The Third Wave – Rise of Lord A.I

Dear humans, You are fired. Signed, Lord A.I

By now, we have all heard about how Artificial Intelligence advances will replace our jobs. Future employment prospects are uncertain. Fear ripples through our eyes and ears... Did you know that robots write articles too? Perhaps even the one you're reading right now…

Building a machine that will eventually, inevitably, supersede its creators seems like a stupendous yet utterly stupid thing to achieve. We'll be slaves of technology! Oh, but not quite so, Sir Pessimistic; surely you're not in league with the tabloid editors - no doubt displaced from The Washington Post by one of Lord AI's subordinates - claiming 'The Matrix' dystopia is upon us? Such Frankensteinish nonsense. In fact, the future simply can't be brighter for natural scientists.

"The Third Wave is for those who think the human story, far from ending, has only just begun." – **Alvin Toffler**, *futurist, revolutionist*

The term 'artificial intelligence' was first coined by the father of this promising new field, John McCarthy, who described AI as the study of "intelligent agents" capable of "doing things that are on the surface of our minds, […] a simulation of human intelligence." Specifically, to a scientist, it is a system's ability to interpret and learn from large datasets, and hence employ those learnings to achieve certain tasks via elastic adaptation. Colloquially, it is when a machine emulates core cognitive functions possessed by humans: learning, reasoning and self-correction.

Let's start with the benefits of AI. Just fifty years ago, scientists used to have to do everything themselves. From data transformations to finding abnormalities, data analysis usually took a long, long, long time. Even after this very long process, checking the distributions quickly became an inefficient use of our time.
Not anymore.
Recall that AI is a 'simulation of human intelligence' with highly advanced 'reasoning' capabilities (i.e. using previous rules to reach approximate or even definite conclusions). Not only that, they can automatically determine the optimal settings for parameters - now we can simply leave the "magic" algorithm to fill in the seemingly incomprehensible gaps in our knowledge.

Surfing on the Third Wave

A prime example of this wonderful sorcery at work is the creation of an organic synthesis robotic AI system, recently deployed by a group of researchers at the University of Glasgow. The challenge was to find the best pathways to unearth new chemical reactions, so as to improve pharmaceuticals and materials' sustainability. However, this was severely limited by top-down approaches plagued with time-consuming and rather erratic processes, where expert knowledge was required in order to target a particular molecule. Not only that, but choosing from among hundreds of promising building blocks and thousands of chemical rules wasn't exactly easy.
With the help of machine learning algorithm, a cost-effective bottom-up approach was implemented, with no specific target in mind. The system quickly compared reactivities based on the reagents' spectra, and was able to effectively determine a reaction's success – thus it was able to prioritise those with greatest potential. Even more astonishingly, by *just* performing 10% of 969 possible reactions from

18 reagents, the autonomous system predicted with 86% accuracy the reactivity of the remaining reactions. Moreover, more experiments were automatically performed based on the reactivity data it had gathered - leading to the discovery of four new reactions. Since these robots performed up to 36 experiments per day – nearly 10 times more than an average human – productive efficiency was noticeably higher. This has huge implications in drug developments; after all, traditional 'brute-force' isn't feasible anymore. These experiments set the milestone for how we go about looking for new reactions, even challenge our current structure and bonding theory. But perhaps the most notable point to take away from this experiment is the fact that the system had no prior information (except the starting reagents' spectra of course). The database was built from scratch. AI had a new recipe in mind.
"This is desirable for reaction discovery since bias's removed," explains Cronin, team leader. "It was tremendously exciting to see the system autonomously navigate for the first time." Indeed, chemical automation systems "aren't normally creative" and would only do what is asked of them. Seed model strategies like those used above are increasingly becoming essential if one is to truly understand the big questions – those involving truly big datasets.

Physicists at Fermilab like working with big datasets. Naturally, their field lends itself to AI, neural networks in this case, since every experiment focuses on finding subtle patterns from an extremely noisy background. Using their electromagnetic calorimeter, they measure the energy of interacting subatomic particles. In a particle detector, photons create what seems like showers of particles. But so do electrons. So do hadrons. And so do many others, some of whom prefer not to wear name tags. Here we are reminded of the needle-in-the-haystack proverb; how to distinguish between the different sprays? Fortunately, machine-learning algorithms tell the difference by identifying correlations among the multiple variables that describe these showers.

Synthetic biologists have caught the wave too. DeepMind raised the bar high last year at the CASP competition, unveiling AlphaFold. It predicted the most accurate structure for 25 out of 43 proteins by anticipating the distances between amino acid pairs, all made possible through neural network trainings on thousands of known proteins - until 3D structures could be predicted from amino acids alone. This is extremely impressive, as the human body makes vast numbers of unique proteins, with estimates ranging from tens of thousands to billions!
What's more surprising is just how fast they can catch up; in the beginning, the program took a fortnight to envisage the first protein structures. Now it takes a couple of hours.

Careful of the Backwash

Our hard-working friends of Lord A.I gives rise to tremendous optimism surrounding the use of AI in our every-day lives. On the other hand, the raging growth in machine-learning tool developments, fed by unrestrained investments by the profit-driven private sectors, calls into question whether it is ethical to endow artificial beings with human-like intelligence. There are many philosophical issues about the nature of the mind without definite right or wrong views. It is time for responsible scientists to pause and ponder upon the ethical risks of incorporating AI into decision-making. And not just in the political context.
Danton Char, Professor of Anesthesiology at Stanford, certainly sees through the outward glamour of this vibrant technology: "From implementations in non-health care areas, there can be ethical problems with algorithmic learning when it's deployed at a large scale." And there are many precedents out there to corroborate this - take Volkswagen's algorithm, which allowed vehicles to pass emissions tests

by reducing their nitrogen oxide emissions during testing.
The worry doesn't end there. Many modern AIs learn to make decisions by being trained on massive datasets. Datasets from where though? If the data itself contains biases, these can be inherited and repeated by these 'amoral' systems.
Last year, AI used for image recognition categorised chefs as women, even when handed images of balding men.

The need for strong independent organisations, with well-informed researchers, that act as watchdogs and hold AI systems accountable is absolutely necessary. And it's incumbent on scientists, as AI innovators, to ensure that they are both safe and fair. Remaining ignorant about the construction of machine-learning systems or allowing them to be constructed as black boxes, as well as overreliance on their guidance, may contribute towards ethically problematic, potentially perilous outcomes.

Surviving Wave Wipeout

Despite the pitfalls, Char says pressure to turn to data for answers is especially intense in fields that are growing at an exponential rate, such as genetic testing and sequencing. Humans simply cannot comprehend the vast volume of data. But with AI-wielding machines, working 24/7, coupled with multi-processing capabilities, patience is a thing of the past.
To quote Father John again, since "our ability to observe our own mental processes is not very good", AI can be used to successfully extrapolate patterns, which may be difficult, or even near impossible, for humans to find.

This idea of AI being an 'experimenter's helper' is comforting. Complementing human intelligence with machine intelligence could have immeasurable positive impacts on continual progress in many fields of natural science, turning our data-rich but chaotic world into a smart planet full of smarter people.

Ginni Rometti, IBM's CEO, is hopeful: "Some people call this artificial intelligence, but the reality is this technology will augment our intelligence." She's right – AI systems could easily become the automatic psychologist, the automatic professor and so on.

Machines are not going to override humans any time soon. Imagine artificial boffins overseeing peer review to validate scientific advances. Can we really trust them?

It seems that Toffler's next big revolution to completely change the way we work is yet to arrive. But remember, machine-learning systems operate as goal-specification systems – they can reconfigure themselves. Unless our minds evolve too, in a creative, ethical manner, AI may one day become our boss.

Dear Lord A.I, Not quite yet.
Signed, humans.