



Harvard Business Review

REPRINT H03857
PUBLISHED ON HBR.ORG
NOVEMBER 09, 2016

ARTICLE **ANALYTICS**


What Artificial Intelligence Can and Can't Do Right Now

by Andrew Ng

ANALYTICS

What Artificial Intelligence Can and Can't Do Right Now

by Andrew Ng
NOVEMBER 09, 2016



Many executives ask me what artificial intelligence can do. They want to know how it will disrupt their industry and how they can use it to reinvent their own companies. But lately the media has sometimes painted an unrealistic picture of the powers of AI. (Perhaps soon it will take over the world!) AI is already transforming web search, advertising, e-commerce, finance, logistics, media, and more. As the founding lead of the Google Brain team, former director of the Stanford Artificial Intelligence Laboratory, and now overall lead of Baidu's AI team of some 1,200 people, I've been

privileged to nurture many of the world’s leading AI groups and have built many AI products that are used by hundreds of millions of people. Having seen AI’s impact, I can say: AI will transform many industries. But it’s not magic. To understand the implications for your business, let’s cut through the hype and see what AI really is doing today.

Surprisingly, despite AI’s breadth of impact, the types of it being deployed are still extremely limited. Almost all of AI’s recent progress is through one type, in which some input data (A) is used to quickly generate some simple response (B). For example:

What Machine Learning Can Do

A simple way to think about supervised learning.

INPUT A	RESPONSE B	APPLICATION
Picture	Are there human faces? (0 or 1)	Photo tagging
Loan application	Will they repay the loan? (0 or 1)	Loan approvals
Ad plus user information	Will user click on ad? (0 or 1)	Targeted online ads
Audio clip	Transcript of audio clip	Speech recognition
English sentence	French sentence	Language translation
Sensors from hard disk, plane engine, etc.	Is it about to fail?	Preventive maintenance
Car camera and other sensors	Position of other cars	Self-driving cars

SOURCE ANDREW NG

© HBR.ORG

Being able to input A and output B will transform many industries. The technical term for building this A→B software is *supervised learning*. A→B is far from the sentient robots that science fiction has promised us. Human intelligence also does much more than A→B. These A→B systems have been improving rapidly, and the best ones today are built with a technology called deep learning or deep neural networks, which were loosely inspired by the brain. But these systems still fall far short of science fiction. Many researchers are exploring other forms of AI, some of which have proved useful in limited contexts; there may well be a breakthrough that makes higher levels of intelligence possible, but there is still no clear path yet to this goal.

Today’s supervised learning software has an Achilles’ heel: It requires a huge amount of data. You need to show the system a lot of examples of both A and B. For instance, building a photo tagger requires anywhere from tens to hundreds of thousands of pictures (A) as well as labels or tags telling you if there are people in them (B). Building a speech recognition system requires tens of thousands of hours of audio (A) together with the transcripts (B).

So what can A→B do? Here’s one rule of thumb that speaks to its disruptiveness:

If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future.

A lot of valuable work currently done by humans — examining security video to detect suspicious behaviors, deciding if a car is about to hit a pedestrian, finding and eliminating abusive online posts — can be done in less than one second. These tasks are ripe for automation. However, they often fit into a larger context or business process; figuring out these linkages to the rest of your business is also important.

AI work requires carefully choosing A and B and providing the necessary data to help the AI figure out the A→B relationship. Choosing A and B creatively has already revolutionized many industries. It is poised to revolutionize many more.

After understanding what AI can and can't do, the next step for executives is incorporating it into their strategies. That means understanding where value is created and what's hard to copy. The AI community is remarkably open, with most top researchers publishing and sharing ideas and even open-source code. In this world of open source, the scarce resources are therefore:

- **Data.** Among leading AI teams, many can likely replicate others' software in, at most, 1–2 years. But it is exceedingly difficult to get access to someone else's data. Thus data, rather than software, is the defensible barrier for many businesses.
- **Talent.** Simply downloading and “applying” open-source software to your data won't work. AI needs to be customized to your business context and data. This is why there is currently a war for the scarce AI talent that can do this work.

Much has been written about AI's potential to reflect both the best and the worst of humanity. For example, we have seen AI providing conversation and comfort to the lonely; we have also seen AI engaging in racial discrimination. Yet the biggest harm that AI is likely to do to individuals in the short term is job displacement, as the amount of work we can automate with AI is vastly bigger than before. As leaders, it is incumbent on all of us to make sure we are building a world in which every individual has an opportunity to thrive. Understanding what AI can do and how it fits into your strategy is the beginning, not the end, of that process.

Andrew Ng is former VP & Chief Scientist of Baidu, Co-Chairman and Co-Founder of Coursera, and an Adjunct Professor at Stanford University.
