




SHAPING  
OUR FUTURE

WILL *AI* STEAL YOUR

JOB?





**BACK IN THE 1950s**, Minsky and McCarthy, the fathers of artificial intelligence (AI) described it as any task performed by a program or a machine that, if carried out by a human, would be considered to have intelligence applied to accomplish the task. That is obviously a fairly broad definition, which is why you will sometimes see arguments over whether something is truly AI or not.

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, manipulation, and, to a lesser extent, social intelligence and creativity.

■ **WHAT ARE THE USES FOR AI?**

AI is ubiquitous today, used to recommend what you should buy next online, to understand what you say to virtual assistants such as Amazon’s Alexa and Apple’s Siri, to recognize who and what is in a photo, to spot spam, or to detect credit card fraud.

■ **WHAT ARE THE DIFFERENT TYPES OF AI?**

At a very high level, artificial intelligence can be split into two broad types: narrow AI and general AI.

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 is 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, coordinating 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 Internet of Things (IoT) devices, the list goes on and on.

■ **WHAT CAN GENERAL AI DO?**

Artificial general intelligence (AGI) is very different from narrow AI, 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. This is the sort of AI more commonly seen in movies, the likes of HAL in *2001* or Skynet in *The Terminator*, but which

does not exist today and AI experts are fiercely divided over how soon it will become a reality.

That said, some AI experts believe such projections are wildly optimistic given our limited understanding of the human brain, and believe that AGI is still centuries away.

■ **WHAT IS MACHINE LEARNING?**

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.

■ **WHAT ARE NEURAL NETWORKS?**

Key to the process of machine learning is neural networks. 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 tasks like speech recognition and computer vision.

Another area of AI research is evolutionary computation, which borrows from Darwin’s famous theory of natural selection,




and sees genetic algorithms undergo random mutations and combinations between generations in an attempt to evolve the optimal solution to a given problem.

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 AI and algorithms are not good at is dealing with ambiguity and gray areas. They do not understand context or nuance, and so they are not good at making judgment calls. That is where humans are much faster and more accurate.**



#### ■ WILL AI KILL US ALL?

Again, it depends who you ask. As AI-powered systems have grown more capable, so the warnings of the downsides have become more dire.

Tesla and SpaceX CEO Elon Musk has claimed that AI is a “fundamental risk to the existence of human civilization.” As part of his push for stronger regulatory oversight and more responsible research into mitigating the downsides of AI he set up OpenAI, a non-profit artificial intelligence research company that aims to promote and develop friendly AI that will benefit society as a whole. Similarly, the esteemed physicist Stephen Hawking has warned that, once a sufficiently advanced AI is created, it will rapidly advance to the point at which it vastly outstrips human capabilities, a phenomenon known as the singularity, and could pose an existential threat to the human race. Yet, the notion that humanity is on the verge of an AI explosion that will dwarf our intellect seems ludicrous to some AI researchers.

Chris Bishop, Microsoft’s director of research in Cambridge, England, stresses how different the narrow intelligence of AI today is from the general intelligence of humans, saying that when people worry about “Terminator, the rise of the machines, and so on, it is utter nonsense. At best, such discussions are decades away.”

#### ■ WILL AN AI STEAL YOUR JOB?

The possibility of artificially intelligent systems replacing much of modern manual labour is perhaps a more credible near-future possibility.

While AI will not replace all jobs, what seems to be certain is that AI will change the nature of work, with the only question being how rapidly and how profoundly automation will alter the workplace.

There is barely a field of human endeavour that AI does not have the potential to impact. As AI expert Andrew Ng puts it: “Many people are doing routine, repetitive jobs. Unfortunately, technology is especially good at automating routine, repetitive work,” saying he sees a “significant risk of technological unemployment over the next few decades.”

The evidence of which jobs will be supplanted is starting to emerge. Amazon has just launched Amazon Go, a cashier-free supermarket in Seattle where customers take items from the shelves and walk out. What this means for the more than three million people in the US who work as cashiers remains to be seen. Amazon again is leading the way in using robots to improve efficiency inside its warehouses. These robots carry shelves of products to human pickers who select items to be sent out. Amazon has more than 100,000 bots in its fulfilment centers, with plans to add many more. But Amazon also stresses that, as the number of bots have grown, so has the number of human





workers in these warehouses. However, Amazon and small robotics firms are working to automate the remaining manual jobs in the warehouse, so it is not a given that manual and robotic labor will continue to grow hand-in-hand.

Fully autonomous self-driving vehicles are not a reality yet, but, by some predictions, the self-driving trucking industry alone is poised to take over 1.7 million jobs in the next decade, even without considering the impact on couriers and taxi drivers.

Yet, some of the easiest jobs to automate will not even require robotics. At present, there are millions of people working in administration, entering and copying data between systems, chasing and booking appointments for companies. As software gets better at automatically updating systems and flagging the information that is important, so the need for administrators will fall.

As with every technological shift, new jobs will be created to replace those lost. However, what is uncertain is whether these new roles will be created rapidly enough to offer employment to those displaced, and whether the newly unemployed will have the necessary skills or temperament to fill these emerging roles.

Not everyone is a pessimist. For some, AI is a technology that will augment, rather than replace, workers. Not only that, but they argue there will be a commercial imperative to not replace people outright, as an AI-assisted worker – think a human concierge with an Augmented Reality headset that tells them exactly what a client wants before they ask for it – will be more productive or effective than an AI working on its own.

Among AI experts, there is a broad range of opinion about how quickly artificially intelligent systems will surpass human capabilities.

Oxford University's Future of Humanity Institute asked several hundred machine-learning experts to predict AI capabilities over the coming decades.

Notable dates included AI writing essays that could pass for being written by a human by 2026, truck drivers being made redundant by 2027, AI surpassing human capabilities in retail by 2031, writing a best-seller by 2049, and doing a surgeon's work by 2053.

They estimated there was a relatively high chance that AI beats humans at all tasks within 45 years and automates all human jobs within 120 years.


The 21st century's longest shadow hanging over the technology industry, workers, and humanity in general, is the effect that artificial intelligence, robots, and automation are going to have on jobs. The popular narrative is generally filled with lots of doom-and-gloom about AI wiping out whole industries and millions of people going unemployed, but let's take a little reality check and look at where humans do a better job than AI, and vice versa, to give you some context for how the workforce of the future is going to change.

#### ■ WHERE AI WINS

It is no secret that algorithms and robots can outperform humans at repetitive tasks and solving problems that involve processing large, well-organized data sets. Humans get bored and distracted doing the same things over and over again. Robots and computers are singularly focused.

Humans are also slow and prone to error when processing and evaluating patterns in big data. Algorithms have long since lapped us. That is why IBM's Deep Blue beat Garry Kasparov in chess in 1997, and it is why Google DeepMind beat Lee Sedol in the Game of Go in 2016. They used pattern analysis on datasets with very clear rules and parameters.

As more and more of the economy gets digitized and automated, it is going to create more opportunities for humans with critical thinking skills.



#### ■ WHERE HUMANS WIN

However, what AI and algorithms are not good at is dealing with ambiguity and gray areas. They do not understand context or nuance, and so they are not good at making judgment calls. That is where humans are much faster and more accurate.

One of the dirty little secrets about AI and big data is that the tech giants working on it are hiring lots of human beings to sort, organize, cleanse, and prepare the data to be analyzed by the algorithms – because humans are better at it.

As more and more of the economy gets digitized and automated, it is going to create more opportunities for humans with critical thinking skills. And those are not just white-collar jobs either. Robots and AI will not be fixing your plumbing or building skyscrapers. But, they will provide those workers with data that can help them work faster, more efficiently, and more safely. 📊