Education in the Age of AI（Excerpt）

https://shc.stanford.edu/arcade/interventions/education-age-ai

I'm currently reading Yuval Noah Harari's latest book, Homo Deus, and the following passage caught my attention:

"**In fact, as time goes by it becomes easier and easier to replace humans with computer algorithms, not merely because the algorithms are getting smarter, but also because humans are professionalizing.** Ancient hunter-gatherers mastered a very wide variety of skills in order to survive, which is why it would be immensely difficult to design a robotic hunter-gatherer. Such a robot would have to know how to prepare spear points from flint stones, find edible mushrooms in a forest, track down a mammoth and coordinate a charge with a dozen other hunters, and afterwards use medicinal herbs to bandage any wounds. However, over the last few thousand years we humans have been specializing. A taxi driver or a cardiologist specializes in a much narrower niche than a hunter-gatherer, which makes it easier to replace them with AI. As I have repeatedly stressed, AI is nowhere near human-like existence. But 99 per cent of human qualities and abilities are simply redundant for the performance of most modern jobs. For AI to squeeze humans out of the job market it needs only to outperform us in the specific abilities a particular profession demands."

This leads to the core argument I'd like to put forth in this post: **the right educational training and curriculum for the AI-enabled job market of the 21st century should create generalists, not specialists.** Intelligent systems will get better and better at carrying out specific activities and specific tasks on our behalf. They'll do them reliably. They won't get sick. They won't have fragile egos. They won't want to stay home and eat ice cream after a breakup. They can and should take over this specialized work to drive efficiencies and scale. But, machines won't be like startup employees any time soon. They won't be able to reliably wear multiple hats, shifting behavior and style for different contexts and different needs. They won't be creative problem solvers, dreamers, or creators of mission. We need to educate the next generation of workers to be more like startup employees. We need to bring back respect for the generalist. We need the honnête homme of the 17th century or Arnheim*** in Robert Musil's Man Without Qualities. We need hunter-gatherers who may not do one thing fabulously, but have the resiliency to do a lot of things well enough to get by.

What types of skills should these AI-resistant generalists have and how can we teach them?

Flexibility and Adaptability

Andrew Ng is a pithy tweeter. He recently wrote: "**The half-life of knowledge is decreasing. That's why you need to keep learning your whole life, not only through college**."

Interdisciplinarity

There is a lot of value in doing the philosophical work to understand just what our methodologies and assumptions are, and how they shape how we view problems and ask and answer questions about the

world. I think one of the best ways to help students develop sensitivities for methodologies is to have them study a single topic, like climate change, energy, truth, beauty, emergence, whatever it may be, from multiple disciplinary perspectives. So understanding how physics studies climate change; how politicians study climate change; how international relations study climate change; how authors have portrayed climate change and its impact on society in recent literature. Stanford's Thinking Matters and the University of Chicago's Social Thought programs approach big questions this way.

Model Thinking

Michael Lewis does a masterful job narrating the lifelong (though not always strong) partnership between Daniel Kahneman and Amos Tversky in The Undoing Project. Kahneman and Tversky spent their lives showing how **we are horrible probabilistic thinkers**. We struggle with uncertainty and have developed all sorts of narrative and heuristic mental techniques to make our world make more concrete sense. **Unfortunately, we need to improve our statistical intuitions to succeed in the world of AI, which are probabilistic systems that output responses couched in statistical terms**. While we can hide this complexity behind savvy design choices, really understanding how AI works and how it may impact our lives requires that **we develop intuitions for how models, well, model the world**. At least when I was a student 10 years ago, statistics was not required in high school or undergrad. We had to take geometry, algebra, and calculus, not stats. It seems to make sense to make basic statistics a mandatory requirement for contemporary curricula.

Synthetic and Analogical Reasoning

There are a lot of TED Talks about brains and creativity. People love to hear about the science of making up new things. **Many interesting breakthroughs in the history of philosophy or physics came from combining together two strands of thought that were formerly separate**: the French psychoanalyst Jacques Lacan, whose unintelligibility is besides the point, cleverly combined linguistic theory from Ferdinand Saussure with psychoanalytic theory from Sigmund Freud to make his special brand of analysis; the Dutch physicist Erik Verlinde cleverly combined Newton and Maxwell's equations with information theory to come to the stunning conclusion that gravity emerges from entropy (which is debated, but super interesting).

As we saw above, AI systems aren't analogical or synthetic reasoners.

Framing Qualitative Ideas as Quantitative Problems

A final skill that seems paramount for the AI-enabled economy is the ability to translate an idea into something that can be measured. Not everyone needs to be able to this, but there will be good jobs—and more and more jobs—for the people who can.

This is the data science equivalent of being able to go from strategy to tactical execution. Perhaps the hardest thing in data science, in particular as tooling becomes more ubiquitous and commoditized, is to figure out what problems are worth solving and what products are worth building. This requires working closely with non-technical business leaders who set strategy and have visions about where they'd like to go. But it takes a lot of work to break down a big idea

into a set of small steps that can be represented as a quantitative problem, i.e., translated into some sort of technology or product. This is also synthetic and interdisciplinary thinking. It requires the flexibility to speak human and speak machine, to prioritize projects and have a sense for how long it will take to build a system that does what it needs to do, to render the messy real-world tractable for computation. Machines won't be automating this kind of work anytime soon, so it's a skill set worth building. The best way to teach this is through case studies. I'd advocate for co-op training programs alongside theoretical studies, as Waterloo provides for its computer science students.