

## Artificial Intelligence and the Future of Education

Forthcoming in *Education for a Changing World*, New South Wales Department of Education.

Catherine Stinson<sup>1</sup>

May 2019

The promise of Artificial Intelligence (AI) to transform both how we think about ourselves and how we live our lives has remained in the realm of science fiction for centuries, since at least as far back as Mary Shelley's *Frankenstein*. Only recently has AI rather suddenly started to deliver on that promise. Robots are caring for the elderly. Self-driving cars are on the roads. Elections have been won and lost due to AI personality profiling. Warehouse workers are being promoted and fired by autonomous algorithms.

A moment of reckoning is happening in the technology sector. Researchers and coders are realizing that their jobs are not purely technical in nature. Governments and corporations are grappling with the social and ethical impact of AI. Regular people are waking up to the hidden effects technology is having on their lives. This chapter projects into the not-so-distant future to explore the societal changes AI is bringing, and to suggest some ways that education systems might adapt in the face of those changes. Students will need new skills and capacities to live in an AI future, and to steer AI in directions they may find more desirable.

## INTRODUCTION

While we can't predict all of the effects that AI will have in the coming years, some likely scenarios are apparent enough that we can begin planning how education systems ought to change in order to prepare society for a future with AI. Three areas where changes to education systems are called for are digital literacy, social responsibility, and the world of work.

**1. Digital Literacy:** The distinction between being offline and online is blurring as our private lives become saturated with technology. Digital literacy needs to evolve beyond knowing how to use a word processor and not divulging personal information to strangers, to include new skills like navigating ubiquitous surveillance and interacting with AI agents.

**2. Social Responsibility:** One by one fields have come to realize that their work has social implications. Tertiary education has responded by introducing applied ethics training in engineering, medicine, science, business, and now computer science. It is time to consider secular moral education as a basic component of secondary schooling.

**3. The World of Work:** Employment conditions are changing rapidly. Gig work has become the norm, globalization is well underway, and the automation of many sectors is expected in the coming years. Instead of preparing students for jobs, education systems need to start preparing students for a world where most people won't have traditional jobs.

---

<sup>1</sup> Email: [hello@catherinestinson.ca](mailto:hello@catherinestinson.ca)

These three areas do not exhaust the changes to education that will be needed to prepare for a future where AI is ubiquitous. Some areas not discussed here are the technical training that will be needed for future AI workers, as well as workers in areas like medicine, which are increasingly making use of AI technology. Another area only discussed in passing is how education in political science and economics will need to adapt. These areas may be expected to shift more organically, rather than requiring major overhauls.

As a way of illustrating the need for changes to educational approaches to digital literacy, social responsibility, and the world of work, the chapter begins with a series of not-so-distant future scenarios that show how current trends in AI might continue along paths that lead to bleak futures. Subsequent sections detail how education systems can respond to and help prevent these outcomes. The future scenarios are then revisited to illustrate how things might turn out differently.

## **FUTURE SCENARIOS**

Among science fiction fans, there's growing sadness that the genre is beginning to die, because the real world is becoming so much like science fiction that there's hardly room left for imaginative speculation about the future. The three future scenarios below take as their starting point technological trends that are currently happening, and project their extension into a not-so-distant future. This is more an exercise in spelling out possible consequences of what is already here than creation of fantasy worlds. The three scenarios focus on privacy, governance, and automation.

### **Scenario 1:**

When you're in a major city, it's hard to escape the gaze of surveillance cameras. In Kashgar, in northwest China, Uighur Muslims are surveilled constantly through tracking software in cell phones, ID checkpoints, and video cameras. A Google affiliate company is proposing to build a prototype "smart" neighbourhood in Toronto, featuring cameras in doorbells and garbage cans. Police forces are investing in facial recognition software, and piloting projects to identify individuals gathered in public spaces, and driving past strategic points like bridges. The public is participating in this surveillance by taking video of altercations they witness, and publicly identifying and shaming perceived wrongdoers. Landlords have been caught hiding cameras in rental units, and families are surveilling one another with nanny cams and spyware.

Cell phones, fitness gadgets, navigation systems, and key finders like Tile are tracking the locations and activities of technology users. Some location data like social media check-ins and face tagging are freely accessible, but much of it is held and sold by tech companies. Electronic payments will soon become mandatory even for small purchases like a cup of coffee or subway fare, allowing corporations to build complete consumer profiles.

Although some cities like San Francisco are banning facial recognition software by police and city agencies, these bans will be isolated to places with wealthy residents (many of whom, ironically, work in AI), and do not apply to other levels of government or private business. Facial recognition's use will continue to expand rapidly, allowing anyone with a

camera to identify individuals, as long as they have access to identity databases, but those databases will not be accessible to all. Governments have databases of faces linked to identities from drivers' licenses and passports. AI companies are building databases from photo apps, and social media tags. Both governments and corporations will soon have the capacity to track individuals' locations and activities in real-time.

Criminal and other undesirable activities will be automatically detected and dealt with by, for example, automatically deducting fines for illegal parking, or apprehending criminals before they act. Which activities will be punished will depend on the priorities of police forces, governments, and AI companies, and these systems will be susceptible to interference by the powerful. Facial recognition software will continue to work unreliably for marginalized groups, particularly black women and trans people. Cases of mistaken identity will occur frequently, and since the systems lack transparency, it will be difficult to prove one's innocence without access to surveillance data controlled by police, governments, and corporations. Deep fakes will frame people for crimes they didn't commit, and access to data will be essential to prove one's innocence.

Personalized advertizing will invade our personal space, popping up on all of the screens we encounter, through ear buds, on wearables, and via personal assistants like Alexa. The only way to opt out of ubiquitous advertizing will be to pay monthly subscription fees, but the more subtle behaviour modification nudges that influence moods and political leanings won't be noticeable or opt-outable. The feeling of constantly being monitored will erode public trust to the point where people do not feel comfortable going out in public, creating social isolation and furthering dependence on social media, wearables, and interactions with AI.

### **Scenario 2:**

The profits of the biggest tech companies exceed the GDP of most countries, giving them power to not only influence regulation and governance on a global scale, but also deeply affect living conditions. This is already apparent in how tech companies alter housing markets, transportation infrastructure, and working conditions in the cities where they operate in ways that affect entire regions. Tech companies are getting into the business of providing community healthcare services, public health monitoring, email and web servers for governments, universities, and non-profits, transportation planning, education delivery, and curriculum design. AI has already shifted public opinion in several political elections, with global effects, and its power to control public opinion will only grow. The EU's General Data Protection Regulation is one measure to regulate the industry, but outside the EU, AI companies are defying national governments by refusing to comply with local privacy laws, and simply paying fines that are too small to affect their profits.

Government revenues are increasingly threatened by these same dynamics of globalization and concentration of wealth in a few large companies, who find ways of avoiding taxation. Without tax revenue, governments are becoming less able to compete with tech companies in their ability to govern. This is leading to the offloading of essential services to technology companies, at the price of relinquishing control over how those services are provided, and where the data goes. These same companies are also spending massive amounts of money lobbying governments and funding candidates' campaigns. We will soon see open corporate ownership of governments, and formerly public services like education and healthcare will

only be available via paid subscriptions to institutions like The Amazon White House, or The Parliament of Australia: powered by Moodle.

### **Scenario 3:**

Tech jobs are becoming increasingly stratified. Tech support has already moved to markets where wages are low and worker protections weak. Many programming jobs have been downgraded to temporary gig work, and tasks like data labeling are being done for extremely low pay through hubs like Mechanical Turk, or in sweatshop conditions in the global south. Power over how AI technology develops is becoming concentrated in a few hands, hampering the ability of users, tech workers and governments to regulate the industry.

As AI and automation destroy job markets in sectors like customer service, care work, education, transportation, and manufacturing, unemployment will rise well above 50%, leaving most people with nothing better to do than spend hours playing video games and surfing social media. These apps will become highly addictive, and monetary incentives to stay online and keep engaging will be woven into the experience, so that tasks like data labeling and amplifying political messages will be performed by uncompensated quasi-workers unaware that they're working. Insofar as democracy still exists, voting will move online, where people's choices can be better manipulated by the forces in control of the technology. Identification verification for everything from opening the front door to your apartment, to voting will use facial recognition, and systematically disenfranchise racialized groups for whom the technology works unreliably.

As unemployment soars and climate refugees migrate en masse, population control measures will be introduced, with algorithms deciding who is fit to reproduce and who is fit to parent. AI applications in genetic engineering will be used to design babies for those who can afford it, further widening the gap in opportunities between rich and poor.

These three not-so-distant future scenarios illustrate some likely next steps in the evolution of AI technology and its effects on societies. None of these scenarios are desirable outcomes, but all of them are preventable through a combination of activism, regulation, and education. Adult and tertiary education will also play an important role steering outcomes away from these future scenarios, but here we focus on what secondary education can do to prepare young people for the future of AI.

## **DIGITAL LITERACY**

Traditionally digital literacy education consists of experience using popular applications like word processors, spreadsheets, presentation tools, and sometimes more specialized software like video editing or design suites. Training in coding skills and website creation is becoming more commonplace. In addition, students typically get training in how to stay safe online by being careful of sharing personal information with strangers, avoiding phishing scams and malware, and responding effectively to online bullying. All of these are useful lessons, but only scratch the surface of what people will need to know to thrive in a world run by AI.

As AI becomes more interactive, we are going to need to know how to speak to it. It may seem harmless to swear in frustration at self-checkout machines when they malfunction, but

as these tools become more human-like, the social and ethical implications become concerning. Already people are noting how disturbing it is that small children shout abuse at personal assistants like Alexa (which are so often coded as feminine). The problem is not just that the conscious Alexa of the future might have her artificial feelings hurt, as that we are getting into the dangerous habit of treating AI like an underclass deserving of abuse. It harms our humanity to behave that way. While it makes sense to drop the pleasantries and simply state a request when emailing a scheduling bot like x.ai, there is nothing to be gained by treating AI in demeaning ways, no matter how frustrating it is to interact with. Just as children used to be taught how to make polite phone calls, they should now be taught to interact respectfully with AI personal assistants.

General familiarity with AI and data analysis techniques is another area where secondary education could be beneficial. The common perception that AI is mysterious and incomprehensible needs to be challenged. Without requiring any advanced mathematics or knowledge of the details of complex algorithms, the basic idea of how AI algorithms work can be communicated. For example, the k-means algorithm could be illustrated by getting students to organize themselves into groups by height. First students are randomly chosen to represent k group means, and each student joins the group whose mean matches their height most closely, then new mean height students are chosen for each group, and other students are reassigned to the groups based on the new means. This repeats until no further corrections are needed. Collaborative filtering could be illustrated by asking a series of questions about which movies the students like. With each question, students who answer the same way stay together, and those who answer differently branch off into a new group. After a few divisions, the students can make predictions about what other movies the people in their group will like based on what they like. The main points to get across are that these algorithms are in essence very simple, but that with enough data and computing power they can make very accurate and detailed predictions and classifications.

The next step is awareness of just how much data is being gathered about us, and what that data is being used for. It used to be that to protect your privacy online you ought not to reveal things like social security numbers, home addresses, or compromising pictures. In the past few years many people learned the hard way that answering questions about which 80s TV character you're most like, or which breed of dog you are can be just as dangerous. These quizzes have been used to build detailed personality profiles, which have been used to manipulate people's beliefs and behaviours. Likewise, commercial DNA sequencing services marketed as ways of finding unknown relatives, or checking whether you're part Cherokee have very broad terms of service that allow them to share and sell incredibly sensitive personal information. The companies also have troubling partnerships with law enforcement. Similarly, apps that match your face to famous paintings, or filter your face to look like a baby are almost certainly being used to build facial recognition databases, which will likely end up being used for surveillance. Protecting your privacy online now requires a lot more savvy about what apparently benign apps are doing behind the scenes.

Privacy policies are often oppressively long and technical to read, but students can learn some of the basic vocabulary needed to understand them, and practice scanning these documents for suspicious content. Navigating the privacy settings of common applications like Google is also a learnable skill. An exercise that can help with awareness of how data is being used is to research the companies making the software students like to use. Who are

the board members? What political contributions have they made? What other companies are they involved with? What patents have they filed? Similar exercises can be done with political memes. Are they being disseminated by bots? Which social media accounts or Reddit threads did they originate from? What other content is posted in those threads? What countries are they logging in from? Can the claims in them be fact checked?

A third component of digital literacy is knowing what can be done to push back when technology companies behave unethically, and when it's important to do so. Knowledge of privacy and consent protections has long been the near exclusive domain of legal scholars, but this knowledge needs to go mainstream. People should know that when stores asks for your email address at the checkout, you can refuse, when a club or service asks for demographic information as part of their registration process, they have no legal basis for requiring it, and that they need your consent to use that data for any additional purposes.

Being put on The Gap's mailing list may seem like a minor inconvenience, but in the age of AI, giving away apparently insignificant data can have significant consequences. It is not difficult to correlate purchase behaviours with personal information like pregnancy status, health problems, age, and political beliefs. Since email addresses are used as website logins, they can be used to track online behaviour, and connect people to locations and social networks. All told, that's a lot of powerful information to give a store in exchange for a coupon.

Students could also benefit from learning more about the history of technology, business regulation and workers' rights. Students could be introduced to case studies illustrating how technology changes society, often in unanticipated ways, and to the common patterns that technological disruptions tend to follow. Students also would benefit from knowing how class action lawsuits, strikes, boycotts, letter writing campaigns, and community advocacy have been successful in the past at changing unethical business practices, and protecting the public from harm. These case studies can be introduced alongside regular curriculum in subject like science, health, social studies, history, mathematics, and physical education.

## **SOCIAL RESPONSIBILITY**

The revelation that Cambridge Analytica's AI-driven personality profiling may have swung both the Brexit and 2016 US Presidential elections, and the looming possibility that self-driving cars will soon be commonplace has provoked a moral reckoning in AI. There has been a shift in the field toward recognizing that technological innovations have social and ethical implications, contrary to the previously popular view that algorithms are politically and ethically neutral. While this realization has come late to many people working in AI, researchers in science and technology studies have been raising alarm bells about AI's potential harm for decades.

The need for ethical oversight of AI is now a hot topic. Some of the proposed oversight measures, like codes of ethics for the field, or ethics officers within tech companies, have been criticized for how ineffectual they're likely to be. Another suggestion—requiring applied ethics courses in tertiary computer science education, as well as more attention to social and ethical implications of technology throughout the curriculum—may be more

effective, as long as it's done well. There is also room for more ethics training earlier, in secondary education.

AI is only the most recent field to have its moment of reckoning. Physics had its moment during the Manhattan project. Biology has faced ethical issues over cloning and genetic engineering. Engineering has long recognized that skyscrapers and bridges need to be built with public safety top of mind. Medicine had its moment during the Nuremberg trials. Business ethics training picked up around the time of the Enron scandal. What all of this illustrates is that moral education is essential in most every field of human endeavour. It may be that the secularization of public schools has left a moral education vacuum.

Two likely reasons why public schools often do not include classes in morality are the conviction that values ought to be taught in the home or places of worship, and a belief in moral relativism: that there are no universal human values. With most parents now working full-time, and fewer people participating in organized religion, formal moral education isn't happening in those other spaces. Although there are certainly areas of disagreement on values across cultures and religions, there is a lot of common ground too. Kindergarten classes already teach universal moral rules like sharing, and keeping one's hands to oneself. Furthermore, agreement about values is not necessary for teaching ethics. The practice of considering ethical implications is often more important than the particular conclusions reached.

Pedagogical research shows that the most effective applied ethics classes teach ethical theory embedded within realistic examples taken from the relevant field, where students are given ample opportunities for discussion and reflection, and time interacting with teachers. Medical ethics classes might illustrate the difference between utilitarian and Kantian ethics using a case study where the benefits and harms to a pregnant person and their fetus are at odds. Business ethics classes might teach similar theoretical content using a case study about oil exploration, where the interests of shareholders and the public are at odds. In such cases, disagreements about values are likely to remain deeply entrenched, but students can develop moral reasoning skills by breaking down the problem to make clear what all the competing interests are, and by defending their choices about which actions are best, based on a careful consideration of those competing interests. Decades of pedagogical research confirm that moral reasoning and empathy can be taught.

There is also evidence that in first year university students, extra-curricular activities promote moral reasoning.<sup>2</sup> This effect was credited to these activities involving thoughtful dialogue with peers about how to solve collective problems, which helps students break out of egocentric thinking about ethical issues. This could be used to argue for the educational value of extra-curriculars in secondary education as well, with the caveat that students need to be given some autonomy in making decisions about their clubs and teams.

---

<sup>2</sup> Mayhew, Matthew J., Seifert, Tricia A., & Pascarella, Ernest T. (2010). A multi-institutional assessment of moral reasoning development among first-year students. *The Review of Higher Education*, 33(3): 357-390.

Moral reasoning skills are not just ends in themselves. In medicine, higher scores on measures of moral reasoning and empathy correlate with higher patient satisfaction and adherence to therapy, decreased medical errors and malpractice claims, and better health outcomes.<sup>3,4</sup> In business, customer and employee loyalty increase when a company is perceived as behaving ethically, which increases profits and reduces staffing costs.<sup>5</sup> Not only are there compelling reasons for teaching ethics, students would benefit from knowing that social responsibility is good for business, in order to counteract the popular view in business and tech circles that the way to get ahead is to be ruthless and break things.

Another proposed corrective to AI's ethical crisis is increasing the diversity of the workforce. However, it is not always made clear how exactly more diversity is supposed to make for more ethical AI. It could be that diversity is seen as a good in itself. It could be that diversity hires are meant to protect the interests of the underrepresented groups to which the people hired belong. A more nuanced set of arguments coming from work in standpoint epistemology suggests that groups of scientists are more likely to achieve objectivity if the group includes members of marginalized communities, because those individuals are in a better position to notice assumptions that are taken for granted by insiders.<sup>6</sup> Others argue that it is the functional diversity of scientific teams that matters for high performance, rather than inclusion of members occupying marginalized positions.<sup>7,8</sup> These may come to much the same thing, as increasing diversity in teams tends to mean including members of marginalized communities.

This sort of message of inclusivity and diversity is already part of liberal education, but could be strengthened in several ways. One would be by modeling diversity practices in the hiring of teachers, principals, and other school staff. Another would be making sure that de facto segregation does not occur due to neighbourhood boundaries, streaming, or specialized programs like second language courses, performing art programs, and gifted programs.

Perhaps the most powerful way of diversifying the future composition of the AI workforce would be to intervene on stereotypes about math and computer science, and to train students in how to enact cultural change. Programs that teach girls to code have largely been unsuccessful in increasing the number of women in the field, not because of a lack of girls'

---

<sup>3</sup> Riess, H., Kelley, J. M., Bailey, R. W., Dunn, E. J., & Phillips, M. (2012). Empathy training for resident physicians: a randomized controlled trial of a neuroscience-informed curriculum. *Journal of general internal medicine*, 27(10), 1280-1286.

<sup>4</sup> Hall, J. A., Roter, D. L., & Katz, N. R. (1988). Meta-analysis of correlates of provider behavior in medical encounters. *Medical Care*, 26(7), 657-675.

<sup>5</sup> Schmidt, C. D., McAdams, C. R., & Foster, V. (2009). Promoting the moral reasoning of undergraduate business students through a deliberate psychological education-based classroom intervention. *Journal of Moral Education*, 38(3), 315-334.

<sup>6</sup> Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3), 575-599.

<sup>7</sup> Longino, H. E. (1994). In search of feminist epistemology. *The Monist*, 77(4), 472-485.

<sup>8</sup> Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385-16389.



interest in technology, but because the pipeline leaks at later points. When those girls enter the workforce, they face stereotype threat, discrimination, and harassment, which lead to attrition. Targeted skill building programs aren't effective if they're not paired with tools for cultural change. All students need to be taught how to act as allies to colleagues facing discrimination, for instance by amplifying voices, putting others forward for rewarding tasks, keeping records of bad behaviour, and speaking out against it. Actively promoting inclusion and making reparations with communities that have suffered systemic injustice are skills that advocacy groups for disability justice, indigenous sovereignty and anti-racism have a lot of knowledge about. Putting people with that experience in charge of developing new curriculum would be the most direct route to making the skills needed to enact cultural change more widespread.

One of the most pressing reasons for including more moral education in secondary schooling, especially embedded within science and technology classes, is that technology is increasingly being automated. It used to be that computer programs had human users running them and making decisions about what actions to take with the guidance of the program. There is a shift happening where the human user is being removed from the equation, and programs are making autonomous decisions about which actions to take, then implementing them directly, without oversight. A problem that gets a lot of research and media attention is whether we'll be able to design moral machines, and how to do so. The problem being faced right now is that we're allowing machines to make decisions of moral import without even trying to build moral decision-making abilities into them. It arguably used to be the case that computer scientists didn't need to know ethics, because their programs had human users who could take responsibility for decisions, but that abdication of responsibility is not acceptable when there is no human user. This means that the people building technology need to follow a very different design process that considers ethical questions about social impact from the earliest stages of the design process. For that new kind of design process to sink in, we need to start much earlier, and invest much more effort than offering a single applied ethics course in university programs.

## **THE WORLD OF WORK**

Schools always play a balancing act, where the pressure to prepare students for jobs or careers competes with the responsibility to raise happy, well-adjusted people. In recent years many school boards have shifted the balance toward spending more time on career readiness skills and less time on things like physical play and the arts. As we transition to economies where many jobs are being automated, this balance will need to be re-adjusted. There is some disagreement and uncertainty over which types of jobs will disappear, which will remain, and which will be created. What does look certain is that there will be far fewer jobs, far more people, yet still enough resources to support the population in the medium term (dependent of course on how climate change plays out).

One way that secondary education can adapt to the changing world of work is to prepare students for the jobs that are most likely to continue to exist, and to be created. Some of the types of jobs expected to remain are technical jobs like writing, debugging and testing code, designing and repairing machinery, maintaining data privacy and security, and doing entrepreneurial work like developing new product ideas. As mentioned already, an additional

layer to these jobs will be the necessity to consider ethical and social impacts of new technologies.

It seems likely that new jobs will be created for people who are both technically competent and trained in fields like sociology or philosophy, so as to act as ethics and social impact advisers within tech companies. In the near term, there will be a need for teachers to re-train current tech workers who did not receive this kind of training in school, and to develop school curriculum and adult education programs for this re-training. AI Ethics programs are already cropping up in universities.

One trend already apparent in the gig economy is the taskification of work, where jobs get broken down into ever smaller tasks. This translates to lower pay for workers who no longer need broad skill sets, and fewer paid hours of work, since transition time between tasks is cut out. Taskification may make many jobs much more specialized and narrow.

But there is also reason to believe that interdisciplinarity will be particularly valued in jobs of the future. The need for people with both technical and sociological or ethical expertise is just one example. In fields like medicine and education, greater digital literacy will be needed as high tech treatments and educational tools become more common, requiring doctors, nurses and teachers to be better trained in the use of technology. Likewise, if there is less work to be done in general, it may be useful to have employees who are more versatile and flexible, so that they can do many kinds of work, depending on what needs to be done. Generalists could therefore become particularly valuable. Recent evidence also suggests that generalists are better at predicting future outcomes than specialists, which may be an important skill in a world where change seems to be accelerating.

Another common prediction about the future of work is that the jobs that will survive will be ones that require creativity and imagination, because these are skills that have proven difficult to automate. While deep learning has created programs that can mimic the artistic styles of famous painters, and write (slightly odd) screenplays, coming up with truly new ideas is still a characteristically human skill. To prepare for this, the return of daily art and music classes, and free play time would be valuable, as well as the addition of creativity training within other disciplines. A biology class on evolution, for example, might include an exercise where students imagine which types of animals and plants might evolve in a fictional environment, or design new hybrid organisms.

Critical thinking is another set of skills projected to increase in importance in coming years. While logical reasoning is easy for computers, the fuzzier skills that make up common sense are more difficult to automate. Picking out the meaning in an allegory or poem, recognizing sarcasm, seeing what is missing or out of place in a picture, and figuring out which information is relevant to a given context are all examples of critical thinking skills at which humans excel, but computers do not. Critical thinking could be taught in combination with digital literacy. Evaluating whether a social media account is spreading propaganda, for example, includes elements of both.

In addition to preparing students for the kinds of work that will be available in the future, a major shift will be needed to prepare students for a future without work. If many jobs are expected to disappear, we will need to shift to a very different economic and social system

where people work much fewer hours per week, people retire at a much younger age, or many people do not work at all. One major task will be to figure out how to enact this rearrangement of economic and social systems. Young people may be the best source of ideas about how to distribute resources in a world without near universal work.

This will require, among other things, a psychological shift away from seeing work as the major source of fulfillment and personal pride. People will need to develop other sources of fulfillment and pride. These could include artistic creation, caring relationships, innovation and invention, volunteer work, sensory experiences, and lifelong learning. Rather than requiring that students follow strict curricula that force competence in a set of basic skills, there could be more freedom to choose educational paths starting as early as students' particular interests appear. If starting to work by the age of 18 or 22 is no longer a financial necessity, there should be much less pressure to follow a linear path through grade levels.

If lifelong learning becomes commonplace, it might also make sense to better integrate education with work, so that student/workers can move more seamlessly back and forth between learning and doing. Instead of having students do major projects that end as a report that only their teacher reads, teachers could be encouraged to give assignments that have real-world impact. Governments and businesses could be encouraged to have their employees work on projects in collaboration with teachers and students.

One of the most difficult shifts will be to the expectations people have about their lives. We will have to stop measuring our worth (and income) based on productivity. If there is less work to be done, it would also make sense to focus more on quality rather than on speed or quantity. To prepare people for this shift, it may be advisable to phase out timed tests and other evaluation methods that valorize speed and productivity over careful work.

There are also more indirect consequences that AI may have on education. One of the most restrictive aspects of schooling is the requirement that students sit quietly at desks for hours a day. That requirement is justified partly by the perceived need to train well-behaved workers, but mostly by the lack of adequate staffing to handle smaller groups where more movement and noise would be manageable. If income is no longer tied strictly to work, and many jobs disappear, teaching may become a much more popular pursuit. With more teachers, it would be feasible to have much smaller class sizes (which would also make teaching more attractive). Students who have trouble sitting still and concentrating could be given more freedom to move around, play, or spend time on physical tasks like building. Math concepts could be taught in much more physical ways involving bodies moving around on a field, balls being thrown or kicked into the air, or through activities like knitting. Biology is very amenable to outdoor, tactile learning. History can be taught through reenactment. The consequences of AI for education may be quite far reaching.

## **REVISED FUTURE SCENARIOS**

Technology is only as powerful as we allow it to be. The pace of technological change is not unstoppable, and the much-discussed singularity is not inevitable. Although AI looks to be moving in sinister directions at the moment, motivated public servants and an educated public have the capacity to change AI's direction. Getting young people on board with the

project of imposing regulation on AI through changes to secondary education is one way of hastening brighter futures. Here we revisit the three future scenarios outlined earlier, and re-imagine how things might turn out given the suggested changes to education.

### **Scenario 1:**

Young people are cluing into the fact that having “nothing to hide” does not keep one safe from ubiquitous surveillance. When adversaries can harm your reputation by slightly slowing down a video of you, taking a private moment out of context, or weaponizing seemingly harmless information to manipulate your behaviour, no person and no behaviour is safe anymore. Ubiquitous location tracking and harvesting of personal data for uses the public neither knows about nor consents to is being uncovered and will be shut down.

Pushback against facial recognition technology is growing, and will turn into a powerful movement. Many jurisdictions are watching San Francisco’s ban on its use by police forces and other public entities. What will follow are more widespread national bans, and regulation of its use in public spaces. It will become so taboo that no company would risk getting caught using facial recognition in public. Countries that continue to use facial recognition will be sanctioned for human rights abuses.

A deeper cultural shift will be necessary before families stop using spyware on one another. Anger is growing over products that are clearly designed for cyberstalking, and a few cybersecurity experts are disseminating tools to disable these apps, and spreading awareness about them. Given the potential dangers, a high profile case where this technology is implicated in violence seems inevitable. The public’s growing digital literacy and awareness of privacy rights will ensure that this sort of case will be followed up with strict regulation.

Surveillance of employees is an issue that labour unions will continue to fight. As income becomes less tied to employment, it will become much less acceptable for employers to exploit their workers’ labour and to impose unreasonable restrictions on workplace behaviour. The rewards for surveilling employees will also diminish as most rote work gets automated. Surveilling creative work has much less of a payoff, since the process is less observable.

Another area where pushback is growing is against overly permissive privacy policies and terms of service agreements online. People are becoming more aware that they need to look closely at the policies in place for the apps they use, and are gaining the tech savvy to either turn off settings that permit rights violations, or choosing to sign off from the offending services. Facebook, for example is losing users, particularly among younger people.

As consumers demand apps that do not track them, harvest their personal data, or serve them constant advertizing, competition will force developers to change their business models. Already alternative platforms are being developed that are premised on personal control of data, or cooperative data governance through local data trusts. The market is ready for solutions to online experience problems like routine data leaks, trolling without consequences, and fake news. The main things standing in the way of innovation in this sector are the monopolies who control most of our online experiences, but there is already talk of the need to break up those companies, and to require platforms to better monitor the content their users post for fakes and slander.

Instead of a future where the thought police fine you even before you've illegally parked, we can instead build a future where AI's uses will benefit the public.

### **Scenario 2:**

GDPR was the first step toward governments taking back control from technology companies. Similar legislation is pending in countries outside the EU. Several governments have legal cases in progress against companies like Facebook and Google, and are considering taking action to break up these companies into smaller ones to encourage competition. Coalitions of countries are also discussing strategies for ensuring that the major technology companies can't continue to avoid taxation on their global earnings.

While the current situation is one where tech companies have a lot of influence over policy, thanks to young people with the digital literacy skills not to be intimidated by tech speak and a greater appreciation for business ethics, the tide will turn. Although continued effort will be needed to keep regulation from lagging behind, the tech giants will largely be reigned in, returning political decision-making as well as tax revenues to governments.

A key piece of regulation will require companies that lay off workers because of automation to provide those workers with pensions. The human costs of automation will have to be weighed into decisions to invest in AI, which will slow its adoption long enough for regulation to catch up. Pushback against austerity measures will help return the provision of social services to the public sector, and continue the global move toward stronger public education, health, and mass transit systems. Attempts by tech companies to get into the business of controlling infrastructure in major cities by setting up smart city pilots will be shut down by movements like #BlockSidewalk.

Only a few nations will switch to holding elections using online ballots. Largely the idea will be abandoned in favour of the higher security of paper ballots filled out and counted by hand. The promised convenience of online elections will fail to materialize due to widespread authentication problems, hacking, and other security issues, as well as the mismatch between the affordances of digital technology, which makes replication, alteration and spying easy, and the mechanics of voting, where authenticated single copies that are unalterable and secure are needed. The UN will step in to reverse the disenfranchisement of citizens denied their right to vote due to facial recognition errors.

### **Scenario 3:**

In a few countries AI automation of jobs will initially roll out in such a way that profits are concentrated in a few hands, the small proportion of workers who keep their jobs are tightly controlled by the bosses, and the large numbers of unemployed are heavily policed or incarcerated. In several powerful countries where social programs are already well developed and natural resources are plentiful, the move to automation of jobs will come more gently, with guaranteed income programs, pensions for laid off workers, and extended subsidized leaves for parenting, elder care, and re-training.

Instead of a few people working under tightly controlled conditions, working hours will be drastically reduced, so that most people who are able to work remain in the workforce, but work what would now be considered part-time hours, with leaves of absence of years at a

time to care for family members, and frequent educational leaves. Businesses will discover that for the work that is best done by humans, it is more cost effective to shorten working hours and allow flexible leaves, because the quality of the work more than makes up for the decrease in working hours. Work hours will become concentrated bursts of highly motivated, efficient work with very few errors. Instead of hanging around water coolers when they're bored, workers will leave the office when they're tapped out for the day. This will allow for much more efficient office environments, with flexible desk assignments, and energy savings.

Because of strong social programs, workers will no longer feel pressured to stay in exploitative jobs, and the balance of power will shift to workers, who can choose to leave en masse if they aren't well treated. Where they are well treated loyalty will improve, which again translates to greater efficiency. Guaranteed incomes will not make people lazy. Instead it will free up people who need to care for family members to do that without pressure to also do other work, allow people with health (including mental health) problems to adjust working conditions to what is manageable, and allow people to find work that they are motivated to do. The reduced stress, better early years childcare, and preventative health care that result will reduce costs in the healthcare, education, and police systems enough to more than make up for increased investments in social programs to cover subsidized leaves and guaranteed incomes. People will find fulfillment from making human connections, contributing high quality work, creative expression, and immersive experiences. The perceived need to build AI that can do care work will be reduced, as will our dependence on social media interactions, and interactions with AI personal assistants.

The move to data governance systems where individuals maintain control of their data, either personally or through local data trusts will ensure a level of trust in technology such that people will feel safe and comfortable freely sharing data for select purposes like evidence-based medical and public policy research. This will allow medical researchers access to high quality health data linked to lifestyle and genetic data, ushering in an age of innovation in healthcare such that real-time tracking and management of contagious diseases will be possible, drastically reducing the number of sick days among workers, and all but eliminating preventable health problems. Social service provision will also improve dramatically both in coverage and efficiency with trusted access to data. As policing becomes much less necessary, arguments for ubiquitous surveillance will become unconvincing.

As self-driving cars become commonplace, their navigation systems will be coordinated with those of mass transit systems, in order to optimize transportation. On popular routes, individual and family sized self-driving pods, as well as cargo vehicles will attach themselves into trains to reduce fuel costs and collision risks. Driving as separated pods will only be necessary on remote routes, at unpopular times, or for final docking into parking spaces. Driverless taxis and busses will also form part of these communal trains, making the only differences between car ownership and public transit use the freedom to make interior design choices, keep one's belongings in the pod between uses, make block payments instead of paying per use, and the need to secure parking space. Car ownership will become uncommon for people without specialized and extremely frequent travel needs. Energy consumption, collisions, and the repair and health costs associated with collisions and poor road maintenance will all reduce dramatically, allowing for further investment into maintenance of and improvements to transportation infrastructure.

Opportunities to contribute piecework in service of the public will become integrated into public infrastructure, for example in the form of gym cars on trains, where commuters can peddle a bike, lift weights, turn a tread wheel, or pull oars so as to contribute energy to the working of the train, offsetting energy needs while also keeping their bodies healthy. This sort of task will be gamified to make it an entertaining alternative to video games that do not contribute useful work or provide benefits to users.

The countries that take this path will be much more productive, efficient, and safe than the few police states. This will lend them economic and political power, and the ability to take in refugees fleeing police states. Eventually economic and political pressure from the outside, combined with internal instability will force the police states to introduce reforms.

## **CONCLUSIONS**

The concentrated power of the big tech companies, combined with their lax ethical standards, and lack of regulation makes for an ominous situation, where change is sorely needed. Current trends suggest that we're headed to a world of ubiquitous surveillance, where facial recognition software is widely used to track activities, and authenticate identities, largely for the purpose of manipulation and marketing, tech companies are encroaching into governance and delivery of essential services, and the world of work is moving toward widespread unemployment and disenfranchisement.

Current trends need not continue unchecked, however. In addition to tighter regulation, the young people of today can be given the skills and knowledge to hold tech giants to account. With expanded digital literacy training, students can gain facility with AI algorithms, and the legalese in terms of service agreements. They can learn how to protect their personal information, how to spot fake news and attempts at manipulation, and how to push back when their labour is being exploited.

The next generation of tech workers can be taught mature ethical reasoning skills and a sense of social responsibility. If teachers model the diversity and allyship needed to enact cultural change in the tech world, today's students will spread that cultural change as they enter the workforce, bringing with them more diversity and more equitable outcomes.

Both young and old will need to change their expectations about work to prepare for automation's shake up of the world of work. Teaching interdisciplinarity, creativity, and the value of careful, high quality work is the best way of preparing students for the future of work. They will need to be prepared to find fulfillment outside of their careers, since available working hours will be greatly reduced.

We can redirect AI toward a future where data is given consensually to be used in evidence-based decision-making toward goals that the public values, where less powerful tech companies contribute tax revenues and pensions to subsidize minimum incomes, and where fewer working hours are used more efficiently.