

13 Forecasts on Artificial Intelligence

Author, Francesco Corea

A Data Science Foundation White Paper

June 2018

www.datascience.foundation

Copyright 2016 - 2017 Data Science Foundation

13 Forecasts on Artificial Intelligence

Where AI is headed

I. Overview

We have discussed some AI topics in the previous posts, and it should seem now obvious the extraordinary disruptive impact AI had over the past few years. However, what everyone is now thinking of is where AI will be in five years time. I find it useful then to describe a few emerging trends we start seeing today, as well as make few predictions around machine learning future developments. The following proposed list does not want to be either exhaustive or truth-in-stone, but it comes from a series of personal considerations that might be useful when thinking about the impact of AI on our world. The interesting aspect of those is that are predictions made one year ago, and many turned out to be true.

II. The 13 Forecasts on AI

1. AI is going to require fewer data to work.

Companies like Vicarious or Geometric Intelligence are working toward reducing the data burden needed to train neural networks. The amount of data required nowadays represents the major barrier for AI to be spread out (and the major competitive advantage), and the use of probabilistic induction (Lake et al., 2015) could solve this major problem for an AGI development. A less data-intensive algorithm might eventually use the concepts learned and assimilated in richer ways, either for action, imagination, or exploration.

2. New types of learning methods are the key.

The new incremental learning technique developed by DeepMind called Transfer Learning allows a standard reinforcement-learning system to build on top of knowledge previously acquired—something humans can do effortlessly. MetaMind instead is working toward Multitask Learning, where the same ANN is used to solve different classes of problems and where getting better at a task makes the neural network also better at another. The further advancement MetaMind is introducing is the concept of dynamic memory network (DMN), which can answer questions and deduce logical connections regarding series of statements.

3. AI will eliminate human biases, and will make us more “artificial”.

Human nature will change because of AI. Simon (1955) argues that humans do not make fully rational choices because optimization is costly and because they are limited in their computational abilities (Lo, 2004). What they do then is “satisficing”, i.e., choosing what is at least satisfactory to them. Introducing AI in daily lives would probably end it. The idea of becoming once for all computationally-effort independent will finally answer the question of whether behavioral biases exist and are intrinsic to the

human nature, or if they are only shortcuts to make decisions in limited-information environment or constrained problems. Lo (2004) states that the satisficing point is obtained through an evolutionary trial and error and natural selection—individuals make a choice based on past data and experiences and make their best guess. They learn by receiving positive/negative feedbacks and create heuristics to solve quickly those issues. However, when the environment changes, there is some latency/slow adaptation and old habits don't fit the new changes—these are behavioral biases. AI would shrink those latency times to zero, virtually eliminating any behavioral biases. Furthermore, learning over time based on experience, AI is setting up as a new evolutionary tool: we usually do not evaluate all the alternatives because we cannot see all of them (our knowledge space is bounded).

4. AI can be fooled.

AI nowadays is far away to be perfect, and many are focusing on how AI can be deceived or cheated. Recently a first method to mislead computer vision has been invented, and it has been called adversarial examples (Papernot et al., 2016; Kurakin et al., 2016). Intelligent image recognition software can indeed be fooled by subtly modifying pictures in such a way the AI software would classify the data point as belonging to a different class. Interestingly enough, this method would not trick a human mind.

5. There are risks associated with AI development.

It is becoming mainstream to look at AI as potentially catastrophic for mankind. If (or when) an ASI will be created, this intelligence will largely exceed the human one, and it would be able to think and do things we are not able to predict today. In spite of this, though, we think there are few risks associated to AI in addition to the notorious existential threat. There is actually the risk we will not be able to understand and fully comprehend what the ASI will build and how, no matter if positive or negative for the human race. Secondly, in the transition period between narrow AIs and AGI/ASI, there will be generated an intrinsic liability risk—who would be responsible in case of mistakes or malfunctioning? Furthermore, there exists, of course the risk of who will detain the AI power and how this power would be used. In this sense, we truly believe that AI should be run as a utility (a public service to everyone), leaving some degree of decision power to humans to help the system managing the rare exceptions.

6. Real general AI will likely be a collective intelligence.

It is quite likely that an ASI will not be a single terminal able to make complex decisions, but rather a collective intelligence. A swarm or collective intelligence (Rosenberg, 2015; 2016) can be defined as “a brain of brains”. So far, we simply asked individuals to provide inputs, and then we aggregated after-the-fact the inputs in a sort of “average sentiment” intelligence. According to Rosenberg, the existing methods to form a human collective intelligence do not even allow users to influence each other, and when they do that they allow the influence to only happen asynchronously—which causes herding biases.

An AI on the other side will be able to fill the connectivity gaps and create a unified collective intelligence, very similar to the ones other species have. Good inspirational examples from the natural world are the bees, whose decision-making process highly resembles the human neurological one. Both of them use large populations of simple excitable units working in parallel to integrate noisy evidence, weigh

alternatives, and finally reach a specific decision. According to Rosenberg, this decision is achieved through a real-time closed-loop competition among sub-populations of distributed excitable units. Every sub-population supports a different choice, and the consensus is reached not by majority or unanimity as in the average sentiment case, but rather as a “sufficient quorum of excitation” (Rosenberg, 2015). An inhibition mechanism of the alternatives proposed by other sub-populations prevents the system from reaching a sub-optimal decision.

7. AI will have unexpected socio-political implications.

The first socio-economic implication usually associated with AI is the loss of jobs. Even if from one hand this is a real problem (and opportunity from many extents), we believe there are several further nuances the problem should be approached from. First, the job will not be destroyed, but they will simply be different. Many services will disappear because data will be directly analyzed by individuals instead of corporations, and of the major impact AI will have is fully decentralizing knowledge. A more serious concern in our opinion is instead the two-fold consequence of this revolution. First of all, using always smarter systems will make more and more human beings to lose their expertise in specific fields. This would suggest the AI software to be designed with a sort of double-feedbacks loop, which would integrate the human and the machine approaches. Connected to this first risk, the second concern is that humans will be devoted to mere “machine technicians” because we will believe AI to be better at solving problems and probably infallible. This downward spiral would make us less creative, less original, and less intelligent, and it will augment exponentially the human-machine discrepancy.

We are already experiencing systems that make us smarter when we use them, and systems that make us feeling terrible when we do not. We want AI to fall into the first category, and not to be the new “smartphone phenomenon” which we will entirely depend on. Finally, the world is becoming more and more robo-friendly, and we are already acting as interfaces for robots rather than the opposite. The increasing leading role played by machines—and their greater power to influence us with respect to our ability to influence them—could eventually make the humans be the “glitches”.

On a geopolitical side instead, we think the impact AI might have on globalization could be huge: there is a real possibility that optimized factories run by AI systems which control operating robots could be relocated back to the developed countries. It would lack indeed the classic economic lowcost rationale and benefits of running businesses in emerging countries, and this is not clear whether it will level out the countries’ differences or incrementing the existing gaps between growth and developed economies.

8. Real AI should start asking “why”.

So far, any machine learning system is pretty good in detecting patterns and helping decision makers in their processes, and since many of the algorithms are still hard-coded they can still be understood. However, even if already clarifying the “what” and “how” is a great achievement, AI cannot understand the “why” behind things yet. Hence, we should design a general algorithm able to build causal models of the world, both physical and psychological (Lake et al., 2016).

9. AI is pushing the limits of privacy and data leakage prevention.

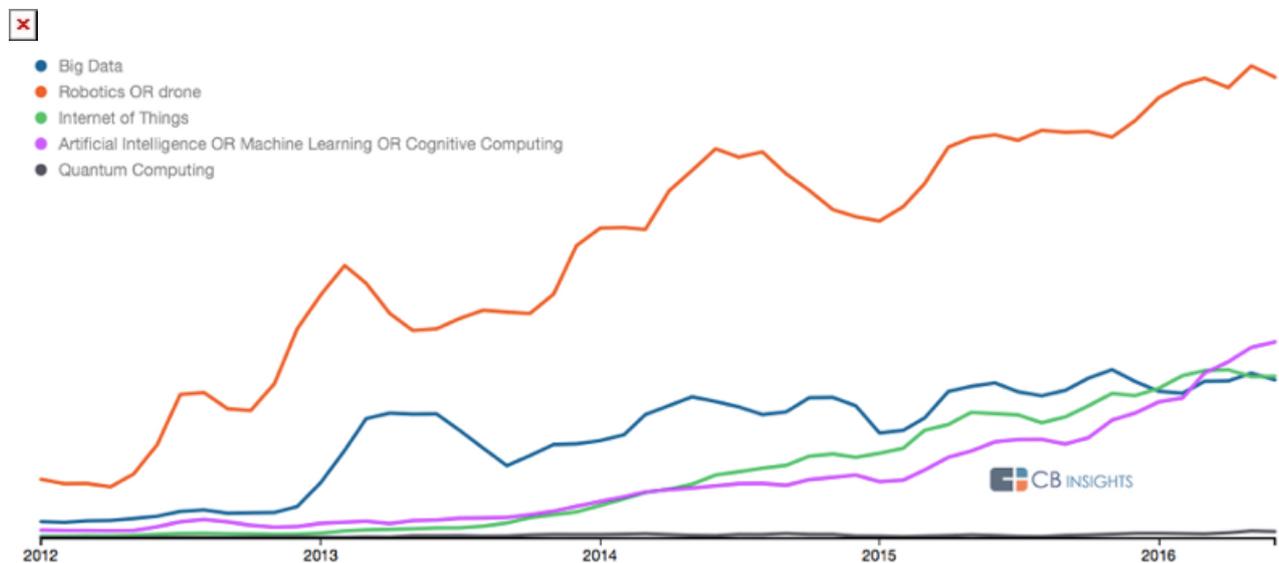
AI is shifting the privacy game on an entirely new level. New privacy measures have to be created and adopted, more advanced than simpler secure multi-party computation (SMPC) or faster than homomorphic encryption. Recent researches show how differential privacy can solve many of the privacy problems we are facing on a daily basis, but there are already other companies looking one step ahead—an example is Post-Quantum, a quantum cybersecurity computing startup.

10. AI is changing IoT.

AI is allowing IoT to be designed as a completely decentralized architecture, where even single nodes can do their own analytics (i.e., “edge computing”). In the classic centralized model, there is a huge problem called server/client paradigm. Every device is identified, authenticated, and connected through cloud servers—that entails an expensive infrastructure. A decentralized approach to IoT networking or a standardized peer-to-peer architecture can solve this issue, reduce the costs, and prevent a single node failure to break down the entire system.

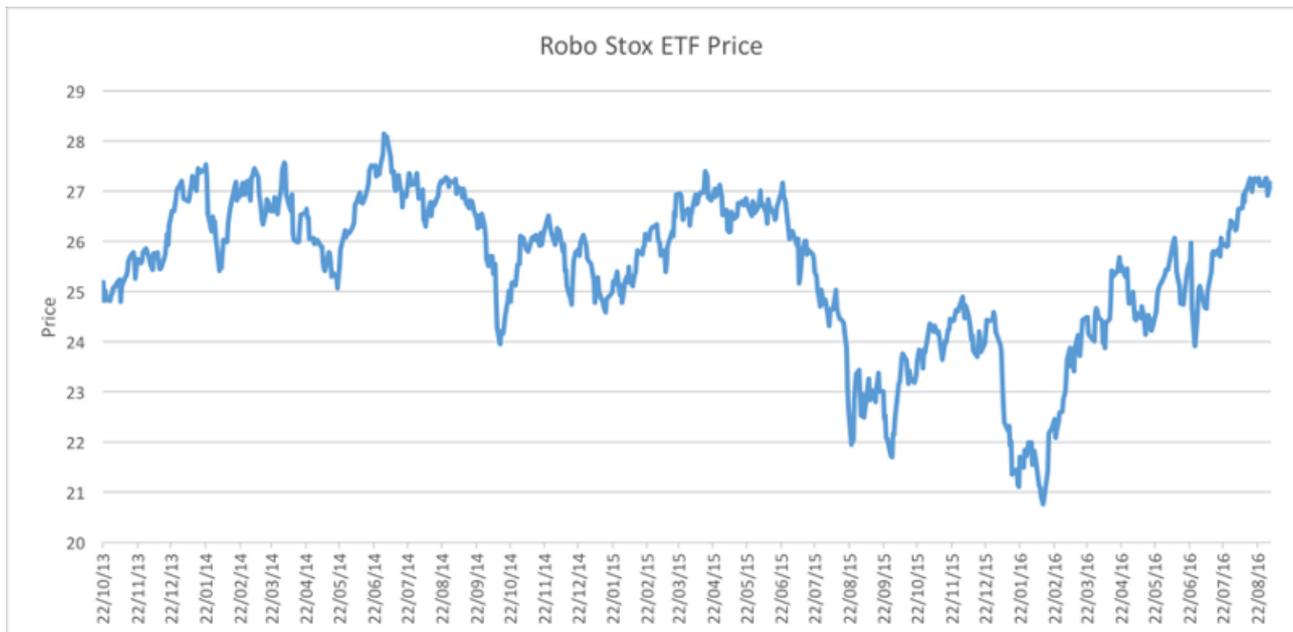
11. Robotics is going mainstream.

I believe that AI development is going to be constrained by advancements in robotics, and I also believe the two connected fields have to go *pari passu* in order to achieve a proper AGI/ASI. Looking at following figure, it is clear how our research and even collective consciousness would not consider an AI as general or super without having a “physical body”.



Search trends for robotics and other fields artificial intelligence alike (created with the CBI Insights Trends tool)

Other evidence that would confirm this trend are: i) the recent spike in robotic patent application, which according to IFI Claims reached more than 3,000 applications in China, and roughly the same number spread across USA, Europe, Japan, and South Korea; ii) the price trend for the Robo Stox ETF, as shown in next figure.



Robo Stox ETF Price trend, for the period 2013–2016.

12. AI might have a real barrier to development.

The real barrier for running toward an AGI today is not the choice of algorithms or data we used (not only at least) but is rather a mere structural issue. The hardware capacities, as well as the physical communications (e.g., the internet) and devices power, are the bottlenecks for creating an AI fast enough—and this is why I believe there exist departments such as Google Fiber. This is also why quantum computing is becoming extremely relevant. Quantum computing allows us to perform computations that Nature does instantly although they would require us an extremely long time to be completed using traditional computers. It relies on properties of quantum physics, and it is all based on the idea that traditional computers state every problem in terms of strings of zeros and ones. The qubits instead identify quantum states where a bit can be at the same time zero and one. Hence, according to Frank Chen (partner at Andreessen Horowitz), transistors, semiconductors, and electrical conductivity are replaced by qubits—that can be represented as vectors—and new operations different from traditional Boolean algebra.

A common way to explain the different approach of traditional vs. quantum computing is through the phonebook problem. The traditional approach for looking for a number in a phonebook proceeds through

scanning entry by entry in order to find the right match. A basic quantum search algorithm (known as Grover's algorithm) relies instead on what is called "quantum superposition of states", which basically analyzes every element at once and determines probabilistically the right answer.

Building a quantum computer would be a scientific revolutionary breakthrough, but it is currently extremely hard to build according to Chen. The most relevant issues are the elevated temperature needed for superconducting materials the computer will be built with; the small coherence time, which is the time window in which the quantum computer can actually perform calculations; the time for performing single operations; and eventually, the energy difference between the right and the wrong answers is so small to be hard to be detected. All these problems shrink the market space to no more than a few companies working on quantum computing: colossus such as IBM and Intel are working on it since some years, and startups such as D-Wave Systems (acquired by Google in 2013); Rigetti Computing; QxBranch; 1Qbit; Post-Quantum; ID Quantique; Eagle Power Technologies; Qubittek; QC Ware; Nano-Meta Technonoliges; and Cambridge Quantum Computing Limited are laying the foundations for quantum computing.

13. Biological robot and nanotech are the future of AI applications.

We are witnesses of a series of incredible innovations lying at the intersection of AI and nanorobotics. Researchers are working toward creating creatures entirely artificial as well as hybrids, and they even tried to develop biowires (i.e., electrical wires made by bacteria) and organs on chips (i.e., functional pieces of human organs in miniature made by human cells that can replicate some of the organ functions - **Emulate** is the most advanced company in this space). Bio-bots research is also testing the boundaries of materials, and soft-robots have been recently created with only soft components. BAE Systems corporation is also pushing the limits of computing trying to create a "chemical computer (the Chemputer)", a machine that would use advanced chemical processes "to grow" complex electronic systems.

References

Kurakin, A., Goodfellow, I. J., Bengio, S. (2016). "Adversarial Examples in the Physical World". Technical report, Google, Inc. Available at arXiv: 1607.02533.

Lake, B. M., Salakhutdinov, R., Tenenbaum, J. B. (2015). "Human-level concept learning through probabilistic program induction". *Science*, 350(6266): 1332-1338.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., Gershman, S. J. (2016). "Building Machines That Learn and Think Like People". Available at arXiv:1604.00289.

Lo, A. W. (2004). "The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective". *Journal of Portfolio Management* 30: 15-29.

Papernot, N., McDaniel, P. D., Goodfellow, I. J., Jha, S., Celik, Z. B., Swami, A. (2016). "Practical black-box attacks against deep learning systems using adversarial examples". CoRR, abs/1602.02697.

Rosenberg, L. B. (2015). "Human Swarms, a real-time method for collective intelligence". Proceedings of the European Conference on Artificial Life: 658-659.

Rosenberg, L. B. (2016). "Artificial Swarm Intelligence, a Human-in-the-Loop Approach to A.I.". Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16): 4381-4382.

Simon, H. A. (1955). "A Behavioral Model of Rational Choice". The Quarterly Journal of Economics, 69 (1): 99-118.

This article appeared in my book "Artificial Intelligence and Exponential Technologies" edited by Springer (2017).

About the Data Science Foundation

The Data Science Foundation is a professional body representing the interests of the Data Science Industry. Its membership consists of suppliers who offer a range of big data analytical and technical services and companies and individuals with an interest in the commercial advantages that can be gained from big data. The organisation aims to raise the profile of this developing industry, to educate people about the benefits of knowledge based decision making and to encourage firms to start using big data techniques.

Contact Data Science Foundation

Email: contact@datascience.foundation
Telephone: 0161 926 3641
Atlantic Business Centre
Atlantic Street
Altrincham
WA14 5NQ
web: www.datascience.foundation

Data Science Foundation

Data Science Foundation, Atlantic Business Centre, Atlantic Street, Altrincham, WA14 5NQ
Tel: 0161 926 3641 Email: contact@datascience.foundation Web: www.datascience.foundation
Registered in England and Wales 4th June 2015, Registered Number 9624670