

How Automation is Going to Redefine What it Means to Work

The time for machines to take over most of humanity's work is rapidly approaching. The world is woefully unprepared to deal with the implications that automation will have over the coming decades. Universal basic income is just beginning to be discussed, and automation has the potential to displace much of the world's workforce. Many decisions have to be made, and quickly, if we hope to keep pace with innovation.

On December 2nd, 1942, a team of scientists led by Enrico Fermi came back from lunch and watched as humanity created the first self-sustaining nuclear reaction inside a pile of bricks and wood underneath a football field at the University of Chicago. Known to history as [Chicago Pile-1](#), it was celebrated in silence with a single bottle of Chianti, for those who were there understood exactly what it meant for humankind, without any need for words.

Now, something new has occurred that, again, quietly changed the world forever. Like a whispered word in a foreign language, it was quiet in that you may have heard it, but its full meaning may not have been comprehended. However, it's vital we understand this new language, and what it's increasingly telling us, for the ramifications are set to alter everything we take for granted about the way our globalized economy functions, and the ways in which we as humans exist within it.

The language is a new class of machine learning known as [deep learning](#), and the "whispered word" was a computer's use of it to seemingly out of nowhere [defeat three-time European Go champion Fan Hui](#), not once but five times in a row without defeat. Many who read this news, considered that as impressive, but in no way comparable to a match against Lee Se-dol instead, who many consider to be one of the world's best living Go players, if not *the* best. Imagining such a grand duel of man versus machine, [China's top Go player predicted that Lee would not lose a single game, and Lee himself confidently expected to possibly lose one at the most](#).

What actually ended up happening when they faced off? Lee went on to lose [all but one](#) of their match's five games. An AI named AlphaGo is now a better Go player than any human and has been [granted the "divine" rank of 9 dan](#). In other words, its level of play borders on godlike. Go has officially fallen to machines, just as Jeopardy did before it to Watson, and chess before that to Deep Blue.

"AlphaGo's historic victory is a clear signal that we've gone from linear to parabolic."

So, what is Go? Very simply, think of Go as Super Ultra Mega Chess. This may still sound like a small accomplishment, another feather in the cap of machines as they

continue to prove themselves superior in the fun games we play, but it is no small accomplishment, and what's happening is no game.

AlphaGo's historic victory is a clear signal that we've gone [from linear to parabolic](#). Advances in technology are now so visibly exponential in nature that we can expect to see a lot more milestones being crossed long before we would otherwise expect. These exponential advances, most notably in forms of artificial intelligence limited to specific tasks, we are entirely unprepared for as long as we continue to insist upon employment as our primary source of income.

This may all sound like exaggeration, so let's take a few decade steps back, and look at what computer technology has been actively doing to human employment so far:

Let the above chart sink in. Do not be fooled into thinking this conversation about the automation of labor is set in the future. It's already here. Computer technology is already eating jobs and has been since 1990.

All work can be divided into four types: routine and nonroutine, cognitive and manual. Routine work is the same stuff day in and day out, while nonroutine work varies. Within these two varieties, is the work that requires mostly our brains (cognitive) and the work that requires mostly our bodies (manual). Where once all four types saw growth, the stuff that is routine stagnated back in 1990. This happened because routine labor is easiest for technology to shoulder. Rules can be written for work that doesn't change, and that work can be better handled by machines.

Distressingly, it's exactly routine work that once formed the basis of the American middle class. It's routine manual work that Henry Ford transformed by paying people middle class wages to perform, and it's routine cognitive work that once filled US office spaces. [Such jobs are now increasingly unavailable](#), leaving only two kinds of jobs with rosy outlooks: jobs that require so little thought, we pay people little to do them, and jobs that require so much thought, we pay people well to do them.

If we can now imagine our economy as a plane with four engines, where it can still fly on only two of them as long as they both keep roaring, we can avoid concerning ourselves with crashing. But what happens when our two remaining engines also fail? That's what the advancing fields of robotics and AI represent to those final two engines, because for the first time, we are successfully teaching machines to *learn*.

I'm a writer at heart, but my educational background happens to be in psychology and physics. I'm fascinated by both of them so my undergraduate focus ended up being in the physics of the human brain, otherwise known as [cognitive neuroscience](#). I think once you start to look into how the human brain works, how our mass of interconnected neurons somehow results in what we describe as the mind, everything changes. At least it did for me.

As a quick primer in the way our brains function, they're a giant network of interconnected cells. Some of these connections are short, and some are long. Some cells are only connected to one other, and some are connected to many. Electrical signals then pass through these connections, at various rates, and subsequent neural firings happen in turn. It's all kind of like falling dominoes, but far faster, larger, and more complex. The result amazingly is us, and what we've been learning about how we work, we've now begun applying to the way machines work.

One of these applications is the creation of [deep neural networks](#) – kind of like pared-down virtual brains. They provide an avenue to machine learning that's made incredible leaps that were previously thought to be much further down the road, if even possible at all. How? It's not just the obvious growing capability of our computers and our expanding knowledge in the neurosciences, but the vastly growing expanse of our collective data, aka big data.

Big data isn't just some buzzword. It's information, and when it comes to information, we're creating more and more of it every day. In fact we're creating so much that a 2013 report by SINTEF estimated that [90% of all information in the world had been created in the prior two years](#). This incredible rate of data creation is even [doubling every 1.5 years](#) thanks to the Internet, where in 2015 [every minute we were liking 4.2 million things on Facebook, uploading 300 hours of video to YouTube, and sending 350,000 tweets](#). Everything we do is generating data like never before, and lots of data is exactly what machines need in order to learn *to learn*. Why?

Imagine programming a computer to recognize a chair. You'd need to enter a ton of instructions, and the result would still be a program detecting chairs that aren't, and *not* detecting chairs that are. So how did we learn to detect chairs? Our parents pointed at a chair and said, "chair." Then we thought we had that whole chair thing all figured out, so we pointed at a table and said "chair", which is when our parents told us that was "table." This is called reinforcement learning. The label "chair" gets connected to every chair we see, such that certain neural pathways are weighted and others aren't. For "chair" to fire in our brains, what we perceive has to be close enough to our previous chair encounters. Essentially, our lives are big data filtered through our brains.

The power of deep learning is that it's a way of using massive amounts of data to get machines to operate more like we do without giving them explicit instructions. Instead of describing "chairness" to a computer, we instead just plug it into the Internet and feed it millions of pictures of chairs. It can then have a general idea of "chairness." Next we test it with even more images. Where it's wrong, we correct it, which further improves its "chairness" detection. Repetition of this process results in a computer that knows what a chair is when it sees it, [for the most part as well as we can](#). The important difference though is that [unlike us, it can then sort through millions of](#)

[images within a matter of seconds](#).

This combination of deep learning and big data has resulted in astounding accomplishments just in the past year. Aside from the incredible accomplishment of AlphaGo, [Google's DeepMind AI learned how to read and comprehend what it read](#) through hundreds of thousands of annotated news articles. [DeepMind also taught itself to play dozens of Atari 2600 video games better than humans](#), just by looking at the screen and its score, and playing games repeatedly. An AI named Giraffe taught itself how to play chess in a similar manner using a dataset of 175 million chess positions, [attaining International Master level status in just 72 hours by repeatedly playing itself](#). In 2015, [an AI even passed a visual Turing test by learning to learn](#) in a way that enabled it to be shown an unknown character in a fictional alphabet, then instantly reproduce that letter in a way that was entirely indistinguishable from a human given the same task. These are all *major* milestones in AI.

However, despite all these milestones, when asked to estimate when a computer would defeat a prominent Go player, the answer even just months prior to [the announcement by Google of AlphaGo's](#) victory, was by experts essentially, "[Maybe in another ten years](#)." A decade was considered a fair guess because Go is a game so complex I'll just let Ken Jennings of Jeopardy fame, [another former champion human defeated by AI](#), describe it:

Go is famously a more complex game than chess, with its larger board, longer games, and many more pieces. Google's DeepMind artificial intelligence team likes to say that there are more possible Go boards than atoms in the known universe, but that vastly *understates* the computational problem. There are about 10^{170} board positions in Go, and *only* 10^{80} atoms in the universe. That means that if there were as many parallel universes as there are atoms in our universe (!), then the *total* number of atoms in *all* those universes combined would be close to the possibilities on a single Go board.

Such confounding complexity makes impossible any brute-force approach to scan every possible move to determine the next best move. But deep neural networks get around that barrier in the same way our own minds do, by learning to estimate what *feels* like the best move. We do this through observation and practice, and so did AlphaGo, by analyzing millions of professional games and [playing itself millions of times](#). So the answer to when the game of Go would fall to machines wasn't even close to ten years. The correct answer ended up being, "Any time now."

Any time now. That's the new go-to response in the 21st century for any question involving something new machines can do better than humans, and we need to try to wrap our heads around it.

We need to recognize what it means for exponential technological change to be entering the labor market space for nonroutine jobs for the first time ever. Machines that can learn mean *nothing* humans do as a job is uniquely safe anymore. From [hamburgers](#) to [healthcare](#), machines can be created to successfully perform such tasks with no need or less need for humans, and at lower costs than humans.

[Amelia](#) is just one AI out there currently being beta-tested in companies right *now*. Created by IPsoft over the past 16 years, she's learned how to perform the work of call center employees. She can learn in seconds what takes us months, and she can do it in 20 languages. Because she's able to learn, she's able to do more over time. In one company putting her through the paces, she successfully handled one of every ten calls in the first week, and by the end of the second month, she could resolve six of ten calls. Because of this, it's been estimated that she can put 250 million people out of a job, *worldwide*.

[Viv](#) is an AI coming soon from the creators of Siri who'll be our own personal assistant. She'll perform tasks online for us, and even function as a Facebook News Feed on steroids by suggesting we consume the media she'll know we'll like best. In doing all of this for us, we'll see far fewer ads, and that means the entire advertising industry — that industry the entire Internet is built upon — stands to be hugely disrupted.

A world with Amelia and Viv — and the countless other AI counterparts coming online soon — in combination with robots like [Boston Dynamics' next generation Atlas](#) portends, is a world where machines can do *all four types of jobs* and that means serious societal reconsiderations. If a machine can do a job instead of a human, [should any human be forced at the threat of destitution to perform that job?](#) Should income itself remain coupled to employment, such that having a job is the only way to obtain income, when [jobs for many are entirely unobtainable](#)? If machines are performing an increasing percentage of our jobs for us, and not getting paid to do them, [where does that money go instead?](#) And [what does it no longer buy?](#) [Is it even possible that many of the jobs we're creating don't need to exist at all](#), and only do because of the incomes they provide? These are questions we need to start asking, and fast.

Fortunately, people *are* beginning [to ask these questions](#), and there's an answer that's building up momentum. The idea is to put machines to work for us, but empower ourselves to seek out the forms of remaining work we as humans find most valuable, by simply providing everyone a monthly paycheck independent of work. This paycheck would be granted to all citizens unconditionally, and its name is [universal basic income](#). By adopting UBI, aside from [immunizing](#) against the negative effects of automation, we'd also be decreasing [the risks inherent in entrepreneurship](#), and [the sizes of bureaucracies](#) necessary to boost incomes. It's for these reasons, it has [cross-partisan support](#), and is even now in the beginning stages of possible implementation in countries like [Switzerland](#), [Finland](#), the [Netherlands](#), and [Canada](#).

The future is a place of accelerating changes. It seems unwise to continue looking at the future as if it were the past, where just because new jobs have historically appeared, they always will. [The WEF started 2016 off by estimating the creation by 2020 of 2 million new jobs alongside the elimination of 7 million](#). That's a net loss, not a net gain of 5 million jobs. In a frequently cited paper, [an Oxford study estimated the automation of about half of all existing jobs by 2033](#). Meanwhile self-driving vehicles, again thanks to machine learning, have the capability of drastically impacting all economies — [especially the US economy as I wrote last year about automating truck driving](#) — by eliminating millions of jobs within a short span of time.

And now even the White House, [in a stunning report to Congress](#), has put the probability at 83 percent that a worker making less than \$20 an hour in 2010 will eventually lose their job to a machine. Even workers making as much as \$40 an hour face odds of 31 percent. To ignore odds like these is tantamount to our now laughable "[duck and cover](#)" strategies for avoiding nuclear blasts during the Cold War.

All of this is why it's those most knowledgeable in the AI field who are now actively sounding the alarm for basic income. During a panel discussion at the end of 2015 at Singularity University, prominent data scientist [Jeremy Howard](#) asked "Do you want half of people to starve because they literally can't add economic value, or not?" before going on to suggest, "If the answer is *not*, then the smartest way to distribute the wealth is by implementing a universal basic income."

AI pioneer [Chris Eliasmith](#), director of the Centre for Theoretical Neuroscience, warned about the immediate impacts of AI on society in an interview with Futurism, "AI is already having a big impact on our economies... My suspicion is that more countries will have to follow Finland's lead in exploring basic income guarantees for people."

Moshe Vardi [expressed the same sentiment](#) after speaking at the 2016 annual meeting of the American Association for the Advancement of Science about the emergence of intelligent machines, "we need to rethink the very basic structure of our economic system... we may have to consider instituting a basic income guarantee."

Even Baidu's chief scientist and founder of Google's "Google Brain" deep learning project, [Andrew Ng](#), during an onstage interview at this year's Deep Learning Summit, expressed the shared notion that basic income must be "seriously considered" by governments, citing "a high chance that AI will create massive labor displacement."

When those building the tools begin warning about the implications of their use, shouldn't those wishing to use those tools listen with the utmost attention, especially when it's the very livelihoods of millions of people at stake? If not then, what about when [Nobel prize winning economists](#) begin agreeing with them in increasing numbers?

No nation is yet ready for the changes ahead. High labor force non-participation leads to social instability, and a lack of consumers within consumer economies leads to economic instability. So let's ask ourselves, what's the purpose of the technologies we're creating? What's the purpose of a car that can drive for us, or artificial intelligence that can shoulder 60% of our workload? Is it to allow us to work more hours for even less pay? Or is it to enable us to choose *how* we work, and to decline any pay/hours we deem insufficient because we're already earning the incomes that machines aren't?

What's the big lesson to learn, in a century when machines can learn?

I offer it's that jobs are for machines, and life is for people.