

The Robot and His Hologram

Blade Runner 2049 was last year's sequel to the classic 1980s sci-fi noir that poked at the question of what it meant to be human in a future in thrall to technology. It stars Ryan Gosling as an android gumshoe whose job is to hunt down and retire other androids. When he is not working, Gosling comes home to an artificially intelligent hologram that follows him around the house, giving the illusion of domesticity and companionship. In a way, this gets at the heart of the double-edged nature of artificial intelligence and our headlong plunge into technological progress: it promises so much, but are they empty promises?

On the one hand, you have an artificially intelligent human facsimile executing its own kind, living a distorted reflection of a human life by availing itself of the services of another technology designed to distract from the emptiness of a shallow existence. Is this the future that technology promises us? To create a fractured world in which our own place is uncertain, and then to soothe our anxieties about that uncertainty with shiny falsehoods? Or does it promise an egalitarian society, where people freed from the burden of work and drudgery are able to engage in higher pursuits that expand the limits of human knowledge?

In wrestling with these questions, I was reminded of my first job working for one of the largest commercial property owners in California. One of the office buildings this company owned was in a large metropolitan downtown area, where space was at a premium and parking a lucrative source of cash flow. At this property the parking generated \$40,000 a month in revenue with minimal overhead. The largest cost was payroll and taxes for two full-time parking attendants earning just above the minimum wage, about \$9/hour, so the all-in cost of employing these two workers came to about \$4,000, or 10% of revenue. Parking alone was thus earning \$36,000 in monthly pre-tax profits, which is a healthy margin by any standard.

Nevertheless, we soon received instructions from above to install automated ticketing machines in our building. These cost around \$10,000 each, but by eliminating at least one full-time attendant they shaved \$2,000 off monthly operating expenses. This was my first experience with automation displacing a real human worker, and I soon came to hate these things. Automated parking machines sound nice in principle, but like any technology they often create unanticipated headaches, such as when people forget to pay for their tickets at the machine and then get stuck at the gate, backing up traffic while a symphony of horns alerts the human attendant that a riot is brewing. People constantly lose tickets or can't pay. Sometimes the machines break. At least one full-time person was thus still required to deal with all the problems created by the automated process, and now he had to work even harder, for the same low hourly wage, while the service provided by this box of circuits served mainly to unite tenants and visitors in their hatred of the property owner.

What always stuck with me was the human cost of this automation, and the breezy almost off-hand way in which it was rolled-out by the executives. A person who was already

making a barely liveable wage had his job taken by a ticketing machine so that the owners of a high-rise office building, already millionaires many times over, could bank an extra \$2,000 a month. What this showed me is that automation, technology, artificial intelligence – these things are not inherently good or bad. They are tools and what determines their impact on society is how we choose to use them. The above example, where owners of capital chose to employ technology to marginally add to their wealth at the expense of low-paid human workers, should not be at all surprising in a society like the United States where the naked pursuit of profit and self-interest is often worn as a badge of honour. In fact, it's probably the most obvious application of labour-saving technologies, and will surely accelerate on a much greater scale in the future.

While it is not surprising that technology is being developed and used as a means to maximize profits, it should be concerning, especially when the unchecked pursuit of wealth ignores downstream and second-order social costs that are often difficult to detect but can rebound in very disruptive ways. I don't think it is any coincidence that many of the poor and rural people who voted for Donald Trump and for Brexit, the people that are most often taken in by the empty promises of populist demagogues, have overwhelmingly seen their wages stagnate and employment prospects decline thanks in part to automation eroding their bargaining power and displacing them in the workforce.¹

This should be especially concerning as technological progress quickens. The anecdote above was drawn from nearly a decade ago. Since then the conversation has moved well beyond automated parking ticket machines to software that is capable of doing daily household tasks and even managing investment funds. Artificial intelligence is no longer at the point where it threatens to displace only low-wage workers in the service industry, like fast food employees or parking attendants. It is quickly reaching a point where it will threaten to displace high-skilled workers, like financial consultants or even music producers. The day that artificial intelligence can do Justin Bieber's job is a day we should fear.

But as the pace of technological development increases, the threat is not just that it will hit and displace a wider range of socio-economic actors. It is that change will start to happen at such a speed and scale that society will struggle to cope. As the recent Facebook scandals have demonstrated, we are already playing catch-up – regulators and even the inventors of technology themselves have been slow to recognize the inherent perils of their own platforms. Mark Zuckerberg envisioned Facebook as a tool for connecting the world. One of the most consistent themes that emerged when Facebook came under Congressional scrutiny earlier this year was that amongst its top executives there was a genuine belief that Facebook was an engine for positive social change, that it was making the world a better place - a sentiment that is widely shared in Silicon Valley.

¹ Pedro Nicolaci da Costa, "There are 2 big reasons your pay isn't going up", *Business Insider*, 1 August 2017.

This is undoubtedly a partial truth. Facebook does connect the world in ways undreamt of just fifteen years ago. But the platform can also be used to spread baseless propaganda, inflame racial and ethnic tensions, and trample on personal privacy. This is because the algorithms and elegant math equations guiding the way these platforms sort and distribute information are unthinking, bounded by code that, for all its sophistication, has no concept of ethics or morality or social values. We don't always fully understand the technologies we are unleashing on the world, and instead push ahead simply because we can without questioning whether we should. These technologies still require human oversight to ensure they are providing the greatest benefit for the maximum number of people, a balance that can take years, even many decades, to get right. And as society struggles to dial-in the right balance, a great many things can go wrong in the interim.

That is the Achilles' heel of rushing headlong into what we imagine to be the bottomless promise of technology, the liability built into Zuckerberg's famous mantra of "Move fast and break things." Things are moving ahead at such a rapid clip now that by the time we figure out there is a problem and go back to diagnose it, it may be too late. In such an environment, where technology races ahead for its own sake without stopping to contemplate its effects on society, there is a distinct possibility that we will end up creating something that we never intended and which cannot be reversed.

The key to avoiding such a fate is to constantly remember that human society is not just about proving what we can do, but about what we can do together. We must privilege the notion that technology is created by human beings for other human beings, to help human society prosper in a way that benefits are shared. This is something that Peter Drucker instinctively understood. He was sceptical that technology would render workers obsolete, but that it would instead require the work force to become more flexible, to continually engage in a process of learning and acquiring new skills so as to keep up with and benefit from technology-driven changes.²

A recent Dell Technologies report took this idea and pushed it to the limit, estimating that 85% of jobs that will exist in the year 2030 haven't even been invented yet.³ It is a headline grabbing figure that suggests we are in an exciting and dynamic process of economic change and evolution, where people will slide effortlessly into the new opportunities generated by this techno-futurist vision of tomorrow. But it is also a terrifying figure. Creating millions and millions of new jobs to off-set displacement in existing jobs in barely over a decade is a crushing burden, and the consequences for failure are potentially disastrous. What if those jobs aren't created fast enough or on the scale envisioned? What if machines continue to slowly peel away little pockets of what used to be exclusively human work? How will people be able to find satisfaction in their lives when conditions are

² Rick Wartzman, "What Peter Drucker had to say about automation", *Harvard Business Review*, 2 November 2015.

³ *Emerging Technologies Impact on Society & Work in 2030*, Report from Dell Technologies and Institute for the Future.

constantly changing according to a schedule dictated by technological advances over which they have no control? The report is silent on these questions.

Of course, large economic restructurings have happened many times before in human history. Agriculture gave way to industrial manufacturing which gave way to service-based economies which are now wrestling with the future of disruptive new technologies. In each instance, when new technology made old professions obsolete, there was turmoil and resistance and confusion. Eventually, these changes were incorporated and society adapted to new and often improved ways of life. The Luddites have not aged well and doomsayers predicting we are on the verge of a technology-induced apocalypse do not have the weight of history behind them.

But history does tell us that as new technologies disrupt existing economic structures and displace the jobs of real human beings, there will be a period of turmoil as people adapt to a redistribution of resources, and some segments win out while others lose. This can get bloody, as any Bolshevik would tell you. A key difference this time is that it is happening faster. In the past, when the world was less connected, major shifts took longer and there was more time for people, industries and feelings to change and adjust. It seems unlikely we will enjoy a long and indulgent transition period when it comes to adapting to the realities of artificial intelligence as the world seems to be accelerating with each passing day. Stories that once would have dominated news headlines for weeks, or even months, are crowded out of my Twitter feed within 24 hours now. There is barely time for anything to sink in, for people to chew over major changes and figure out how to feel about them or what the long-term consequences might be.

This acceleration and information overload is bad news for major social changes because while their disruptive outward effects are often quick and messy, their underlying causes tend to brew in the background for a long time. They have long tails, lumbering through the echo chamber of history until they suddenly boil over, intruding into everyday life all at once and the world is never the same. It's hard to dial the clock back on those kinds of changes, so the best thing to do is anticipate and understand them before they happen.

Yet this is difficult to do when technology that we do not even fully understand bombards us daily with an excess of information that we are ill-equipped to process. When we talk about artificial intelligence, about machines getting ever closer to performing tasks at the level of a human, we need to be sure that we don't just talk up the achievements, but that we also think hard about potential downsides. Disruption creates opportunities as it destroys them, so it is critical that benefits are shared and not simply concentrated in the hands of the rich and the powerful, that people who lose jobs are still able to find meaningful ways to contribute to society.

But if you go to a tech conference, few if any people will be talking about these things. In their zeal to show off the latest bit of tech that can make an appointment at a hair salon, no one is asking how a generation of coal miners are going to be re-trained to be IT specialists or the impact such a restructuring will have on their communities. Nobody is even

asking them if they want that, or who they might turn to when the lofty promises made in Silicon Valley turn out to be hollow for the auto worker in Michigan or the McDonald's employee in Ohio.

I will be the first to admit I don't know the answer to these questions. I don't know how to re-train people two or three times within a single lifetimes to do different jobs as their old ones are displaced. I don't know how to take people raised on manufacturing jobs and plug them into a tech ecosystem and make them feel good about it. I don't know how to prevent the upper crust of society from leveraging technology to benefit itself at the expense of everyone else. I don't know what the future holds. But I know that technology alone does not have all the answers, that we cannot place our faith blindly in the idea that sleek gadgets and technological miracles will inevitably help us arrive at the future we think we deserve.

Which is ultimately what makes the story of the robot and his hologram so poignant. Technology helps us do amazing things – go to the stars, communicate instantly across oceans, create complex and engrossing fantasy worlds. But all that glossy tech is merely a tool, a means to an end. As a society, the burden is on us to decide what that end is now and how we want to use these tools to get there. Blade Runner 2049 depicts a society where humans raced ahead with technology without stopping to ask the hard questions. It is a world where a wonder of technology, an android programmed to think and act like a human, is exploited by his makers to hunt down creations just like himself. When he comes home, he indulges in a lie created by yet more technology, the lie that he is something like a human, that he is not alone, that everything is OK.

Hopefully this is not the future we are making for ourselves, where our own leaps into technology we aren't prepared for force us to seek solace in the glossy shine of yet more distractions. But if we continue to speed ahead building computer programs and robots without considering how these advances can be slotted into the overall arc of human progress in a way that ensures benefits are shared and disruptions are minimized, we may find ourselves living in our own dystopian film someday. As we race to cover ourselves in our own genius, we must take steps to protect the most vulnerable members of society, to make sure their jobs and their wages are not eroded by machines just so the wealthiest members of society can improve their bottom line and show off an impressive toy. We need to ensure that machines help humans live better lives, rather than snatching their lives away from them and strangling them of meaning.

The drive to innovate is a very human instinct - to create something that mirrors us and by doing so perhaps capture a glimmer of what it is that makes us special. But it also requires careful thought and planning. It requires a central authority to consider the long-term ethical implications of short-term scientific achievements. It requires a careful consideration of whether these cutting-edge tools are helping us to create a just and equitable society, or something much different. It forces us to ask just what kind of society we even want. The answers are elusive and complicated and hard, but few enough people

are even asking the questions. And by the time we realize which questions we should have been asking, we may find it is too late to go back and answer them.