

Will AI Save Or Destroy Us?

Consider this...

A Grandma walks hand in hand through the park with her four-year-old granddaughter, as dappled golden light plays along the edges of their shadows. They laugh as they watch a squirrel scamper up a tree when suddenly grandma falls to the ground. Frightened, the little girl doesn't know what to do as strangers rush towards her.

But grandma wears an Apple watch with fall detection. It's heartbeat monitor noticed when grandma's heart started to race and detected the fall. So, the watch's artificial intelligence calls for help. A few minutes later, an ambulance arrives. AI saved grandma.

And consider this, too...

In a dark basement, a Super Villian types furiously into a computer. He's writing an algorithm that will collect all the public "Likes" millions of people click on Facebook. From that data, he'll create an AI system able to predict a person's gender, sexual orientation, political persuasion, and a host of other things. He intends to use that information to corrupt the voting system, destroy democracy, and demand one millllllion dollars...

What is Artificial Intelligence?

AI saturates modern life. Every time you buy something on Amazon, robots scurry to fill your order and [Amazon's sophisticated AI system, Flywheel](#), gets to work figuring out what you should buy next. When you talk to Siri, you're talking to AI. Netflix, Pandora, even Nest all use AI to figure out what you want before even you know it.

So, what is AI? Put simply, AI uses input data to predict future events. And the rules governing an AI system don't have to be that complex. In fact, it's possible to build a functioning [AI system in just a few lines of code](#). The power of AI comes from the quality and amount of input data delivered to the system.

```
1 from numpy import exp, array, random, dot
2 training_set_inputs = array([[0, 0, 1], [1, 1, 1], [1, 0, 1], [0, 1, 1]])
3 training_set_outputs = array([[0, 1, 1, 0]]).T
4 random.seed(1)
5 synaptic_weights = 2 * random.random((3, 1)) - 1
6 for iteration in xrange(10000):
7     output = 1 / (1 + exp(-(dot(training_set_inputs, synaptic_weights))))
8     synaptic_weights += dot(training_set_inputs.T, (training_set_outputs - output) * output * (1 - output))
9 print 1 / (1 + exp(-(dot(array([1, 0, 0]), synaptic_weights))))
```

AI program written in nine lines of code.

Like all human inventions, AI is amoral. It functions according to the parameters set by the software. It never thinks about how its results will be used. Just like for grandma in the story above, AI can make our lives safer, and healthier. In the hands the unscrupulous, however, AI can do a lot of damage.

The Good – AI Improves Our Health

Already, [AI has been put to good use helping doctors provide better care](#) for their patients. Many forms of information can serve as input for medical AI, from health history to genetic testing. Even the wearable devices coming on the market, [like the Apple watch](#), aid people who need extra monitoring.

Also, people who live in rural or remote areas can benefit from new AI technology. AI provides a link between rural patients and big city doctors, allowing people to get the quality care they couldn't get any other way. Used properly, AI offers doctors a powerful mechanism to improve the lives of millions, while also saving money through greater efficiency.

The Bad – AI Aids In Deception

The Super Villian in the story above is actually a company called Cambridge Analytics. There's no way those guys hang out in dark basements, or only want a million dollars. It's even possible they don't believe they're villainous or on a quest to destroy democracy. But as they say, "the proof is in the taste of the pudding." And this pudding tastes like... well, [democracy destruction](#).

Don't get me wrong! The technology they use is awesome. Less awesome is their willingness to lie and mislead. Cambridge Analytics used a technique created [by Dr. Michal Kosinski](#), whose team spent years



Dr. Michal Kosinski maintains [he never worked with Cambridge Analytics](#).

researching the correlations between the things people "Like" on Facebook and standard personality questionnaires used by Psychologists. The model Dr. Kosinski

created with his team is both groundbreaking and beautiful. And when the good doctor was approached to license his model to a company hoping to manipulate elections, he said no. However, one of his colleagues said yes... And that's how we got here.

Cambridge Analytics gathered up the data people leave online or stole what they couldn't get otherwise. Then this information served as an input to sophisticated AI systems targeting individuals with the sole purpose of swaying their opinion through targeted Facebook "dark posts," which only the person being targeted can see. It didn't matter whether the information in the posts was true, so long as it got the result of pushing a person in the direction Cambridge Analytics chose. So yeah, this technique can be as evil as it is effective.



A deepfake image of Steve Buscemi's face on Jennifer Lawrence.

It doesn't end with Cambridge Analytics. New AI technology creates convincing [fake people](#). And, deepfake face swapping is a thing now, too. Watching [Jennifer Lawrence talk with Steve Buscemi face](#)... kind of cracks me up... But Scarlett Johansson raising awareness for deepfake porn is no laughing matter.

The Awesome – AI Powers the Future

AI is powerful. No doubt about it. And it holds the promise of unlocking some of humanity's biggest technological challenges. Right now, for instance, quantum computers mostly only inhabit science fiction. However, Dr. Irfan Siddiqi, a professor of

physics at the Quantum Nanoscience Laboratory and the Department of Physics at the University of California Berkeley, recently proposed a new way to [make quantum computing a reality using AI.](#)

Computers nowadays use binary switches to perform calculations: it's all bits made from 1 and 0. Such switches are stable and easy. A quantum bit, or qubit, however, can occupy vastly more states than 1 or 0. A qubit state is also maddeningly hard to measure. The thing is, you can't measure at a qubit without changing it. When you put enough qubits together to be useful, the little buggers entangle with each other, making measurement even more difficult.

To get around this conundrum, Dr. Siddiqui and his colleagues proposed training a recurrent neural network (RNN) to predict the state of the qubit without actually having to measure it. RNNs learn by example, which makes them useful for translating languages. When trained to filter qubits, the RNN can translate a noisy signal into quantum information. If this process works, super powerful quantum computers are on the horizon.

The Future is Almost Here...

The point here is, with evolving AI and with the exponential increase of computing power, more grandmas can be saved, and more Super Villians will try an con us. Forwarned is forearmed, and knowing the dangers helps you avoid them. We can't stop the future. We can try to use emerging technologies in a way that realizes the benefits while mitigating the potential damage. We'll always have super villains. But we don't have to give them power.