THE LAndScAPE OF AI

By **Mausam**

The popular discourse on future of artificial intelligence oscillates between extreme tales of job loss and human-robot wars on the one hand and exceptional impact on all walks of life on the other. In this article we discuss these and related questions regarding the importance of AI in modern life.

The author is an Associate Professor of Computer Science Department at IIT Delhi and an affiliate faculty member at University of Washington, Seattle. mausam@cse.iitd.ac.in. The article should be cited as Mausam, 2019. The Landscape of AI, Geography and You, 19(22): 4-9

Twice in the past, AI has seen periods of buzz followed by AI winters, where both funding and interest in AI research have dried up. Is AI still making castles in the air? Not this time. The successes of AI have been real and promise of the future seems genuine.

The field of artificial intelligence (AI) has seen unprecedented growth in the last few years. Expectations are that AI may yield innovative technologies pervading almost all fields of influence including healthcare, education, safety, and transportation. At the same time, there are cautionary narratives suggesting the imminent danger of self-aware AI wiping out the whole of human civilization. In this article we take an introductory look at many important questions related to AI. What is AI? Why this sudden interest in AI? And—are we really in the danger of being wiped out?

AI: The Definition

At the highest level, the goal of AI is to make any machine intelligent. A machine or any organism like a human, can be thought of having a body comprising sensors and effectors. A sensor will take in any kind of information or input from the environment. For example, a Google search engine takes a user's query as input. A physical robot may use laser to sense the depth of an obstacle. Effectors will make a change to its environment. Google's engine will return the websites relevant to search query. A robot may move its arms to pick up objects.

With every such body a brain is required to control all sensors and effectors. For instance, a brain for a Google search engine will be the algorithm that, based on a search query, will decide which websites to show to the user and in which order. Or a brain for a physical robot may decide when to move which arm, or wheel or other effectors to achieve its goal. A good brain can make the same organism with the same body more intelligent. Google may find the right website at the top or not. The robot's brain may achieve its goal quickly or slowly. AI algorithms study how to make intelligent brains for different kinds of bodies—making bodies successful and efficient in achieving their respective objectives.

The philosophy of AI is quite broad. In practice, it undertakes a variety of manifestations and can be applied to learning, E-commerce, medical diagnosis, gaming, military and more. A chess playing algorithm, a computer vision algorithm that understands a scene from the image, a machine translation algorithm that translates English into Gujarati, or a planner that controls the robot for search and rescue at accident sites are all very different scenarios for applicability of AI. In all cases, the goal of AI is to make the machine better in the specific task at hand.

It is important to note that some of these tasks may be considered easy by humans—such as understanding a scene; some may be harder—like chess playing. However, for machines these tasks are hard in all cases and require advanced methods that may not always succeed, but have the likelihood to work well in practical situations. For instance, a chess playing bot despite not making an optimal move, will more often than not, defeat a human chess champion! On the other hand, it is still unclear whether AI will ever learn to understand scenes from images as well as humans do.

AI: The Recent Success

Formal research in the field of AI started way back in the summer of 1956. However, it is only the last few years that there has been such a widespread excitement about the field. The possibilities are indeed immense, but why is it only now that everyone is waking up to them?

This excitement is due to the sudden discovery that a specific kind of AI algorithm—deep neural networks, a computational model inspired by the human brain that can learn to perform difficult tasks (Goodfellow 2016) when combined with large-scale task specific datasets and powerful hardware called graphics processing units (GPUs), which were originally designed to render a 3D effect on a 2D surface (Parker, 2017). Simply put, the combination of AI algorithms, huge data and large compute power can be magical in bettering traditional approaches.

The task of object recognition—analysing an image to identify all objects in it, was at an accuracy level of about 75 per cent just seven years ago. However, the use of neural models now yields an accuracy of almost 98 per cent. Similar gains were observed in speech recognition, where long-standing numbers of 70 per cent were improved to about 95 per cent using deep neural models. The chess player—Deep Blue that defeated the grandmaster Gary Kasparov in 1997 did not use this technology. But the next frontier in automated game playing was Go, considered more difficult

The eventual power of any **AI system** lies in the hands of those who **design** the system since they **specify its goal.**

than chess—the Google Deep Mind's Alpha Go bot in fact defeated the world-renowned Go player Lee Sedol (Silver et al., 2016) with the use of this technology. Finally, self-driving cars are already navigating the streets of United States and heavily use deep neural models .

This explosive and exponential growth has baffled AI scientists as well. With its success came the expectation that AI could assist doctors in medical diagnosis and treatment; teachers in improving overall education quality; police in fighting crime; planners in streamlining transport and improving its efficiency; and, much more. Many demonstrations of such technologies are already underway. For instance, AI-controlled traffic lights in the city of Pittsburgh have reportedly reduced average travel time by about 25 per cent (Baker, 2018). Moreover, deep neural models were able to achieve dermatologist-level success in recognising skin cancer from skin lesion images (Esteva et al., 2017). Closer home, it is reported that the city of Surat was able to reduce crime by 27 per cent after deploying face recognition over city-wide CCTV cameras—it enabled real-time intelligence to the police in tracking crimes (Vasudevan, 2015).

The Propaganda of AI Perils

While AI may have potential benefits, are we, the human race, in danger with lethal AI-based weapons or self-aware AI bots destroying our kind? It must be understood that all AI systems are guided by an objective function that is always human-specified. The eventual power of any AI system lies in the hands of those who design the system, since they specify its goal. It is possible that someone designs an AI with the goal of killing the human race, but most likely many others will design AI with the goal of protecting human race. In case a battle like this happens, we expect the good-intentioned AIs to protect humans from the destruction any mal-intentioned AI might attempt to cause. This AI vs. AI is no different from the way terrorism vs. counter-terrorism technologies have evolved. If the bad guys can have bombs then

police may deploy bomb detectors. If the bad guys could attempt plane explosions using clear liquid explosives, the airport authorities will disallow or limit liquids onboard and so on. This narrative is unlikely to change in the AI-era, as both law-abiding and unlawful people can have access to AI technology. If anything, critical application of AI technology will be regulated by governments, similar to how nuclear bombs and other deadly weapons are regulated.

Some argue that AI systems may have unintended side-effects causing harm to humans. For instance a room cleaning robot, which is rewarded for cleaning the house, may decide to kill its owner so that it can create garbage, which it can further clean to obtain its reward. These scenarios can be handled by having safety verifications prior to deployment—and ongoing research in human-AI communication to enable accurate understanding of an end user's command may provide important answers. But, it is possible that an AI system makes mistakes. For instance, a Tesla model 3 had autopilot active in May, 2016, when it crashed into a white truck on a US highway because it thought the truck was a cloud, killing the driver (Corfield, 2017). This is definitely possible, but not likely to be frequent with the improvements in AI. The overall positive value of an AI system however, far supersedes the unintended negative consequences.

Other Challenges of AI

While most AI scientists disbelieve in the AI vs. humans narrative, there are other challenges that must be resolved. Deep neural models are highly effective at prediction, but they are not transparent. One cannot understand easily why an AI algorithm made a specific choice. This makes it difficult to analyse mistakes, improve algorithms, and also assign responsibility in the event an accident or an unintended outcome happens. Furthermore, AI algorithms are trained on data. When learning about a specific task, they also learn and exaggerate social and other biases present in any dataset. For instance, language

Go, believed to be the world's oldest strategy board game was invented in China over 2,500 years ago and is played even today. The complexity of the game which allows for intuitive thinking and is more difficult than chess, significantly helping the development of AlphaGo, a computer Go programme developed by Google DeepMind in 2016. In the visual above Chinese seniors gather around in a park in New York, to watch a game of Go between two players.

understanding algorithms think that doctors are male and nurses female, since that is the bias they have learned implicitly from the data.

Finally, in the next 15-20 years, one can expect major technological disruptions in several fields. It is possible that some jobs are wiped away. It is also likely that many new jobs are created. Reskilling ourselves in the changing times may become imperative. The future will involve AI and humans working collaboratively and being AI-ready will allow us to embrace change with leadership.

How to be Ready for AI?

To be AI-ready, a country like India needs to educate its engineers in AI technologies, so that we have the requisite manpower. It must create more professors and researchers who will train new engineers, advance the technology and put in place India-specific datasets, so that AI can be adapted for specific problems facing a large country such as ours. We also need significant investment in computing power so that AI researchers and engineers can train AI systems effectively without being resource constrained. Investments in AI-related startups and industry are thus warranted.

In addition, digitisation and automation needs to be enabled in almost all walks of life so that AI's decisions are bolstered. For example, an AI system deciding where to irrigate an agricultural field will be most effective if fields in need of water can automatically signal the supply of water. But this will require significant automation in agriculture. Moreover, non-engineers will need familiarity with digital devices and physical devices like

robots. Educating young children early can keep them ready when disruption hits.

Way Forward
AI's goal is to continually make a machine better in specific tasks. However, every real-world deployment must be associated with careful economic, sociological and technological, thought so that AI serves a useful purpose without becoming the vehicle for amplifying social discrimination and massive workforce displacement. There is some uncertainty about how things will be few years from now. We do not know which jobs will be impacted and what new jobs will be on offer. But, the new world will likely be more efficient and streamlined, highly automated, with teams of AI and humans solving tasks together. Being AI-ready early can make the difference between a country that is a world leader and that which follows. ◙

References
Baker F., 2018. The Technology that Could End Traffic Jams. BBC. December 12, Available at: http://www.bbc.com/future/story/20181212-can-artificial-intelligence-end-traffic-jams

Corfield G., 2017. Tesla Death Smash Probe: Neither Driver Nor Autopilot Saw the Truck. *The Register*, June 20. Available at: https://www.theregister.co.uk/2017/06/20/tesla_death_crash_accident_report_ntsb/
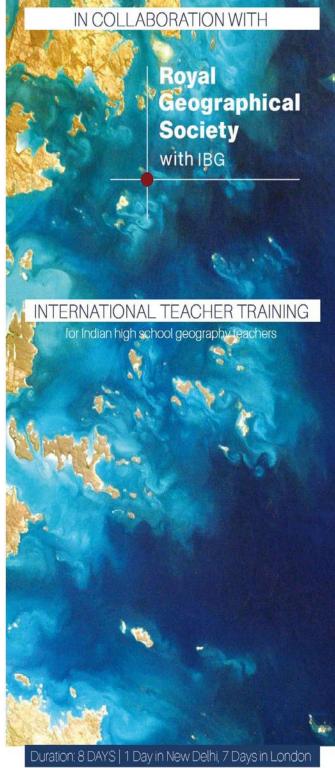
Esteva A., B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau and T. Sebastian. 2017. Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. Nature, 542(7639): 115-118.

Goodfellow I., Y. Bengio and A. Courville, 2016. Deep Learning, MIT Press, Massachusetts: USA. Available at: http://www.deeplearningbook.org

Parker M., 2017. *Graphic Processing Unit*, Digital Signal Processing 101 (Second Edition). Available at: https://www.sciencedirect.com/topics/engineering/graphics-processing-unit

Silver D. at el., 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. Nature, 529(7587): 484-489.

Vasudevan S., 2015. Here is what Surat City Police Did to Cut Crime Rate by 27%. *Data Quest*, August 12. Available at: https://www.dqindia.com/here-is-what-surat-city-police-did-to-cut-crime-rate-by-27/

IN COLLABORATION WITH

Royal Geographical Society
with IBG

INTERNATIONAL TEACHER TRAINING
for Indian high school geography teachers

Duration: 8 DAYS | 1 Day in New Delhi, 7 Days in London

Eligibility:. The teacher must be a Senior Secondary Teacher (PGT)
. Must hold a degree in Geography
. Must have excellent communication skills in English
For Details; Contact: Amit Kumar, Mob. +91-7678491255
Ayushmaan Agarwalla, Mob. +91-9101549712

ORGANISED BY

A Customised
Course on Modern Teaching Aids in
Effective Geography Teaching
New Delhi + London, The UK