The Mystery Go Master

Last March, AlphaGo shocked the world by beating Lee Sedol at Go, a game in which humans were once thought to be unbeatable. In the first few days of 2017, a mystery online player named "Master" appeared and proceeded to beat every top professional including China's Ke Jie, the world #1 ranked player. In a matter of days, it amassed a record of 60 wins, no losses, and one draw. In case you were wondering, the draw was because the opponent's internet connection dropped, and the system called it a draw by default. Online mockeries like "haul this one away, next victim" seemed rather appropriate, considering the best score against Master was a bad internet connection. After defeating China's Gu Li, Google finally confirmed that AlphaGo was the actual player behind "Master".

Arguably what was more impressive than the wins were the highly unconventional moves that AlphaGo made. Those moves seemingly made little sense to humans, but turned out to be highly effective - so much so that after his defeat, Ke Jie said "Humans have evolved in games in thousands of years—but computers now tell us humans are all wrong. I think no one is even close to knowing the basics of Go". This poignant reflection from the #1 human player in the world is both a commentary on human's imagined superiority, and more importantly, the unimagined depths of the game yet to be explored.

Artificial intelligence is after all, artificial. We humans created it. We don't feel bad because we cannot outrun a car, or out-calculate a computer; yet somehow this time, it's different. Our sense of awe is tinged with sadness, and our remaining pride is subdued by resignation. But why does it feel so different this time around?

Great Go players develop an intuitive "sense" of the game from years of playing. This intuition, much like art, cannot be described or taught, and we mistakenly believe to be beyond the grasp of a computer. And because Go is considered a game that combines strategy and creativity, there is something deeply unsettling about losing at a game that embodies the very essence of what makes us human.

I believe that this is a misguided view. Despite our intuitions, Go is still fundamentally a game of calculation (albeit a very complex one), and advances in computer science have shown us that we should not confuse unteachable with unlearnable. AlphaGo is still "Narrow AI", based on an imperfect approximation of our own brains, the actual workings of which we still do not fully understand. In fact, here's a handy little secret, courtesy of Dr. Winston of MIT: when it comes to questions in neurobiology, you too can give the same answer as actual neuroscientists 80% of the time – just by answering "I don't know".

Human ingenuity has taken us from a hunting and gathering society to a fully connected world, and fundamentally altered the way we work. From machinery that relieve us of physical labor to computers that free us of basic mental labor, technology has not only made our jobs easier, but more intellectually demanding. It has freed us from the banalities and allowed us to focus on ideas and creativity. Historically this has happened relatively slowly, and we have had time to adjust our workforce and prepare our upcoming generation for the new environment. With rapid improvements in AI, that may no longer be the case.

For example, in the not distant future, ride sharing and autonomous vehicles will likely free us from having to drive, and change the way we think about car ownership. Unlike previous changes, this will happen in years instead of generations., will present a social problem for which nobody seems to have a good solution. When these professional drivers are made redundant in massive numbers and not easily retrained, they lose more than their egos – they lose their livelihoods and more sadly, their sense of purpose. Yet since there is no good solution to this social problem, it is easier to sweep it under the rug and pretend it is someone else's problem.

What about the upcoming generation? Not everyone has the wherewithal or disposition to become a coder, nor are there unlimited positions to absorb them. Careers not easily impacted by technology, such as arts, are simply not growing enough. What will happen when many of us are no longer needed for physical or mental labor? Will we end up in a "Wall-E" world? Perhaps we will end up participating in activities that cannot be replaced by technology, like drinking and gambling. OK, maybe it won't be so bad after all.

Another philosophical question worth pondering is this: have humans stopped evolving? Evolution by natural selection is the greatest driver behind life, but it works on an extremely long time scale. What works on a much shorter time scale is artificial selection – just look at the many dog breeds we have today, achieved through a process of untold cruelties. We can already edit genes with CRISPR, eradicate diseases, and pass on genetic material artificially; it seems that humans, to an extent, have already beat evolution.

So where do we go from here? We already enhance our bodies with technology (for example, robotic limbs or Cochlear implants), and it is not unreasonable to think that our future will involve artificial enhancements to ourselves, progressing from physical to cognitive. This is called transhumanism, and if you don't see what could possibly go wrong, try watching a few episodes of *Black Mirror*.

There is no slowing of AI and technological progress, and the stakes could not be higher. We live in a world built for the past, and rather than letting the latest technology blindly lead us into the future, we need to ask the tough social, ethical, and philosophical questions to make sure we are not led astray. If we don't? Well, I guess there's always drinking and gambling.

References:

1. AlphaGo versus Lee Sedol, *Wikipedia*
2. DeepMind's AlphaGo is secretly beating human players online, *New Scientist*
3. Google's AlphaGo AI secretively won more than 50 straight games against the world's top Go players, *Quartz Media*
4. Master 颠覆了围棋传统 这更证明了藤泽秀行那句话, *乐视体育*
5. Artificial Intelligence, Lecture 12a: Neural Nets (6.034), *MIT Open Courseware*
6. Patrick Winston, *MIT Computer Science and Artificial Intelligence Laboratory*
7. WALL-E, *Wikipedia*
8. Black Mirror, *Wikipedia*