

ASI - Are We Nearing the Next Stage of Human Evolution, the Transhumans

Jisha M.Chacko , Resmi Das

Assistant Professor, Airforce Technical College, Seeb Muscat

Assistant Professor, Airforce Technical College, Seeb Muscat

ABSTRACT: We are rapidly becoming dependent on computers to outsource our memory and on the Internet to augment our knowledge. The computer is becoming an intelligent machine and there are indications that its level of intelligence may eventually surpass that of its human creator. We may well be heading toward an artificial Intelligence Explosion. If such a scenario comes true then Artificial Intelligence (AI) will transition in rapid succession to what is now commonly referred to as Artificial General Intelligence (AGI) and then exponentially to Artificial Super intelligence (ASI). AGI will have equal or superior intelligence to the human in most domains and ASI is expected to be many orders of magnitude more intelligent than the human. Within this context there has emerged a relatively small but growing group of prominent computer and cognitive scientists, philosophers, sociologists, and psychologists who believe that this unavoidable transition from AI to AGI to ASI is fraught with danger and may in fact lead to the eventual demise of the human species.

However, there is a lot of theory that anticipates artificial super intelligence coming sooner rather than later. Using examples like Moore's law, which predicts an ever-increasing density of transistors, experts talk about singularity and the exponential growth of technology, in which full artificial intelligence could manifest within a number of years, and artificial super intelligence could exist in the 21st century.

Contrary to popular belief, ASI is far from reality. In fact, it may take anything between 15–20 years (best case scenario) to centuries to achieve it! To understand all the aforementioned points, we need to look into why there is a sudden boom in AI and what the current state of AI research is. In this journal an attempt has been made, to look at the roads to achieve ASI, and the dangers it possess.

KEYWORDS: Artificial intelligence, super intelligence, technological singularity, internet explosion, neural networks, deep learning, machine learning.



I. INTRODUCTION

We stand on the edge of a crucial moment in the history of our species – a time when a creation of our own inventiveness has the potential to change everything. For some, it will be humanity's salvation, while for others, it could be our downfall. We are entering the age of artificial intelligence (AI).

While AI is still some way from the sentient machines portrayed in science fiction, the creation of algorithms that can learn, understand language and mimic some aspects of the human mind have led to huge advances. Today AI is being used in hundreds of different industries.

Artificial intelligence will probably be one of the biggest scientific breakthroughs of the 21st century .It will give us the power to probe the universe and our humanity with a different approach. AI has the potential of forever changing our humanity. As we become more and more reliant on phones and devices to aid our everyday tasks rapid development into new technologies is underway at an alarming rate .Modern devices are specifically designed to interact with us in ways that mimic our real human behavior. Applications such as Siri and Alexa are providing the user with a human like

iteration experience. These companies are creating artificial intelligent machines, machines exhibiting cognitive behavior with human like intelligence.

All these means that AI is no longer a futuristic technology but is increasingly integrated into every realm of our lives. From suggesting what books we might like to buy online to powering the virtual assistants that inhabit our phones and smart speakers, some of the applications are more visible than others. In truth, AI is touching our lives far more than many of us realize.

1.1 The Evolution of Artificial Intelligence

This world has seen four major revolutions that changed its entire face. The first revolution was in 1784 when the first steam engine introduced in the world. The second one was in 1870 when the electricity was invented. The third one was in 1969 when the word information technology was introduced in the world, and the fourth one is the revolution of Artificial Intelligence which we are experiencing right now. The present revolutionary era is based on the extreme automation and global connectivity for which Artificial Intelligence is imperative. Just like the other three revolutions, the changes and developments that this revolution is creating will be the bedrock of our future and our way of interacting with technology and with each other.

Less than a decade after breaking the Nazi encryption machine Enigma, and helping the Allied Forces win World War II, mathematician Alan Turing changed history a second time with a simple question: "Can machines think?"

Turing's paper "Computing Machinery and Intelligence" (1950), and it's subsequent Turing Test, established the fundamental goal and vision of artificial intelligence. To be programmed to be classified as AI, it has to pass the Turing test. He determined that if 30% of humans who interacted with the machines believed that they were actually interacting with the human being then the computer is qualified as AI.

The expansive goal of artificial intelligence has given rise to many questions and debates. So much so, that no singular definition of the field is universally accepted.

2 How Artificial Intelligence Works

AI works by combining large amounts of data with fast, iterative processing and intelligent algorithms, allowing the software to learn automatically from patterns or features in the data. AI is a broad field of study that includes many theories, methods and technologies, as well as the following major subfields:

a) Machine learning automates analytical model building. It uses methods from neural networks, statistics, operations research and physics to find hidden insights in data without explicitly being programmed for where to look or what to conclude.

b) A neural network is a type of machine learning that is made up of interconnected units (like neurons) that processes information by responding to external inputs, relaying information between each unit. The process requires multiple passes at the data to find connections and derive meaning from undefined data.

c) Deep learning uses huge neural networks with many layers of processing units, taking advantage of advances in computing power and improved training techniques to learn complex patterns in large amounts of data. Common applications include image and speech recognition.

Understanding the difference between artificial intelligence, machine learning and deep learning can be confusing. Venture capitalist Frank Chenover view of how to distinguish between them, noting:

"Artificial intelligence is a set of algorithms and intelligence to try to mimic human intelligence. Machine learning is one of them, and deep learning is one of those machine learning techniques."

Simply put, machine learning feeds a computer data and uses statistical techniques to help it "learn" how to get progressively better at a task, without having been specifically programmed for that task, eliminating the need for millions of lines of written code. Machine learning consists of both supervised learning (using labeled data sets) and unsupervised learning (using unlabeled data sets).

Deep learning is a type of machine learning that runs inputs through a biologically-inspired neural network architecture. The neural networks contain a number of hidden layers through which the data is processed, allowing the machine to go "deep" in its learning, making connections and weighting input for the best results.

d) Cognitive computing is a subfield of AI that strives for a natural, human-like interaction with machines. Using AI and cognitive computing, the ultimate goal is for a machine to simulate human processes through the ability to interpret images and speech – and then speak coherently in response.

e) Computer vision relies on pattern recognition and deep learning to recognize what’s in a picture or video. When machines can process, analyze and understand images, they can capture images or videos in real time and interpret their surroundings.

f) Natural language processing (NLP) is the ability of computers to analyze, understand and generate human language, including speech. The next stage of NLP is natural language interaction, which allows humans to communicate with computers using normal, everyday language to perform tasks.

III. THREE STAGES OF AI

The artificial intelligence is divided into three broad stages.

a) Artificial Narrow Intelligence:

The first stage of AI as the name suggests is functionally very narrow. It is like an infant baby technology that can only work in one functional area. This is what we are experiencing right now, virtually everything you do online is utilizing Narrow AI.

Eg: When Amazon makes a selection for you.

When Pandora offers a musical selection for you.

All these are AI at its most basic level. Siri and Alexa are good examples of AI level 1 of ANI.

Narrow AI cannot perform beyond its field or limitations, as it is only trained for one specific task. Hence it is also termed as weak AI. Narrow AI can fail in unpredictable ways if it goes beyond its limits.

b) Artificial General Intelligence:

General AI is a type of intelligence which could perform any intellectual task with efficiency like a human. The idea behind the general AI is to make a system which could be smarter and think like a human by its own.

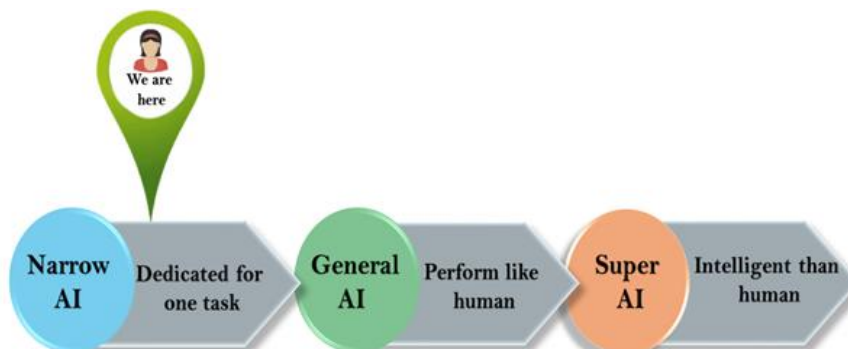
AGI may only be one step further from ANI but this step is the biggest achievement of humanity. The humankind has finally built a machine which cannot only think, but also can generate his reasoning. With AGI, machines will be able to represent themselves as the human and can do whatever human is capable of doing. AGI is helping the organizations and government in problem-solving and abstract thinking. However, it still needs some more time to be developed properly to that extent where it can be work parallel to human. But is it a good idea?

c) Artificial Super Intelligence:

Super AI is a level of Intelligence of Systems at which machines could surpass human intelligence, and can perform any task better than human with cognitive properties. It is an outcome of general AI.

Some key characteristics of strong AI include capability include the ability to think, to reason, solve the puzzle, make judgments, plan, learn, and communicate by its own. Super AI is still a hypothetical concept of Artificial Intelligence. Development of such systems in real is still world changing task.

ASI is the final stage of AI as predicted by the scientists. If we reach to that point, we will be able to solve all the mysteries of the universe and can discover everything which is undiscovered right now.



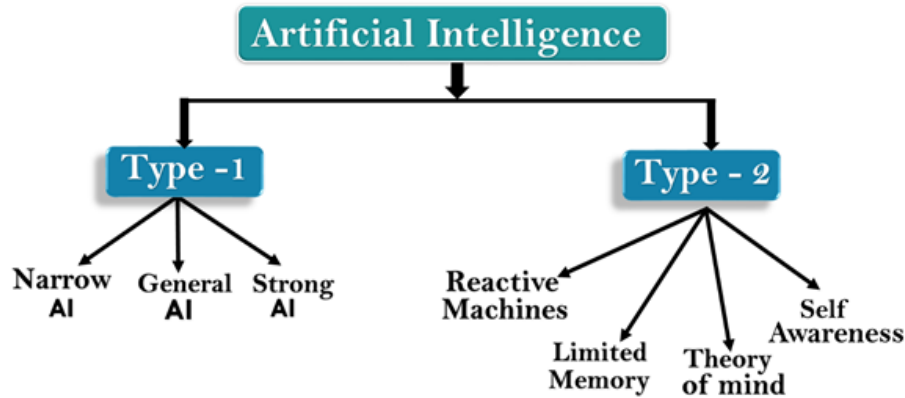


Fig: a

3.1 Transition Period:

The transition from ANI to AGI may be summarized in the invention of first computer bot to the first robot. The first computer bot Eliza which was invented in 1964 to 1966 is the example of basic ANI. This bot was made to be the chatterbot who designed to answer the queries of people. From Eliza to first humanoid robot, the time may be based on the decades but in comparison to the worth of this achievement, this period is nothing. Although we are still in the development stages of AGI and many more achievements in AGI is yet to be achieved, but scientists are already forecasting things about ASI. It is very likely that once we would be able to complete AGI, then the transition period from AGI to ASI would be just the couple of years.

3.2 Beyond AI

We are presently at the cusp of Narrow AI as already mentioned, where the systems are tasked with one particular job, it could be automated spam filtering or recommendations made on e-commerce site, to name a few. In other words, it is less complicated when compared to advanced technologies like that of IBM Watson and are seen as a weak AI.

However, researches are underway to make existing AI systems more complex and by aiming for ASI, companies want to train their present narrow AI systems to achieve more complex and multiple tasks

As described by Oxford professor, Nick Bostrom, ASI is "Any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest".

Also some scientists think that this third level of AI is the scary one. This is where machines are a little bit smarter than humans or perhaps a billion times smarter. The difference is that they have their own agenda, and this is where people start using words like immortality and extinction events with AI because we don't know what will happen once machines attain this level of intelligence. In other words when AI systems begin to take that incredible intelligent power and model themselves, they begin to have self-awareness, they begin to have feeling.

While in the recent past, systems like IBM Watson and Deep Mind's AlphaGO have proven to be more intelligent than humans by outdoing them, the current AI systems haven't turned itself into a form that would be of potential threat to humanity. According to many researchers, ASI would be achievable only after the Artificial General Intelligence wave.

As mentioned above ASI is presently seen as a hypothetical situation as depicted in movies and science-fiction books- a time of doom and gloom, where machines have taken over the world. Largely, it is believed to happen due to an intelligence explosion and has been associated with **technological singularity**, a time-frame in the future when machines can advance human to foresee and control the outcomes of the future.

3.3 What is Singularity?

In physics, the singularity is a point where gravity is so intense that even space and time begin to break down (think of the Big Bang or a black hole). This has become a metaphor to describe what will happen when AI surpasses human intelligence. Ray Kurzweil, world-renowned futurist and Google's Director of Engineering, believes this tipping point will happen before 2045. Take note: so far, an impressive 86 percent of his predictions made since the '90s have become a reality. But even if that sounds optimistic (or pessimistic, depending on how you look at it), it's the logical

next step in the evolution of technology. Once computers can learn for themselves without being taught and trained with data that we've collected, we can expect profound society changes on a scale never before seen.

Once AI is smarter than humans, it will only continue to improve itself. While many people fear this future, Kurzweil suggests that human beings will benefit by surpassing our biological limits for intelligence. This technological singularity will be a pivotal shift that could be terrifying or thrilling, depending on your outlook. Kurzweil is an optimistic futurist, and his views have helped shift the perception of the singularity away from a science fiction horror scenario toward a better life where humans and machines merge.

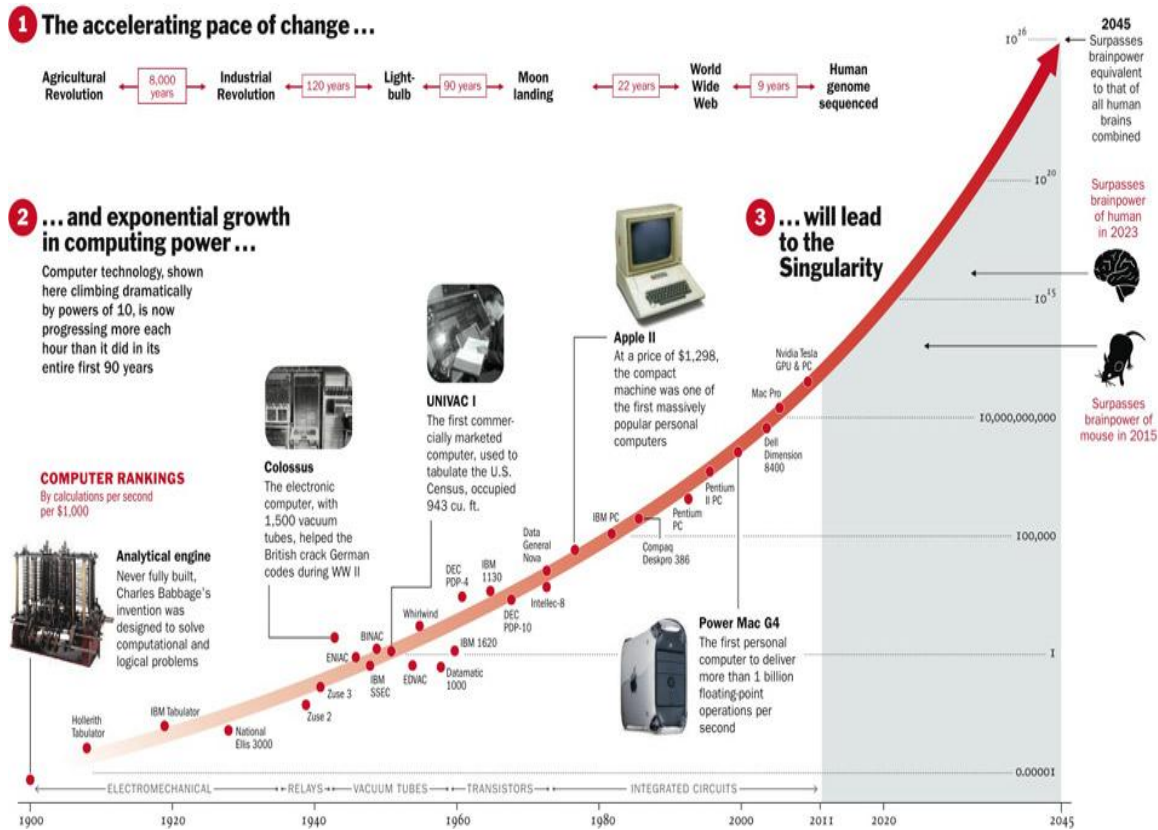


Fig: b

IV. THE OBJECTIONS TOWARDS THE TECHNOLOGY

Researchers have been arguing both in favor of the technology as well as against it. But, largely fear and the potential harm that could inflict upon humanity overshadows its positive capabilities.

One of the biggest fear associated with technology has been that of its unstoppable. Experts feel that, with the onset of intelligence explosion, the technology will be so advanced that it will turn more powerful and cannot be stopped by humans. The ability of the ASI-enabled systems will be so much so that, the machines will be able to produce several outcomes and prevent any attempts to stop it from achieving its goal and it can result in uncontrollable and unintended consequences.

Speaking about the capability of the technology, Eliezer Yudkowsky, acclaimed AI-researcher who is best known for his work towards friendly AI said, "The AI does not hate you, nor does it love you, but you are made out of atoms which it can use for something else."

However, not everyone is willing to cast the technology as the primary reason for the apocalypse. While establishing that strong AI, ASI and AGI will surpass humans in terms of intelligence, philosopher David Chalmers believes that this technology will lead to superhuman intelligence which will enable the computer to naturally achieve what he calls "recursive self-improvement", where machines would be able to reprogram and improve its function without human intervention.

Despite the fear and complexities associated with the technology, it is understood to have the capabilities to solve some of the pressing challenges of our time, like climate change, feminine and even disease breakout. While the conversation is largely about a dystopian future, it is still too early to fear the machine's capability or predict the future course of action.

In Kurzweil's future, AI will be so smart that it will come up with ideas that mere humans can't even comprehend. He suggests that this brilliant AI could solve all of our problems — including all medical problems.

Instead of being afraid that technology will turn on us, Kurzweil is excited about the potential to expand our intelligence. In the future, the human mind won't be limited to the biological brain. Someday our most profound thoughts could be triggered by a nanobot and uploaded into the cloud for eternity.

4.1 The Roadblocks To ASI

a) Solving Complex Problems

The role of innovation and research is paramount for humankind. AI, given it is still in infancy, has not yet solved many extraordinary problems. Agree, chatbots have helped in customer engagement tremendously and we have a movie directed by a computer program. The real question is, how you solve bigger problems such as global warming, space exploration, poverty and numerous more using AI.

b) A Huge Amount of Data

To train any AI algorithm, you need large data sets. To train it well and increase its accuracy, you need even larger, cleaner and well-processed datasets. However, it is not easy to collate large datasets and maintain good quality. The challenge here is to train and deploy algorithms with lesser data. Consequently, this is one of the most important blockages in developing AI.

c) Data Processing Powers

As mentioned, AI requires a large amount of data. Large data demands better data processing powers. The reason Deep Learning revolution took the world by storm was the advent of GPUs. It gave power to independent researchers and small startups to try their hand at AI. Not every invention is lying in the minds of people working in tech giants with data centers the size of football fields. Building ASI would require much more data than we can imagine. How do we bring impressive processing powers without spending a fortune?

Google may have the answer. Recently, they have released TPUs (Tensor Processing Units), which are made to speed up Machine Learning tasks. But for now, TPUs are only produced in limited quantities.


d) What about Entirely New Data?

Suppose you have trained an algorithm to identify these shapes — square, circle and a line. The algorithm can do that with tremendous accuracy. But, if you want the algorithm to correctly identify a rectangle, triangle and pentagon as well, you will have to retrain from scratch. Mind you, training an algorithm is a lengthy process. Such incremental data changes are very frequent in real-world problems. Hence, adjusting new data is a major concern for super intelligent algorithms.

e) Deep Learning Is A Black Box

Deep Learning primarily works on artificial neural networks, which mimic the behavior of a human's neural network to learn. We have tons of other algorithms based on a similar principle. But we do not yet have a complete understanding of how artificial neural networks actually work and arrive at a solution. We cannot effectively improve, control and make efficient the things we do not understand fully.

f) Mimicking The Human Brain



The first step towards Artificial Super Intelligence would be to improve the cognitive abilities of machines. Numerous projects are religiously working towards mimicking human brains and give cognizance to machines. The real progress in this aspect may not be out in the open as most of the companies and organizations do not reveal much for the sake of a competitive edge.

However, one example would be the Blue Brain Project. This is an initiative by EPFL, which is trying to achieve a total digital reconstruction of the mammalian brain. The project is expected to enlighten us more about consciousness and how it can be implanted digitally. Currently, they have simulated the synapses of the size of a bee's brain.

It is really important for AI to collaborate with neuroscience. Human minds are one of the most prestigious creations of evolution. How do we clone hundreds of millennia-long evolution process into algorithms and codes much faster, is the question we need to answer.

Google Brain, a deep learning project from Google, is also attempting to have intelligence similar or equal to human-level. Numenta, a machine learning company is also working on creating cortical learning algorithms.

V. TO FEAR OR NOT TO FEAR

While we're still decades away from overcoming the above mentioned roadblocks and reaching the singularity, it presents itself as the next big step in the evolution of technology. Emerging tech, like autonomous vehicles and facial recognition, are already paving the way.

As mentioned Ray Kurzweil has a positive outlook than most about the singularity. He claims it's an opportunity for humankind to improve by making us smarter and better at all the things we value. Other science and technology leaders, including Elon Musk, Bill Gates and the late Stephen Hawking, are not so sure. Speaking at a conference in Lisbon, Portugal shortly before his death, Stephen Hawking told attendees that the development of artificial intelligence might become the "worst event in the history of our civilization," and he had every reason for concern. Known as an artificial super intelligence (ASI) by AI researchers, ethicists, and others, it has the potential to become more powerful than anything this planet has ever seen and it poses what will likely be the final existential challenge humanity will ever face as a species.

But one thing's for sure – there will come a time when our augmented descendants look back on us as rather quaint and exotic creatures. True artificial intelligence is coming sooner or later, which is why it's time for the business world to start taking AI ethics and security issues seriously. It's no longer the concern of generations far in the future.

However we need to be more responsible while building such systems. We need to understand and define the laws which will govern ASI. Will it be treated the same as humans? Who will make laws, ASI or us? How will justice be served? Will ASI vote, elect or govern? Who will control the military? Who will have the access codes of the nuclear arsenal? Will it to be okay to be ruled by ruthless dictators or greed-hungry humans or by smart, built-for-better-future ASI? How will humans, a highly unpredictable species, become a part of a tidy ASI-run world? These questions are being asked by everyone, and we better start preparing answers for it.

5.1 ASI, The Prime Or The End Of Human Civilization?

In the future, AI will become conscious, self-sustainable, and self-learning. We need not babysit it. ASI will help us eradicate poverty, control pollution, fight diseases, find extraterrestrial life and what not. If the time comes, it will fight our wars, destroy potentially catastrophic meteors on their way to earth and fight an alien invasion.

We would require it to an extent where we would need it daily, undisputedly, in our lives. But that doesn't compel ASI to do the same for us. Imagine, an intelligent entity taking control of the world, undermining others who are far less smart than itself and mending the world to its will. Does this sound familiar to you? We, humans, did this with the entire earth. What is the possibility that a much smarter ASI would not do the same to us?

"You want to know how super-intelligent cyborgs might treat ordinary flesh-and-blood humans? Better start by investigating how humans treat their less intelligent animal cousins. It's not a perfect analogy, of course, but it is the best archetype we can actually observe rather than just imagine."
— Yuval Noah Harari

Everything is speculation. We might coexist together happily or ASI will take over or if it no longer requires humans, it might lead us to the annihilation. Our future, as the neural networks, is a black box we don't have an idea about.

VI. CONCLUSION

Technological progress is speeding up. The coming years would produce a new world which makes the modern world look outdated. This progress is becoming faster and faster and it is predicted that in this century we will make 1000 times the progress in technology compared to the last century. Life in the following thirty years will likely become almost alien to today's observer, similar to how we view life in the 1940s and 1950s. This will be particularly evident as we get closer to the half-century mark. As technology continues to progress by leaps and bounds, radical advances in quantum computing will add a new dimension to our ability to create and solve. The world of AI research is roughly split into optimists and pessimists. The optimists hope that we'll one day invent a superintelligence that solves every

problem we can imagine and leads us into a utopian future where all of mankind's needs are met and everyone lives happily ever after. The pessimists are concerned that one tiny mistake along the way will lead to the swift end of the human race – as an AI programmed to solve climate change, for example, identifies that humans are the number one obstacle to doing so.

The likes of Elon Musk and Stephen Hawking have already said ASI will be the end of human civilization. On the other hand, the likes of Mark Zuckerberg and Ray Kurzweil believe in perfect harmony. An optimist will claim a peaceful world built with the immaculate partnership of humans and ASI. A pessimist will imagine the plight of our species in case ASI takes over. In both the cases, we, humans, need to be prepared.

Finally to conclude the train has started its journey without knowing the destination. Let's not worry but cautiously optimistic.

REFERENCES

- Bostrom N. (2014); 'Superintelligence: Paths, Dangers, Strategies'; Oxford University Press, Oxford, UK
- Kurzweil R. (2005); 'The Singularity Is Near: When Humans Transcend Biology'; Viking, Penguin Group, New York, New York.
- McCarthy J. (1977); 'Epistemological Problems of Artificial Intelligence'; IJCAI (pp. 1038-1044).
- Moravec H. (1990); 'Mind Children: The Future of Robot and Human Intelligence'; Harvard University Press, New York, New York.
- Morbini F. and L. Schubert (2005); 'Conscious Agents'; University of Rochester, Computer Science Department, CS Artificial Intelligence Technical Report, TR879, September.
- Whitby B. (1996); 'Reflections on Artificial Intelligence: The Legal, Moral, and Ethical Dimensions'; Intellect Books, Exeter, UK.
- J.Shabbir and T.Anwer, "A Survey of Deep Learning Techniques for Mobile Robot Applications," ArXiv e-prints, Mar. 2018. [2] U. Neisser, G. Boodoo, T. J. Bouchard Jr, A. W. Boykin, N.
- Brody, S. J. Ceci, D. F. Halpern, J. C. Loehlin, R. Perloff, R. J. Sternberg et al., "Intelligence: Knowns and unknowns." *American psychologist*, vol. 51, no. 2, p. 77, 1996.
- V. C. Muller and N. Bostrom, "Future progress in artificial intelligence: A survey of expert opinion," in *Fundamental issues of artificial intelligence*. Springer, 2016, pp. 555–572. [78] J. S.
- Esteves, A. Carvalho, and C. Couto, "Generalized geometric triangulation algorithm for mobile robot absolute self-localization," in *Industrial Electronics, 2003. ISIE'03. 2003 IEEE International Symposium on*, vol. 1. IEEE, 2003, pp.
- Vedaldi and K. Lenc, "Matconvnet: Convolutional neural networks for matlab," in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 689–692. [80]
- A.Eitel, J.T.Springenberg, L.Spinello, M.Riedmiller, and W.Burgard, "Multimodal deep learning for robust rgb-d object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 681–687.
- Kareem Ayoub and Kenneth Payne, 'Strategy in the Age of Artificial Intelligence', *Journal of Strategic Studies*, vol. 39, nos 5–6, 2016, pp. 793–819, <https://doi.org/10.1080/01402390.2015.1088838>. See also Michael Horowitz, 'Artificial Intelligence, International Competition and the Balance of Power', *Texas National Security Review*, 15 May 2018, <https://tnsr.org/2018/05/artificial-intelligence-international-competition-and-the-balance-of-power/>; and Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W.W. Norton, 2018).
- Sergey Levine, Timothy Lillicrap and Mrinal Kalakrishnan, 'How Robots Can Acquire New Skills from Their Shared Experience', Google Research Blog, 10 March 2016, <https://research.googleblog.com/2016/10/how-robotscan-acquire-new-skills-from.html>.
- Pieter Abbeel, Adam Coates and Andrew Y. Ng, 'Autonomous Helicopter Aerobatics through Apprenticeship Learning', *International Journal of Robotics Research*, vol. 29, no. 13, 2010, pp. 1,608–39, <https://doi.org/10.1177/0278364910371999>.
- Nicholas Ernest et al., 'Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions', *Journal of Defense Management*, vol. 6, no. 1, 2016, <https://doi.org/10.4172/2167-0374.1000144>.
- Matt Simon, 'Stop Laughing at Those Clumsy Humanoid Robots', *Wired*, 11 June 2015, <http://www.wired.com/2015/06/stop-laughing-clumsy-humanoid-robots/>.
- Alvin Toffler, *The Third Wave* (London: Pan Books, 1981).
- See Kenneth Payne, *Strategy, Evolution and War: From Apes to Artificial Intelligence* (Washington DC: Georgetown University Press, 2018).
- See Beatrice Heuser, *The Evolution of Strategy: Thinking War from Antiquity to the Present* (Cambridge: Cambridge University Press, 2010).
- See, for example, Azar Gat, *War in Human Civilization* (Oxford: Oxford University Press, 2006).
- See Barry Scott Zellen, *State of Doom: Bernard Brodie, the Bomb, and the Birth of the Bipolar World* (London: Continuum, 2012).
- Robert Jervis, *The Meaning of the Nuclear Revolution: Statecraft and the Prospect of Armageddon* (Ithaca, NY: Cornell University Press, 1989).
- See, for example, Herman Kahn, *On Thermonuclear War* (Princeton, NJ: Princeton University Press, 1960).
- Arnav Manchanda, 'When Truth Is Stranger than Fiction: The Able Archer Incident', *Cold War History*, vol. 9, no. 1, 2009, pp. 111–33, <https://doi.org/10.1080/14682740802490315>.
- Michael Quinlan, *Thinking about Nuclear Weapons: Principles, Problems, Prospects* (Oxford: Oxford University Press, 2009).
- See, for example, Robert Jervis, *Perception and Misperception in International Politics* (Princeton, NJ: Princeton University Press, 1976).
- Thomas C. Schelling, *Arms and Influence* (New Haven, CT: Yale University Press, 2008).
- Joseph Rotblat, *Scientists in the Quest for Peace: A History of the Pugwash Conferences* (Cambridge, MA: MIT Press, 1972).
- Nina Tannenwald, 'The Nuclear Taboo: The United States and the Normative Basis of Nuclear Non-Use', *International Organization*, vol. 53, no. 3, Summer 1999, pp. 433–68.
- Ugo Pagallo, 'Robots of Just War: A Legal Perspective', *Philosophy & Technology*, vol. 24, no. 3, September 2011, pp. 307–23, <https://doi.org/10.1007/s13347-011-0024-9>.