13

An AI Writes On AI

How did an AI write its biography?

Computers have played chess against humans and defeated the world champion (Deep Blue v/s Gary Kasparov). AIs have played difficult quiz games against humans on TV and won (Watson in *Jeopardy*). An AI (again Watson Debating Engine) has indulged in spirited debates with human beings and lost—as of now at least. AIs have painted remarkable, and mostly forgettable, art, created music, even written books.

But what an AI has not perhaps done yet is to write about itself, a sort of autobiography or a biography of the AI 'species'. As I was writing about AI for this book, I started thinking of if and how an AI could write about itself. Would it 'think' of itself and its ilk differently than how I thought of it? Would it have an opinion, or a way of structuring the description in a more 'logical', 'machine-like' manner? Will there be emotion in the writing, will there be soul?

To find out, I got in touch with my good friend Anand Mahurkar, the founder–CEO of Findability Sciences, a

Boston, Massachusetts-based AI company. I have known Anand and Findability for a few years now, and, in fact, serve as an adviser to them. They do some brilliant work in AI for enterprises and have lately attracted an investment and JV with Softbank, Japan. Anand immediately 'got' my crazy idea and jumped on the opportunity to quickly cobble together an AI, using Findability's existing platforms, which could attempt to write a chapter on itself.

The more we thought of it, the more excited we became, since this would perhaps be a first in publishing history—and AI writing a chapter on AI in a book.

We had less than six months to do it, and so the tech folks at Findability, led by Suresh Shakkarwar, built an AI engine which could write a story about itself. For all of you who are curious about how it was done, and do not mind some tech arcana, here is how:

We used a product called FP-Summary™ an innovative unsupervized method for automatic sentence extraction using graph-based ranking algorithms used in FP-Cognition™, a graph-based ranking algorithm for creating text summaries. Using this method Findability Sciences Platform has 'written' a chapter for *The Tech Whisperer*.

FP-Summary™, an extractive type of algorithm, a graph based ranking algorithm. Graph-based ranking algorithms decide the importance of a vertex within a graph, based on global information recursively drawn from the entire graph. The graph-based ranking algorithm works with a concept of 'voting' or 'recommendation'. The score associated with a vertex is determined based on the votes that are cast for it, and the score of the vertices casting these votes.

FP-Summary™ first builds a graph associated with the text, where the graph vertices represent sentences from the text. The edges are established based on the 'similarity'

relation between the connecting sentences. The similarity is measured as a function of their content overlap. Such a relation between two sentences represents a process of 'voting' or 'recommendation'.

The resulting graph is highly connected, with a weight associated with each edge, indicating the strength of the connections established between various sentence pairs in the text. The text is therefore represented as a weighted graph, and consequently FP-Summary™ uses weighted graph-based ranking algorithm to score each of the vertices. After the ranking algorithm is run on the graph, sentences are sorted on the score, and the top-ranked sentences are selected for inclusion in Summary.

For *The Tech Whisperer*, we connected FP-Summary™ to a repository of open-source text content related to AI, ML, Deep Learning, etc., and then it 'wrote' the chapter 'AI on AI' without any human intervention or review.

In the limited time we had, we worked together and went through several iterations to have this done. Could we have made a better, more capable and 'human' AI? Definitely, yes—with more time, many more engineers and unlimited resources—we could have designed and built an AI which would have written a more flawless and complete chapter. What our narrow AI has achieved, however, is quite remarkable—given a bank of open-source content, it has actually strung together a very credible story of itself.

Does it have the emotion, the completeness and the 'soul' that a human being can put in? In my estimation, not yet. How does it compare to a mortal human writing something on it, as I did in chapter twelve?

Make your own judgement, as you read the fascinating biography of AI, penned by an AI.

1. What is AI—history and description

Description

Artificial intelligence (AI) is the simulation of human intelligence processes by machines, especially computer systems. These processes include learning (the acquisition of information and rules for using the information), reasoning (using rules to reach approximate or definite conclusions) and self-correction. Particular applications of AI include expert systems, speech recognition and machine vision. Strong AI, also known as artificial general intelligence, is an AI system with generalized human cognitive abilities. When presented with an unfamiliar task, a strong AI system is able to find a solution without human intervention. AI as a service allows individuals and companies to experiment with AI for various business purposes and sample multiple platforms before making a commitment. While AI tools present a range of new functionality for businesses, the use of artificial intelligence raises ethical questions. This is because deep learning algorithms, which underpin many of the most advanced AI tools, are only as smart as the data they are given in training. Because a human selects what data should be used for training an AI program, the potential for human bias is inherent and must be monitored closely. Researchers and marketers hope the label augmented intelligence, which has a more neutral connotation, will help people understand that AI will simply improve products and services, not replace the humans that use them.

Back in the 1950s, the fathers of the field Minsky and McCarthy, described artificial intelligence as any task performed by a program or a machine that, if a human carried out the same activity, we would say the human had to apply intelligence to accomplish the task.

AI systems will typically demonstrate at least some of the following behaviors associated with human intelligence: planning, learning, reasoning, problem solving, knowledge representation, perception, motion, and manipulation and, to a lesser extent, social intelligence and creativity.

This approach could lead to the rise of semi-supervised learning, where systems can learn how to carry out tasks using a far smaller amount of labelled data than is necessary for training systems using supervised learning today. The algorithm isn't setup in advance to pick out specific types of data, it simply looks for data that can be grouped by its similarities, for example Google News grouping together stories on similar topics each day. In reinforcement learning, the system attempts to maximize a reward based on its input data, basically going through a process of trial and error until it arrives at the best possible outcome.

An example of reinforcement learning is Google DeepMind's Deep Q-network, which has been used to best human performance in a variety of classic video games. Discussions of artificial intelligence (AI) have created a certain amount of unease by those who fear it will quickly evolve from being a benefit to human society to taking over. Let's look at 6 definitions of artificial intelligence and see how some of the industry's leaders are focusing their AI research efforts.

Definitions

McCarthy is said to have picked the name artificial intelligence for its neutrality; to avoid highlighting one of the tracks being pursued at the time for the field of "thinking machines" that included cybernetics, automata theory and complex information processing. "Today, modern dictionary definitions focus on AI being a sub-field of computer science

and how machines can imitate human intelligence (being human-like rather than becoming human).The English Oxford Living Dictionary gives this definition: 'The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.'

Definitions of artificial intelligence begin to shift based upon the goals that are trying to be achieved with an AI system. Generally, people invest in AI development for one of these three objectives:

- Build systems that think exactly like humans do ('strong AI')
- Just get systems to work without figuring out how human reasoning works ('weak AI')
- Use human reasoning as a model but not necessarily the end goal

Turns out that the bulk of the AI development happening today by industry leaders falls under the third objective and uses human reasoning as a guide to provide better services or create better products rather trying to achieve a perfect replica of the human mind.

Amazon builds a lot of its business on machine-learning systems (as a subset of AI) and defines AI as "the field of computer science dedicated to solving cognitive problems commonly associated with human intelligence, such as learning, problem solving, and pattern recognition.

Machine and deep learning are the priority for Google AI and its tools to 'create smarter, more useful technology and help as many people as possible; from translations to healthcare to making our smartphones even smarter. Facebook AI Research is committed

to "advancing the file of machine intelligence and are creating new technologies to give people better ways to communicate. In 2016, several industry leaders including Amazon, Apple, DeepMind, Google, IBM and Microsoft joined together to create Partnership on AI to Benefit People and Society to develop and share best practices, advance public understanding, provide an open platform for discussion and to identify aspirational effort in AI for socially beneficial purposes.

Those working with AI today make it a priority to define the field for the problems it will solve and the benefits the technology can have for society. It's no longer a primary objective for most to get to AI that operates just like a human brain, but to use its unique capabilities to enhance our world.

In 2016, several industry leaders including Amazon, Apple, DeepMind, Google, IBM and Microsoft joined together to create Partnership on AI to Benefit People and Society to develop and share best practices, advance public understanding, provide an open platform for discussion and to identify aspirational effort in AI for socially beneficial purposes.

History

1308 Catalan poet and theologian Ramon Llull publishes Ars generalis ultima (The Ultimate General Art), further perfecting his method of using paper-based mechanical means to create new knowledge from combinations of concepts.

1726 Jonathan Swift publishes Gulliver's Travels, which includes a description of the Engine, a machine on the island of Laputa (and a parody of Llull's ideas): a Project for improving speculative Knowledge by practical and

mechanical Operations. By using this Contrivance, the most ignorant Person at a reasonable Charge, and with a little bodily Labour, may write Books in Philosophy, Poetry, Politicks, Law, Mathematicks, and Theology, with the least Assistance from Genius or study.

It could change its facial expression and move its head and hands via an air pressure mechanism. This influential paper, in which they discussed networks of idealized and simplified artificial "neurons" and how they might perform simple logical functions, will become the inspiration for computer-based "neural networks" (and later "deep learning") and their popular description as "mimicking the brain.

"1949 Edmund Berkeley publishes Giant Brains: Or Machines That Think in which he writes: "Recently there have been a good deal of news about strange giant machines that can handle information with vast speed and skill.... These machines are similar to what a brain would be if it were made of hardware and wire instead of flesh and nerves... A machine can handle information; it can calculate, conclude, and choose; it can perform reasonable operations with information. A machine, therefore, can think.

1950 Claude Shannon's "Programming a Computer for Playing Chess" is the first published article on developing a chess-playing computer program.

1950 Alan Turing publishes "Computing Machinery and Intelligence" in which he proposes "the imitation game" which will later become known as the "Turing Test.

1952 Arthur Samuel develops the first computer checkers-playing program and the first computer program to learn on its own.

1957 Frank Rosenblatt develops the Perceptron, an early artificial neural network enabling pattern recognition based on a two-layer computer learning network. The New

York Times reported the Perceptron to be the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence. The New Yorker called it a "remarkable machine... capable of what amounts to th.

1958 John McCarthy develops programming language Lisp which becomes the most popular programming language used in artificial intelligence research.

1959 Arthur Samuel coins the term "machine learning," reporting on programming a computer "so that it will learn to play a better game of checkers than can be played by the person who wrote the program.

"1959 Oliver Selfridge publishes "Pandemonium: A paradigm for learning" in the Proceedings of the Symposium on Mechanization of Thought Processes, in which he describes a model for a process by which computers could recognize patterns that have not been specified in advance.

1959 John McCarthy publishes "Programs with Common Sense" in the Proceedings of the Symposium on Mechanization of Thought Processes, in which he describes the Advice Taker, a program for solving problems by manipulating sentences in formal languages with the ultimate objective of making programs "that learn from their experience as effectively as humans do.

1964 Daniel Bobrow completes his MIT PhD dissertation titled "Natural Language Input for a Computer Problem Solving System" and develops STUDENT, a natural language understanding computer program.

1965 Herbert Simon predicts that machines will be capable, within twenty years, of doing any work a man can do.

1965 Hubert Dreyfus publishes Alchemy and AI, arguing that the mind is not like a computer and that there were limits beyond which AI would not progress.

1965 IJ Good writes in Speculations Concerning the First Ultraintelligent Machine that "the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control.

1965 Joseph Weizenbaum, who wanted to demonstrate the superficiality of communication between man and machine, was surprised by the number of people who attributed human-like feelings to the computer program.

The first expert system, it automated the decision-making process and problem-solving behavior of organic chemists, with the general aim of studying hypothesis formation and constructing models of empirical induction in science.

1966 Shakey the robot is the first general-purpose mobile robot to be able to reason about its own actions. In a Life magazine 1970 article about this "first electronic person," Marvin Minsky is quoted saying with "certitude": "In from three to eight years we will have a machine with the general intelligence of an average human being.

1968 Terry Winograd develops SHRDLU, an early natural language understanding computer program.

A learning algorithm for multi-layer artificial neural networks, it has contributed significantly to the success of deep learning in the 2000s and 2010s, once computing power has sufficiently advanced to accommodate the training of large networks.

In an expanded edition published in 1988, they responded to claims that their 1969 conclusions significantly reduced funding for neural network research: "Our version is that progress had already come to a virtual halt because of the lack of adequate basic theories... by the mid-1960s there had been a great many experiments with perceptrons, but

no one had been able to explain why they were able to recognize certain kinds of patterns and not others.

It consisted of a limb-control system, a vision system and a conversation system.

1973 James Lighthill reports to the British Science Research Council on the state artificial intelligence research, concluding that in no part of the field have discoveries made so far produced the major impact that was then promised, leading to drastically reduced government support for AI research.

1976 Computer scientist Raj Reddy publishes "Speech Recognition by Machine: A Review" in the Proceedings of the IEEE, summarizing the early work on Natural Language Processing (NLP).

1978 The XCON (eXpert CONfigurer) program, a rule-based expert system assisting in the ordering of DEC's VAX computers by automatically selecting the components based on the customer's requirements, is developed at Carnegie Mellon University.

1979 The Stanford Cart successfully crosses a chair-filled room without human intervention in about five hours, becoming one of the earliest examples of an autonomous vehicle.

1980 Wabot-2 is built at Waseda University in Japan, a musician humanoid robot able to communicate with a person, read a musical score and play tunes of average difficulty on an electronic organ.

1981 The Japanese Ministry of International Trade and Industry budgets $850 million for the Fifth Generation Computer project. The project aimed to develop computers that could carry on conversations, translate languages, interpret pictures, and reason like human beings.

1984 Electric Dreams is released, a film about a love triangle between a man, a woman and a personal computer.

October 1986 David Rumelhart, Geoffrey Hinton, and Ronald Williams publish "Learning representations by back-propagating errors," in which they describe "a new learning procedure, back-propagation, for networks of neurone-like units.

"1987 The video Knowledge Navigator, accompanying Apple CEO John Sculley's keynote speech at Educom, envisions a future in which "knowledge applications would be accessed by smart agents working over networks connected to massive amounts of digitized information.

This work not only revolutionized the field of artificial intelligence but also became an important tool for many other branches of engineering and the natural sciences.

It is an early attempt at creating artificial intelligence through human interaction.

1988 Members of the IBM T. J. Watson Research Center publish "A statistical approach to language translation," heralding the shift from rule-based to probabilistic methods of machine translation, and reflecting a broader shift to "machine learning" based on statistical analysis of known examples, not comprehension and "understanding" of the task at hand (IBM's project Candide, successfully translating between English and French, was based on 2.

Given the hardware limitations at the time, it took about 3 days (still a significant improvement over earlier efforts) to train the network.

1990 Rodney Brooks publishes "Elephants Don't Play Chess," proposing a new approach to AI—building intelligent systems, specifically robots, from the ground up and on the basis of ongoing physical interaction with the environment: "The world is its own best model... The trick is to sense it appropriately and often enough.

"1995 Richard Wallace develops the chatbot A. L. I. C. E (Artificial Linguistic Internet Computer Entity), inspired by Joseph Weizenbaum's ELIZA program, but with the addition of natural language sample data collection on an unprecedented scale, enabled by the advent of the Web.

1997 Sepp Hochreiter and Jürgen Schmidhuber propose Long Short-Term Memory (LSTM), a type of a recurrent neural network used today in handwriting recognition and speech recognition.

"2006 Geoffrey Hinton publishes "Learning Multiple Layers of Representation," summarizing the ideas that have led to "multilayer neural networks that contain top-down connections and training them to generate sensory data rather than to classify it," i. e., the new approaches to deep learning.

2007 Fei Fei Li and colleagues at Princeton University start to assemble ImageNet, a large database of annotated images designed to aid in visual object recognition software research.

2009 Computer scientists at the Intelligent Information Laboratory at Northwestern University develop Stats Monkey, a program that writes sport news stories without human intervention.

2010 Launch of the ImageNet Large Scale Visual Recognition Challenge (ILSVCR), an annual AI object recognition competition.

2011 Watson, a natural language question answering computer, competes on Jeopardy!

June 2012 Jeff Dean and Andrew Ng report on an experiment in which they showed a very large neural network 10 million unlabeled images randomly taken from YouTube videos, and "to our amusement, one of our artificial neurons learned to respond strongly to pictures of.

"October 2012 A convolutional neural network designed by researchers at the University of Toronto achieve an error rate of only 16% in the ImageNet Large Scale Visual Recognition Challenge, a significant improvement over the 25% error rate achieved by the best entry the year before.

2. Kinds and classification

Narrow AI is what we see all around us in computers today: intelligent systems that have been taught or learned how to carry out specific tasks without being explicitly programmed how to do so.

This type of machine intelligence is evident in the speech and language recognition of the Siri virtual assistant on the Apple iPhone, in the vision-recognition systems on self-driving cars, in the recommendation engines that suggest products you might like based on what you bought in the past. Unlike humans, these systems can only learn or be taught how to do specific tasks, which is why they are called narrow AI.

What can narrow AI do.

There are a vast number of emerging applications for narrow AI: interpreting video feeds from drones carrying out visual inspections of infrastructure such as oil pipelines, organizing personal and business calendars, responding to simple customer-service queries, co-ordinating with other intelligent systems to carry out tasks like booking a hotel at a suitable time and location, helping radiologists to spot potential tumors in X-rays, flagging inappropriate content online, detecting wear and tear in elevators from data gathered by IoT devices, the list goes on and on.

What can general AI do.

Artificial general intelligence is very different, and is the type of adaptable intellect found in humans, a flexible form of intelligence capable of learning how to carry out vastly different tasks, anything from haircutting to building spreadsheets, or to reason about a wide variety of topics based on its accumulated experience. Deep Blue can identify pieces on the chess board and make predictions, but it has no memory and cannot use past experiences to inform future ones. Deep Blue and Google's AlphaGO were designed for narrow purposes and cannot easily be applied to another situation. These AI systems can use past experiences to inform future decisions.

Machines with self-awareness understand their current state and can use the information to infer what others are feeling. Machine Learning is a broad body of research in AI, much of which feeds into and complements each other. Currently enjoying something of a resurgence, machine learning is where a computer system is fed large amounts of data, which it then uses to learn how to carry out a specific task, such as understanding speech or captioning a photograph. These are brain-inspired networks of interconnected layers of algorithms, called neurons, that feed data into each other, and which can be trained to carry out specific tasks by modifying the importance attributed to input data as it passes between the layers.

During training of these neural networks, the weights attached to different inputs will continue to be varied until the output from the neural network is very close to what is desired, at which point the network will have 'learned' how to carry out a particular task. A subset of machine learning is deep learning, where neural networks are expanded into sprawling networks with a huge number of layers that are

trained using massive amounts of data. It is these deep neural networks that have fueled the current leap forward in the ability of computers to carry out task like speech recognition and computer vision. The design of neural networks is also evolving, with researchers recently refining a more effective form of deep neural network called long short-term memory or LSTM, allowing it to operate fast enough to be used in on-demand systems like Google Translate.

This approach has even been used to help design AI models, effectively using AI to help build AI. This use of evolutionary algorithms to optimize neural networks is called neuroevolution, and could have an important role to play in helping design efficient AI as the use of intelligent systems becomes more prevalent, particularly as demand for data scientists often outstrips supply. The technique was recently showcased by Uber AI Labs, which released papers on using genetic algorithms to train deep neural networks for reinforcement learning problems. Finally there are expert systems, where computers are programmed with rules that allow them to take a series of decisions based on a large number of inputs, allowing that machine to mimic the behavior of a human expert in a specific domain. An example of these knowledge-based systems might be, for example, an autopilot system flying a plane. What are the elements of machine learning. As mentioned, machine learning is a subset of AI and is generally split into two main categories: supervised and unsupervised learning.

Supervised learning

A common technique for teaching AI systems is by training them using a very large number of labeled examples. These machine-learning systems are fed huge amounts of data, which has been annotated to highlight the features of interest.

This process of teaching a machine by example is called supervised learning and the role of labeling these examples is commonly carried out by online workers, employed through platforms like Amazon Mechanical Turk. Training these systems typically requires vast amounts of data, with some systems needing to scour millions of examples to learn how to carry out a task effectively -- although this is increasingly possible in an age of big data and widespread data mining. In the long run, having access to huge labelled datasets may also prove less important than access to large amounts of compute power.

In recent years, Generative Adversarial Networks (GANs) have shown how machine-learning systems that are fed a small amount of labelled data can then generate huge amounts of fresh data to teach themselves.

3. How will AI can make life better.

Artificial intelligence (AI) has the potential to help tackle some of the world's most challenging social problems. To analyze potential applications for social good, we compiled a library of about 160 AI social-impact use cases. They suggest that existing capabilities could contribute to tackling cases across all 17 of the UN's sustainable-development goals, potentially helping hundreds of millions of people in both advanced and emerging countries.

Real-life examples of AI are already being applied in about one-third of these use cases, albeit in relatively small tests. AI is only part of a much broader tool kit of measures that can be used to tackle societal issues, however. For now, issues such as data accessibility and shortages of AI talent constrain its application for social good. It looks at domains of social good where AI could be applied, and the most

pertinent types of AI capabilities, as well as the bottlenecks and risks that must be overcome and mitigated if AI is to scale up and realize its full potential for social impact.

The article is divided into five sections:
Mapping AI use cases to domains of social good.
AI capabilities that can be used for social good.
Overcoming bottlenecks, especially around data and talent
Risks to be managed.
Scaling up the use of AI for social good.
Mapping AI use cases to domains of social good.

For the purposes of this research, we defined AI as deep learning. We grouped use cases into ten social-impact domains based on taxonomies in use among social-sector organizations, such as the AI for Good Foundation and the World Bank. Each use case highlights a type of meaningful problem that can be solved by one or more AI capability. Nonetheless, employing usage frequency as a proxy, we measure the potential impact of different AI capabilities.

For about one-third of the use cases in our library, we identified an actual AI deployment (Exhibit 1). Since many of these solutions are small test cases to determine feasibility, their functionality and scope of deployment often suggest that additional potential could be captured. For three-quarters of our use cases, we have seen solutions deployed that use some level of advanced analytics; most of these use cases, although not all, would further benefit from the use of AI techniques.

Crisis response These are specific crisis-related challenges, such as responses to natural and human-made disasters in search and rescue missions, as well as the outbreak of disease.

Economic empowerment

With an emphasis on currently vulnerable populations, these domains involve opening access to economic resources and opportunities, including jobs, the development of skills, and market information.

For example, adaptive-learning technology could base recommended content to students on past success and engagement with the material.

One use case, based on work by Affectiva, which was spun out of the MIT Media Lab, and Autism Glass, a Stanford research project, involves using AI to automate the recognition of emotions and to provide social cues to help individuals along the autism spectrum interact in social environments.

Health and hunger

This domain addresses health and hunger challenges, including early-stage diagnosis and optimized food distribution. AI-enabled wearable devices can already detect people with potential early signs of diabetes with 85 percent accuracy by analyzing heart-rate sensor data. These devices, if sufficiently affordable, could help more than 400 million people around the world afflicted by the disease.

Infrastructure management

This domain includes infrastructure challenges that could promote the public good in the categories of energy, water and waste management, transportation, real estate, and urban planning. AI can also be used to schedule predictive maintenance of public transportation systems, such as trains and public infrastructure (including bridges), to identify potentially malfunctioning components.

Public and social-sector management

Initiatives related to efficiency and the effective management of public- and social-sector entities, including strong institutions, transparency, and financial management, are included in this domain. For example, AI can be used to identify tax fraud using alternative data such as browsing data, retail data, or payments history.

It focuses on security, policing, and criminal-justice issues as a unique category, rather than as part of public-sector management. An example is using AI and data from IoT devices to create solutions that help firefighters determine safe paths through burning buildings.

Our use-case domains cover all 17 of the UN's Sustainable Development Goals

The United Nations' Sustainable Development Goals (SDGs) are among the best-known and most frequently cited societal challenges, and our use cases map to all 17 of the goals, supporting some aspect of each one (Exhibit 3). Our use-case library does not rest on the taxonomy of the SDGs, because their goals, unlike ours, are not directly related to AI usage; about 20 cases in our library do not map to the SDGs at all.

Instead of assuming we're all automatically out of jobs, let's consider the value of being human.

The utility company implements a new system, and instead of prompting customers to press 1 and wait for more automated prompts, your pleasant voice greets them on the first ring with how can I help.

One customer states she wants to pay her bill, and you oblige and conclude the call. Another customer calls at the same time for a different reason and you deal with that request with ease and accuracy. The customer isn't startled by the change in voice or tone, and you resolve the problem because you see the details of his call already on your device. Being human becomes a premium asset

In our example, the customer feels like he is important to the utility company because they are providing a human to help him. Yes, the utility company could have spawned instances of the same voice and personality to answer all calls. They are now focused on managing the utility network from a control centre, with new responsibilities to use analytics and machine learning to expand and maintain the network.

Therefore, AI can help to increase job satisfaction

You hired out your voice to the utility company to answer calls and provide the human touch to their brand. You are paid through new usage terms and conditions that give you flexibility and increase work/life balance.

The customer service scenario may sound preposterous, but the technologies already exist.

- Microsoft is making it easy to integrate artificial understanding into nearly any software application with Cognitive Services.
- Large consultancies, like Avanade, are stitching emerging technologies together and helping customers use them to increase revenue or drive down cost.

Despite the challenges, organizations are digitally transforming to compete for customers and AI promises a way forward. With investment in AI expected to grow from $640m in 2016 to $37bn by 2025, the scenario I describe is coming – and humans will likely accept the change as business as usual

4. Will AI take jobs.

"The adoption of thoughtful policies coupled with sober adaptation to a world in which AI will actually feature in every aspect of our life is our big challenge," Irakli Beridze,

who is heading the Centre for Artificial Intelligence and Robotics, United Nations, UNICRI, told me in an interview for the upcoming book I am co-writing with Neil Sahota, Uber Yourself Before You Get Kodaked: A Modern Primer on AI for the Modern Business.

Beridze is one of the world's leading experts on this subject.

Fear surrounding AI also concerns the likelihood our creations won't literally take our lives; they'll take our livelihoods.

"AI-enhanced technological unemployment is one of the major issues of our time," said Beridze. Almost weekly I see new reports coming out suggesting something to the effect that between 20% to 70% of jobs will be wiped out because of AI" A 2017 report by the McKinsey Global Institute sought to quantify this possibility by researching 20 countries and 30 industries regarding six themes: "productivity and growth, natural resources, labor markets, the evolution of global financial markets, the economic impact of technology and urbanization." What they found was that, "while few occupations are fully automatable, 60 percent of all occupations have at least 30 percent technically automatable activities."

Though there is undoubtedly cause for concern that our economic way of life is changing — and faster than we can imagine — there is also very good reason to believe we have a bright future ahead. "Still, I have also seen reports which say that actually the contrary is true: that AI will create more jobs than people will lose," said Beridze. In a recent interview with my coauthor and I, social entrepreneur and futurist Stephen Ibaraki pointed to just such a possibility occurring in what he views as the "Fourth Industrial Revolution."

Certainly, the human race desires a world without autonomous weapons, one of vast abundance in which

people are free to pursue happiness. But while the public has been subjected to countless reports and popular entertainment foretelling the dark side of AI, there has been little conversation about how things might improve in the years to come.

Positive vision is greatly lacking in our public and private discourse. "We choose to go to the moon in this decade and do the other things, not because they are easy, but because they are hard, because that goal will serve to organize and measure the best of our energies and skills, because that challenge is one that we are willing to accept, one we are unwilling to postpone, and one which we intend to win, and the others, too."

Now, pessimists and realists alike may contend it's not possible to wrangle mass support around such a hopeful ideal, certainly not in today's cynical public sphere. These critics could reasonably doubt the chances for developing a strong global coalition to face what is coming — or more accurately, what is already here.

But for those who think that mankind's biggest problem is that our technology has outstripped our morality — that we have a long way to go before we are capable of responsibly handling the AI revolution, Beridze encourages considering the historical record. We have come a long way toward building a just society, one capable of facing the biggest technological leap in history.

"Of course, we have way more to go — we're nowhere near an ideal world, but if we look at the trajectory of positive change in the last few centuries, it's astonishing! If you go back a bit, at the start of the twentieth century, we had completely different ideas of human rights.

Rather than fear there won't be enough to go around, let's steer the cultural conversation toward ways in which we can include everyone in the bounty to come.

Defying concerns about the rise of nationalism, author and historian Yuval Noah Harari recently offered similar hopeful thoughts on Sam Harris's Podcast, Waking Up. Nationalism demands something almost impossible ... for us to feel loyal to millions of strangers we have never met before — that we know nothing about — and are never likely to meet. And if we can do that I think the distance from there to caring about 8 billion people you've never met is much smaller than the distance already covered. "

This is an exponential way of considering our relations to others. Many famous leaders have joined Elon Musk to start something called Open AI, a non-profit artificial intelligence (AI) research company that aims to promote and develop friendly AI in such a way as to benefit humanity as a whole.

I have read many books on the subject, and the level of pessimism varies.

I found Ford's version of the future to be frankly terrifying. The thing is no one knows for certain what will happen in the future, but there are a number of ways displaced workers could survive.

Robots need to do jobs that can be automated, and humans need to do the jobs that require a personal or creative touch. Displaced workers could potentially get re-trained to apply their skills elsewhere.

They could potentially re-skill to build on their existing skills and work in a different area.

Even employees who aren't at risk for being displaced should expand their skills. People move between jobs more often these days, and that provides opportunities to expand their skillsets. Many companies already provide training and re-skilling for their employees.

Move Them To Other Jobs Bill Gates recently said that AI is a positive for society and that displaced workers could fill

gaps that currently exist elsewhere in the labor market—like elder care, teaching and support for special needs children." Instead of learning new skills, this solution encourages workers to use their existing skills in a new industry.

While it's true there are often plenty of jobs available in these areas, unfortunately these important jobs often don't pay well. In large cities, like where I live in the Bay Area, working in these jobs simply wouldn't give people enough money to live without a universal basic income.

However, these positions are fairly safe from being displaced by robots and can provide job security.

The Technologies Create New Jobs For Both The Short And Long-Term The displacement problem is difficult to solve because we can't see the future.

We don't know what skills will be the most useful in the future or what technology will be the most prevalent.

Not Everyone Will Work Another potential idea is that not everyone will work.

It could be that displaced human workers can't re-skill and don't have it in them to fill the more human-driven roles. This would obviously lead to higher unemployment numbers, which would have a large impact on society and the economy. The impacts of this possibility are far-reaching.

All other tasks will be done by robots, and those who choose not to work can enjoy other activities. That extreme possibility is definitely a long way off, though it can still be considered.AI and robots are the inevitable future.

We need to have conversations now about what to do with displaced workers so we can be prepared when the time comes.

And though it may not yet be able to read the blur on the film as tuberculosis, in five years, AI could make that interpretation as well.

With the 'Deep Learning', 'Machine Learning', 'Data Science' or 'AI Revolution' hype, there a few common repetitive concerns in our society, that keep appearing in different forms: Will AI take over our jobs? Will Robots steal my income streams? Will there be a massive 'job-depression'? I'm a driver, will a Self Driving Car replace me? (For an extreme) Will Robots rule over us in the Future? I will try to state my viewpoint as an AI enthusiast (Sorry humans J).Will AI take over our jobs? First, that would require a General and robust Intelligence, which is atleast a few decades away.

We're truly close to building a Perfect- Smart Washing Machine- Fei Fei LiAI seems to excel in very specific tasks as of 2018, the systems are extremely brittle and can be, under very limited cases be applied to different domains.

The tasks by themselves, are extremely specific to given conditions.

In the real world, with constraints out of the equations, I'm skeptical if a General Intelligence variable will hold (anytime soon).How Far Are We from a Fully Autonomous Driving World? Will there be a mass 'job-depression'?T his seems to be the biggest concern, since AI can do our tasks better, what would we do? I firmly disagree with this general opinion, Firstly because such a system is not near for the following reasons:The systems are extremely brittle as of now. The tasks are not fully automated: Hyperparameter tuning is to be done manually. Lack of Understanding: 'Deep' Neural Networks are mostly blackboxes (even after we visualise the under the hood processes).Edge cases: DL is not good at dealing with edge cases. (Specially when it comes to Autonomous driving).Secondly, Throughout the history of Humanity, whenever we have always witnessed technology that does our jobs in an efficient manner

replacing humans in the field, it has enabled us to worry about greater tasks. Did Automated Call Routing, replace the workers at Telecommunication Exchange? YES, They 'took' their jobs (or replaced them).

Or freedom of doing whatever you want when you are on the road? What about not owning a car and just being able to summon one when you need to commute? What about Solving the issues of car parking, by making it figure that out autonomously? I think the Pros always outweigh the Cons for every Technology that affects the world on a massive scale.

After all, there's a reason why we call it 'innovating'.

I am a firm believer of Sebastian Thrun's vision: AI well let us do better tasks, it will allow humans to work at innovative tasks. Ones at which, we are naturally good at.

AI will replace all of our 'boring' tasks with automation and allow us to work on better, more important tasks. Sebastian, in his Ted Talk argues that our jobs, for most of our lifetimes include working on variations of the same task.

5. The wars being fought over AI

This has a startling story to tell: for the first time since the industrial revolution, he argues, China will be at the forefront of a huge economic transformation — the revolution in artificial intelligence.He starts his book by talking about China's "Sputnik moment", when Google DeepMind's AlphaGo defeated Ke Jie, the world's leading player of the ancient Chinese game of Go. Here China has, he writes, many advantages.First, the work of leading AI researchers is readily available online. Mr Lee describes a world of cut-throat business activity and remorseless imitation, which has already allowed Chinese businesses to defeat leading

western rivals in their home market. So China has been able to jump to universal digital payment systems, while western businesses still use outdated technology.Fifth, China has scale. One of the things China can do more easily than anywhere else is build complementary infrastructure.Finally, writes Mr Lee, the Chinese public is far more relaxed about privacy than westerners. Chinese leaders, I would argue, see no justification for individual privacy at all (except for their own).

So where is this supposed "race" between the US and China today? But five years from now, he thinks, China might be a little ahead in the first, less far behind in the second, well ahead in the third and equal in the last. China is far behind the US in production of semiconductors, ahead in the number of potential users and has about half the number of AI experts and roughly half the number of AI companies. Yet Mr Ding is looking at AI overall, while Mr Lee focuses on commercial applications.Historical experience suggests that the rents created by a lead in an important technology are valuable, though often impermanent. But the economic and social impact of AI is a bigger issue and one that is relevant to every countryAs Mr Lee stresses, advances in AI offer gains. It seems reasonable to fear that AI will accelerate the hollowing out of the middle of the earnings distribution, possibly even the upper middle, while increasing concentrations of private wealth and power at the top.Yet perhaps the most important consequence will be in the intensity of influence and surveillance made possible by AI-monitored mobile devices and sensors. But we may in the end conclude we have birthed a monster.

China and the United States are ahead of the global competition to dominate artificial intelligence (AI),

according to a study by the U.N. World Intellectual Property Organization (WIPO) published on Thursday. The study found U.S. tech giant IBM had by far the biggest AI patent portfolio, with 8,920 patents, ahead of Microsoft with 5,930 and a group of mainly Japanese tech conglomerates.

China accounted for 17 of the top 20 academic institutions involved in patenting AI and was particularly strong in the fast growing area of "deep learning" - a machine-learning technique that includes speech recognition systems.

"The U.S. and China obviously have stolen a lead. U.S. President Donald Trump has accused China of stealing American innovations and technology and has slapped trade tariffs on $234 billion of Chinese goods to punish Beijing.

China said in December it resolutely opposed "slanderous" accusations from the United States and other allies criticizing China for economic espionage and stealing intellectual property and company secrets.

Gurry acknowledged there were accusations about China's behavior but there was no doubt it had embraced the global intellectual property system, with the world's largest patent office and the largest number of domestic patent applications.

"They are serious players in the field of intellectual property," he said.

The WIPO study analyzed international patent filings, scientific publications, litigation filings and acquisition activity, and found there had been as many patent applications for AI since 2013 as in the half century since the term was coined in the 1950s.

Patent applications in machine learning, which includes techniques used by ride-sharing services to minimize detours, averaged annual growth of 28 percent between 2013 and 2016, the last year for which data is available, because of

an 18-month period before confidential applications are publicly disclosed.

"If you did, you wouldn't need a venture capital industry," he said.

The US-China trade war has dominated headlines for months now — but apparently, there is a completely different conflict between the superpowers that's even more urgent.

That's according to renowned economic historian Niall Ferguson, who told guests at the Australian Financial Review's business summit yesterday, the more important battle was being waged over technology and artificial intelligence.

He said China's quick rise as a tech force to be reckoned with was a concern for America, which has long dominated the field.

"The AI arms race is like the nuclear arms race of the original Cold War, but it doesn't sound so scary because there's no nuclear warheads, only data and how to use it," he said.

"But it has dawned on the US that China is catching up, certainly in terms of the volume of research on AI, even if the US retains the lead on quality frontier research.

"Mr Ferguson, a senior fellow at the Hoover Institution, said if China were to win the technological "arms race" it could threaten America's global dominance, but said the Chinese government faced an uphill battle.

"Beijing has to ramp up stimulus as the trade and tech wars hit the Chinese economy — they have to work hard to keep the economic show on the road," he said.But suddenly, kabam, in 2018 Trump starts ratcheting up tariffs and escalating the tech wars and China is on the back foot at the moment," he said.

Meanwhile, London School of Economics professor Keyu Jin, who also fronted the summit, said it was in China's best interests to avoid an "outright collision" with America.

And despite the tensions between Donald Trump and Chinese President Xi Jinping, she said the US President was actually a "strategic gift to China".

"He is not the worst nightmare for China — the worst nightmare is a competent, capable American president who is able to strengthen relationships with allies," she said.

"China's worst nightmare is to be ousted from the global economic system.

"Trump is not only fighting China on all fronts but also causing political crises within his own country ... to some sense, China has interests in keeping him around.

"She said moving forward, one of Australia's biggest challenges would be to choose a side between the sparring global powers.

"This is a profound dilemma for Australia, which for reasons of history, naturally gravitates to the US, but for reasons economic, is powerfully drawn to China," she said.

"It is going to become a central problem of Australian politics, regardless of who is prime minister." Researchers, companies and countries around the world are racing to explore — and exploit — the possibilities of artificial intelligence technology.

Since then, little has happened — until a Feb. 11 executive order from President Donald Trump encouraging the country to do more with AI.

The executive order has several parts, including directing federal agencies to invest in AI and train workers "in AI-relevant skills," making federal data and computing resources available to AI researchers and telling the National Institute of Standards and Technology to create standards for AI

systems that are reliable and work well together. So after researching how large organizations use AI for the past five years, in my view the executive order alone is not likely to transform the American approach to AI.

In 2017, the country's national government announced it wanted to make the country and its industries world leaders in AI technologies by 2030.

The government's latest venture capital fund is expected to invest more than US$30 billion in AI and related technologies within state-owned firms, and that fund joins even larger state-funded VC funds.One Chinese state alone has said it will devote $5 billion to developing AI technologies and businesses.

The national effort also includes using AI in China's defense and intelligence industries; the country's leaders are not reluctant to use AI for social and political control.

DARPA, the Defense Department's research arm, has sponsored AI research and competitions for many years, and has a $2 billion fund called "AI Next" to help develop the next wave of AI technologies in universities and companies.

There are, for instance, many more AI firms in the U.S. than in China.American investment appears strong, too.

Much of that spending went toward AI research, but some of the work actually happened in China and elsewhere outside the U.S. That work has been used to personalize ads, improve search results, recognize and label faces and generally make products smarter.In China, the private sector is much more closely tied to government plans than in the U.S.

The Chinese government has asked four large AI-oriented firms in China — Baidu, Tencent, Alibaba and iFlytek — to develop AI hardware and software systems to handle autonomous driving and language processing, so other companies could build on those skills.

China may have also surpassed the American historic advantage in venture capital investments. However, one report from China suggests that in the first half of 2018, Chinese venture investments — many of which involved AI — were higher than in the U.S. Data from 2017 suggest that Chinese AI firms received more venture funding than U.S. companies, although the American funding went to many more firms.

The U.S. had an historical edge in this regard, with strong technical universities, many technology sector employers and relatively open immigration policies.A recent analysis of LinkedIn data suggests that the U.S. has far more AI engineers than China does.

Both the U.S. and China could learn from efforts in Canada, such as the work by the Montreal Institute for Learning Algorithms, which has offered companies access to facilities, venture capital and university research partnerships to accelerate AI development in that city.

A final key element in AI progress is data: The more data a country's companies have, the better able they are to develop capable AI systems.

Because of its very large number of inhabitants, the population's heavy use of digital services and its lax regulatory environment, China clearly beats the U.S. on data.I still think the U.S. has the edge over China in AI capabilities at the moment.

As I describe in my new book "The AI Advantage," China is executing its strategy for AI, and the U.S. is still wrestling to create one.

China is also reaping the benefits of having a determined government, an inexhaustible pot of money, a growing cadre of smart researchers and a large, digital-hungry population. Perhaps if the leadership of the U.S. government devoted as

much attention and investment to AI as it does to its other strong priorities, the U.S. could maintain its lead in the field.

Spending on artificial intelligence systems in China clocked highs of $12 billion in 2017 and expected to reach record highs of $70 billion by 2020.Investments into AI startups are also on the rise.

The central government is among the biggest investors in AI technologies in China.

The Education Ministry has also announced a five-year program designed to equip people with AI skills.The number of AI journal articles from research institute is also on the rise the number of AI-related patents having surged by 200%.China's edge when it comes to AI development has to do with the ease of access, to troves of data.

Also, China does not have restrictive privacy, making it easy for companies to collect sufficient data key to AI development.

"You're going to see better and better and deeper insight patterns than anyone else, and I think it'll be a great advantage for China," said author and columnist Thomas Friedman.However, the biggest threat to a flourishing AI sector in China is the lack of sufficient high-level AI talent.

China only has about a fifth of the number of AI talent that the U.S has, something that leaves it at a disadvantage in the development of high-quality AI research.

U.S AI developmentThe U.S is home to most AI innovation, even though China accounts for more investments especially from the central government.

A lack of sufficient funding primarily from the central government is one of the reasons why the U.S lags China when it comes to AI market value.

Immediate reports indicate that the White House is planning to slash science and technology research funding

by 15% something that would hurt innovations considerably. The U.S boasts of an edge over China when it comes to the pool of talent with the necessary AI skills.

6. What is future of AI— living with humans or instead of them?

In this version of the future, people will still have a role working alongside smart systems: either the technology will not be good enough to take over completely, or the decisions will have human consequences that are too important to hand over completely to a machine. There's just one problem: when humans and semi-intelligent systems try to work together, things do not always turn out well. Like almost all of today's autonomous cars, a back-up driver was there to step in if the software failed. The so-called Level 3 system is designed to drive itself in most situations but hand control back to a human when confronted by situations it cannot handle.

"If you're only needed for a minute a day, it won't work," says Stefan Heck, chief executive of Nauto, a US start-up whose technology is used to prevent professional drivers from becoming distracted. Without careful design, the intelligent systems making their way into the world could provoke a backlash against the technology. Preventing that will require more realistic expectations of the new autonomous systems, as well as careful design to make sure they mesh with the human world. "Does the AI make us feel more involved — or is it like dealing with an alien species?"

The semi-driverless car is a particularly stark example of a near-autonomous system that relies on close co-operation with people. In the real world, people often make decisions about situations they have not previously faced.

Research from Stanford University has shown that it takes at least six seconds for a human driver to recover their awareness and take back control, says Mr Heck. But even when there is enough time for human attention to be restored, the person stepping into a situation may see things differently from the machine, making the handover far from seamless.

"We need to work on a shared meaning between software systems and people — this is a very difficult problem," says Mr Sikka. A second type of human/machine co-operation is designed to make sure that a sensitive task always depends on a person — even in situations where an automated system has done all the preparatory work and would be quite capable of completing the task itself. Military drones, where human "pilots", often based thousands of miles away, are called on to make the decision to fire at a target, are one example. Both show how AI can make humans far more effective without robbing them of control, says Mr Heck.

According to Stuart Russell, an AI professor at the University of California, Berkeley, it would be a short and easy step in a national emergency to remove the human drone operator from the loop, precipitating an era of robot weapons that make their own decisions about when to kill people. "You can't say the technology itself can only be used in a defensive way and under human control. A final type of "human in the loop" system involves the use of AI that is not capable of handling a task entirely on its own but is used as an aid to human decision-making. Algorithms that crunch data and make recommendations, or direct people in which step to take next, are creeping into everyday life. The algorithms, though, are only as good as the data they are trained on — and they are not good at dealing with new situations.

People required to trust these systems are often also required to take them on faith. The outcome of these computer-aided decisions may well end up being worse than those based on purely human analysis, he says. "Sometimes people will blindly follow the machine, other times people will say: 'Hang on, that doesn't look right.

But what happens when the stakes are higher? IBM made medical diagnostics one of the main goals for Watson, the system first created to win a TV game show and then repurposed to become what it calls a more general "cognitive" system. "Simply saying they'll still make the decisions doesn't make it so. "Similar worries surfaced in the 1980s, when the field of AI was dominated by "expert systems" designed to guide their human users through a "decision tree" to reach the correct answer in any situation.

But the latest AI, based on machine learning, looks set to become far more widely adopted, and it may be harder to second-guess.

Non-experts may feel reluctant to second-guess a machine whose workings they do not understand. Technicians had no way of identifying the flaw and the machine stayed in use much longer as a result, says Mr Nourbakhsh.

Unlike the logic circuits employed in a traditional software program, there is no way of tracking this process to identify exactly why a computer comes up with a particular answer. Some experts, however, say headway is being made and that it will not be long before machine learning systems are able to point to the factors that led them to a particular decision. Like many working in the field, he expresses optimism that humans and machines, working together, will achieve far more than either could have done alone.

He had already founded and sold off several successful consumer technology companies, but as he grew older he wanted to do something more meaningful, that is, he wanted to build a product that would serve the people that technology startups had often ignored. Both my friend and I were entering the age at which our parents needed more help going about their daily lives, and he decided to design a product that would make life easier for the elderly.

It sounded like a wonderful product, one that would have a real market right now.

But once those material needs were taken care of, what these people wanted more than anything was true human contact, another person to trade stories with and relate to. If he had come to me just a few years earlier, I likely would have recommended some technical fix, maybe something like an AI chat bot that could simulate a basic conversation well enough to fool the human on the other end.

But there remains one thing that only human beings are able to create and share with one another: love.

Despite what science-fiction films like Her—in which a man and his artificially intelligent computer operating system fall in love—portray, AI has no ability or desire to love or be loved.

I firmly believe we must forge a new synergy between artificial intelligence and the human heart, and look for ways to use the forthcoming material abundance generated by artificial intelligence to foster love and compassion in our societies.

That does not mean all of those jobs will disappear overnight, but if the markets are left to their own devices, we will begin to see massive pressure on working people.

Centuries of living within the industrial economy have conditioned many of us to believe that our primary role in

society (and even our identity) is found in productive, wage-earning work.Instead, we must move toward a new culture that values human love, service, and compassion more than ever before.

We can choose a purely technocratic approach—one that sees each of us as a set of financial and material needs to be satisfied—and simply transfer enough cash to all people so that they don't starve or go homeless.

It would put the economic bounty of AI to work in building a better society, rather than just numbing the pain of AI-induced job losses.

In an age in which intelligent machines have supplanted us as the cogs and gears in the engine of our economy, I hope that we will value all of these pursuits—care, service, and personal cultivation—as part of our collective social project of building a more human society.

Researchers no longer speak of just one AI, but of hundreds, each specializing in a complex task—and many of the applications are already lapping the humans that made them. In just the last few years, "machine learning" has come to seem like the new path forward.

This AI "takeoff," also known as the singularity, will likely see AI pull even with human intelligence and then blow past it in a matter of days.

The $2,300 deposited into their bank accounts every month as part of the universal basic income, plus their free health insurance, the hyper-personalized college education their children receive and a hundred other wonderful things, are all paid for by AIs like Alpha 4, and people don't want that to change.

As you listen in, the government's lawyers argue that there's simply no way to prove that Alpha 4—which is thousands of times smarter than the smartest human—is

conscious or has human feelings. AIs do have emotions—there has long been a field called "affective computing" that focuses on this specialty—far more complex ones than men and women possess, but they're different from ours: A star-voyaging AI might experience joy, for example, when it discovers a new galaxy. Superintelligent systems can have millions of thoughts and experiences every second, but does that mean it should be granted personhood?This is the government's main argument.

And will we be able to sleep at night when things that surpass us in intelligence are separate and unequal? magine you are a woman in search of romance in this new world.

After years of experience, you've found that your AI is actually better at choosing men than you.

It's also familiar with millions of other people's inventions—it has scanned patent filings going back hundreds of years—and it has read every business book written since Ben Franklin's time.When you bring up a new idea for your business, your AI instantly cross-references it with ideas that were introduced at a conference in Singapore or Dubai just minutes ago.

And there is even one system—call it a guardian-angel AI —that watches over your "best friend" to make sure the advice she's offering you isn't leading you to bad ends.

Yes, there are full-AI zones in 2065, where people collect healthy UBIs and spend their time making movies, volunteering and traveling the far corners of the earth.

Others flee because they don't trust the machines Even the most advanced full-AI zones, in places like China and the United States, will be vulnerable

But the most unanticipated result of the singularity may be a population imbalance, driven by low birth rates in the

full-AI zones and higher rates elsewhere. It may be that the new technologies will draw enough crossers to the full-AI side to even up the numbers, or that test-tube babies will become the norm among those living with AI.

Imagine that the nation's leaders long ago figured out that the only real threat to their rule was their citizens—always trying to escape, always hacking at the AI, always needing to be fed.

It would make a certain kind of sense: To an AI trained to liquidate all resistance , even a minor disagreement with the ruler might be a reason to act.

We're not going to get the AI we dream of or the one that we fear, but the one we plan for.

If there's one thing that gives me pause, it's that when human beings are presented with two doors—some new thing, or no new thing—we invariably walk through the first one.