

Intelligent Machines vs. Human Intelligence



Guide to Critical Analysis

This guide is intended to help you evaluate the topics presented in EBSCO Information Services' *Points of View Reference Center*[™]. Using this analysis will enhance your ability to read critically, develop your own perspective on the issue of intelligent machines, and write or debate an effective argument on the topic.

Before You Read

- In a single sentence, state what you currently think is/are the argument(s) in the intelligent machines vs. human intelligence debate.
- On the opening screen, Browse by Category, click on "People & Society," and then click on "Intelligent Machines vs. Human Intelligence."
- View the Result List of four records on the Points of View tab that includes this guide.
- Each topic in *Points of View Reference Center*[™] has three main Points of View essays. While there may be many opinions on the issue, these essays are designed to provide a base from which you may engage in further analysis. The Overview provides general background information on the issue. The Point and Counterpoint essays present opposing perspectives on the issue, with talking points that support those positions.

As You Read

Reading all three essays will enable you to better understand the controversy, providing a foundation from which you can broaden your knowledge of the topic.

Read the Overview, Point and Counterpoint essays. You will be able to move between the Overview, the Point and the Counterpoint essays by clicking on the colored rectangles in the sidebar featuring the Points of View icon.

Take notes on the main ideas. Start by printing out the essay and highlighting the main ideas and circling key words.

Review the terms in the Understanding the Discussion section of the Overview.

Organize your ideas. Go to the Research Guide "Organize Your Ideas Visually." You can access it from the Research Guides link in the Reference Shelf section on the home page. Use one of the graphic organizer templates provided to break the essay down into its main points.

Judge Fact and Opinion

Review the research guide article titled "Judge Fact vs. Opinion."

FACTS

- Facts are statements that can be proved true or false.
- Facts tell what actually happened.
- Facts tell what is happening now.
- Facts state something that can be easily observed or verified.

OPINIONS

- Opinions are statements that cannot be proved true or false because they express a person's thoughts, beliefs, feelings, or estimates.
- Opinions express worth or value.
- Opinions express what the author or speaker think should or should not be thought or done.
- Opinions are based on what seems true or probable.

Note the difference between the following facts and opinions drawn from the Overview, Point or Counterpoint articles, or bibliographic sources referenced in those articles:

Example 1:

Fact: We use AI in many ways today--from computer games to digital personal assistants to self-driving cars.

Opinion: Greater use of AI will be even more beneficial to humanity.

In this example, the opinion speculates about a future outcome that cannot yet be known.

Example 2:

Fact: Some jobs have been lost through automation in the past.

Opinion: The use of intelligent machines will replace human jobs and drive down wages for human workers.

In this example, the author is generalizing without substantiating evidence.

Example 3:

Fact: The speed, performance, and capacity of computers have increased at an exponential rate.

Opinion: The speed, performance, and capacity of computers will continue to increase at an exponential rate.

In this example, the author is extrapolating from a current trend without additional evidence to support it.

Example 4:

Fact: Intelligent machines have remained under the control of humanity.

Opinion: Superintelligent machines are the future of humanity.

In this example, the opinion expresses the author's personal belief.

Based on the criteria listed above, select:

- Three fact statements from the Overview, Point, or Counterpoint articles (or one of the bibliographic sources within those articles) on intelligent machines vs. human intelligence.
- Three opinion statements from the Overview, Point, or Counterpoint articles (or one of the bibliographic sources within those articles) on intelligent machines vs. human intelligence.

Intelligent Machines vs. Human Intelligence in Depth

You will see several links in the Related Items box on the right hand side of the screen. Exploring these added features will enhance your understanding of the topic you're studying.

- Click on one of the photos in the Related Items (if applicable) and record your reaction to the picture you chose. Did viewing this picture change your current view on Uber? How?
- Under the "What Experts Say" tab, click on the link and read the article, if one is provided. Does the author of this article make any of the same arguments outlined in the Talking Points sections of the Point or the Counterpoint essays? If the writer is supporting a new viewpoint or Talking Point, what is it?

- The Result List also offers many other kinds of research articles, such as Periodicals, Newspapers, Reference Books, etc. Click on one of these tabs and read one of the articles. What source tab did you use and what is one benefit of using this type of source? What is one drawback? What new information did you learn from reading this article?

You Decide

Based on the readings and review of the Related Items offered, consider the Point and Counterpoint essays in light of your own set of personal values and understanding of the topic.

- Reflect on which arguments about intelligent machines vs. human intelligence you accept, those outlined in the Point, Counterpoint, or a completely different argument. For you, what is the single most compelling argument regarding intelligent machines, and why?
- Write an 800- to 1,000-word essay response to the question above; include a thesis statement and 2-3 supporting arguments. See "Write a Thesis Statement" and "Write a Persuasive Essay" research guides for tips on how to write your essay.
- Use facts from other resources found in Points of View Reference Center™, other than the Overview, Point, and Counterpoint essays. See "Judge Fact vs. Opinion" in the Research Guide section to help determine the value of the information you find.
- Investigate some of the sources cited in the Bibliography at the end of each Point of View essay to review the kinds of sources those authors consulted.
- For website information, first review the research guide article titled "Evaluate a Website." Use the information you find there to guide your decisions on which sites provide the most credible information and would add value to your essay.
- Suggest ways in which the controversy surrounding intelligent machines could be resolved.

See these other *Points of View Reference Center*™ "How To" topics for additional help. They can be found using the Research Guides link on the home page.

- Write a Topic Sentence
- Take Good Notes
- Go From Notes to Outline
- Organize Your Ideas Visually
- Write a Conclusion
- Cite Sources
- Create a Bibliography
- Write a Research Paper

Debate It

- Debate the issue. To create an effective debate, refer to the tips offered in the research guides - "Use Strategy in a Debate," "Give an Oral Presentation," and "Start a Debate Club."
- Have someone else read an excerpt from a bibliographic source in either the Point or the Counterpoint essay on intelligent machines vs. human intelligence or within an article from the "What the Experts Say" in the Related Items section of a Point or Counterpoint article, if available. The excerpt should clearly express an argument for one side or other of the discussion. Imagine you are facing this person in a debate setting. Take notes on the points (contentions) you would make to oppose the arguments (rebuttals).
- When the speaker is finished, respond to the arguments presented with opposing arguments, pointing out any contradictions you may have heard.

Overview

Introduction

Artificial intelligence (AI) has penetrated many aspects of contemporary life, from computer games to autonomous vacuum cleaners to digital personal assistants. However, these AI programs are built for a specific purpose (to play card games, to clean the house, to direct the user to the nearest sushi restaurant, and so on), and not intended to or able to demonstrate general intelligence. Will machines one day be as intelligent, or much more intelligent, than human beings?

Some scientists in the AI field subscribe to the idea that at some point in the not-so-distant future, AI technology will reach and surpass human intelligence. The term "the singularity" refers to this point in time. There are many theories as to the likelihood of the singularity, what form it might take, and what impact it might have on human society and economy. Supporters of the singularity concept tend to point to the accelerating rate of technology as putting it within reach as early as 2030 or 2045. They argue that it would be beneficial for humanity to have intelligent and superintelligent machines to work on its behalf, and indeed people might see these machines as the future of humanity, rather than as a threat to it. Skeptics, on the other hand, doubt both the realizability and desirability of the singularity. They contend that the human mind is too complex and human intelligence too nuanced to be imitated or replaced by machines. Nor is it clear that the singularity, were it to happen, would be unambiguously beneficial to humanity: intelligent machines might begin with replacing human workers and end up replacing human beings altogether.

Understanding the Discussion

Artificial general intelligence (AGI): AI that possesses intelligence that is comparable to human-level intelligence. AGI is distinct from more narrowly focused AI programs, such as chess-playing programs or stock-trading programs.

Artificial intelligence (AI): Defined by John McCarthy, computer scientist and AI pioneer, as the scientific methods employed to create intelligent machine, particularly through the use of intelligent computer programs.

Moore's law: The observation, made by Gordon Moore (cofounder of Intel), that the number of transistors that can be fit onto an integrated circuit doubles roughly every two years.

Singularity: A point in the future when AI reaches and then surpasses human intelligence.

Superintelligence: A term coined by philosopher Nick Bostrom, to refer to intelligence that possesses cognition that significantly and consistently outstrips human cognition.

History

The idea of intelligent machines goes back at least Samuel Butler's 1872 satire of Victorian society, *Erewhon*, and his musings on the evolution of consciousness in machines as a parallel to the evolution of consciousness in human beings. Another significant literary exploration of the world of intelligent machines is the 1920 Czech play *RUR* (standing for "Rossum's Universal Robots" in English), which gave the term "robot" to the English language. The science-fiction play explores the benefits and dangers of using robots; in the play, robots begin by taking over human jobs and end by taking over the world.

English mathematician Alan Turing began discussing intelligent machines or, as he referred to them, "thinking machines," in the 1950s. Turing considered whether machines were capable of thought in his famous 1950 essay "Computing Machinery and Intelligence." He suggested that the question should be considered as answered in the affirmative if a human judge could not accurately distinguish between a machine and a man, after several minutes of conversation with both, at least 30 percent of the time. This process has become known as the Turing test and is still considered by many to be one of the milestones in the development of AI. The test has also been contested, and other researchers have developed more in-depth or alternative tests. One of the first mentions of a technological singularity came in a tribute to mathematician and physicist John von Neumann published by mathematician Stanislaw Ulam, in which he summarized a conversation with von Neumann in which they discussed the possibility that technological advancement could change the human race irrevocably. Ulam referred to this point as "some essential singularity." Von Neumann and Ulam's remarks suggest that the exponential rate of technological progress puts it within reach in the foreseeable future and that humanity does not have the ability to foresee the future beyond the singularity. These elements are still considered an integral part of singularity hypotheses. As other thinkers have put it, it is impossible for humans, with their limited intelligence, to foresee what a superintelligent machine might do.

The term itself, the "technological singularity" (generally just known as "the singularity"), was coined by computer scientist and science-fiction writer Vernor Vinge. Vinge first used it at a conference on AI in 1982, where he observed that creating human-level AI would be paramount and that superhuman-level AI would render all that came after wholly different on a fundamental level. He expanded on his ideas in a seminal paper, "The Coming Technological Singularity: How to Survive in the Post-Human Era," in 1993. In the abstract of the paper, Vinge theorizes that the technology for superhuman intelligence will exist within thirty years and that the creation of such technology will end the era of human dominance. This characterization of the singularity offers an element not present in von Neumann and Ulam's remark, that of the creation of superhuman levels of intelligence. It also offers another element present in many of the characterizations of the singularity as a likely eventuality--that it is only a few decades away. In an interview in 2008 in *IEEE Spectrum*, Vinge said he stood by his original prediction that the singularity would happen, barring catastrophes, by 2030.

What does it mean for a machine to have human or superhuman levels of intelligence? The question might be asked another way: what would it mean for a machine to have general intelligence? This is also known as strong AI or artificial general intelligence (AGI), to differentiate it from "narrower" or more specific forms of AI, such as programs designed to play chess or trade on the stock market or direct users to the nearest sushi restaurant. In his book *Superintelligence: Paths, Dangers, Strategies*, Oxford philosopher Nick Bostrom has defined superhuman levels of intelligence, or "superintelligence," as intelligence which that possesses cognition that significantly and consistently outstrips human cognition. It is not sufficient, then, to consider a machine as an AGI if it can beat a world champion at chess, but that same machine cannot hold a simple conversation or tie its shoelaces.

The singularity is somewhat of a specialist interest but has gained more attention in the media and general public with developments in AI. It has been popularized by enthusiasts such as the futurist and inventor Ray Kurzweil, who argues in his 2005 book *The Singularity Is Near*, that the accelerating pace of technological change, particularly in the fields of genetics, nanotechnology, and robotics, puts the singularity within reach by 2045. The singularity, in Kurzweil's conception, is "a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed." The transformation might include the end of human society, or even of humanity, as we know it, but this may not be a bad thing. The singularity, Kurzweil explains, would allow humans to transcend the "limitations of [their] biological bodies and brains" to extend their intelligence and gain power over their own mortality.

Intelligent Machines vs. Human Intelligence Today

There are several possible paths to the singularity, including the creation of intelligent machines (which may enable humans to create superintelligent machines, or may enable machines to redesign their own programming to enhance their intelligence, in what is sometimes known as an "intelligence explosion"); whole brain emulation (also known as "mind uploading"); enhancement of humans' own biological intelligence (through genetic engineering, for example); and the development of more sophisticated brain-computer interfaces. The first two paths are generally considered to be more likely than the others, or are at least the ones which attract more popular attention.

The development of AI has come a long way since the term was first coined by computer scientist and AI pioneer John McCarthy for a summer conference at Dartmouth University in 1956. Computers are now able to do things that used to be considered the exclusive province of human activity, including playing chess, vacuuming the floor, driving a car, and appearing on the game show Jeopardy. In 1997, the IBM supercomputer Deep Blue defeated world chess champion Garry Kasparov in a series of matches publicized as the ultimate battle between machine and man. Fourteen years later, another IBM creation, the computer system Watson, defeated two champions on Jeopardy to win \$1 million. Watson's victory meant that it was able to understand and answer questions in natural language and to analyze and execute a winning strategy. In 2014, Eugene Goostman, a computer program developed to simulate a thirteen-

year-old boy with a quirky sense of humor and a pet guinea pig, successfully convinced a third of its (human) judges that it was human at a Turing test in London. Experts, however, disagree on the significance of the event: organizers claim that it was the first time a computer program passed the Turing test, while other scientists point out that people do not really think of Eugene Goostman as possessing human-level intelligence. The disagreement points, perhaps, to a deeper confusion about what it means to possess human-level intelligence: even as computer programs achieve milestones that used to be considered essential tests of human intelligence (playing chess, conversing with a human being), more milestones appear to be added (such as recognizing objects, using natural language, or understanding a scene). The project of building AI is teaching us not just what it is to be intelligent, but what it is to be human.

Bibliography

Books and Reports

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford UP, 2014. Print.

Kurzweil, Ray. *The Singularity Is Near: When Humans Transcend Biology*. New York: Penguin, 2005. Print.

Vinge, Vernor. "The Coming Technological Singularity: How to Survive in the Post-Human Era." VISION-21 Symposium. Holiday Inn, Westlake. 1993. Address.

Periodicals

Turing, A. M. "Computing Machinery and Intelligence." *Mind* 59 (1950): 433-60. Print.

Point: The Singularity Is Realizable and Beneficial for Humanity

Thesis: Scientific and technological advances make it more likely that humans will be able to create AGI. Having human-level and superhuman-level intelligence will be valuable for humanity.

Talking Points

- The accelerating rate of technological change makes it likely that people will have both the computational capacity and scientific knowledge to create AGI in the next few decades.
- Having AGI will be beneficial to humanity. For example, it may enable humans to mitigate global problems such as climate change.
- It may also lead to an increase in the rate of economic growth and level of production.
- Intelligent and superintelligent machines are the future of humanity.

Summary

Supporters often point to the exponential pace of technological change as one reason to take the singularity seriously. In order to develop, AGI would require both hardware (computational capacity) and software (algorithms). Kevin Kelly, founding editor of *Wired*, wrote in 2014 that three developments had introduced AI into more mainstream technology: inexpensive parallel computation (computing that performs multiple large functions at once), big data (with which to train AI), and better algorithms (including renewed interest in machine learning techniques). The first point suggests that humans will be able to meet the hardware requirements of AGI, assuming that Moore's law can be generalized beyond the number of transistors on an integrated circuit to computational speed and capacity. To meet the software requirements, however, may require not just better algorithms but, in all probability, a better understanding of human cognition. In his 2005 book *The Singularity Is Near*, Ray Kurzweil argued that three simultaneous revolutions in genetics, nanotechnology, and robotics would give humans the understanding and technology needed for AGI: genetics for a better understanding of the human brain and body, and nanotechnology and robotics for the ability to model, emulate, and improve on the human brain.

Having more intelligence at its disposal should, on the face of it, be beneficial to humanity. It may enable humans to eradicate certain diseases or to mitigate climate change and, in those ways, to raise the standard of human life. For example, after IBM's computer system Watson was victorious on the game show *Jeopardy*, IBM started work training it to recognize medical images, as a first step to helping doctors diagnose and treat particular illnesses.

It may also enable humans to be more productive and so increase the rate of economic growth. Economist Robin Hanson has argued in his essay, "Economics of the Singularity," that sharp increases in the historic rate of world economic growth have been the result of singularities.

One past singularity was the Industrial Revolution, which saw a sixtyfold increase in the world economic growth rate. The "second machine age," as Massachusetts Institute of Technology (MIT) professors Andrew McAfee and Erik Brynjolfsson call it, may result in an even larger increase in the growth rate and usher in an age of greater prosperity.

Rather than see intelligent and superintelligent machines as threats to humanity, people could see them as the future of humanity. Proponents argue that AI makes it possible for humans to surpass their own biological limitations. In a 1993 essay on the "Age of Robots," roboticist Hans Moravec predicted that as computing power increased, so would the ability of robots to learn, model their world, and reason as humans do. Moravec noted in the essay that "depending on your point of view, humanity will then have produced a worthy successor, or transcended inherited limitations and transformed itself into something quite new."

Ponder This

- The author has presented the fundamental positions for this perspective in the debate. Outline the strengths and weaknesses of each perspective.
- If asked to begin forming an argument for this position, what sources would you need to build your case? What fundamental information do you need? What opinion leaders in this debate would you look to in solidifying your argument?
- What are the weakest aspects of the position outlined by the author? How might those weaker arguments help you prepare a counterargument?
- What additional Talking Points could you add to support this position?

Bibliography

Books and Reports

Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: Norton, 2014. Print.

Periodicals

Hanson, Robin. "Economics of the Singularity." *IEEE Spectrum*. *IEEE Spectrum*, 1 June 2008. Web. 16 Nov. 2015.

Kelly, Kevin. "The Three Breakthroughs That Have Finally Unleashed AI on the World." *Wired*. Condé Nast, 27 Oct. 2014. Web. 16 Nov. 2015.

Websites and Digital Files

Moravec, Hans. "The Age of Robots." Field Robotics Center. Carnegie Mellon Inst., June 1993. Web. 18 Nov. 2015.

Counterpoint: The Singularity Is Neither a Realistic nor a Desirable Future for Humanity

Thesis: The human mind is too complex and nuanced to be replaced by AI. Nor is it clear that such a development would be beneficial for humanity.

Talking Points

- Contemporary technology and understanding of neuroscience is insufficient to create AI that could replace the human brain.
- It is unclear that the effects of superintelligence would be positive.
- A superintelligent entity may not be friendly to human beings.

Summary

As John Horgan, a science journalist, has put it, humans do not know how a brain creates a conscious mind. Some of the complexity is technical: a healthy adult brain has a 100 billion neurons and quadrillions of connections between neurons--not to mention the fact that some of these synaptic connections may weaken or be dissolved, and new ones may be formed. Even if scientists had an exact model of all the neurons and connections of the brain, they do not as yet know the neural code by which signals are governed or, indeed, whether there is a neural code. It may be, as neuroscientist Christof Koch has observed, that there is no universal neural code, "above and beyond the insight that brains are amazingly adaptive and can extract every bit of information possible, inventing new codes as necessary."

Some of the complexity is conceptual: what is it for a machine to be conscious? English mathematician Alan Turing considered whether or not machines could think. In a 2008 report on the singularity by IEEE Spectrum, the journal of the Institute of Electrical and Electronic Engineers, Massachusetts Institute of Technology (MIT) computer scientist Rodney Brooks suggested that an AGI worthy of its name would also have the "object-recognition capabilities of a 2-year-old child ☐ the language capabilities of a 4-year-old ☐ the manual dexterity of a 6-year-old ☐ [and] the social understanding of an 8-year-old." Brooks stated that creating AI with these capabilities may be further off than most singularity enthusiasts suggest and that an eventual utopia brought on by the AI is unlikely.

It is not clear that having intelligent and superintelligent machines would be beneficial for humanity. Skeptics fear that even if the use of such machines increases the rate of economic growth, the rewards of such growth might not extend to all workers, or even, in the extreme case, to human beings. A 2013 study by the Oxford Martin School considered over seven hundred occupation types and concluded that some 47 percent of total US employment was at risk of being replaced by automation, including jobs in transportation, logistics, and office and administrative support.

Nor is it clear that superintelligent machines would necessarily be friendly to human beings. As Oxford philosopher Nick Bostrom has pointed out, a superintelligent entity may not share the values people associate with human intelligence, such as scientific curiosity or benevolent concern for others. It may or may not be possible to create a superintelligent machine that does value such things. In his book *Superintelligence*, Bostrom suggested that "it is no less possible--and probably technically easier--to build a superintelligence that places final value on nothing but calculating the decimal expansion of pi." Other thinkers have warned of the dangers of AI. In May 2014, for example, a group of prominent scientists, including Stephen Hawking, warned that AI could be "potentially the best or worst thing to happen to humanity in history" and that nothing much was being done to ensure one over the other. Later the same year, entrepreneur Elon Musk said in a symposium at MIT that AI was probably the biggest threat to human existence; he recommended implementing an international code to regulate AI creation because of the uncertainty of whether or not humans could control future AI.

Ponder This

- The author has presented the fundamental positions for this perspective in the debate. Outline the strengths and weaknesses of each perspective.
- If asked to begin forming an argument for this position, what sources would you need to build your case? What fundamental information do you need? What opinion leaders in this debate would you look to in solidifying your argument?
- What are the weakest aspects of the position outlined by the author? How might those weaker arguments help you prepare a counterargument?
- What additional Talking Points could you add to support this position?

Bibliography

Books and Reports

Barrat, James. *Our Final Invention: Artificial Intelligence and the End of the Human Era*. New York: St. Martin's, 2013. Print.

Periodicals

Hawking, Stephen, Stuart Russell, Max Tegmark, and Frank Wilczek. "Stephen Hawking: 'Transcendence Looks at the Implications of Artificial Intelligence--But Are We Taking AI Seriously Enough?'" *Independent*. Independent Print, 22 May 2014. Web. 16 Nov. 2015.

Horgan, John. "The Consciousness Conundrum." *IEEE Spectrum*. *IEEE Spectrum*, 1 June 2008. Web. 16 Nov. 2015.

Marcus, Gary. "The Brain in the Machine." *New Yorker*. Condé Nast, 20 Nov. 2012. Web. 16 Nov. 2015.

~~~~~

All content written by Tsin Yen Koh

Tsin Yen Koh is a doctoral student at Harvard University, with research interests in eighteenth- and nineteenth-century political thought, contemporary political theory, and social policy.

~~~~~

These essays and any opinions, information, or representations contained therein are the creation of the particular author and do not necessarily reflect the opinion of EBSCO Information Services.

Copyright of Points of View: Intelligent Machines vs. Human Intelligence is the property of Great Neck Publishing and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.

Artificial intelligence and the human mind are totally different constructions that complement each other's shortcomings. Think of AI as augmented intelligence. Let's stop comparing AI with human intelligence. None of this means that AI is superior to the human brain, or vice versa. They point is, they're totally different things. AI is good at repetitive tasks that have clearly defined boundaries and can be represented by data, and bad at broad tasks that require intuition and decision-making based on incomplete information. In contrast, human intelligence is good for settings where you need common sense and abstract decisions, and bad at tasks that require heavy computations and data processing in real time. Everyone agrees that the human brain is an intelligent system; in fact it is the only system everyone agrees is intelligent. We believe that by studying how the brain works we can learn what intelligence is and what properties of the brain are essential for any intelligent system. For example we know the brain represents information using sparse distributed representations (SDRs), which are essential for semantic generalization and creativity. We use the term "machine intelligence" to refer to machines that learn but are aligned with the Biological Neural Network approach. Although there still is much work ahead of us, we believe the Biological Neural Network approach is the fastest and most direct path to truly intelligent machines.

Reflect on which arguments about intelligent machines vs. human intelligence you accept, those outlined in the Point, Counterpoint, or a completely different argument. For you, what is the single most compelling argument regarding intelligent machines, and why?

Write an 800- to 1,000-word essay response to the question above; include a thesis statement and 2-3 supporting arguments.