

Artificial intelligence: Mute, Blind, and Boring?

Author

Pekka Pekkala

PublishedJune 20th 2017**Topic**

Artificial Intelligence

—• “Everybody takes for granted that a bulldozer is stronger than a human but the idea of a machine being smarter than a human is harder to grasp. But from the engineering perspective, there is no fundamental reason that would prevent machines being smarter than men.”

Jaakko Särelä, data scientist at Reaktor, describes the potential of artificial intelligence¹.

However, if you think closer about a bulldozer, you realize that it is great at one thing—pushing sand or demolishing buildings. If you need to build something instead of ripping it down, a human would do a better job. The horsepower and agility of a bulldozer might be superhuman but it is no match to a human when you need to build a house or sandcastle. The powers of a machine are magnificent but limited.

Narrow Artificial Intelligence²: Do One Thing Well

Like a bulldozer, current AI solutions have very limited talent. If artificial intelligence was a human, it would be described as a savant, a person with extraordinary skills in certain areas, such as mathematics. At the same time, it would have trouble with everyday tasks, such as cooking or crossing a road.

1

Artificial Intelligence

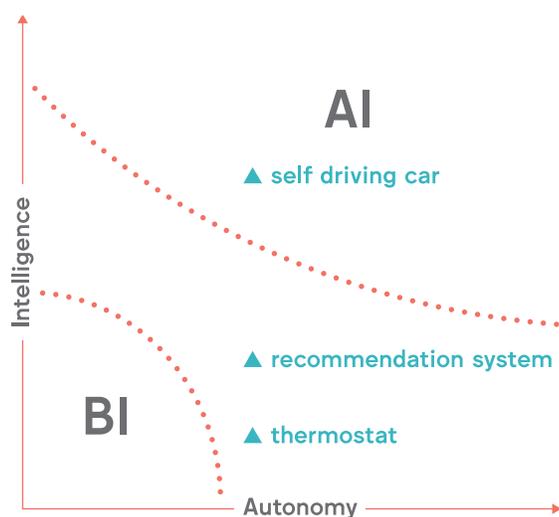
A broader term for the idea that machines are showing intelligent behavior by reacting to their environment and making successful decisions based on that information. Best known examples of AI are the successful algorithms winning strategic games like chess or Go. Speech and image recognition are some of the most advanced forms of AI as well.

The artificial intelligence inside Google’s game-winning DeepMind, Tesla’s self-driving autopilot, and Apple’s voice-assisting Siri are completely different machines. Tesla’s autopilot is not able to play Go and Siri has no idea how to drive. In a sense, they are different machines, just like a bulldozer is different from an assembly robot in a car manufacturing plant. They have no means to modify their behavior and use their skills in new situations.

General Artificial Intelligence: Asimov Class Robots

The artificial, adaptive intelligence that we see in Hollywood movies is called **general artificial intelligence**³. The sensing and thinking robots that have multiple human-like skillsets are still far in the future. However, the little helpers are already here. From a business perspective, it is important to understand that what we consider a “smart” or “intelligent” machine changes over time. When you saw the first flatbed scanner that could recognize letters from a printed paper, it felt like a miracle. Now, character recognition is commonplace and it is not considered “intelligent” anymore.

Another important factor to consider is the difference between intelligence and autonomy. Both of them are needed to create AI. A thermostat keeping your room temperature comfortable is not considered smart, even when it makes something effortless, that would require at least some skill from a human counterpart. When a self-driving car drives as fluently as a thermostat does setting the temperature, it becomes autonomous. In this sense, Tesla’s autopilot is just a smarter version of a thermostat.



We can classify systems by their general intelligence or how autonomous a system is. The more intelligent, or autonomous a system is, the more easily we call it an AI system. In contrast, systems with less intelligence and autonomy fall often in the category of Business Intelligence.

2

Narrow Artificial Intelligence

Current state of AI, the narrowly-skilled but superb programs that recognize cats from pictures or beat humans in Go. Even the technology needed to create self-driving cars is considered narrow, in spite of the multiple technologies it needs for the task.

3

General Artificial Intelligence

This is the stuff from science fiction movies and books: a human-like robot like C-3PO from Star Wars or the voice of Samantha from the movie Her. This is not the AI of today and it might be easily 50 years from now before we see any type of General AI.

The current wave of artificial intelligence is popularized with bots that have human-like communication skills and, thus, appear a bit like a human. Apple's Siri, Amazon's Alexa, and Google's Assistant are amazing, useful, and entertaining chatbots. But, how useful they are for business cases?

Industrial Intelligence: Machine to Machine Communication

- “What about AI bots that do not communicate with humans but with other machines? They might be thousand times smarter than a chatbot but since they don't have visible or audible signs of human intelligence, they are not considered as AI.”

Johan Himberg, Reaktor data scientist

There is also a lot of practical AI innovation made in this field today. Image or sound recognition has industrial uses as well and all the data that we need to recognize—including the pictures and sounds—might be produced by machines. Computer vision does not care if a picture is an adorable cat or a sound is harsh. It is all just data to be processed.

These kind of AI applications might listen to a sound of a certain machine at a manufacturing plant or record a video of a furnace. The idea is to do artificial intelligence in industrial setting: gather as much data as possible (audio or video) and recognize patterns in it. These patterns might show a difference in the way a certain machine sounds before it needs serviced or how a burning process produces too much visible soot and creates pollutants in the process.

The data created from these sensors might not make any sense to a human but machines can interpret them and send commands to other machines, ask for more oxygen for the furnace, or alert the service bot to order new parts next week.

Computers do not get tired watching the flames or listening to the monotonous sound of an engine. They gather data, process it, and look for patterns unrecognizable to human eye or ear. The same basic principle of artificial intelligence applies to everything from industrial solutions to personalization algorithms in Netflix recommendations or Amazon.com store. Gather data; recognize patterns; make decisions based on the given rules; apply them; repeat a million times.

Machine Learning: Going Deep Layer by Layer

Most AI technologies have been around for decades. A major reason for the current AI boom is the number of sensors, data, and processing power that we have available today. This creates endless possibilities for **machine learning**⁴ and even the much-hyped **deep learning**⁵.

Google search is probably the best-known example of machine learning. It has been trained by different feedback systems for 20 years to give us the most relevant search results available. Feedback systems include the number of links pointing to a certain page, the trustworthiness of a certain domain, or what you have clicked before in your search results. All these feedback loops, created by the millions of Google users every second, make the algorithm better.

In industrial applications, just imagine machines instead of people and data generated by those machines instead of search queries. The same basic AI principles apply to both: gather data, analyze patterns, make decisions, get feedback, analyze again, make better decisions.

When we enter the realm of deep learning and ways to build **neural networks**⁶ with it, AI gets more complicated to understand. However, you don't need to worry, even scientists have trouble understanding the AI systems that they have created. Despite the human-like term "neural network", computers still make decisions very differently from humans. The programmers have trouble understanding the errors that AI makes. They literally make no sense, in human terms.

Perhaps the best example of this is Google Translate. When it was upgraded with Neural Machine Translation, the results improved by 25% and even the creators of the system failed to understand why it happened. To make things even more mysterious, it seemed that Google Brain created its own language, a kind of interlingua, to help with the translations. This, with other advances in **natural language processing**⁷, will keep the human-like communication interfaces and outputs in the limelight of artificial intelligence.

Boring Applications, Awesome Results

However, as intriguing as these human-like advances in AI are, they are just one side of the revolution created by the amount of available

4

Machine Learning

Instead of programming machines to do certain a task directly, they are being "trained" to create programs for specific tasks by themselves. Training means giving computers large amounts of data, models to analyze it, and a chance to give feedback.

5

Deep Learning

The current cutting-edge of "AI", an advanced model of machine learning made possible by the increase of computing power and the amount of available data this is fed to a neural network. When the computer has enough data and parallel processing power, it can teach itself and narrow the margin of error in the probabilities that it calculates, to resemble human intelligence, going from "there is a cat in this picture with 94% probability" to "99.9% probability". This is the technology that you need for self-driving cars. It was also used to train the Google Brain that beat humans in the game of Go.

6

Neural Networks

A network of different algorithms that work together to find classification or prediction models from data. If you want the computer to recognize a cat picture, you have different algorithms to recognize borders and patterns that are understood as cat whiskers, paws, claws, or tails. From these individual clues, the network can decide that, yes, this is indeed a cat picture. The network model is called neural because it resembles the model of the neural network in human brain. Neural networks make deep learning possible.

data and processing power. Artificial intelligence is great at doing repetitive, boring tasks that you or your business would love to outsource to someone else.

Computers, even the artificially intelligent ones, don't get tired or bored and they work around the clock. AI likes repetition because it makes it smarter. Google's AlphaGo studied 30 million positions from previous games and played endlessly against itself before beating a human. Computer vision and sound recognition gather data around the clock, making better predictions about upcoming production equipment failures. Algorithms made by Reaktor give personalized recommendations to every user in different situations, such as gaming, shopping and entertainment. The applications for AI are endless and the time to use it is right now.

If you need help to figure out how Reaktor could help your company with artificial intelligence and data science, check out our offerings at reaktor.ai.

7

Natural Language Processing

The idea of computers communicating with people with natural language instead of commands has been around since 1950s, when Alan Turing developed Turing test. This has been one of the fastest developing fields in AI. For example, Google Translate can now translate over 100 languages. There are multiple services offering speech-to-text services and semantic text analysis. Speech and language are easily recognized as "human" skills, making NLP one of the easiest ways to demonstrate the power of AI to the public.