Scenario Exploration


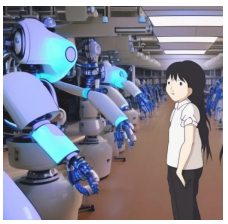
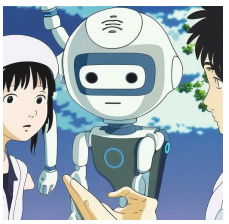
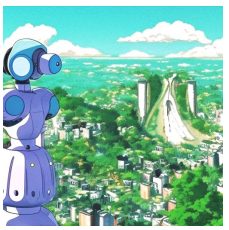

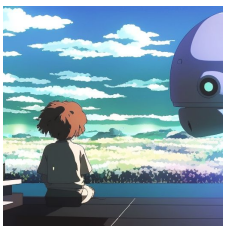
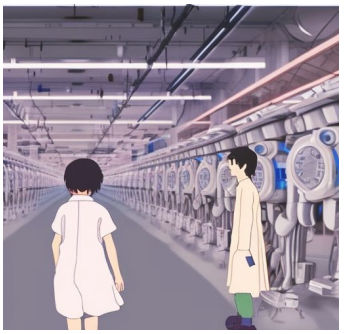
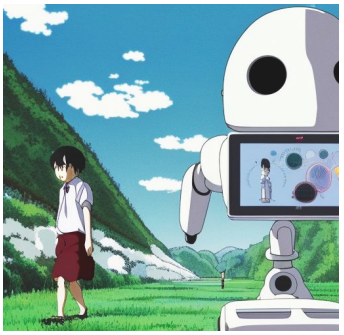

GTP-3 Prompt: "Humans jobs have been replaced by artificial intelligence and automation"

GTP-3 Prompt: "AI Job Displacement creates a utopian society"

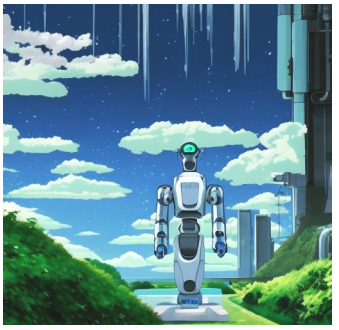



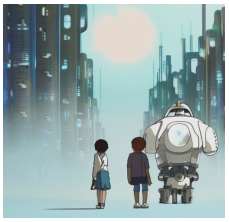
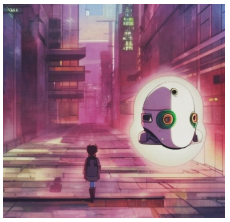
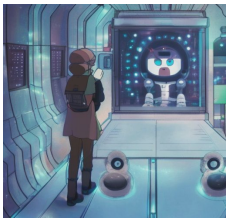

GTP-3 Prompt: "But whilst humans automate more and more of human decision making artificial intelligence is becoming a better artist and music creator than humans"

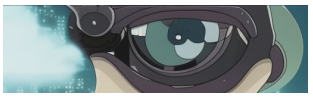

GTP-3 Prompt: "Humans have more time for leisure because of automation by artificial intelligence"

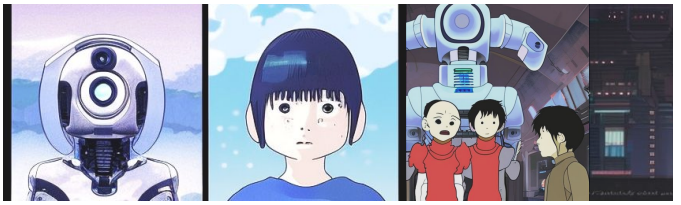

CONTENTS

Introduction
- Automating manual labour
- Automating human decision making
- Where we are today: ANI
- Exploring ANI limitations

Roadmap to our scenario
- Artificial General Intelligence
- Predictions for AGI arrival
- The Roads to AGI
- Computing Hardware
- Understanding human hardware to develop HLMI
- Debate: Could our scenario happen without AGI
- Debate 2: AGI is inevitable
- Trilemma impacts

Indicators of our scenario happening
- When will we be able to tell when we've attained AGI?

Impacts of our scenario
- Likelihood of automation

Adapting to the future of work and education
- Modelling the future of work
- Upskilling to a knowledge-based economy
- Delivering AI and IA

Conclusion

References

Appendix


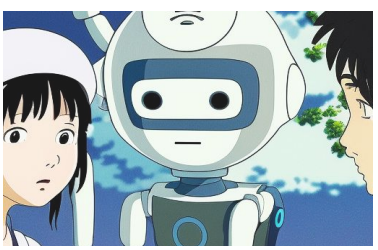

Introduction

Artificial Intelligence (AI) has already displaced human labour in domains that were previously deemed to be uncomputerizable. We asked the question: if AI/ML development continues, what impact would it have on employment, and ultimately society?

Our scenario depicts a future where AI displaces over ten million jobs, greatly affecting entry-level jobs that do not need complicated decision making, creativity, or leadership. It takes the perspective of a student entering the workforce and dealing with emotions of obsolescence in a future where universal basic income (UBI) has been introduced and work is no longer a necessity. Humans' place in this world seems to be in areas which require empathy and compassion, but the scenario hints that this too could one day be challenged as a chat-GTP bot takes on the job of a careers counsellor to help a human decide where to find his purpose.

The following report examines the future of AI/ML development, offering a roadmap to where such conditions may become a reality over the next decade, analysing the repercussions, and suggesting remedies for the future of technological unemployment.

Automating manual labour

Aristotle in 350 B.C.E warned against innovation as something that crept innocuously into society until the entire state was altered, and he predicted that machines would one day abolish the necessity for human labour.(1) We've spent the last 1000 years building tools to decrease physical labour. Machines prove to be more precise and cost efficient than humans, from ploughs to tractors in agriculture to the mechanical spinning machine, which displaced labourers in the 1700s and prompted the rise of the Luddites, who attempted to prohibit or destroy new technology in order to preserve their old ways of life.(2)

Automating human decision making

Earlier industrial revolutions saw machines replacing humans in manual labour. The AI revolution will see machines replacing humans in cognitive functions, including human decision making. Automating human decision making differs from physical automation, It requires greater complexity to achieve and carries much broader implications(3).

In an age in which machines maximise our productivity by increasingly performing tasks only humans used to be capable of, what will constitute our identity as human beings? Driven by economic incentives, AI will be adopted to enhance efficiency and productivity. The Tetrad model(4) (Figure 1) depicts a trade-off between efficiency and productivity improvements, which will render human intellect obsolete. Consequently it is likely to redefine our idea of 'intelligence' to prioritise emotional intelligence which requires the retrieval of human empathy and compassion. Empathy and compassion will be crucial as the range of possible outcomes for humanity expands – for better or worse – as AI is pushed to its limits. (5)

Where we are today: ANI This is currently achieved to various extent with narrow AI (ANI), which is highly specific to one task, and requires significant redesign to be adapted for other tasks. In certain sectors, ANIs provide results that are comparable to, and occasionally exceed, human intelligence.(6)

ANI applications are firmly ingrained in the infrastructure of every industry and can outperform humans in terms of speed, as well as the ability to manage and examine thousands of variables simultaneously.(7) Such as in detecting cancer from CT scans, translating languages or mastering any game without knowing the rules.(8,9,10) Integral these ANI capabilities are machine learning (ML) and natural language processing (NLP).(11) ML enables systems to learn patterns from data, make predictions, and then improve future experience through applying the discovered patterns to situations absent in their initial design.(12)

Exploring ANI limitations

ANI systems are really good at particular tasks, but doing them together in a coordinated way seems to be a lot harder. Something appears to be lacking in ANI's ability to generalise and conceptualise. An ANI can recognise patterns in data and distinguish between dogs and cats, but it falls short when faced with a more abstract challenge. They are unable to generalise outside of data that they have not seen.

Bayesian reasoning determines the probability of an event with uncertain knowledge. Humans also struggle with such concepts, hence quantitative algorithms serve to remove bias in investment decisions and respond to indicators much faster than humans, making algorithmic trading responsible for 60-73% of all U.S. equity trading.(13) But during the COVID-19 pandemic, quantitative AI models did not perform so well. Overall human decision making outperformed quant funds in 2020.(14) Practical AI systems shall have to cope with uncertainty — that is, they shall have to deal with incomplete evidence leading to beliefs that fall short of knowledge, with fallible conclusions and the need to recover from error, called non-monotonic reasoning.(15) Until this is achieved, human decision making remains central, and it is more likely that ANI deployment increases tools utilised by humans to inform their decisions, than it is to result in wide scale job displacement.

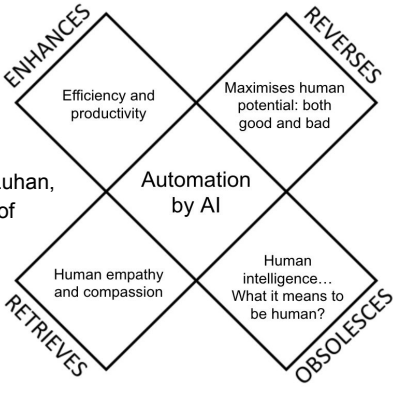

Figure 1: The Tetrad (McLuhan, 1988) to model the effects of automation

Roadmap to our scenario

Artificial General Intelligence

For AI progress to surpass human levels in all domains, it would need to overcome limitations – including handling uncertainty. Many stipulate that this would require artificial general intelligence (AGI).[16-18] The following section explores expectations for AGI arrival. Vernor Vinge proposed the concept of a "technological singularity" in 1993, believing that mankind would possess reach the technological capability to create superhuman intelligence within thirty years.[19] AI development has since seen inflection points such as the application of GPUs in 2012.[20] These advanced microprocessors allowed parallelisation of computation – to run hundreds or thousands of computing operations in parallel, millions of times every second. This freed AI/ML development from the constraints of Moore's law, providing the speed for ML deep neural networks for the ANI of today.

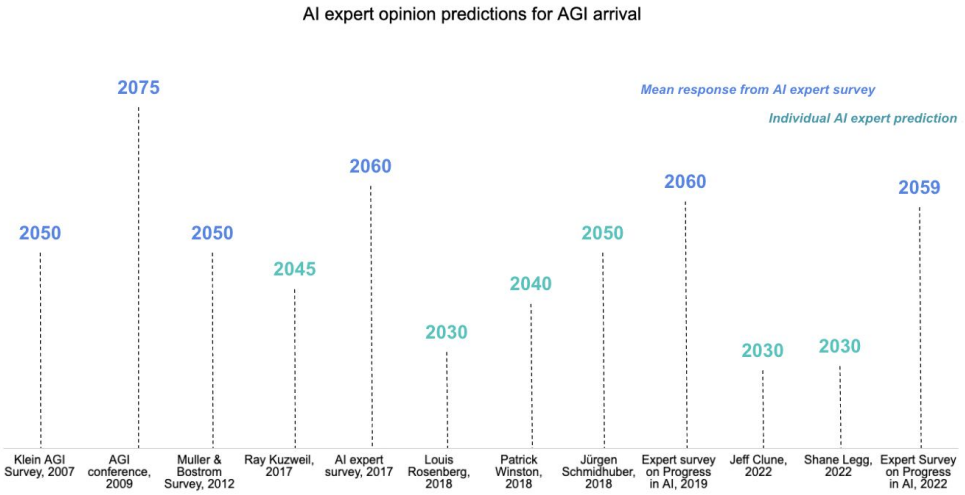

Figure 2: AI expert opinion predictions for AGI arrival, sources Table in Appendix

Predictions for AGI arrival Despite rapid increases in computation, there are hurdles to developing generalised intelligence and the expectations for AGI varies greatly amongst experts. Analysing 6 surveys conducted amongst AI experts as well as individual expert predictions indicates that there is no clear consensus for AGI arrival (Figure 2) and highlights the potential fallibility of our own predictions in our scenario.

In 2018, The Future of Humanity Institute at the University of Oxford surveyed the views of 352 academics and industry experts in machine learning; In the next ten years, we should have A.I. do better than humans in translating languages (by 2024), writing high-school level essays (by 2026), writing top 40 songs (by 2028) and driving trucks (by 2027). But it is not until 2049 that the consensus could envision A.I. writing New York Times best sellers, or until 2053 to be performing surgeries.[21] The average believed overall, A.I. should be better than humans at pretty much everything by 2061. More pessimistic forecasts assert that programs still require an exceptional amount of structured data to learn from and cannot adapt to problems such as those seen with the performance of autonomous vehicles, despite billions of dollars invested into their development.[22] Individual predictions from figures in the field tend to be more optimistic and there has been a reduction in the time expected to develop AGI between the 2019-2022 expert surveys on progress in AI.

Over the past 2 months, the Metaculus (collective community) predictions for estimated years to AGI has fallen by 70%, predicting AGI arrival within 6 years and implies superintelligence by 2040 at median.[23,24]

The Roads to AGI

The road to AGI is marked by mastery of the most significant and powerful trait of human civilization, perhaps of our planet's whole evolutionary history: intelligence. Which technical developments will enable AGI and will they cause an intelligence explosion, or a gradual but steady growth?

Computing Hardware Over the previous decade, computation has grown at an exponential rate, with maximal compute doubling every 3.4 months (Figure 4). Price to train to obtain a given accuracy on ImageNet has fallen by a factor of 200 over four years, reflecting software and hardware price performance gains.[25] In 1998 Hans Moravec rated entities by the computational power and memory of the smallest universal computer needed to mimic their behavior.[26] Extrapolating computer hardware price trends to these assumptions suggests computers suitable for human-level machine intelligence (HLMI) will appear around 2025 (Figure 3) and that today's most powerful supercomputers are within a factor of a hundred from required capacity. Projections are based on extrapolating semiconductor technology; as Moore's law nears its end (about 2025), circuits will most likely integrate an increasing number of quantum interference components. Before 2030, quantum computers, DNA computers, carbon semiconductors, and other developing technologies might approach 100 million MIPS to equal brain capacity.[27] However, the amount of computation necessary to achieve HLMI has been widely contested.[28]

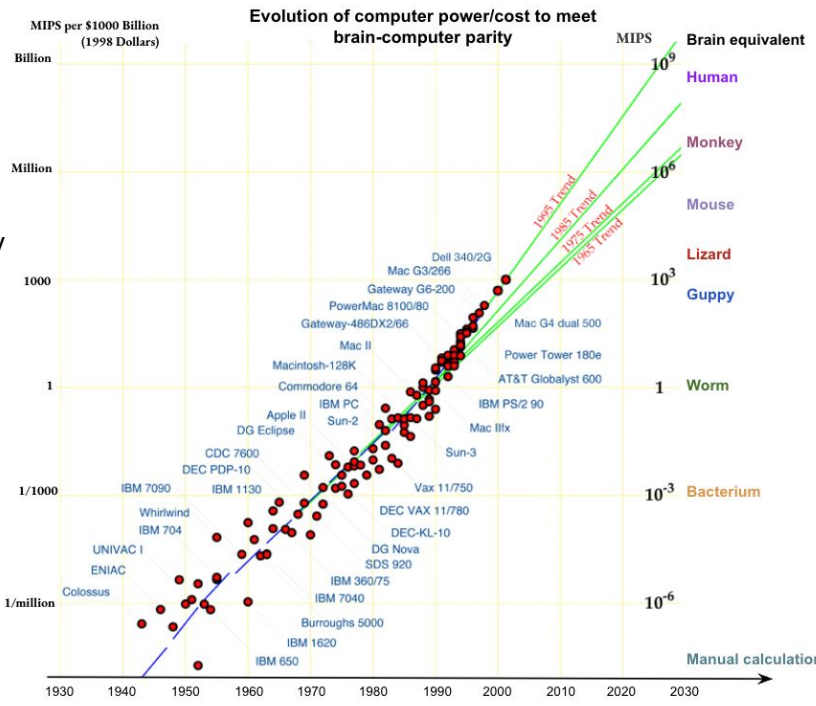

Figure 3: Evolution of computer power / cost, drawing adapted from Hans Moravec (26)

Roadmap to our scenario

ML deep learning models emulate connected brain neurons with several layers of a neural network mimicking synapses.(29) Lower image processing layers, for example, may recognise boundaries, but higher levels may discover significant data such as numbers or faces. When these notions are appropriately recognised, the transistor connections in the firing pathways are strengthened. Neural networks may have 100 billion parameters (layers), however estimates for the human brain are around 86 Billion Neurons and 150 trillion Synapses.(30) Japan's K computer (formerly the world's fastest supercomputer) utilised 82,000 processors to replicate brain function.(31) It took 40 minutes for the entire supercomputer to mimic just one second of brain activity. Rain Neuromorphics are attempting to replicate 150 trillion synapses by connecting enough of its neuromorphic chips together, which have been designed to increase efficiency over 1000x to today's standards.(32) Even with chip architecture breakthroughs, hardware would need to gain processing capacity to perform more total computations per second (cps).(33) The human brain is estimated to operate at one exaflop (a billion billion cps). In 2022 The Frontier supercomputer became the first to break the exascale barrier by achieving 1.102 exaFlops.(34) Quantum computers use quantum mechanics to process exponentially more data than normal computers and may facilitate faster hardware for to develop AGI.(35)

Understanding human hardware to develop HLMI

The hurdles to develop AI up until this point have met the limitation of computation, but to reach AGI the greatest breakthroughs necessary seem to be in understanding the human brain, and essentially: intelligence.

"Innovation is nothing more than a modest reverence for the past and mastery of its achievements."
– René Girard on the Philosophy of Innovation.

It is in the nature of technology to grasp a phenomenon and then build systems to drastically amplify it. Therefore development of AGI is predicated on the premise that a computer should emulate how the human brain functions. A survey of expert opinion on AI progress asked which research approaches would contribute the most to the development of HLMI: Cognitive science and integrated cognitive architecture gained 48% and 42% of respondents' votes.(36)

Meanwhile techniques to understand the human brain are becoming more sophisticated – the cost performance and temporal and spatial resolution of brain scans, the quantity of data and information accessible about brain function, and the sophistication of simulations.(37)

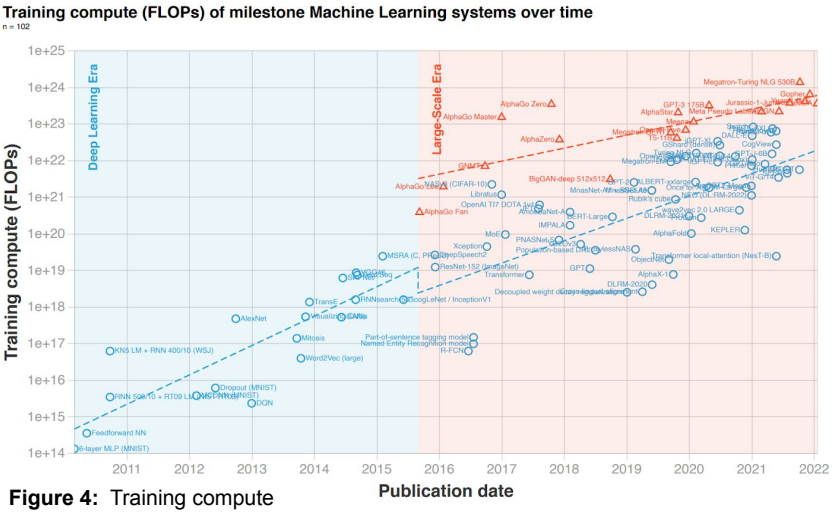

Figure 4: Training compute of 102 milestone ML systems between 2010 and 2022. Source: (38)

Numerous projects are attempting to imitate neocortical function in order to give cognizance to machines; Google Brain, Numenta and EPFL Blue Brain Project, are aiming to build a complete digital recreation of the mammalian brain. Scientists have so far been able to reproduce the brain of a 1-millimetre flatworm, which has 302 neurons.(39-42)

Debate: Could our scenario occur without AGI

Could our scenario occur in the absence of an inflection point toward AGI, such as the theorised deeper understanding of neurobiology?

AGI may not be necessary for our scenario since ANIs could be trained to specific tasks and consequently replace many human jobs. Generality as an indicator of AI sophistication may be a flawed concept since no intelligent system with finite geographical, temporal, and energetic resources can be entirely general. Humans have greater 'general intelligence' than any known AI system, yet they aren't particularly adept at high speed processing of computational problems simultaneously.

Every time our ability to access information and to communicate it to others is improved, we have achieved “Intelligence Amplification” (IA). If the scenario were to take place without development of AGI, it would likely arrive much slower and require more human contact to make decisions on the basis of deploying a wider array of ANIs. Human-machine collaboration is critical because people give advice by labelling data from which AI machines may learn. So far, AI has augmented rather than replaced human talents. In this IA systems point of view, Vernor Vinge in 1993 developed assumptions for possible future scenarios with IA that have been realised in the 21st century: (19)

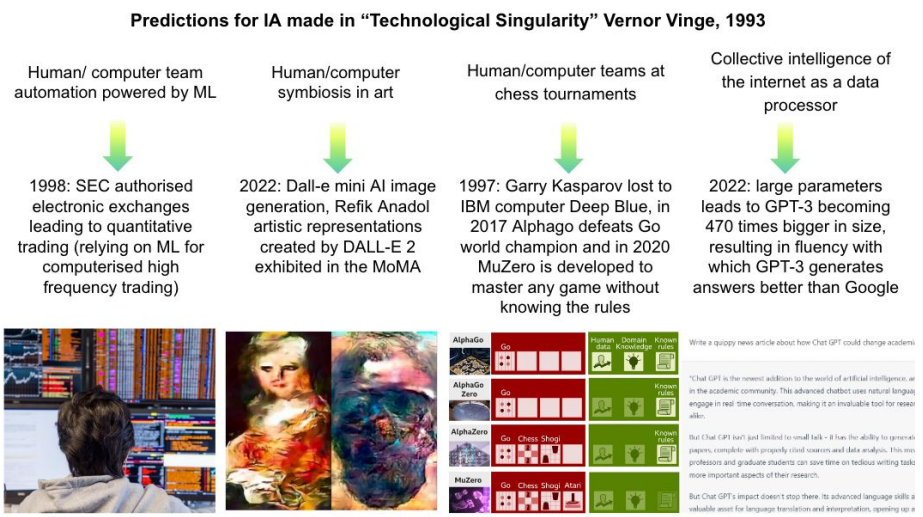

Figure 5: AI advances that brought 1993 IA predictions (19) to reality

Roadmap to our scenario

Is AGI needed for our scenario This presents a potential re-routing of our scenario; where human intelligence is amplified by AI/ML tools. What would Human Purpose look like if we were utilising AI tools rather than seeing ourselves as displaced by them?

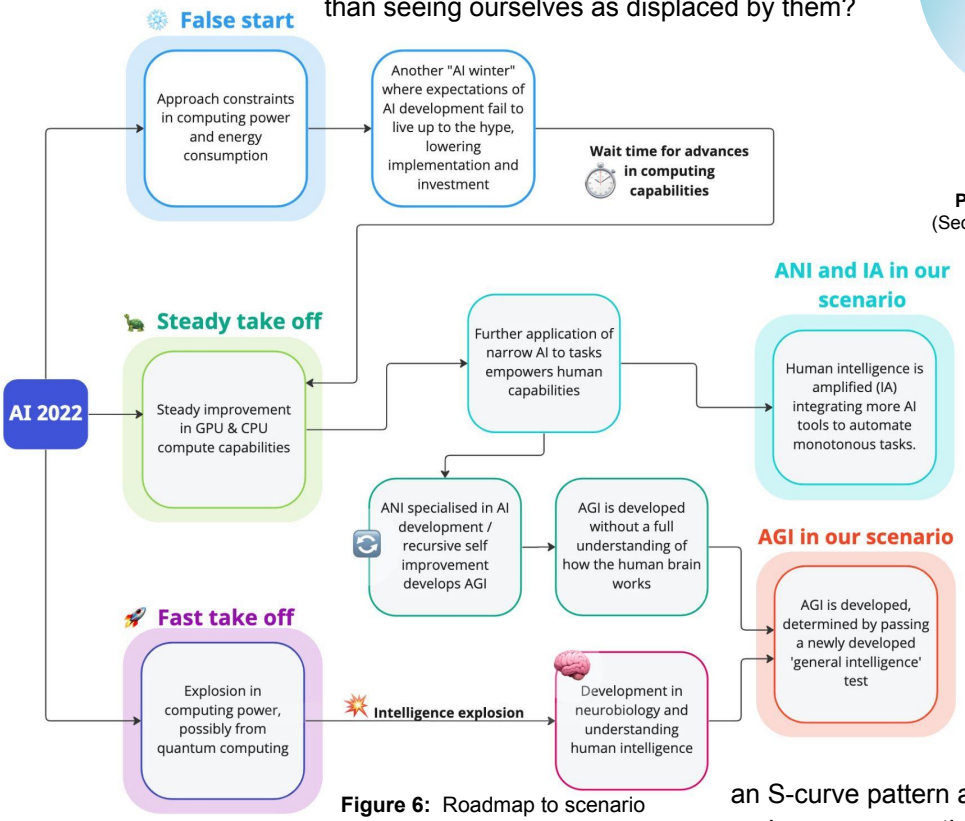

Figure 6: Roadmap to scenario

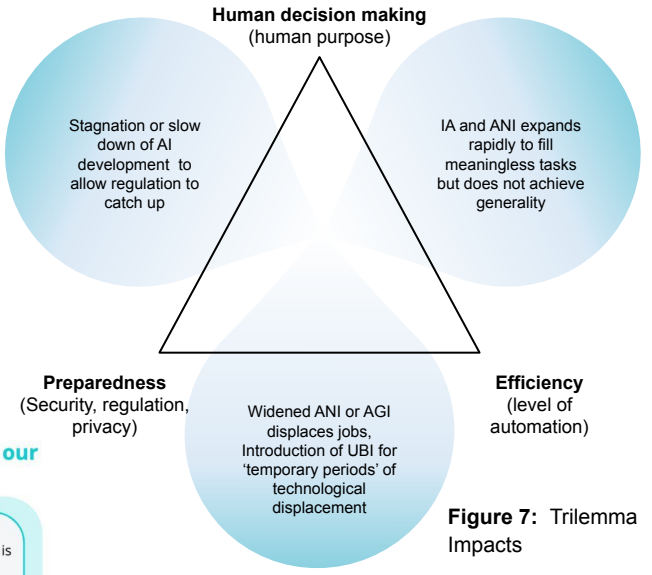

Figure 7: Trilemma Impacts

Trilemma Impacts AGI could take many multiple different forms: it could be a service much like how we use Google, housed within robots or within the infrastructure of a city. It is anticipated that the implementation of AGI would result in GDP doubling every two years, as compared to 4.3% per year or doubling every sixteen years.(44) Economic impacts could be dependent on adoption timelines, a slow takeoff scenario for example would present an S-curve pattern as the adoption of increasingly sophisticated ANI tools grows over time. Slower adoption allows for regulation to catch up to the speed of AI development, but sacrifices efficiency. The Trilemma model depicts the trade-off between the level of human-decision making that remains vs efficiency and preparedness.

Counter Argument: AGI is inevitable

There is little to convince against the fact that ANI would be able to apply the same level of rapid learning and self-teaching to generalised programming and algorithm development. Yudkowsky suggests that if an AI became capable of directly editing itself, this could spark a rapid increase in intelligence, as the actual process causing increases in intelligence could itself be improved upon.(43) Although a false-start or steady-takeoff scenario (figure 6) seem to permit human interaction during the ascent to AGI, they might result in AGI being developed without the understanding of intelligence that would be gained from focused development of AGI – such as a breakthrough in cognitive neuroscience.

Either route could enable a speed explosion whereby intelligent machines develop ever-faster devices. There might also be an intelligence explosion, in which one AGI discovers how to make a qualitatively smarter AGI, and that AGI exploits its newfound intelligence to generate even more clever AGIs, and so on, leaving human intellect well behind.(43) In this sense it seems AGI is inevitable, wider application of ANI might be an indicator of AGI arrival rather than a sign that the world will forever take the route of IA, as once such a tool becomes available it would be impossible to stop its escalation.

When will we be able to tell when we've attained AGI?

With the above roadmap to our scenario, key indicators would be widened application of ANI tools or when AGI is achieved. The Turing test, devised in 1950, is one of the most well-known assessments of human and computer intelligence but has been achieved by an AI in 2014.(45) Hector Levesque proposed multiple-choice Winograd schemas as a substitute to the Turing test in 2011. By 2019, a handful of AI systems based on massive pre-trained transformer-based language models that had been fine-tuned for these types of challenges reached greater than 90% accuracy.(46)

It is difficult to quantify intelligence; Intelligence conceptions based on goal attainment or reward optimization fail to capture important properties of human intelligence.(6) Steve Wozniak proposed the coffee test for determining AGI; suggesting that a programme should be able to enter any kitchen, identify the required supplies, and accomplish the task of preparing a coffee. This requires a system to be multi-modal, capable of generalising across tasks and orchestrating a sequence of operations.(47) Figure 8 proposes a comprehensive system designed to determine the strength of AGI.

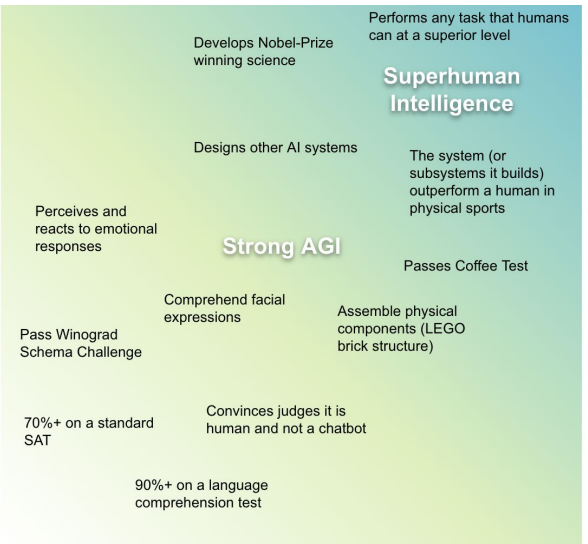

Figure 8: System comprising of multiple tests to determine AGI

Assessing the impacts

Likelihood of automation (LOA)

Studies estimate the likelihood of occupations being automated find that automation could displace 850,000 UK jobs by 2030, 47% of US jobs before 2033, and 137 million workers in Southeast Asia by 2036.[48-50] In 2013 Frey and Osborne applied a ML algorithm to assess how easily 702 different kinds of jobs in America could be automated - concluding 47% of jobs could be within 10-20 years.[51] Based on this study, an OECD paper[52] estimated that 14% of employment in 32 countries are highly susceptible (70% LOA), and 32% were slightly less vulnerable (50-70% LOA). At current employment rates, this puts 210 million jobs at danger across the 32 nations studied. However these estimates are based on O*NET variables that define jobs involving perception, dexterity, creative intelligence and social intelligence are unlikely to be automated.[53] These assumptions may be suitable in the instance of widened ANI and IA in our scenario, whereas AGI may be able to outcompete in these tasks depending on its sophistication

Tacit knowledge defines tasks that can be undertaken without thinking about it such as driving a car – and cannot be represented by a simple formula. Explicit knowledge is information can be relayed by a set of instructions. Therefore boundary between explicit and tacit knowledge is used to define limits of automation and thus the exclusive precincts of humans.The advanced tasks described as being beyond the grasp of automation by O*NET, tacit vs explicit knowledge or degree of cognition (Figure 10) therefore would include driving, medical diagnosis and education — which involve a combination of mental and manual labour, but relied on tacit knowledge. This has been challenged by emerging AI capabilities such as Google's car and medical diagnostic algorithms – resetting the boundaries to the limits of automation.[54,55]

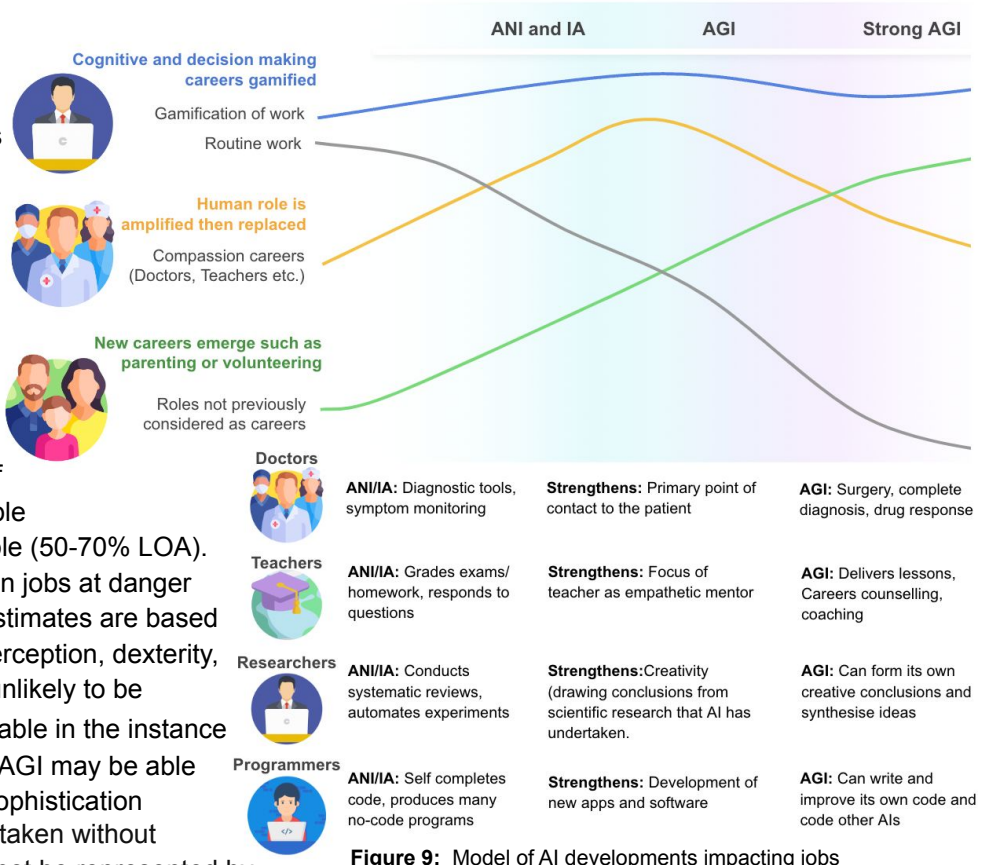

Figure 9: Model of AI developments impacting jobs

Modelling the future of work: Predicting the capabilities AGI will fall short of is important; In order to understand which jobs may be displacement-resistant, and train more people for them. As outlined in this report, three key skills prove hard to automate: Creativity, Empathy and Dexterity. Until AGI reaches superintelligence, it cannot feel or interact with feelings like empathy and compassion. Figure 9 visualises how certain careers might change over time in our scenario – As ANI/AGI become more prevalent in our situation, it will do routine tasks alongside humans, strengthening human roles where warmth and compassion are required. As stronger AGI advances, this trend could reverse since AGI will become more capable than humans, even in those skills we associate as uniquely human. Economic growth and prosperity would invariably result in more jobs, including new ones we cannot yet predict. For example, mapping human games to real-world tasks may be an emerging technology. In the long-term, AGI might envelop all human occupations resulting in jobs that become novel games between humans like chess. In this case work may be optional, and will likely involve the introduction of UBI as seen in our scenario. We may reconsider what constitutes paid work — volunteering, or caretaking duties, such as foster parenting, home-schooling, or caring for a dependent family member, may be considered full-time occupations in the future.

Categorising LOA of by the degree of repetition and cognition involved deems Blue-collar work more at risk since repetitive jobs fall inside the band of tacit knowledge. These assumptions result in large variance among countries: locations with lower GDP per person appear to have a larger percentage of total occupations at risk of automation.[52] Resultantly, the productivity increases seen by businesses that can invest in the emerging ANI/ AGI technologies may expand the wealth disparity gap. These assumptions also suggest creativity and interpersonal abilities are more robust. But much like how chess was once perceived as a uniquely human creative skill, it becomes challenged by recent advances such as GTP-3 and DALL-E mini, which have successfully generated musical compositions and art exhibitions.[56-58] Similarly, occupations in AI/ML development that are expected to flourish will most likely be displaced - especially if AI development follows the path of recursive self improvement.[59]

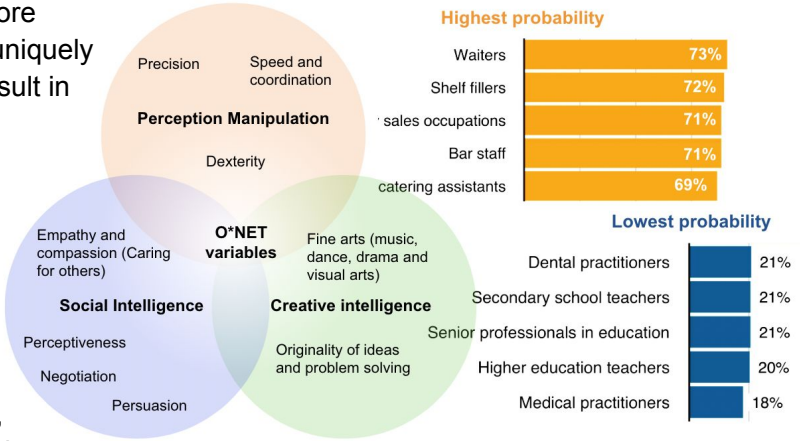

Figure 10: O*NET variables contributing to % probability of automation. Created in google drawings, data source: (53)

Adapting to the future of work and education

Upskilling to a knowledge-based economy

According to the World Economic Forum, by 2025, 50% of all employees will require reskilling from job displacement.[60] Japan and China have made significant investments in AI education and worker training as the cornerstone of long-term national AI ambitions.[61] In contrast the United States invests just around 0.1 percent of its GDP on initiatives that assist workers in adapting to job changes.[62] However most upskilling proposals as a solution to focus on training individuals in AI/ML computer science. Since these fields are themselves not safe from automation, education should be broader. Studies imply that automation propensity is inversely correlated with the level and length of education[64] signifying that education in general would ease the transition to temporary periods of technical displacement.

The Trifecta for Innovation model (figure 11) investigates the motivating drivers behind the various ANI/AGI products that would form our scenario. The feasibility factors have been investigated as the likely developments to enable AGI, while the viability factors include economic and business incentives. Finally, desirability is influenced by societal views and level of knowledge about AI.

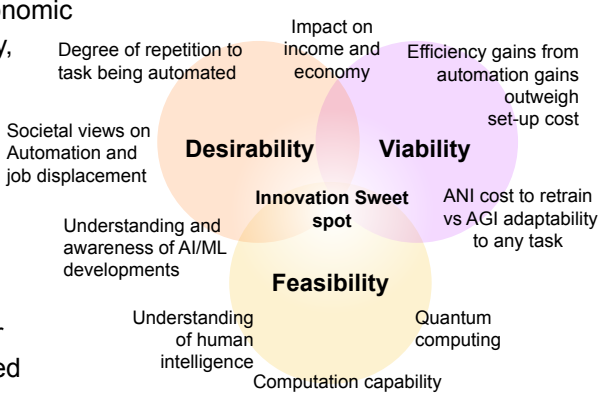

Figure 11: Trifecta for innovation in ANI and AGI

Delivering AI and IA

The path of development toward AI in our scenario – whether developed internally or in an open-source democratised manner – would result in vastly different worlds. Developing AGI will necessitate immense computing power, which will most likely be generated by a few well-funded organisations. Is democracy conceivable in a society dominated by a few "genius" systems run by a few organisations? Who controls AGI? Who grants access to it? By making AI advances available as open-source code, we may accelerate innovation and reduce the impact on groups with lower AI knowledge and power. For example, Stability AI offered a $200,000 reward in a competition to discover who could create the greatest deep fake detection solution and release it for free.[64]

The Chord diagram (figure 12) has been used to illustrate relationships between those that build and use the AI developments; businesses and governments would likely hold the most power as they develop AI to be used by individuals (either as products they are aware of or unseen entities). Businesses will also develop AI for other organisations to use internally which will have second-order effects on individuals if these tools serve to displace their roles. Developers may release developments as open source for anyone to access which will democratise the process. Scientists also have power over the developments that will enable developers/business to build AI.

Groups with little technological advancement will feel most unfulfilled Optimising the precision and efficiency of decision making is excellent for society, but for the individual, meaning is derived from autonomy and the capacity to explain the world around them. In most countries, a small but rising group of people comprehend AI – but, a larger number of people may not comprehend why AI does what it does, reducing their feeling of autonomy and capacity to attach meaning to the environment. By incorporating AI into consumer products (including those that automate tasks at work), it will enable benefits of the technology to be widely distributed. AI will also manage networks and systems that are beyond the control of any one user. In many circumstances, experiences with AI may be unsettling or disempowering, especially for people who lack technical competence. Those with little knowledge of AI or power over it may be more prone to dismiss it. This requires a reform of the labour system to a digital knowledge-based economy. AI tools themselves could be utilised to transform education in this world.

In 2020 the USC initiative used AI to build holograms of Holocaust survivor accounts that audiences may engage with for years to come.[62] GTP-3 can already respond to prompts in the style of a particular author or fictional character; these developments could be combined to deliver knowledge from any domain. Such a tool could spark a new renaissance based on human expression and creativity.

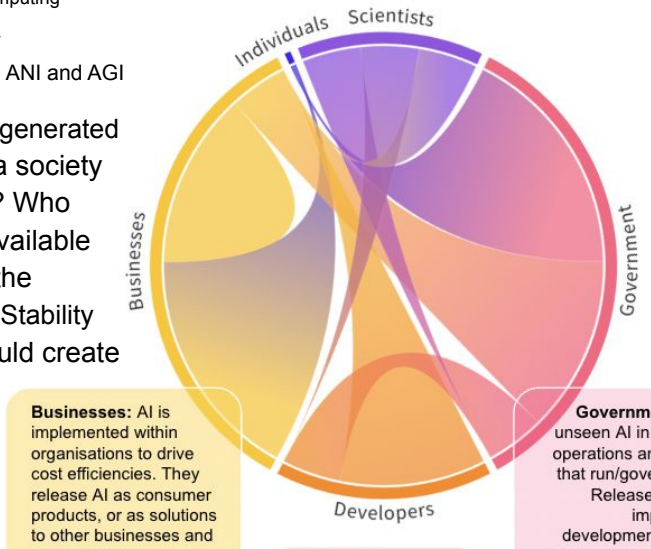

Businesses: AI is implemented within organisations to drive cost efficiencies. They release AI as consumer products, or as solutions to other businesses and organisations.

Government: Builds unseen AI in underlying operations and systems that run/govern society. Release regulation impacting the development of AI and individuals affected by it.

Developers: Produce AI/ML for organisations and government. And release open-source AI developments for any individuals to access.

Individuals: AI is implemented within organisations to drive cost efficiencies. They release AI as consumer products, or as solutions to other businesses and organisations.

Scientists: deliver developments in computation and cognitive neuroscience to develop AGI for businesses, governments, and inform developers.

Figure 12: Chord diagram (created in python) showing builders/users of AI

Conclusion: Interventions on the risk of AI and Job displacement thus far, have focused on upskilling programs largely focused around preparing workers in AI/ML programming skills. In our scenario we describe a world that challenges the view of coding capabilities being in high demand and more importantly, characterise the problem not of job loss to automation, but loss of human meaning and purpose. In preparation for this world, we propose a redirection of educational and upskilling programs to be focused towards human meaning and compassion.

References

All diagrams and figures have been generated in google drawings or in python unless otherwise stated. Data sources used to generate these figures have been referenced.

Title page comic book images generated with GTP-3

1. The Human Brain is Not Obsolete Vol. 38, No. 2 - [Internet]. RBC. 1957 [cited 2023 Jan 11]. Available from: https://www.rbc.com/en/about-us/history/letter/february-1957-vol-38-no-2-the-human-brain-is-not-obsolete/
2. Klein C. The Original Luddites Raged Against the Machine of the Industrial Revolution [Internet]. HISTORY. 2019. Available from: https://www.history.com/news/industrial-revolution-luddites-workers
3. Wagner B. Liable, but Not in Control? Ensuring Meaningful Human Agency in Automated Decision-Making Systems. Policy & Internet. 2019 Jan 24;11(1):104–22.
4. McLuhan M, McLuhan E. Laws of Media: The New Science [Internet]. Google Books. University of Toronto Press; 1988 [cited 2023 Jan 11]. Available from: https://books.google.co.uk/books/about/Laws_of_Media.html?id=4lkNAQAAMAAJ&redir_esc=y
5. Bostrom N. The Future of Humanity [Internet]. Future of Humanity Institute Faculty of Philosophy & James Martin 21st Century School , Oxford University. New Waves in Philosophy of Technology, eds. Jan-Kyrre Berg Olsen, Evan Selinger, & Soren Riis (New York: Palgrave McMillan, 2009): 186-216]; 2007. Available from: https://nickbostrom.com/papers/future
6. Korteling JE (Hans)., van de Boer-Visschedijk GC, Blankendaal RAM, Boonekamp RC, Eikelboom AR. Human- versus Artificial Intelligence. Frontiers in Artificial Intelligence. 2021 Mar 25;4.
7. Dormehl L. Thinking machines : the quest for artificial intelligence--and where it's taking us next. New York: Tarcherperigee; 2017.
8. Hunter B, Hindocha S, Lee RW. The Role of Artificial Intelligence in Early Cancer Diagnosis. Cancers. 2022 Mar 16;14(6):1524.
9. New AI Model Translates 200 Languages, Making Technology Accessible to More People [Internet]. Meta. 2022 [cited 2023 Jan 11]. Available from: https://about.fb.com/news/2022/07/new-meta-ai-model-translates-200-languages-making-technology-more-accessible/
10. MuZero: Mastering Go, chess, shogi and Atari without rules [Internet]. www.deepmind.com. Available from: https://www.deepmind.com/blog/muzero-mastering-go-chess-shogi-and-atari-without-rules
11. Khurana D, Koli A, Khatter K, Singh S. Natural language processing: state of the art, current trends and challenges. Multimedia Tools and Applications. 2022 Jul 14;
12. Sarker IH. Machine Learning: Algorithms, Real-World Applications and Research Directions. SN Computer Science [Internet]. 2021 Mar 22;2(3). Available from: https://link.springer.com/article/10.1007/s42979-021-00592-x
13. Global Algorithmic Trading Market to Surpass US$ 21,685.53 Million by 2026 [Internet]. www.businesswire.com. 2019. Available from: https://www.businesswire.com/news/home/20190205005634/en/Global-Algorithmic-Trading-Market-to-Surpass-US-21685.53-Million-by-2026#:~:text=Algorithmic%20trading%20is%20responsible%20for
14. Human-run hedge funds outperform robot quants in pandemic year [Internet]. amwatch.com. 2020 [cited 2023 Jan 11]. Available from: https://amwatch.com/AMNews/Fund_Management/article12659643.ece
15. Bayesian Reasoning 1.1 Reasoning under uncertainty [Internet]. [cited 2023 Jan 11]. Available from: https://bayesian-intelligence.com/publications/bai/book/BAI_Chapter1.pdf
16. Eberding LM, Belenchia M, Sheikhlar A, Thórisson KR. About the Intricacy of Tasks. Artificial General Intelligence. 2022;65–74.
17. Özkural E. Measures of Intelligence, Perception and Intelligent Agents. Artificial General Intelligence. 2022;174–83.
18. Rathi S. Approaches to Artificial General Intelligence: An Analysis [Internet]. [cited 2023 Jan 11]. Available from: https://arxiv.org/pdf/2202.03153.pdf
19. Vinge V. TECHNOLOGICAL SINGULARITY by Vernor Vinge [Internet]. frc.ri.cmu.edu. 1993. Available from: https://frc.ri.cmu.edu/~hpm/book98/com.ch1/vinge.singularity.html
20. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. Communications of the ACM. 2012 May 24;60(6):84–90.
21. Grace K, Salvatier J, Dafoe A, Zhang B, Evans O. Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts. Journal of Artificial Intelligence Research [Internet]. 2018 Jul 31;62(62):729–54. Available from: https://arxiv.org/pdf/1705.08807.pdf
22. Fjelland R. Why general artificial intelligence will not be realized. Humanities and Social Sciences Communications [Internet]. 2020 Jun 17;7(1). Available from: https://www.nature.com/articles/s41599-020-0494-4
23. When will the first general AI system be devised, tested, and publicly announced? [Internet]. www.metaculus.com. 2020. Available from: https://www.metaculus.com/questions/5121/date-of-artificial-general-intelligence/
24. After a (weak) AGI is created, how many months will it be before the first superintelligent AI is created? [Internet]. www.metaculus.com. 2021 [cited 2023 Jan 11]. Available from: https://www.metaculus.com/questions/9062/time-from-weak-agi-to-superintelligence/
25. Amodei D. AI and Compute [Internet]. OpenAI. OpenAI; 2018. Available from: https://openai.com/blog/ai-and-compute/
26. Hans Moravec. When will computer hardware match the human brain? [Internet]. Robotics Institute Carnegie Mellon University, editor. Journal of Evolution and Technology. Vol. 1; 1998. Available from: https://jetpress.org/volume1/moravec.htm
27. Fan D, Wang J, Wang E, Dong S. Propelling DNA Computing with Materials' Power: Recent Advancements in Innovative DNA Logic Computing Systems and Smart Bio‐Applications. Advanced Science. 2020 Nov 9;7(24):2001766.
28. Collins H. Hubert L. Dreyfus, Forms of Life, and a Simple Test for Machine Intelligence. Social Studies of Science [Internet]. 1992 [cited 2023 Jan 11];22(4):726–39. Available from: https://www.jstor.org/stable/285462
29. Lotter W, Kreiman G, Cox D. A neural network trained for prediction mimics diverse features of biological neurons and perception. Nature Machine Intelligence [Internet]. 2020 Apr 20 [cited 2023 Jan 11];2(4):210–9. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8291226/
30. Herculano-Houzel S. The Human Brain in numbers: a Linearly scaled-up Primate Brain. Frontiers in Human Neuroscience [Internet]. 2009;3(31). Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2776484/
31. January 2014 TG 14. Supercomputer Takes 40 Minutes To Model 1 Second of Brain Activity [Internet]. livescience.com. Available from: https://www.livescience.com/42561-supercomputer-models-brain-activity.html

References

32. Koetsier J. 1000X More Efficient Neural Networks: Building An Artificial Brain With 86 Billion Physical (But Not Biological) Neurons [Internet]. Forbes. [cited 2023 Jan 11]. Available from: https://www.forbes.com/sites/johnkoetsier/2022/03/25/1000x-more-efficient-neural-networks-building-an-artificial-brain-with-86-billion-physical-but-not-biological-neurons/?sh=25b99316408e
33. Perez CE. Revisiting the AGI Capability Roadmap [Internet]. Intuition Machine. 2021 [cited 2023 Jan 11]. Available from: https://medium.com/intuitionmachine/revisiting-the-agi-capability-roadmap-a89a533e5e4c
34. ORNL's Frontier First to Break the Exaflop Ceiling | TOP500 [Internet]. www.top500.org. Available from: https://www.top500.org/news/ornls-frontier-first-to-break-the-exaflop-ceiling/
35. IBM. What is Quantum Computing? | IBM [Internet]. www.ibm.com. 2022. Available from: https://www.ibm.com/topics/quantum-computing
36. Müller VC, Bostrom N. Future progress in artificial intelligence. AI Matters [Internet]. 2014 Sep 17;1(1):9–11. Available from: https://nickbostrom.com/papers/survey.pdf
37. XUE G, CHEN C, LU Z-L, DONG Q. Brain Imaging Techniques and Their Applications in Decision-Making Research. Acta Psychologica Sinica [Internet]. 2010 Feb 3;42(1):120–37. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2849100/
38. Sevilla J, Heim L, Ho A, Besiroglu T, Hobbhahn M, Villalobos P. COMPUTE TRENDS ACROSS THREE ERAS OF MACHINE LEARNING [Internet]. [cited 2023 Jan 11]. Available from: https://arxiv.org/pdf/2202.05924.pdf
39. The Thousand Brains Theory of Intelligence [Internet]. Numenta. [cited 2023 Jan 11]. Available from: https://www.numenta.com/technology/research/thousand-brains-theory/
40. Blue Brain's Scientific Milestones [Internet]. www.epfl.ch. Available from: https://www.epfl.ch/research/domains/bluebrain/blue-brains-scientific-milestones/
41. Brain Team – Google Research [Internet]. Google Research. Available from: https://research.google/teams/brain/
42. Fessenden M. We've Put a Worm's Mind in a Lego Robot's Body [Internet]. Smithsonian Magazine. Available from: https://www.smithsonianmag.com/smart-news/weve-put-worms-mind-lego-robot-body-180953399/?no-ist
43. Yudkowsky E. Artificial Intelligence as a Positive and Negative Factor in Global Risk [Internet]. 2008. Available from: https://intelligence.org/files/AIPosNegFactor.pdf
44. Hanson R. Economic Growth Given Machine Intelligence [Internet]. [cited 2023 Jan 11]. Available from: https://mason.gmu.edu/~rhanson/aigrow.pdf
45. Turing AM. Computing Machinery and Intelligence. Mind. 1950;LIX(236):433–60.
46. The Winograd Schema Challenge [Internet]. cs.nyu.edu. Available from: https://cs.nyu.edu/~davise/papers/WinogradSchemas/WS.h
47. Shick M. Wozniak: Could a Computer Make a Cup of Coffee? [Internet]. Fast Company. 2010. Available from: https://www.fastcompany.com/1568187/wozniak-could-computer-make-cup-coffee
48. Experts State Robots Will Take Over Additional 850,000 Jobs By 2030 [Internet]. Futurism. Available from: https://futurism.com/experts-state-robots-will-take-over-additional-850000-jobs-by-2030
49. Frey CB, Osborne MA. The Future of employment: How Susceptible Are Jobs to computerisation? Technological Forecasting and Social Change [Internet]. 2013 Sep 17;114(1):254–80. Available from: https://www.oxfordmartin.ox.ac.uk/downloads/academic/T
50. 137 Million Southeast Asian Workers Could Lose Their Jobs to Automation in the Next 20 Years [Internet]. Futurism. [cited 2023 Jan 11]. Available from: https://futurism.com/137-million-southeast-asian-workers-could-lose-their-jobs-to-automation-in-the-next-20-years
51. Frey CB, Osborne MA. The Future of employment: How Susceptible Are Jobs to computerisation? Technological Forecasting and Social Change [Internet]. 2013 Sep 17;114(1):254–80. Available from: https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf
52. Automation, skills use and training. OECD Social, Employment and Migration Working Papers. 2018 Mar 8;
53. Work Context: Degree of Automation [Internet]. www.onetonline.org. Available from: https://www.onetonline.org/find/descriptor/result/4.C.3.b.2
54. Press A. Google's Waymo to offer driverless ride-hailing service in San Francisco [Internet]. the Guardian. 2022. Available from: https://www.theguardian.com/technology/2022/mar/30/waymo-self-driving-ride-hailing-service-san-francisco-alphabet-google
55. Longoni C, Morewedge CK. AI Can Outperform Doctors. So Why Don't Patients Trust It? [Internet]. Harvard Business Review. 2019. Available from: https://hbr.org/2019/10/ai-can-outperform-doctors-so-why-dont-patients-trust-it
56. Dhariwal P, Jun H, Payne C, Kim JW, Radford A, Sutskever I. Jukebox: A Generative Model for Music. arXiv:200500341 [cs, eess, stat] [Internet]. 2020 Apr 30; Available from: https://arxiv.org/abs/2005.00341
57. Ramesh A, Pavlov M, Goh G, Gray S, Voss C, Radford A, et al. Zero-Shot Text-to-Image Generation. arXiv:210212092 [cs] [Internet]. 2021 Feb 26; Available from: https://arxiv.org/abs/2102.12092
58. Caulfield B. MoMA Installation Marks Breakthrough for AI Art [Internet]. NVIDIA Blog. 2022 [cited 2023 Jan 11]. Available from: https://blogs.nvidia.com/blog/2022/11/17/moma-ai-art/
59. AI learns to write computer code in "stunning" advance [Internet]. www.science.org. Available from: https://www.science.org/content/article/ai-learns-write-computer-code-stunning-advance
60. These are the top 10 job skills of tomorrow – and how long it takes to learn them [Internet]. World Economic Forum. [cited 2023 Jan 11]. Available from: https://www.weforum.org/agenda/2020/10/top-10-work-skills-of-tomorrow-how-long-it-takes-to-learn-them/#:~:text=For%20those%20workers%20who%20stay
61. Managing Japan's Shrinking Labor Force With AI and Robots - IMF F&D Magazine - June 2018 | Volume 55 | Number 2 [Internet]. IMF. Available from: https://www.imf.org/en/Publications/fandd/issues/2018/06/japan-labor-force-artificial-intelligence-and-robots-schneider
62. Lee K. Artificial Intelligence, Automation, and the Economy [Internet]. whitehouse.gov. Whitehouse; 2016. Available from: https://obamawhitehouse.archives.gov/blog/2016/12/20/artificial-intelligence-automation-and-economy
63. Albuquerque PHM, Saavedra CAPB, Morais RL de, Peng Y. The Robot from Ipanema goes Working: Estimating the Probability of Jobs Automation in Brazil. Latin American Business Review [Internet]. 2019 [cited 2023 Jan 11];20(3):227. Available from: https://www.academia.edu/39760365/The_Robot_from_Ipanema_goes_Working_Estimating_the_Probability_of_Jobs_Automation_in_Brazil
64. Jeffries D. Why the Future of Open Source AI is So Much Bigger Than Stable Diffusion 1.5 and Why It Matters to You [Internet]. danieljeffries.substack.com. Available from: https://danieljeffries.substack.com/p/why-the-future-of-open-source-ai

Appendix

Year of prediction	Respondant average prediction for AGI arrival	Number of respondants	Source
2007	Klein AGI Survey, 2007 **2050**	888 Respondants	Klein AGI Survey [Internet]. AI Impacts. 2014 [cited 2023 Jan 11]. Available from: https://aiimpacts.org/klein-agi-survey/
2009	AGI conference, 2009 **2075**	21 Respondants	Baum SD, Goertzel B, Goertzel TG. How long until human-level AI? Results from an expert assessment. Technological Forecasting and Social Change. 2011 Jan;78(1):185–95.
2012	Muller & Bostrom Survey, 2012 **2050**	550 Respondants	Müller VC, Bostrom N. Future progress in artificial intelligence. AI Matters [Internet]. 2014 Sep 17;1(1):9–11. Available from: https://nickbostrom.com/papers/survey.pdf
2017	Ray Kuzweil, 2017 **2045**	Individual expert opinion	Reedy C. Kurzweil Claims That the Singularity Will Happen by 2045 [Internet]. Futurism. Futurism; 2017. Available from: https://futurism.com/kurzweil-claims-that-the-singularity-will-happen-by-2045
2017	AI expert survey, 2017 **2060**	325 Respondants	Grace K, Salvatier J, Dafoe A, Zhang B, Evans O. Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts. Journal of Artificial Intelligence Research [Internet]. 2018 Jul 31;62(62):729–54. Available from: https://arxiv.org/pdf/1705.08807.pdf
2018	Louis Rosenberg, 2018 **2030**	Individual expert opinion	Rosenberg L. METAVERSE 2030 [Internet]. Predict. 2022. Available from: https://medium.com/predict/metaverse-2030-ee59e4d4010d
2018	Patrick Winston, 2018 **2040**	Individual expert opinion	Separating science fact from science hype: How far off is the singularity? [Internet]. Futurism. Available from: https://futurism.com/separating-science-fact-science-hype-how-far-off-singularity
2018	Jürgen Schmidhuber, 2018 **2050**	Individual expert opinion	The “father of artificial intelligence” says singularity is 30 years away [Internet]. Futurism. Available from: https://futurism.com/father-artificial-intelligence-singularity-decades-away
2019	Expert survey on Progress in AI, 2019 **2060**	32 Respondants	Testbanking D. When Will We Reach the Singularity? - A Timeline Consensus from AI Researchers (AI FutureScape 1 of 6) [Internet]. Emerj. Available from: https://emerj.com/ai-future-outlook/when-will-we-reach-the-singularity-a-timeline-consensus-from-ai-researchers/
2022	Jeff Clune, 2022 **2030**	Individual expert opinion	Machines that think like humans: Everything to know about AGI and AI Debate 3 [Internet]. ZDNET. [cited 2023 Jan 11]. Available from: https://www.zdnet.com/article/ai-debate-3-everything-you-need-to-know-about-artificial-general-intelligence/
2022	Shane Legg, 2022 **2030**	Individual expert opinion	The road to AGI - DeepMind: The Podcast (S2, Ep5) [Internet]. www.youtube.com. [cited 2023 Jan 11]. Available from: https://www.youtube.com/watch?v=Uy4OYU7PQYA&ab_channel=DeepMind
2022	Expert Survey on Progress in AI, 2022 **2059**	738 Respondants	2022 Expert Survey on Progress in AI [Internet]. AI Impacts. 2022 [cited 2023 Jan 11]. Available from: https://aiimpacts.org/2022-expert-survey-on-progress-in-ai/

Source Table for AI expert AGI predictions analysis