^ "Max More and Ray Kurzweil on the Singularity". KurzweilAI. Archived from the original on 21 November 2018. Retrieved 9 September 2011.
59. ^ "Concise Summary | Singularity Institute for Artificial Intelligence". Singinst.org. Archived from the original on 21 June 2011. Retrieved 9 September 2011.
60. ^ a b "Bostrom, Nick, The Future of Human Evolution, Death and Anti-Death: Bibcode:2008gcr..book..303Y, ISBN 978-0-19-857050-9, archived from the original (PDF) on 7 August 2008
88. ^ "The Uncertain Future". *theuncertainfuture.com; a future technology and world-modeling project*. Archived from the original on 30 April 2019. Retrieved 17 August 2010.
89. ^ Anders Sandberg and NickBostrom (2008). "Global Catastrophic Risks Survey (2008) Technical Report 2008/1" (PDF). Future of Humanity Institute. Archived from the original (PDF) on 16 May 2011.
90. ^ "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards". *nickbostrom.com*. 2002. Archived from the original on 27 April 2011. Retrieved 25 January 2006.
91. ^ a b c Hawking, Stephen (1 May 2014). "Stephen Hawking: 'Transcendence looks at the implications of artificial intelligence – but are we taking AI seriously enough?' ". *The Independent*. Archived from the original on 25 September 2015. Retrieved 5 May 2014.
92. ^ Nick Bostrom, "Ethical Issues in Advanced Artificial Intelligence" Archived 2018-10-08 at the Wayback Machine, in *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, Vol. 2, ed. I. Smit et al., Int. Institute of Advanced Studies in Systems Research and Cybernetics, 2003, pp. 12–17
93. ^ Eliezer Yudkowsky: Artificial Intelligence as a Positive and Negative Factor in Global Risk Archived 2012-06-11 at the Wayback Machine. Draft for a publication in *Global Catastrophic Risk* from August 31, 2006, retrieved July 18, 2011 (PDF file)
94. ^ "The Stamp Collecting Device, Nick Hay". Archived from the original on 17 June 2012. Retrieved 21 August 2010.
95. ^ "Andart: Why we should fear the Paperclipper". *www.aleph.se*. Retrieved 14 June 2023.
96. ^ "Omohundro, Stephen M., "The Basic AI Drives." Artificial General Intelligence, 2008 proceedings of the First AGI Conference, eds. Pei Wang, Ben Goertzel, and Stan Franklin. Vol. 171. Amsterdam: IOS, 2008". 30 November 2007. Archived from the original on 17 September 2018. Retrieved 20 August 2010.
97. ^ de Garis, Hugo. "The Coming Artilect War". *Forbes*. Retrieved 14 June 2023.
98. ^ Coherent Extrapolated Volition, Eliezer S. Yudkowsky, May 2004 Archived 2010-08-15 at the Wayback Machine
99. ^ Hibbard, Bill (2012), "Model-Based Utility Functions", *Journal of Artificial General Intelligence*, **3** (1): 1, arXiv:1111.3934, Bibcode:2012JAGI....3....1H, doi:10.2478/v10229-011-0013-5, S2CID 8434596.
100. ^ a b Avoiding Unintended AI Behaviors. Archived 2013-06-29 at the Wayback Machine Bill Hibbard. 2012 proceedings of the Fifth Conference on Artificial General Intelligence, eds. Joscha Bach, Ben Goertzel and Matthew Ikle. This paper won the Machine Intelligence Research Institute's 2012 Turing Prize for the Best AGI Safety Paper Archived 2021-02-15 at the Wayback Machine.
101. ^ Hibbard, Bill (2008), "The Technology of Mind and a New Social Contract", *Journal of Evolution and Technology*, **17**, archived from the original on 15 February 2021, retrieved 5 January 2013.
102. ^ Decision Support for Safe AI Designl. Archived 2021-02-15 at the Wayback Machine Bill Hibbard. 2012 proceedings of the Fifth Conference on Artificial General Intelligence, eds. Joscha Bach, Ben Goertzel and Matthew Ikle.
103. ^ a b c Kemp, D. J.; Hilbert, M.; Gillings, M. R. (2016). "Information in the Biosphere: Biological and Digital Worlds". *Trends in Ecology & Evolution*. **31** (3): 180–189. doi:10.1016/j.tree.2015.12.013. PMID 26777788. S2CID 3561873. Archived from the original on 4 June 2016. Retrieved 24 May 2016.
104. ^ Hilbert, Martin. "Information Quantity" (PDF).
105. ^ a b Markoff, John (26 July 2009). "Scientists Worry Machines May Outsmart Man". *The New York Times*. Archived from the original on 1 July 2017.
106. ^ Robinson, Frank S. (27 June 2013). "The Human Future: Upgrade or Replacement?". *The Humanist*. Archived from the original on 15 February 2021. Retrieved 1 May 2020.
107. ^ Eliezer Yudkowsky. "Artificial intelligence as a positive and negative factor in global risk." Global catastrophic risks (2008).
108. ^ Bugaj, Stephan Vladimir, and Ben Goertzel. "Five ethical imperatives and their implications for human-AGI interaction." Dynamical Psychology (2007).
109. ^ Sotala, Kaj, and Roman V. Yampolskiy. "Responses to catastrophic AGI risk: a survey." Physica Scripta 90.1 (2014): 018001.
110. ^ Naam, Ramez (2014). "The Singularity Is Further Than It Appears". Archived from the original on 17 May 2014. Retrieved 16 May 2014.
111. ^ Naam, Ramez (2014). "Why AIs Won't Ascend in the Blink of an Eye – Some Math". Archived from the original on 17 May 2014. Retrieved 16 May 2014.
112. ^ Hall, J. Storrs (2008). "Engineering Utopia" (PDF). *Artificial General Intelligence, 2008: Proceedings of the First AGI Conference*: 460–467. Archived (PDF) from the original on 1 December 2014. Retrieved 16 May 2014.
113. ^ Goertzel, Ben (26 September 2014). "Superintelligence — Semi-hard Takeoff Scenarios". h+ Magazine. Archived from the original on 25 October 2014. Retrieved 25 October 2014.
114. ^ More, Max. "Singularity Meets Economy". Archived from the original on 28 August 2009. Retrieved 10 November 2014.
115. ^ K. Eric Drexler, *Engines of Creation*, 1986
116. ^ Feynman, Richard P. (December 1959). "There's Plenty of Room at the Bottom". Archived from the original on 11 February 2010.
117. ^ Ray Kurzweil, The Singularity Is Near, p. 215. Penguin Group, 2005.
118. ^ *The Singularity Is Near*, p. 216.
119. ^ Lanier, Jaron (2010). *You Are Not a Gadget: A Manifesto*. New York, New York: Alfred A. Knopf. p. 26. ISBN 978-0307269645.
120. ^ Prasad, Mahendra (2019). "Nicolas de Condorcet and the First Intelligence Explosion Hypothesis". *AI Magazine*. **40** (1): 29–33. doi:10.1609/aimag.v40i1.2855.
121. ^ Intelligent machines: How to get there from here and What to do afterwards by Hans Moravec, 1977 (wikidata)
122. ^ a b c Smart, John (1999–2008), *A Brief History of Intellectual Discussion of Accelerating Change*
123. ^ Today's Computers, Intelligent Machines and Our Future by Hans Moravec, 1979, wikidata
124. ^ Solomonoff, R.J. "The Time Scale of Artificial Intelligence: Reflections on Social Effects", Human Systems Management, Vol 5, pp. 149–153, 1985.
125. ^ Vinge, Vernor (1 October 2004). *Marooned in Realtime*. Macmillan. ISBN 978-1-4299-1512-0.
126. ^ David Pringle (28 September 1986). "Time and Time Again". *The Washington Post*. Retrieved 6 July 2021.
127. ^ Vinge, Vernor (1988). *Threats and Other Promises*. Baen. ISBN 978-0-671-69790-7.

Two Hundred Years After Kant, Fifty Years After Turing, ed. Charles Tandy, pp. 339–371, 2004, Ria University Press". Archived from the original on 28 August 2018. Retrieved 20 August 2010.

61. ^ Shulman, Carl; Sandberg, Anders (2010). Mainzer, Klaus (ed.). "Implications of a Software-Limited Singularity" (PDF). *ECAP10: VIII European Conference on Computing and Philosophy*. Archived (PDF) from the original on 30 April 2019. Retrieved 17 May 2014.
62. ^ Muehlhauser, Luke; Salamon, Anna (2012). "Intelligence Explosion: Evidence and Import" (PDF). In Eden, Amnon; Søraker, Johnny; Moor, James H.; Steinhart, Eric (eds.). *Singularity Hypotheses: A Scientific and Philosophical Assessment*. Springer. Archived (PDF) from the original on 26 October 2014. Retrieved 28 August 2018.
63. ^ Dreyfus & Dreyfus (2000). *Mind Over Machine*. p. xiv. ISBN 9780743205511.: 'The truth is that human intelligence can never be replaced with machine intelligence simply because we are not ourselves "thinking machines" in the sense in which that term is commonly understood.'
64. ^ John R. Searle, "What Your Computer Can't Know", *The New York Review of Books*, 9 October 2014, p. 54.: "[Computers] have, literally ..., no intelligence, no motivation, no autonomy, and no agency. We design them to behave as if they had certain sorts of psychology, but there is no psychological reality to the corresponding processes or behavior. ... [T]he machinery has no beliefs, desires, [or] motivations."
65. ^ Hawking, Stephen (2018). *Brief Answers to the Big Questions*. New York: Bantam Books. p. 161. ISBN 9781984819192. "Some people say that computers can never show true intelligence whatever that may be. But it seems to me that if very complicated chemical molecules can operate in humans to make them intelligent then equally complicated electronic circuits can also make computers act in an intelligent way. And if they are intelligent they can presumably design computers that have even greater complexity and intelligence."
128. ^ Dooling, Richard. *Rapture for the Geeks: When AI Outsmarts IQ* (2008), p. 89
129. ^ "Will Robots Inherit the Earth?". *web.media.mit.edu*. Retrieved 14 June 2023.
130. ^ *Colonizing the Galaxies* by Graham Oppy
131. ^ *Episode dated 23 August 2006* at IMDb
132. ^ Sandberg, Anders. "An overview of models of technological singularity." Roadmaps to AGI and the Future of AGI Workshop, Lugano, Switzerland, March. Vol. 8. 2010.
133. ^ Singularity University Archived 2021-02-15 at the Wayback Machine at its official website
134. ^ Guston, David H. (14 July 2010). *Encyclopedia of Nanoscience and Society*. SAGE Publications. ISBN 978-1-4522-6617-6. Archived from the original on 15 February 2021. Retrieved 4 November 2016.
135. ^ "Nanotechnology: The Future is Coming Sooner Than You Think" (PDF). Joint Economic Committee United States Congress. March 2007. Archived (PDF) from the original on 15 February 2021. Retrieved 29 April 2012.
136. ^ Treder, Mike (31 March 2007). "Congress and the Singularity". Archived from the original on 7 April 2007. Retrieved 4 November 2016.
137. ^ Dadich, Scott (12 October 2016). "Barack Obama Talks AI, Robo Cars, and the Future of the World". *Wired*. Archived from the original on 3 December 2017. Retrieved 4 November 2016.

Sources [edit]

- Kurzweil, Ray (2005). *The Singularity Is Near*. New York, New York: Penguin Group. ISBN 9780715635612.
- William D. Nordhaus, "Why Growth Will Fall" (a review of Robert J. Gordon, *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton University Press, 2016.ISBN 978-0691147727, 762 pp., $39.95), *The New York Review of Books*, vol. LXIII, no. 13 (August 18, 2016), pp. 64, 66, 68.
- John R. Searle, "What Your Computer Can't Know" (review of Luciano Floridi, *The Fourth Revolution: How the Infosphere Is Reshaping Human Reality*, Oxford University Press, 2014; and Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, 2014), *The New York Review of Books*, vol. LXI, no. 15 (October 9, 2014), pp. 52–55.
- Good, I. J. (1965), "Speculations Concerning the First Ultraintelligent Machine", in Franz L. Alt; Morris Rubinoff (eds.), *Advances in Computers Volume 6*, Advances in Computers, vol. 6, Academic Press, pp. 31–88, doi:10.1016/S0065-2458(08)60418-0, hdl:10919/89424, ISBN 9780120121069, archived from the original on 27 May 2001, retrieved 7 August 2007
- Hanson, Robin (1998), *Some Skepticism*, Robin Hanson, archived from the original on 28 August 2009, retrieved 19 June 2009
- Berglas, Anthony (2008), *Artificial Intelligence will Kill our Grandchildren*, archived from the original on 23 July 2014, retrieved 13 June 2008
- Bostrom, Nick (2002), "Existential Risks", *Journal of Evolution and Technology*, **9**, archived from the original on 27 April 2011, retrieved 7 August 2007
- Hibbard, Bill (5 November 2014). "Ethical Artificial Intelligence". arXiv:1411.1373 [cs.AI].

Further reading [edit]

- Krüger, Oliver, *Virtual Immortality. God, Evolution, and the Singularity in Post- and Transhumanism.*, Bielefeld: transcript 2021. ISBN 978-3-8376-5059-4
- Marcus, Gary, "Am I Human?: Researchers need new ways to distinguish artificial intelligence from the natural kind", *Scientific American*, vol. 316, no. 3 (March 2017), pp. 58–63. *Multiple* tests of artificial-intelligence efficacy are needed because, "just as there is no single test of athletic prowess, there cannot be one ultimate test of intelligence." One such test, a "Construction Challenge", would test perception and physical action—"two important elements of intelligent behavior that were entirely absent from the original Turing test." Another proposal has been to give machines the same standardized tests of science and other disciplines that schoolchildren take. A so far insuperable stumbling block to artificial intelligence is an incapacity for reliable disambiguation. "[V]irtually every sentence [that people generate] is ambiguous, often in multiple ways." A prominent example is known as the "pronoun disambiguation problem": a machine has no way of determining to whom or what a pronoun in a sentence—such as "he", "she" or "it"—refers.
- Scaruffi, Piero, "Intelligence is not Artificial" (2016) for a critique of the singularity movement and its similarities to religious cults.

External links [edit]

- singularity | technology, britannica.com
- The Coming Technological Singularity: How to Survive in the Post-Human Era (on Vernor Vinge's web site, retrieved Jul 2019)
- Intelligence Explosion FAQ by the Machine Intelligence Research Institute
- Blog on bootstrapping artificial intelligence by Jacques Pitrat
- *Why an Intelligence Explosion is Probable* (Mar 2011)
- *Why an Intelligence Explosion is Impossible* (Nov 2017)
- *How Close are We to Technological Singularity and When?*

Listen to this article (52 minutes)

52:26

This audio file was created from a revision of this article dated 3 November 2018, and does not reflect subsequent edits.

(Audio help · More spoken articles)

V·T·E	**Existential risk from artificial intelligence**	[show]
V·T·E	**Emerging technologies**	[show]
V·T·E	**Global catastrophic risks**	[show]
Authority control: National	United States	

Categories: Singularitarianism | Existential risk from artificial general intelligence | Philosophy of artificial intelligence

This page was last edited on 11 July 2023, at 08:39 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License 4.0; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.

Privacy policy About Wikipedia Disclaimers Contact Wikipedia Code of Conduct Mobile view Developers Statistics Cookie statement