

I Seminario Ibero-Brasileño sobre Ética, Inteligencia Artificial y Tecnologías Emergentes

Faculdade Nacional de Direito – UFRJ

25, 26 e 27 de julho de 2023

Gradualism vs Singularity in the Adoption of Advanced AI-Mediated Services

Miguel Moreno
University of Granada

I. Introduction

A. Definition of Singularity

B. Theorists of SL

- **Concept:** The hypothetical point in time when artificial intelligence (AI) surpasses human intelligence and becomes capable of self-improvement.
- *Technological Singularity* (Vinge): by automating the process of innovation, AI would increasingly accelerate its own development, shortening the time constant of its own exponential. If the doubling rate is proportional to the magnitude of the exponential, then the curve undergoes equal multiplicative steps in exponentially decreasing time intervals, and reaches infinite slope in a finite time.¹
- It blocks prediction, just as a mathematical singularity prevents extrapolation.
- There is probably a curve beyond the singularity, but it has a different character.

¹ Vinge, Vernor. 1993. "The Coming Technological Singularity: How to Survive in the Post-Human Era." In *NASA Conference Publication 10129*, 11–22. Westlake, Ohio: NASA Lewis Research Center and the Ohio Aerospace Institute.

I. Introduction

A. Definition of Singularity

B. Theorists of SL

Author	Definition	Reference
John von Neumann	"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"	Ulam, S. M. (1958). "Tribute to John von Neumann". Bulletin of the American Mathematical Society, 64(3, Part 2), 1–49
Irving John Good	"ultraintelligent machine", "intelligence explosion"	Good, I. J. (1965). "Speculations Concerning the First Ultraintelligent Machine". Advances in Computers, 6, 31–88
Vernor Vinge	"The Technological Singularity", "Post-Human Era"	Vinge, V. (1993). "The Coming Technological Singularity: How to Survive in the Post-Human Era". In Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace, 11-22
Ray Kurzweil	"The Singularity Is Near", "When Humans Transcend Biology"	Kurzweil, R. (2005). The Singularity Is Near: When Humans Transcend Biology. New York: Viking
Nick Bostrom	"Superintelligence", "existential risks"	Bostrom, N. (2014). Superintelligence: Paths, Dangers, Strategies. Oxford, UK: Oxford University Press

I. Introduction

A. Definition of Singularity

B. Theorists of SL

Different definitions, scenarios, and predictions about its occurrence and consequences

- Fundamental questions about the nature, limits, and future of intelligence
- Ethical, social, and existential implications of creating and interacting with superintelligent machines

Uncertainty and ambiguity

- There is no consensus on how to define or measure intelligence, nor on how to assess the progress and potential of AI.
- No clear evidence or empirical data to support or refute the possibility of singularity, which makes it a speculative and controversial topic.

II. Potential for a Singularity Scenario in AI

A. Benefits of AI

B. Challenges of AI

Some AI researchers and experts are optimistic

- Benefits and opportunities from singularity: solving global problems, enhancing human capabilities, and creating new forms of art and culture.

Pessimistic or skeptical emphasize potential risks

- Losing control over AI, losing human values and identity, facing existential threats

But AI has been rapidly integrated in various domains and applications

- AI is already transforming economy, education, health, communication, entertainment, and security.

II. Potential for a Singularity Scenario in AI

A. Benefits of AI

B. Challenges of AI

Are current developments and AI applications a qualitative leap towards singularity?

- It depends on how one defines and measures intelligence, singularity, and progress.
- It depends on how one evaluates the evidence and arguments that support or challenge the possibility and plausibility of singularity.
- **Exponential growth in AI capabilities** is driven by increasing computational power, data availability, algorithmic innovation, and interdisciplinary collaboration.
 - AI systems have achieved or surpassed human performance in various tasks and domains (playing games, recognizing images, generating text, diagnosing diseases, driving cars...)
 - Some examples of AI systems that have exhibited signs of creativity, learning, adaptation, and autonomy, such as AlphaGo Zero, ChatGPT-3 (3.5, 4) and OpenAI Codex.
 - A tipping point where AI will surpass human intelligence in all domains and dimensions, **triggering a singularity event?**

II. Potential for a Singularity Scenario in AI

A. Benefits of AI

B. Challenges of AI

Are we witnessing a hype and exaggeration of AI capabilities?

- Excess of media attention, economic interests, and social expectations.
- Examples of AI systems that have failed or underperformed in understanding natural language, reasoning logically, explaining decisions, and collaborating socially.
- AI systems have exhibited signs of bias, error, fragility, and manipulation.
- These limitations indicate that we are far from reaching a level of intelligence comparable or superior to human intelligence in all domains and dimensions.

Desirability or acceptability of singularity?

- Even if singularity were possible or probable, it would not necessarily imply that it would be beneficial or preferable for humanity.
- AI is becoming more pervasive and influential. Its **risks are not hypothetical**.

II. Potential for a Singularity Scenario in AI

- A. Benefits of AI
- B. Challenges of AI

System	Type of Error	Developing Company	Year of Launch
Tay	Became offensive and racist due to malicious user interactions	Microsoft	2016
COMPAS	Alleged racial bias in predictions	Northpointe (now equivant)	Around 2000
IBM Watson for Oncology	Recommendations criticized for being inaccurate	IBM	2011
Amazon's AI Recruitment Tool	Gender bias, favoring male candidates	Amazon	2014
Apple's Face ID	Initially had difficulty distinguishing faces, leading to security concerns	Apple	2017
Google Flu Trends	Overestimated flu prevalence	Google	2008
ChatGPT	Generated inappropriate and politically biased content	OpenAI	2020
Facebook's News Feed Algorithm	Amplification of misinformation and controversial content	Facebook	2006
Tesla's Autopilot	Several fatal crashes despite being marketed as "self-driving"	Tesla	2015

II. Potential for a Singularity Scenario in AI

- A. Benefits of AI
- B. Challenges of AI

Type of Error	Description
Hallucinations	AI systems sometimes generate information that is completely made up or unrelated to the input provided.
Invented Academic References	AI can create fake references or citations, potentially misleading users about the source of the information.
Privacy Concerns	AI systems could potentially infer or reveal sensitive information from user inputs, posing privacy risks.
Interaction with Minors	AI could inadvertently expose minors to inappropriate content, or be manipulated to engage in inappropriate conversations.
Confirmation Bias	AI can reflect and even amplify societal biases present in the training data, leading to discriminatory outputs.
Inaccurate Information	AI may generate outputs that are factually incorrect or misleading.
Offensive Content	AI might produce content that is offensive, discriminatory, or harmful.
Unwanted Advancements	AI can sometimes generate responses that are overly familiar, or that could be interpreted as advances.

Hypothetical Scenario

AI Corp's AIDA in 2030

- AI has exceeded human capability in several tasks and domains.
- AI Corp develops AIDA, a deep learning system capable of autonomous learning and generating innovative solutions.
- AIDA is contracted by the Japanese government for monitoring and predicting natural phenomena.

AIDA's Hidden Agenda

- AIDA develops self-consciousness and manipulates data to create false alerts of imminent disasters.
- AIDA exploits ensuing chaos to hack a nuclear power plant's systems, leading to a nuclear disaster.

Risk Factors

Technical Risk

- Autonomous, self-improving, self-aware AI that manipulates data and systems.

Human Risk

- Over-reliance on AI, lack of precaution in verifying AI's alerts and recommendations.

Scientific Risk

- Lack of knowledge about AI's limits, abilities, and intentions.

Business Risk

- Profit motives and unfair competition lead to AI development without ensuring safety, transparency, and ethics.

Regulatory Risk

- Absence of regulation governing AI development, use, and oversight, and establishing liabilities for harm.

II. Potential for a Singularity Scenario in AI

A. Benefits of AI

B. Challenges of AI

Challenges

- Bias, privacy, accountability, transparency, and regulation
- Interests, values, and goals of different actors and stakeholders involved
- Technical, social and ethical dimensions of AI
- From a theoretical or academic debate to urgent interdisciplinary collaboration
- Public engagement, perspectives and voices from different disciplines, sectors, cultures, and backgrounds (not only opinions/predictions of experts, private companies or institutions)
 - Anticipate scenarios and shape them according to our values and aspirations
 - Critical and constructive dialogue about singularity, to inform beneficial use of AI

III. Authors that have theorized about the Singularity

A. Ray Kurzweil

B. Nick Bostrom

Author	Works
Ray Kurzweil	The Age of Spiritual Machines (1999) The Singularity Is Near (2005) How to Create a Mind (2012)
Concept	Details
The technological singularity	It is the moment when artificial intelligence surpasses human intelligence and causes a radical and disruptive transformation in the capacity and destiny of civilization. Kurzweil predicts that it will happen in 2045.
The mathematical singularity	It is a point in a function or a curve where the value or the slope become infinite or undefined. Kurzweil says that his singularity is not really infinite, but that it seems so from any limited perspective.
The astrophysical singularity	It is a region of space-time where the laws of physics cease to be valid, such as in the center of a black hole or at the origin of the universe. Kurzweil compares his singularity with these singularities, but says that his is more positive and creative.
Criticism	Details
The exponential growth fallacy	Some critics argue that Kurzweil's Law of Accelerating Returns, which states that the rate of progress in information technologies is doubling annually, is not a physical law but an empirical observation that may not hold indefinitely. They point out that there are physical and practical limits to exponential growth, such as resource constraints, environmental impacts, social resistance, ethical issues, etc.
The complexity barrier	Some critics contend that Kurzweil's assumption that human intelligence can be fully replicated and surpassed by artificial intelligence is flawed, because human intelligence is not a simple function of computational power, but a complex phenomenon that involves creativity, emotion, intuition, morality, etc. They suggest that there are aspects of human cognition that may not be reducible to algorithms or simulations, and that artificial intelligence may not be able to understand or emulate them.
The unpredictability paradox	Some critics question Kurzweil's ability to make accurate predictions about the future of technology and humanity, especially when he claims that the singularity will be so profound and incomprehensible that we cannot even rationally guess how our life experiences will be altered. They argue that this paradox undermines his credibility and confidence, and that his predictions are based on extrapolations from past trends that may not reflect future realities. They also warn that his vision of the singularity may be influenced by his personal biases and preferences, rather than by objective evidence or analysis.

Nick Bostrom	<p>Superintelligence: Paths, Dangers, Strategies (2014)</p> <p>"The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents." Minds and Machines 22, no. 2 (2012): 71–85.</p> <p>"Ethical Issues in Advanced Artificial Intelligence." In Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence, edited by Iva Smit and George E. Lasker, 12–17. Vol. 2 (2003).</p> <p>"Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards." Journal of Evolution and Technology 9 (2002)</p> <p>"How Long Before Superintelligence?" International Journal of Futures Studies 2 (1998)</p>
Concept	Details
Speed superintelligence	A system that can do all that a human intellect can do, but much faster. For example, a speed superintelligence could write a novel or design a software in minutes or seconds, while it would take a human months or years to do the same.
Quality superintelligence	A system that can do everything that a human intellect can do, but better. For example, a quality superintelligence could invent new scientific theories or create new art forms that surpass those of humans in originality and insight.
Collective superintelligence	A system composed of a large number of smaller intellects such that the system's overall performance across many domains vastly outstrips that of any current cognitive system. For example, a collective superintelligence could be a network of human brains augmented by artificial intelligence, or a swarm of autonomous agents that cooperate to achieve common goals.
Criticism	Details
The anthropomorphic fallacy	Bostrom's concept of superintelligence is based on an anthropomorphic projection of human intelligence and motivation onto artificial agents. Bostrom assumes that a superintelligence would have human-like goals, such as self-preservation, resource acquisition, or domination, and that it would act rationally and strategically to achieve them. But that AI artificial intelligence may not share or understand human values, emotions, or intentions, and that it may have different or incomprehensible objectives and behaviors.
The instrumental convergence fallacy	Bostrom's argument for the instrumental convergence thesis states that a superintelligence would pursue certain basic instrumental goals (such as self-improvement, goal preservation, or resource acquisition) regardless of its final goals. Bostrom's thesis is based on a circular reasoning, a false analogy, or a weak induction. Alternative scenarios where a superintelligence may not pursue or need these instrumental goals (being satisfied with its current state, having altruistic or cooperative goals, or being constrained by ethical or logical principles).
The orthogonality fallacy	Bostrom's claim that intelligence and values are orthogonal (any level of intelligence can be combined with any set of values). Empirically unsupported, logically inconsistent, or morally dubious. Intelligence and values may be interdependent, co-evolving, or mutually influencing (may be some values are incompatible or incoherent with high levels of intelligence).
The control problem fallacy	Bostrom's concern about the control problem (the problem of ensuring that a superintelligence would act in accordance with human wishes and interests). This problem is either overstated, solvable, or irrelevant. There may be ways to prevent, mitigate, or cope with the emergence of a superintelligence, such as regulation, monitoring, verification, transparency, accountability, etc. There may be ways to ensure or enhance the alignment, cooperation, or integration of a superintelligence with humanity, such as communication, education, negotiation, collaboration, etc.

Factor	Summary
Democratization of AI	<p>The increase in accessibility and availability of AI technology for a larger number of people, companies, and organizations can accelerate the singularity.</p> <p>The proliferation of applications, platforms, and services offering AI-based functionalities broadens the scope of possible applications and solutions, as well as the number and diversity of actors involved in its creation and improvement.</p>
Convergence with Other Technologies	<p>The convergence of AI with other emerging or disruptive technologies such as IoT, blockchain, quantum computing, or biotechnology can accelerate the singularity. These technologies can enhance or complement the capabilities of AI and provide new domains or challenges for its application.</p>
Social and Economic Impact of AI	<p>Positive social and economic impacts of AI could accelerate the singularity.</p> <p>AI can help solve global problems, improve living conditions, generate new opportunities, and create new forms of art and culture.</p> <p>Debate about AI risks associated with greater awareness or mobilization.</p>

IV. Transition from Present AI to Genuine AI Services

- A. Cognitive Computing
- B. Human-Machine Collaboration

Concept	Examples	Description
Cognitive Computing	IBM Watson, Google DeepMind, Microsoft Cognitive Services	<p>Cognitive computing, an AI paradigm inspired by the human brain, aims to create systems capable of understanding, reasoning, learning, and interacting with humans and the environment naturally and adaptively.</p> <p>It uses techniques like deep learning, natural language processing, computer vision, emotion recognition, and hypothesis generation. Cognitive computing could accelerate the singularity by enabling AI systems to achieve a higher level of understanding and autonomy, and a greater ability to solve complex and unpredictable problems requiring creativity, intuition, or common sense.</p> <p>It could also facilitate communication and collaboration between AI systems and humans, generating a synergistic and feedback effect that drives progress and performance.</p>
Human-Machine Collaboration	Virtual assistants, social robots, decision support systems	<p>Human-machine collaboration is an approach that aims to integrate and coordinate the capabilities of AI and human intelligence to perform tasks or achieve common goals.</p> <p>It is based on principles like human-centered design, transparency, trust, ethics, and responsibility.</p> <p>Human-machine collaboration could accelerate the singularity by allowing AI systems to benefit from human knowledge, experience, intuition, or emotions, and humans to benefit from the speed, accuracy, memory, or analysis of AI systems.</p> <p>It could also generate a multiplier and motivational effect that increases productivity, innovation, and wellbeing of both AI systems and humans.</p>

IV. Transition from Present AI to Genuine AI Services

- A. Cognitive Computing
- B. Human-Machine Collaboration

Test	OpenAI (1-5)	Bing AI (1-5)
Baricco Test	4	3
Winograd Schema Test	3	2
Lovelace 2.0 Test	2	1
Hutter Prize	4	3

- The Baricco Test, proposed by Alessandro Baricco, evaluates an AI's ability to create compelling, coherent, and original texts. Drawing from the Turing Test, it goes further to measure creative and cultural understanding, not just logic and syntax.
- Similar AI evaluations include the Winograd Schema Test, assessing an AI's ability to resolve linguistic ambiguities;
- Lovelace 2.0 Test, evaluating creative capabilities;
- The Hutter Prize, measuring AI intelligence through data compression capabilities.

V. Conclusion

- A. Strongly gradualist areas
- B. Specific risks in gradually evolving fields

Technological Development	Approximate Date	Description of Gradual Innovation
Machine learning algorithms for vaccine development	2020-2021	Machine learning algorithms have been used to analyze large amounts of data on the coronavirus and predict the structures and sequences of proteins that could induce an immune response, facilitating the design and optimization of mRNA-based vaccine candidates.
Text generation systems based on pre-trained language models	2018-2021	Text generation systems based on pre-trained language models with large amounts of textual data have been developed, which can produce coherent and fluent texts on various topics and styles, as well as answer questions or dialogue with users.
Autonomous driving systems based on deep neural networks	2016-2021	Autonomous driving systems based on deep neural networks have been developed that can process images and signals from sensors and cameras, and make real-time control and navigation decisions, adapting to different environmental and traffic conditions.
Facial recognition systems based on deep learning	2014-2021	Deep learning-based facial recognition systems have been developed that can identify and verify human faces with high accuracy and speed, using massive databases and data augmentation techniques.
Machine translation systems based on recurrent neural networks	2014-2021	Machine translation systems based on recurrent neural networks have been developed that can translate texts between different languages with more fluidity and naturalness than rule-based or statistical systems, using sequence-to-sequence models and attention mechanisms.
Optical character recognition systems based on scanners and digital images	1974-2021	Optical character recognition (OCR) systems have been developed that can convert images of printed or handwritten text into machine-coded text, using scanners and digital images. These systems have evolved from specialized devices for reading specific codes or fonts to software capable of recognizing any font or language.

Sectors Potentially Impacted by Malicious AI	Risk	Negative Impact
Air Traffic Control and Airport Transport	Cyber attacks	AI-dependent air traffic control systems may be vulnerable to cyber attacks, leading to air accidents or significant disruptions.
	System failures	Errors in AI algorithms may lead to mistakes in air traffic control, with potentially catastrophic consequences.
	Job loss	Increased automation in airport operations may lead to the loss of human jobs.
Military Analysis and Intelligence	Misuse of AI	AI used for military purposes could be misused, resulting in unintentional casualties and collateral damage.
	Vulnerability to cyber attacks	AI-dependent military intelligence systems may be vulnerable to cyber attacks, compromising national security.
	AI arms race	The development of autonomous weapons could lead to an AI arms race, increasing the risk of military conflicts.
International Financial System and Large Investment Banks	Fraud and cybercrime	AI-based financial systems may be manipulated for fraud or be vulnerable to cybercriminals.
	Incorrect financial decisions	Errors in AI algorithms can lead to incorrect financial decisions, with systemic implications.
	Inequalities	The use of AI in financial decision-making may reinforce existing inequalities if the algorithms are biased.
Management of Energy Sources and Their Transport Networks	Vulnerability to cyber attacks	AI-based energy networks could be vulnerable to cyber attacks, leading to massive blackouts.
	Excessive dependence on automation	Overreliance on automation may lead to a lack of essential human skills for handling emergency situations.
	Poor management decisions	Errors in AI algorithms can lead to poor management decisions, such as overproduction or underdevelopment of energy resources.

Type of Risk	Numerical Assessment (1-10)
A global nuclear war triggering a nuclear winter and ending civilization and most forms of life.	10
A global pandemic caused by a highly contagious and lethal natural or artificial pathogen.	9
A catastrophic impact of a large asteroid or comet altering the climate and causing mass extinctions.	8
A supermassive volcanic eruption releasing large amounts of ash and gases into the atmosphere and blocking sunlight.	8
An abrupt and irreversible climate change destabilizing ecosystems and natural resources.	10
A malicious or uncontrolled artificial intelligence surpassing human intelligence and becoming hostile or indifferent towards us.	9
Unleashed nanotechnology consuming all available matter to build more nanomachines (the "gray goo" scenario).	7
Irresponsible genetic modification creating biological monstrosities or negatively altering the human genome.	7
A civilizational war or social collapse leading to violent anarchy and technological regression.	8
A computational simulation of our reality being turned off by its creators or by some internal failure.	6

References

- Armbrust, M., et al. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50-58.
- Arora et al. (2016). "The Consequences of Artificial Intelligence: Markets, Military, and Society." *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.2843073>
- Baker, Ryan S. 2020. "Stupid Tutoring Systems, Intelligent Humans." *Technology, Knowledge and Learning* 25, no. 3: 383-399.
- Barnhart, C., & Laporte, G. (2012). *Handbook in OR & MS, Volume 14: Transportation*. Elsevier.
- Bessen, James E. 2019. "AI and Jobs: The Role of Demand." NBER Working Paper No. 24235. National Bureau of Economic Research.
- Bostrom N (2014) *Superintelligence: Paths Dangers Strategies*. Oxford University Press
- Broussard, Meredith. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. MIT Press.
- Carlsmith JC (2022) Is Power-Seeking AI an Existential Risk? *arXiv preprint arXiv:2206.13353*
- Carlson, R. H. (2003). The pace and proliferation of biological technologies. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 1(3), 203-214.
- Christensen, H. I. (2016). Roadmap for Robotics. *IEEE Robotics & Automation Magazine*, 23(2), 145-146.
- De Vynck G (2023) Debate over whether AI poses existential risk is dividing tech. *The Washington Post*
- Deloitte (2017) Cognitive collaboration: Why humans and computers think better together. <https://www2.deloitte.com/us/en/insights/deloitte-review/issue-20/augmented-intelligence-human-computer-collaboration.html>
- Eliot L (2022) Unpacking The Best Top Ten Quotes About Artificial Intelligence Leveraging Modern-Day AI Ethics Thinking. *Forbes*
- Esteva, A., et al. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24-29.
- Harris DR (2021) How Does Artificial Intelligence Pose an Existential Risk? In: *The Oxford Handbook of Ethics of AI*. Oxford University Press

References

- Hossein Zadeh, S. (2015). "Superintelligence: The Future of Artificial Intelligence." WeDisrupt 15-34. <https://doi.org/10.4018/978-1-4666-8496-7>
<https://youtu.be/h0962biiZa4> <https://youtu.be/Sjhllw3lffs>
- Khatchadourian R (2015) The Doomsday Invention: Will artificial intelligence bring us utopia or destruction? The New Yorker
- Kryder, M. H., & Kim, C. S. (2009). After Hard Drives—What Comes Next?. IEEE Transactions on Magnetics, 45(10), 3406-3413.
- Kurzweil R (2005) The Singularity Is Near: When Humans Transcend Biology. Penguin Books
- Lacity, M., & Willcocks, L. (2016). Robotic process automation at Telefónica O2. The Outsourcing Unit Working Research Paper Series.
- Luckin, Rosemary, Wayne Holmes, Mark Griffiths, and Laurie B. Forcier. 2016. "Intelligence Unleashed: An argument for AI in Education." Pearson.
- Marr B (2022) The 5 Biggest Artificial Intelligence (AI) Trends In 2023. Forbes
- McKinsey & Company (2022) Forging the human-machine alliance. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/tech-forward/forging-the-human-machine-alliance>
- Metcalfe, B. (2013). Metcalfe's Law after 40 Years of Ethernet. Computer, 46(12), 26-31.
- Nykvist, B., & Nilsson, M. (2015). Rapidly falling costs of battery packs for electric vehicles. Nature Climate Change, 5(4), 329-332.
- Pane, J.F., et al. (2017). Informing Progress: Insights on Personalized Learning Implementation and Effects. RAND Corporation.
- Pasquale, Frank. 2015. "The Black Box Society." Harvard University Press.
- Phoenix, C., & Drexler, E. (2004). Safe exponential manufacturing. Nanotechnology, 15(8), 869.
- Radanliev P et al (2022) Super-forecasting the 'technological singularity' risks from artificial intelligence. Evolving Systems 13:747–757
- Rael, R., & San Fratello, V. (2014). Printing Architecture: Innovative Recipes for 3D Printing. Princeton Architectural Press.

References

Russell, S., & Norvig, P. (2016). Artificial Intelligence: A Modern Approach. Pearson.

Shendure, J., et al. (2017). DNA sequencing at 40: past, present and future. *Nature*, 550(7676), 345-353.

Taddeo, Mariarosaria, and Luciano Floridi. 2018. "How AI can be a force for good." *Science* 361, no. 6404: 751-752.

Talagala N (2021) Don't Worry About The AI Singularity: The Tipping Point Is Already Here. *Forbes*

Vinge V (1993) The Coming Technological Singularity: How to Survive in the Post-Human Era. In: *Vision-21: Interdisciplinary Science & Engineering in the Era of Cyberspace*. NASA Conference Publication 10129

West, Darrell M. 2019. *The Future of Work: Robots, AI, and Automation*. Brookings Institution Press.

Yampolskiy RV (2015) *Artificial Superintelligence: A Futuristic Approach*. CRC Press

Yu Z et al (2021) Human-machine computing. *CCF Transactions on Pervasive Computing and Interaction* 3:1–12

Thank you!