


WHEN MACHINES START THINKING

HOW AI IS
SHAPING
OUR DAY
AFTER
TOMORROW

by Peter Hinssen
in partnership with  nexworks

EDITION

1

ANY SUFFICIENTLY
ADVANCED
TECHNOLOGY IS
INDISTINGUISHABLE
TO MAGIC

ARTHUR C CLARKE




If you really want to understand Artificial Intelligence, you must understand its roots. They date back to a pivotal conference that was organized on the campus of Dartmouth College – the smallest university in the Ivy League – in the glorious summer of 1956. For eight weeks in a row, a group of dedicated people worked diligently together on the top floor of the Dartmouth Mathematics Department to create an entirely new discipline in the history of computer science.

The Dartmouth Conference was organized by Marvin Minsky, John McCarthy and Claude Shannon. These scientists firmly believed that “every aspect of learning or any other feature of intelligence can be so precisely described, that it is possible to perform it by machines”.

This conference is now widely considered to be the birth of Artificial Intelligence. It is where the field got its name, received its mission, and where the major players gathered for the first time to address the challenges ahead. Those who attended would become the leaders of AI-research for decades to come.

The 1956 Dartmouth conference might have been pivotal, but others had already been working on AI before that. In the 1940s and 1950s, research into neurology had unveiled that our brain was essentially a vast electrical network of neurons and synapses which used all-or-nothing electrical pulses. This insight created great excitement amongst scientists and engineers in the emerging electronics field about how they could simulate or emulate this.

Claude Shannon, the father of information theory, described how digital (binary) signals could be used to transfer information. Norbert Wiener coined the phrase ‘Cybernetics’ and showed how control theory could help us use electrical networks to drive and steer robots. And the brilliant British mathematician Alan Turing devised a theory of computation, showing that any ‘mathematical logical problem’ could be described digitally. The coming together of these disciplines made the scientists dream that, one day, it would be possible to design and build an electronic brain.



THINKING MACHINES

Turing is most known for developing a computer that broke the German encryption codes during World War II in the Enigma project. He was deeply enamored with the huge potential of computing power to solve logical challenges. In 1950, he wrote a landmark paper in which he speculated about the possibility of creating machines that could think, arguing that 'thinking' is difficult to actually define. In order to address that problem he devised the now famous Turing Test: "If a machine is capable of carrying on a conversation that is indistinguishable from a conversation with a human being, it is reasonable to say that the machine is 'thinking'. The Turing Test was probably the first serious proposal in the philosophy of artificial intelligence.


Companies like Disney employ algorithms to engage with users on social media who want to travel to Disney theme parks, and the customers have no idea that they are talking to a computer instead of a human. I'm not sure Turing would have thought that his ideas would lead to people booking hotel rooms, spa treatments and Mickey Mouse photoshoots at the Magic Kingdom. But hey, there's progress I guess.

One of the students deeply inspired by Turing's papers was the young Marvin Minsky. This AI-rock star built the first randomly wired neural network learning machine in 1951 when he was just 24 years old.

He called it the SNARC. Neural networks are fascinating: the basic idea behind them is to simulate lots of densely interconnected brain cells inside a computer so you can get it to learn things, recognize patterns and make decisions in a very humanlike way. The amazing thing about a neural network is that you don't have to program it to learn explicitly: it learns all by itself, just like a brain. Minsky would become one of the most important leaders and innovators in the field of AI for the next fifty years.

Claude Shannon was the father of information theory. John McCarthy developed the first programming language for Artificial Intelligence: LISP. And Marvin Minsky built the first neural network machine. These three men organized the first Dartmouth conference, and a new field was born. And they were WAY too optimistic. Almost naive.

<http://www.explainthatstuff.com/introduction-to-neural-networks.html>



SLOWER THAN EXPECTED

They predicted that a digital computer could become chess champion of the world by 1970. Nice try. The founders of AI were off by almost 30 years.

They also predicted that by 1985 "machines will be capable of doing any work a man can do". Marvin Minsky predicted in Life magazine that by the end of the 20th century "We will have a machine with the general intelligence of an average human being."

Well. Not really. The majority of us humans are still baffled by how incredibly stupid our computers are, how horribly slow they are to understand our demands. Most of us are frustrated by their ignorance when we have to instruct automated systems in a bank or an airline what we want to do.

After the Dartmouth conference, money poured into the field. The government, the military, the industrial players all loved the idea of a computer that could think, machines that could be like human beings, or better, or cheaper, and they started sending truckloads of cash to the researchers, engineers and universities.

They wanted to believe. They really wanted to recreate mankind, in every aspect. In 1979 McCarthy wrote an article called "Ascribing Mental Qualities to Machines." In it, he stated that "Machines as simple as thermostats can be said to have beliefs, and having beliefs seems to be a characteristic of most machines capable of problem-solving performance."

WINTER IS COMING

But then came the disappointment. It was the advent of the first 'AI-winter' in the late seventies. The reason was simple: there was simply not enough computer power and memory to run all the ideas and concepts of the AI-researchers. These were the days when companies like Atari and Commodore started building home computers that had 64Kilobytes of memory, and the systems that the AI-researchers had were not that much more powerful. Tackling the challenges of AI with the computers of that era was like trying to get to space with the use of sticks and stones. No go.

Hans Moravec, one of the leading researchers at the time, stated, in 1976, that 'computers were still millions of times too weak to exhibit intelligence'. Computers are measured in MIPS (million instructions per second). An Apple II at the time would have had 1 MIPS. The fastest computer back then (a Cray-1 supercomputer) would have an awesome 130 MIPS. In order for machines to display the beginnings of actual intelligence, they would need 1,000,000 MIPS.

So, being nowhere near this number of 1,000,000 MIPS, the first winter of AI began. The press was jumping all over the failed promises of the AI-community. Governments, military and corporates alike felt misled.

They were disappointed that the AI-researchers had grossly over-promised and massively under-delivered. So funds started to dry up. Science fiction visions of computers smarter than humans were tucked away, and many of the researchers in the field went on to do other things. Artificial Intelligence was regarded as a pipe dream, to be stored in the same category as the search for the philosopher's stone or the fountain of youth.



HIGH HOPES

It was the Japanese government that rekindled the hope of Artificial Intelligence in the early 1980s. At that time, companies like SONY and Toyota had shown the world that Japanese industry was not just about copying the West and making it cheaper. They were actually redefining electronics and automotive. Japan decided to show to the world that they were leading in technological innovation. So, as of 1982, Japan's Ministry of International Trade and Industry funded its 'Fifth Generation Computer Project' which aimed to create an "epoch-making computer" with supercomputer-like performance and to provide a platform for future developments in artificial intelligence¹. The result, among other things, was massive enthusiasm around the concept of 'Expert Systems'.

The latter are programs that answer questions, or that can solve problems about a particular domain of knowledge, based on logical rules that come from the pooled knowledge of 'experts'. Instead of a 'dumb' database of facts, expert systems have content, knowledge and rules that help solve questions, search for knowledge, and present solutions.

Soon universities offered expert system courses and two-thirds of the Fortune 500 companies applied the technology in daily business activities. Instead of the old AI-approach that was a lot broader, the expert systems focused on a very small domain of specific knowledge. And, for that, the limited horsepower of the computers of the time was powerful enough to achieve results.

Expert systems were used by oil companies to find new oil wells, by mining companies to understand where to drill, and by hedge funds to understand where to invest.

But they were expensive. Not just to build, and to run, but also to maintain. The systems did not 'learn' like humans; they had to be regularly fed with new rules. Likewise, their intelligence did not 'evolve' or grow. By the mid-eighties, the initial spring revival and enthusiasm for expert systems had started to fade, and gave way to another AI-winter.

¹https://en.wikipedia.org/wiki/Fifth_generation_computer

THE MOTHER OF SEARCH

But the knowledge of how to query information in these expert systems and the capability to find solutions in large amounts of information gave way to an incredibly lucrative new industry: search. When the worldwide web exploded onto the scene in 1995, several companies were able to develop their search algorithms thanks to the enormous research and knowledge that came out of this expert systems' era. It made many of them very wealthy. But it did not change the fact that, at the end of the eighties, AI was put into hibernation again.

Journalist John Markoff remarked in an article in the New York Times in 2005: "At its low point, some computer scientists and software engineers avoided the term artificial intelligence for fear of being viewed as wild-eyed dreamers.

But Moore's law was on their side. When the first AI-winter was caused simply because there was not enough computing horsepower to go around, it was just a matter of time. This was a game just like the second half of the chess board. Computers were getting more powerful every day, computer chips started to become incredibly strong and we could combine more and more to create 'parallel' computers that worked in harmony. Big computer manufacturers had not given up, and had genius teams working on the dream of computer intelligence. And they had their mind set on the initial promise of Dr. Marvin Minsky: beat the best chess player in the world.

DEEP BLUE AND BEYOND


The most pivotal moment in the history of Artificial Intelligence came on the 11th of May 1997, when the first computer chess-playing program defeated the reigning world chess champion, Garry Kasparov.

But it did not stop there, it was like the uncorking of a Champagne bottle that unleashed an avalanche of new discoveries, new breakthroughs and new exciting headways into the evolution of Artificial Intelligence.

A mere 8 years later, in 2005, a Stanford Robot drove a car autonomously over an unrehearsed desert trail for more than 130 miles as part of the Grand DARPA Challenge.

Two years later, a team from Carnegie Mellon University drove a car autonomously through an urban environment for more than 50 miles, tackling traffic hazards and adhering to all traffic laws.

In 2011, a computer defeated the two all-time best human Jeopardy Players, Brad Rutter and Ken Jennings, by a striking margin. It was capable of listening to the Jeopardy questions and reason and respond much faster than the human contestants. It was clear that the AI-winter was coming to an end.



The question is “was there ever an AI-winter”? Artificial Intelligence never really went away. As a matter of fact, many AI-researchers claim that it is quite the opposite: AI is everywhere. Rodney Brooks, one of the prominent researchers in the field complained in 2002 that “There seems to be this stupid myth out there that AI has failed, but AI is around you every single second of the day.”

To understand his claim, it’s necessary to understand that it’s not only artificial intelligence if it emulates the working of the human brain. That’s the ambition, but not the only type of AI. Basically, there are three different levels of artificial intelligence: Artificial Narrow Intelligence, Artificial General Intelligence and the fairest of them all: Artificial Superintelligence.

Artificial Narrow Intelligence specializes in just one area. It can beat the world chess champion at chess but it won’t be able to play monopoly. This is the type Rodney Brooks is talking about. Our world is pervaded with it: it’s in the maps on our phones, fights off spam, regulates our thermostats, helps Spotify recommend tunes and arrange your feed. It’s everywhere. But basically, it’s really good at doing one thing. That’s why they call it Weak AI.

As we step into the next room of the AI museum, we arrive at the phase of Artificial General Intelligence. This type aims to be as smart as a human and able to perform the same kind of intellectual tasks that we can. We are moving in this direction, thanks to self-learning systems, as we will see later, but we haven’t succeeded yet in this domain.

This type is also referred to as Strong AI, or Human-Level AI. And then there’s the scary type of Artificial Superintelligence – which surpasses that of humans in an exponential manner – which will (probably) not be here for quite a few years².

Artificial Narrow Intelligence is everywhere. Like in the case of smart speakers, which are essentially an AI interface to the internet. You can communicate to it in natural language, ask it simple questions in English like: “What is the weather forecast for tomorrow?”, or more complex challenges like: “Book me an Uber to get to the cinema”. Soon you will be able to throw complex tasks to these AI-interfaces like: “Check the timing of my upcoming flight to Houston, and make sure there is an Uber to pick me up in time”. The AI-device will not only understand your question, but will make sure to calculate the traffic on the road on the way to the airport. Marvin Minsky would be proud.

Smart speakers will bring AI towards a mainstream audience. On the one hand of the spectrum you have the Self-Driving Ubers and Autonomous Drones and, on the other hand, you will see that AI becomes a mainstream interface towards consumers. In just a short period of time we have seen ‘digital first’ – as the way to communicate to customers – become ‘mobile first’ as mobile becomes the dominant gateway to consumers. But very soon ‘AI first’ will replace ‘mobile first’ as the preferred consumer-facing interaction.

²<http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>


WHERE IS HAL?

Timing is a bit of an issue in Artificial Intelligence. I became fascinated with the field when I watched '2001: A Space Odyssey'. Stanley Kubrick's magnificent film was based on the novel by science fiction writer Arthur C. Clarke. The HAL – short for Heuristically programmed Algorithmic computer – 9000 computer was brilliantly visualized by a throbbing red light that observes the activities in the Discovery One spacecraft. Described as having become operational on the 12th of January 1997, it interacts with the astronauts and speaks in a soft and calming voice. It is clear that HAL is extremely smart. Yet it also goes completely cuckoo and tries to murder the astronauts.

The brilliant book by Arthur C. Clarke explains that HAL is caught in a moral dilemma. It was forced to lie to the crew about the true nature of their mission, and this causes enormous internal conflicts inside the HAL circuits. Unable to resolve the conflict between his general purpose to relay information accurately, and his direct orders to withhold the true purpose of the space mission, HAL comes to the conclusion that if the crew dies, he would not need to lie to them anymore.

In the end, the commander of the mission, Bowman, manages to enter the central core system of HAL and starts shutting it down.

The movie was, and still is, fascinating, both in the magnificent splendor of the visuals and in the serenity of the dialogues and interactions between HAL and the crew-members. But it was pure science fiction, of course. No factory in the world could have produced a HAL in 1997.



But '2001: A Space Odyssey' raises fundamental questions about how to control AI when it reaches HAL's level of intelligence, which it almost certainly will in our lifetimes. How can we ensure the stability of systems, the integrity of logic, and how can we make sure that AI systems can resolve such conflicts or dilemmas?

Stewart Russel is a professor of computer science at the University of Berkeley, just outside of San Francisco. He has spent his life on exactly this problem: how can we marry the future of AI with the future of the human race? He believes we must integrate 'human' values into the AI systems of the future. And he has very outspoken and no-nonsense examples: "If you want to have a domestic robot in your house, it has to share a pretty good cross-section of human values. Otherwise it's going to do pretty stupid things, like put the cat in the oven for dinner because there's no food in the fridge and the kids are hungry."

Well, we certainly don't want that!

That's why Stewart Russel believes that implementing 'values' in AI will be crucial. "It only takes one or two things like a domestic robot putting the cat in the oven for dinner for people to lose confidence and not buy them."

Before we get to 'HAL-9000-level' intelligence we will probably see plenty of strange and absurd things happen as we start to utilize and bring AI into service. And some of it will grossly malfunction. We will have people getting injured or losing their life when a machine learning autopilot on their self-driving car malfunctions. Or we will see irrational and nonsensical behavior in algorithms.



A ROBOT FOR THE ELIMINATION OF TEDIOUS TASKS

Next door to the office of Dr. Stewart Russel at Berkeley, a group of researchers has been working for a long time on the creation of BRETT: the 'Berkeley Robot for the Elimination of Tedious Tasks'. The project was the brainchild of Russel's colleague Dr. Pieter Abbeel, who runs the robotics lab at Berkeley and, in 2016, became the right-hand man of Elon Musk on the Open AI Initiative.

I had the pleasure to visit Pieter Abbeel many times in his lab in Berkeley, and see the progress of BRETT. The project started out as a joke when he was still studying to get his PhD at Stanford in robotics and AI. If there was one thing that Pieter hated more than anything in the world, it was folding laundry. He felt that the time spent on the mindless task of folding up shirts, pants and socks was just a terrible waste of anyone's mental capabilities.

So he set off to build a robot that could fold the laundry. His attempts made him a viral star on the internet. He and his team of students embarked on the journey of teaching BRETT – who was based on a standard PR2 robot from Willow Garage – how to fold a shirt. It turned out to be more complex than they initially thought. The robot has to 'see' where there is a cuff, an elbow, a shoulder, and then reason how (in 3D) to figure out the folding pattern. It learned a lot, and eventually succeeded, although BRETT's first incarnation took about 2 hours to fold a complete shirt.

Pieter Abbeel is a Belgian born scientist, who went to Stanford to study under Sebastian Thrun who developed the first driverless car. Pieter is an absolute techno-optimist, who seems less concerned about a robot putting a cat in the oven. He is absolutely convinced that we are turning the corner in AI and are leaving the AI-winters behind us.



"The huge breakthrough came around 2012", he recalls, "When we made amazing quantum leaps forward in computer vision." He would know, as his BRETT robot had to 'see' the shirts and pants in order to fold them.

According to Abbeel, that is exactly where the melting of the glacial AI-winter started. "2012 saw the creation of AlexNet, essentially taking the concept of a huge neural network, that was trained specifically for computer vision. Think of it as a huge flexible block of computation, an 8-layer neural network with more than 60 million parameters to learn. We proceeded to feed the network images and pictures, to train the system.'

By 'showing' the neural network millions of pictures, and 'training' that network to recognize objects like a 'cat' or a 'tree', the system started to learn very quickly. AlexNet spawned an enormous amount of research, and today not only can computer-vision AI networks recognize cats and trees, people and cars, locations and weather conditions, they can now be used to analyze complex situations on images and pictures.

You can show an image to a neural network, ask it to observe and then get responses like: "A woman holding a camera in a crowd.", or "A guy on a skateboard on the side of a ramp".

THE TIME IS (W)RIGHT

A vintage biplane is shown in flight against a solid golden background. The aircraft is a two-winged plane with a prominent tail boom and a tail wheel. The wings are dark with a grid-like pattern, and the fuselage is a light color. The plane is angled upwards and to the right, suggesting forward motion.

In this intoxicating springtime excitement for AI we have arrived at a pivotal moment where the biggest network players are investing a lot in order to be at the forefront of this new revolution.

THAT TIME UBER TOOK OVER (A PART OF) CARNEGIE MELLON

Carnegie Mellon is one of the world's top research universities. It was founded in Pittsburgh in 1900 by the steel magnate Andrew Carnegie who wanted to create a top-class engineering school in order to boost the steel industry around Pittsburgh. It evolved into a top engineering college that has some of the world's best researchers in robotics, AI and autonomous systems.

In September 2015, Uber surprised the world when it announced that it had poached no fewer than 49 top researchers from Carnegie Mellon's National Robotics Engineering Center, the NREC. Uber knew exactly what they wanted. They wanted the world's best engineers and researchers working on the future of driverless cars, and they happened to be the ones working at the NREC.

Uber came in and made NREC's researchers an offer they could not refuse: Silicon Valley king-sized salaries, and a chance to build the greatest fleet of self-driving cars in the world. The group was settled in Pittsburgh and rebranded to Uber ATC: Advanced Technology Center.

That is precisely the reason why, in the fall of 2016, the first fleet of self-driving Ubers did not start in San Francisco, the headquarters of Uber, but in Pittsburgh.

There seems to be a new kind of exodus of talent towards the new technology platform players and AI is the name of the game. Artificial Intelligence will become big bucks in the next few years and could reshape the industrial landscape like never before. And if you're a top researcher, you'll want to be where the action is.



THE MONEY AND THE POWER

The new technology giants, investing heavily in AI, have the computational horsepower to really make a dent in the universe. They have the money, the resources, and the infrastructure to make things happen.

When you visit the Robotics lab in Berkeley and see the Postdoc and PhD students working on the Berkeley Robot for the Elimination of Tedious Tasks, you feel that there is a huge gap with the real world out there. The BRETT robot is a PR2 robot, that was probably really hot in 2007 when it first came out, but today is hopelessly outdated. Actually the company that built this robot, Willow Garage, went out of business in 2014, and the university students spend more time repairing the damn thing than doing real breakthrough work.

You feel this tension between academics and the real world when you talk to Pieter Abbeel.

That's probably why Pieter decided to take a sabbatical to join Elon Musk in his Open AI initiative in 2016. It was founded by Musk and Sam Altman (the president of 'Y combinator' – one of the leading incubators in Silicon Valley) to address the challenges of bringing AI into the open. It seeks to promote the benefits of AI by doing research and making its patents and research open to the public.

The Open AI initiative has been given a \$1 billion endowment to get it going. A cool billion. I can fully understand why Pieter Abbeel spends four days a week in the Open AI offices in San Francisco, and only one day a week with his students trying to patch up good old laundry-folding BRETT.

You need money to power AI, because AI needs power.



The analogy I really like is the story of the Wright Brothers. The Wright Brothers were the first humans to achieve controlled, human-operated and sustained flight with a heavier-than-air machine. Their Wright Flyer made its first successful trip on the 17th of December 1903, at Kitty Hawk, North Carolina.

For most people, that is the moment when we humans learned how to fly. Not really. The theory behind the heavier-than-air flight had been developed as far back as 1738, by Daniel Bernoulli. That is 165 years before the Wright Brothers made it happen on the sunny beaches of North Carolina. Yes, 165 years.

If you're an engineer, you're bound to have studied Bernoulli's principle. It explains how the flow of air over the shape of a wing can create enough lift to make an airplane take off. But when Bernoulli wrote it down in his book *Hydrodynamica* in 1738 it was just a theory: pure and beautiful mathematics and physics.

A long time would pass before anyone could prove that it was correct. Bernoulli did not conceive the airplane. But he did describe the fluid dynamics principles that would be needed to build carburetors and airplane wings. It took the brilliance and guts of the Wright Brothers to apply it.

The Wright brothers had been trying to make airplanes work for years. The problem was that they needed an engine – a very powerful one – to drive the propeller that could 'pull' the plane forward fast enough to create enough airflow over the wings and generate enough lifting power to get it off the ground. Gasoline engines were available thanks to the burgeoning automobile industry. But powerful engines were still too heavy at the time and would make the contraption impossible to get off the ground.



The real breakthrough that was key to the Wright brothers' success on the 17th of December 1903 was of their own doing. They had built an extremely powerful, efficient and lightweight custom-built engine out of aluminum by themselves. The smashing 12 horsepower output that it produced was just enough to make the Wright Flyer take off. When the Wright brothers achieved their success, they had no idea what they'd unchained. Today we take a plane like we take a bus. I'm writing this chapter on an Airbus A-380. It's amazing to observe how something of this sheer size and mass takes off. Neither Bernoulli nor the Wright brothers, could have ever anticipated how flight would evolve like this.

THE CLOUD AS THE MISSING PIECE OF THE PUZZLE

In the world of Artificial Intelligence, I believe we are exactly at that 'Kitty Hawk, North Carolina, 1903' moment. For the last 70 years, we have had the mathematics of AI worked out for us. The works of Alan Turing, Marvin Minsky, Norbert Wiener and John McCarthy laid the mathematical foundations of machine learning and artificial intelligence. But for 70 years, they lacked the equivalent of the 12 HP aluminum engine to make it work.

Until now that is... With the advent of cloud computing – where we can stitch together the power of thousands, hundreds of thousands, of machines – we are finally arriving at producing the tools that are powerful enough to make machines that think. A computer was able to beat Lee Sedol in 2016 because – for the first time – we had enough computing power to make it happen.

We knew HOW we could do it for a long time, just like Bernoulli had given us the theory of flying 165 years before we could fly. And now we are at this exciting point where AI is gathering the power to really take off.



ONE LAST OBSTACLE

On one of my visits at Pieter Abbeel's lab in Berkeley and after a long discussion, he took a marker and started to scribble on a flipchart. "We don't have to actually worry about anything yet", he says. We had just been discussing superintelligence, the work by the Future of Humanity institute, and the dangers of AI "Us humans are still way too cheap to be replaced".

As a true scientist, he starts making comparisons to the world of computing today, and the complexity of brainpower in the animal world. He sketches a column where he writes down the number of neurons, and number of synapses in animals, and tries to calculate how much computing capacity they have.

And then he starts calculating how much computing power you would need to buy on one of the largest cloud providers on the planet, to get to the equivalent of a human brain. The scary thing is that you could. After a rough calculation, he comes up with a number.

According to Abbeel, to rent enough capacity in the cloud to have the equivalent of a human brain, you would need to shell out about \$5,000 per hour. "So, we could do that, but today a regular human brain is still cheaper", is his conclusion.

So, we definitely DO have the computing power to create computers that are as smart as a human.

What we need is there, right around the corner, waiting for its prices to drop. And they WILL drop. They always do. So we're not there YET. But with the evolution of the second half of the chess-board, it is merely a matter of time.



UTOPIA OR DYSTOPIA

Machine Learning and Artificial Intelligence could definitely become the 'new electricity', but we still cannot know what the outcome on society will be. It could be utopian, an autonomous world to benefit all of humanity. Being a bit of a techno-optimist myself, I'm inclined to believe in its positive value. But we cannot just stand by and wait to see what happens. Because – even if we disregard the dystopian Superintelligence type of scenarios – many, many jobs will disappear if AI evolves just a little bit further along the spectrum towards Artificial General Intelligence.

There are still many out there living in denial – when it comes to employment. One of the arguments people love to use is how AI can only be employed for dull and routine jobs. If a task has anything to do with innovation, creativity and emotion, only humans can and ever will be able to perform it. It sounds really reassuring, doesn't it? Like a little blanket against this cold world of AI. Well I hate to burst your bubble, but the major tech giants are tackling this last beacon of humanity domain as well.

It really might not be a bad idea to find out how we can control AI before we bring it into our midst. And, at the very least, train and prepare ourselves, and our children, for its arrival. Because more likely than not, the actual accomplishment of Artificial General Intelligence will happen very suddenly and evolve faster than we will be able to comprehend.

So it's a good idea to start exploring the possibilities of AI NOW. Not in Q4. Not after the next board meeting. Not even in two weeks. Now. Because we are on the brink of a major shift that will completely transform how we work, learn, live and even think. And WHEN it happens, it will be huge. And it will move so fast that laggards will no longer be able to catch up. So don't miss that train (or self-driving car, if you prefer).



WHAT ARE YOU
DOING TO INNOVATE
AND THRIVE IN YOUR
DAY AFTER
TOMORROW?



LaLiga

BUILDING THE NEXT
GENERATION SPORTS
EXPERIENCE

[READ STORY HERE](#)



HOW TO FEED THE
WORLD WITHOUT
WRECKING THE PLANET

[READ STORY HERE](#)



AUTONOMOUS
VEHICLES

[READ STORY HERE](#)



ENERGISING
THE PLANET

[READ STORY HERE](#)



ADENTS

DATA
SERIALIZATION

[READ STORY HERE](#)



Microsoft

GLOBAL IMPACT
ACROSS INDUSTRIES

[READ STORY HERE](#)

microsoft.com/ai

As machine learning becomes a fundamental ingredient to helping organizations transform, it becomes incumbent on drivers of transformation to responsibly create and own AI and infuse it into digital systems. The focus on AI development means that organizations will have to mature in specific ways to successfully develop, train and own AI components that integrate into larger digital experiences.

Digital transformation means “doing the same things better”, or “doing new things that no one has ever seen before”. These represent fundamental opportunities for organizations to realize new revenue streams, disrupt industries and create new opportunities for themselves and their customers. AI is critical to making smart, fast and helpful digital experiences.

However, organizations must be aware of how to select the right technologies that they can support in a ownership mode. If not, organizations may stumble and fall as they attempt to infuse AI into their organization that they are not ready to own or operate.

Microsoft believe the real power of AI rests in its ability to holistically transform the enterprise and redefine business in ways that move beyond our imagination.

To make this happen, organizations need a long-term strategy and a technology partner that goes beyond providing single-shingle solutions and acts as a strategic thought partner.

As AI implementation continues to expand, this partnership must be capable of meeting the needs and concerns of the enterprise, such as security and scalability. It must also ensure that all employees, regardless of technical expertise, are able to benefit.

The following section provides an overview on quick approaches organizations can use to track, adapt and select the right technologies.

Microsoft is focused on developing AI in a way that it is human-centric and augments human abilities, especially humankind's innate ingenuity. Developing AI technology that leverages the unique strengths of computers—such as probabilistic reasoning and pattern recognition—with the creativity, ingenuity, and capacity for meaning-making of humans. Innovating AI to enable better decision-making across organizations, amplify the tools and processes employees already use, and tear down knowledge siloes—so people can do more, together. Empowering and accelerating the impact that people around the world can have in solving some of the society's biggest challenges:



AI for Good

AI can be a powerful tool for increasing access to information, education, employment, government services, and social, and economic opportunities. There are no limits to what people can achieve when technology reflects the diversity of everyone who uses it. Enterprises should play an active role to ensure that these new technologies are applied responsibly and inclusively.



AI for Accessibility

Promotes inclusion through intelligent technology. Focused on empowering organizations and developers to harness AI to amplify human capabilities for people with disabilities. The program focuses on driving breakthroughs that make the workplace more inclusive, providing equal access to information through innovations in vision, speech, and machine reading, and helping people with disabilities gain more independence to perform daily tasks.



AI for Humanitarian Action

Harness the power of AI to support disaster response and recovery, help ensure the safety and wellbeing of children around the world, protect refugees and displaced people, and promote respect for human rights.



AI for Earth

Empowers people and organizations to create breakthrough innovations in the way we monitor, model, and ultimately manage Earth's natural systems. Focused on four key areas vital to creating a sustainable future – agriculture, water, biodiversity, and climate change.



Ethics

As we look to a future powered by a partnership between computers and humans, we address ethical challenges head-on. Designing trustworthy AI requires creating solutions that reflect ethical principles deeply rooted in important and timeless values—fairness, reliability and safety, privacy and security, inclusivity, transparency, and accountability—to guide the cross-disciplinary development and use of artificial intelligence.



Get Started

Every enterprise is unique and have their own path to transforming their organization. To help organization to take their first step towards their AI transformation, Microsoft has created an AI ready assessment tool that helps evaluate your organizational readiness for adopting AI-based systems and provides customized recommendations around AI implementations for your business.

Microsoft build and test their AI capabilities in their own internal processes, so they can get a better grasp on how to make the most useful tools for their customers across industries. These are just a few of the internal Microsoft projects that are leveraging AI to create better outcomes:

Customer support virtual agent

Created a conversational virtual agent to support customer queries on a variety of Microsoft products, including Windows, Office, Xbox, and more. The agent resulted in a 2x increase in self-help success and a massive decrease in agent-to-agent transfers.

Revamped sales processes

Simplified complex sales processes, increased the accuracy of sales data, and enabled an individualized customer experience by creating a new sales process built on Dynamics 365 and Azure Cloud Services.

Smart buildings powered by data analytics

Leveraged data analytics, IoT, and Azure Machine Learning for predictive maintenance climate control, and HVAC optimization—keeping buildings comfortable while minimizing our environmental footprint.

It really might not be a bad idea to find out how we can control AI before we bring it into our midst. And, at the very least, train and prepare ourselves, and our children, for its arrival. Because more likely than not, the actual accomplishment of Artificial General Intelligence will happen very suddenly and evolve faster than we will be able to comprehend.

So it's a good idea to start exploring the possibilities of AI NOW. Not in Q4. Not after the next board meeting. Not even in two weeks. Now. Because we are on the brink of a major shift that will completely transform how we work, learn, live and even think. And WHEN it happens, it will be huge. And it will move so fast that laggards will no longer be able to catch up. So don't miss that train (or self-driving car, if you prefer).



AI MATURITY IS CRITICAL TO SUCCESS



FOUNDATIONAL

Questioning what AI is and how to apply it

Wrong expectations or disappointment
Low digitization
Basic analytical capabilities



APPROACHING

Hopeful on AI and its promise

Digitization underway
Looking to increase or optimize processes
Cautious about disruption



ASPIRATIONAL

Experimented and applied AI

High digitization
Desires new business models
Achieved a data culture



MATURE

Emerging data science and operational capability

Understands model lifecycle and management
Building a foundational data architecture

EVOLVING ARTIFICIAL INTELLIGENCE



More general intelligence



Better human and AI collaboration



Address opportunities and challenges ethically



Transformative digital experiences



9 RULES TO INNOVATE AND THRIVE IN THE DAY AFTER TOMMORROW

1

BREAK THE RULES

Rule-followers won't survive the future. Disruptors will. Make your own rules, then bend, break, and renew them, and never stop.

2

SPEND 10% ON BEING RADICAL

More than that might weaken the 'today' and 'tomorrow' business that funds your innovation. Less will mean you'll get left behind.

3

TRAVEL BEYOND THE LIMITS

Fight the status quo, push the boundaries and break down barriers. The impossible is just a possibility waiting to be born.

4

GROW A PAIR

Innovation is messy and chaotic. Avoidance of risk isn't safe, it keeps you from evolving fast enough.

5

CULTURE BEFORE STRUCTURE

Hire anyone, so long as they are passionate and committed to your customers. Your people are your culture, so choose them carefully.

6

THOU SHALT NOT MISTRUST

The age of disruption is about trust. Trust in empowered employees and the opinions of your customers. Trust that shared ideas will grow. Suspicious minds will miss the Day After Tomorrow.

7

FOLLOW THE VEXERS

Listen to the troublesome customers that demand the impossible. They will make you see things you didn't see before and push your company beyond its limits.

8

MOVE FAST AND BREAK THINGS

Move before your company peaks, before you think you should be moving, then keep moving. That's the secret of eternal youth.

9

UN-BECOME YOURSELF


You have to reinvent yourself constantly. There are so many tools, ideas, platforms, and people waiting to be connected. If you are willing to learn, un-learn, and take radical new directions, you will find your Day After Tomorrow.

The engine that drives innovation is simple:




DREAM IT. DO IT.

Watch the movie



Read the book



Author: Peter Hinssen

Contributors: Devon Baker, Eric Charran, Brad Clayton, John Emrick, Steve Franklin, Samantha Greensmith, Beatrix Guimaraes, Donal Higgins, Johan Klut, Laura Longcore, Brian Marple, John Nisi, Xavier Pereira, Jaime Perena, Kim Schulze, Roy Sharples, Blessing Sibanyoni, Steve Sweetman, Tuan Jean Tee, Ian Webster

In partnership with nexxworks

Nexxworks inspires companies to kickstart their innovation. By helping them regain the flexibility to radically innovate, that they lost as they grew. Provoking action through innovation tours, bootcamps, tailored programs and keynotes with the latest insights in business and technology, that encourage companies to combine a long-term vision with short-term initiatives. Brought to life through a network of international and local like-minded entrepreneurs, innovators, experts and speakers. nexxworks.com



In partnership with **nexxworks**

microsoft.com/industry

