48

Beyond the end of work / 48th St. Gallen Symposium / 2–4 May 2018
Janis Goldsmith – The robot apocalypse is techno-
logical illteracy in disguise

The robot apocalypse is technological illiteracy in desguise

Janis Goldschmidt is one of the top five contributors to this year's St. Gallen Wings of Excellence Award. He studies at the University of Potsdam and will attend the 48[th] St. Gallen Symposium as a Leader of Tomorrow.

Janis Goldschmidt (DE), Ph.D. Candidate in Non-Linear Dynamics, University of Potsdam

Richard Janis Goldschmidt is a Ph.D. researcher in nonlinear dynamics, and Co-Founder of a neurotech start-up. He builds and simulates mathematical models to gain insight into how cognition comes about, how our brain works, and how to build technologies that influences it. He also does research in computational sociology, where he applies natural language processing to investigate patterns in the language used by online news media. Since September 2017, he is the proud dad to a little daughter Leontine. Janis is currently on parental leave.

I am a PhD student and cofounder in a neuroscience startup. I create and use the kinds of technologies against which people want to augment themselves. But take it from me: autonomous robots are not real and there is no intelligence in Artificial Intelligence (AI). What is real is supercharged automation . Automation is making machines perform mundane tasks previously done by humans, a process which has been with us from the days of the Industrial Revolution. The difference is that with computational power exploding, tasks are being automated at a speed unheard of in human history. In a world where nothing runs without computers and where whole industries are being revolutionized, those who do not understand how and why this is done will necessarily be left behind. Seeing this, the real danger is not AI, but technological illiteracy. In fact, the overwhelming majority of people is not able to use even a fraction of the computational power

available at their finger tips. The rules of progress in the form of free market capitalism are going to drag society into the future kicking and screaming. If we want to continue functioning as a society and actively shape our future, the question to be answered is not how to stay economically relevant, but how to become politically mature and understand how supercharged automation changes the world as we know it. Being a productive member of society is just a side-effect of that.

In order to step and not be dragged into this new world, being technologically illiterate has to become socially unacceptable. Just as it became unacceptable not to be able to read 100 years ago, statistics and information processing need to become the foundational corner stones the rest of our learning is structured around. Until this long term goal is achieved and to keep society from buckling under the load of mass unemployment and under-

employment, we have to help those who cannot adapt in the short term and provide them with jobs beneficial to society even if their economic value is automated away. A feeling of self-sufficiency and being a useful and needed member of society lie at the heart of human dignity, which has to be the ultimate goal of societal and technological progress. Otherwise the legitimacy of our economic and political system will erode.

As a software engineer and PhD student in computational physics, I belong to the group of people which is at the forefront of developing our autonomous future. The danger I see is not one of making the human race useless at large, but of people staying wilfully ignorant of what our current technologies are able and not able to do. When talking to friends and family — and even colleagues in physics — I observe that people are not able to harness the power that the machines they interact with every day are

capable of performing. These are highly educated people with degrees in law, economics, chemistry, and physics, yet they both exaggerate the upcoming robot revolution while simultaneously underestimating what is currently happening. What they do not understand is that not only menial labourers need to adjust. Knowledge-workers who think themselves save from disruptions should be worried, too.

The limits of machine cognition or AI

To understand why robots will not take over everything, we need to know the limitations of AI. As Andrew Ng, the founder of the Google Deep Brain project, puts it[1], "we don't see a realistic path for our software to become sentient." My own Ph.D. research is about how and why biologically inspired neural networks obtain their computational power. When taken by themselves, we understand the single neurons making up this network perfectly. But when many neurons act together, they create an emergent collective behaviour which is impossible to predict and hard to understand. Using massive amounts of data and computational power, we hammer these networks into shape to act in accordance with our expectations. For example, thousands of hours of humans driving under various scenarios is being recorded, and then the software extracts a blueprint of what the machine should do given a certain scenario. Outside this scenario, the machine cannot act appropriately. The human brain, on the other hand, is able to infer very complex conclusions in a relatively information-constrained environment, all the while performing a multitude of other tasks. Humans can do this because we have a higher idea of what it is that we are doing. We can ask "why?" and answer with "because!".

It follows that what we understand by AI should not be cognition or sentience,

but the automation of information processing. To use a more applied example than my research, AI can be several cameras pointing at the street and telling a car in real-time that it is straying from its lane, and by how much it should steer to the right. But taking the car off-road already poses insurmountable challenges to AI, as the machine does not understand why and to what purpose a car is being driven.

The easiest way to think about this is by recognizing that we are able to automate only those tasks that we do not actively think about while performing them. Of course, learning to drive a car is mentally exhausting, but once you get the gist of it, you no longer think about it. Making sure not to run the car off the road or into the car ahead requires a lot of attention, but not a lot of thought. Doing this for 8 hours straight we feel mentally and physically drained. Similarly, scanning the pages of an unknown text for a keyword is incredibly tiresome. It would be even more tiresome to read every single page to make sure not to miss anything, so we start skimming over. It is a silly proposition that a human should do this task, if a computer could.

How automation enhances human productivity

In fact, this is what I did in a research collaboration with a sociologist. I automated textdiscovery, and my colleague could focus on doing sociology. We automatically collectedmillions of news stories and extracted a few thousand relevant ones. To put this effort into perspective: the United Nations commissioned a study[2] to assess how European media reacted to the refugee crisis, with as many as 14 researchers having spent a year collecting the information. In our case, automating and running the process took two weeks and provided a reproducible sample of data, free from

errors that occur when the reader is tired, grumpy, or otherwise biased. As a result we were able to move quickly from data collection to analysis and knowledge generation, which are the purpose of our research. Analyzing the remaining text is much harder, which is exactly the point: we are not even close to teaching a computer figures of speech and how to understand context. Sarcasm alone is something even most humans struggle with because it requires a subtle understanding of intent, which a computer has no concept of.

Even though automation does not imply machine intelligence, it is a great boon to those who can harness its potentials, but a challenge to those who cannot. I am a cofounder in a neuroscience startup. We want to support scientists to produce better research with higher replicability. The functioning of our company as well as its future success are predicated on modern tools of automation . With our team being located in the US, Brazil, and Germany, our work would be impossible without these tools. They automatically check logical errors in our software code, merge the work of several people into one product, test it, and then ship it to our production system. They make each single one of us wildly productive and we are not even located on the same continent.

Automation is also what drives our idea of creating reproducible clinical experiments: we store every bit of information, from the timing of each electrical pulse, to the impedance of the subject's skin. This way, researchers can pinpoint how and why the effectiveness of their study was influenced. Cross-correlating this data with other studies, also using tools of machine learning and big data, will allow them to find new treatments, while shrinking clinical trials to a minimum. We automate the tedious parts of a researcher's work, from her having to

keep track of everything she is doing, to her peers having to fine tune each knob to figure out how to reproduce the results. This way, everyone can focus on thinking about their research and not on cleaning up their workbench.

How ignorance of automation is dangerous

But rather than making full use of the potentials automation offers, our society is backward facing. We keep hearing about the good old days in the 1960s and 1970s, when we had full employment and everyone made a good living from hard, manual work. In school, we teach classical literature and study the Roman Empire. Don't get me wrong, I myself had Latin in school and value my learning in this regard. Yet considering the relative importance of Latin versus computer science for our modern world, it is not entirely comprehensible why the latter would receive no treatment at all.

This is a failure of our role models, too: politicians go on television and proclaim that they speak five languages, all the while making light of being bad with numbers. I recently attended a major policy conference, where the CEO of an international human resource consulting firm haughtily explained that people need to adapt and learn the skills that companies require of them. Next he failed to start his slide show presentation.

Our smartest people fail to understand the modern world. How can they then make good choices in shaping it? The wonders of modernity are built around the scientific method and logic. Mathematics and science are the foundations of our prosperity. High speed information processing and automatic statistical inference are changing society faster than any other human invention in history. Not understanding how the very world we live in works is outright dangerous, irrespective of economic considerations. As

Eben Moglen, Law Professor at Columbia University and founder of the Software Freedom Law Center, puts it[3], "surveilling and predicting human behaviour is the new economy". And at his recent keynote in Davos, George Soros noted[4] that "social media companies are inducing people to give up their autonomy." Even if you consider these statements exaggerated, the fact remains that only a minority of people can form an educated opinion on these topics because they have a notion of the sophisticated methods of automatic data processing that lie at their heart. Being tech illiterate thus becomes dangerous.

How to move forward

People like me are coming up with ways to make skills redundant faster than they can be taught. To have everyone stay economically relevant is hence unrealistic, because economic success is entirely predicated upon making more with less (which in this case means less employees). Alternatives like the unconditional basic income suffer from the misconception that people will automatically find something to do with their free time. For one, the paradox of choice[5] really does create anxiety, but also the notion of getting something for free without putting in work appears so deeply shunned by society that fighting it is futile. Instead we need to make people feel needed and embedded in society, and this needs to be reflected in how they are able to share in the fruits of our technological progress. The legitimacy of our system rests on people not feeling disconnected and disenfranchised.

One solution is to elevate jobs that don't outrightly produce anything, yet are beneficial to society. These exist, but are currently woefully underpaid. They are, for example, taking care of children, as well as the sick and old. As Bill Gates notes[6], "if you can take the labor that

used to do the thing automation replaces, and financially and training-wise and fulfillment-wise have that person go off and do these other things, then you're net ahead." Human empathy will not be automated away, and neither will be those jobs that lack a profit motive.

Automation is going to free up a lot of resources . We need to make sure that ultimately all of humanity can share in its benefits — both by actively harnessing it and by being freed from mundane tasks to engage in work that is truly meaningful.

Endnotes

1 https://www.theregister.
co.uk/2015/03/19/andrew_ng_baidu_
ai/

2 http://www.unhcr.org/56bb369c9.pdf

3 https://arstechnica.com/information-
technology/2016/05/eben-moglen-gpl-
online-advertising-isbecoming-
a-perfect-despotism/

4 https://eandt.theiet.org/content/artic-
les/2018/01/tech-giants-are-a-menace-
to-society-says-georgesoros-
in-blistering-speech/

5 https://www.ted.com/talks/barry_
schwartz_on_the_paradox_of_choice

6 https://qz.com/911968/bill-gates-the-ro-
bot-that-takes-your-job-should-pay-taxes/