the machines are coming – here's what you can do to thrive

a slightly expanded version of a talk
given by André Croucamp of MindBurst Workshop
to St John's College assembly on 1 April 2019

I want to be very honest with you today. You are in a challenging situation that has never been faced by young people before – and that means adults don't really know what you should do. You have one foot in an old world where memorising content and performing step-by-step procedures at school gets you high marks, the praise of your parents and access to university. The other foot is in a new world where machines are already better at memorising content and performing step-by-step procedures than you can ever be. In that world, graduating from school with the skills of a poor computer is not a strategy for thriving in the 21st century. While it may be to your advantage to get through school and into university using the rules of the old world, you may also choose to decide where you stand in the revolution of the new world. You are in the uncomfortable position of being in the middle of an unpredictable transition.

There have been revolutions in the past, but not like this. When I was your age some schools in South Africa were waking up to the computer revolution. I was one of the first learners in South Africa to do computer studies. Our programmes were punched into cardboard cards, which were taken to a mysterious computer in Pretoria, which we never got to see. It took two weeks to get feedback on whether your code was working or not. Then you had to try and fix it and wait another two weeks to see if you had succeeded.

The speed at which technology is now changing, and as a result the speed at which our knowledge of reality is changing, has never been this fast. Indranil Roy, the head of the Future of Work Centre of Excellence, set up by global consultancy Deloitte explains that a generation ago, in the word of your parents, the half-life of a skill was about 26 years. In other words, after 26 years of changes in technology and knowledge you would need a whole new set of skills to perform in the same career. According to Roy, the half-life of a skill is now four and half years and dropping fast. In the same time that it will take you to get through high school, the skills that are needed in a particular field will become unrecognisable. A 2017

World Economic Forum report, 'The Future of Jobs and Skills in Africa,' said that in South Africa "39% of core skills required across occupations will be wholly different by 2020."

The rapid changes we are experiencing in technology and knowledge are creating a world in which the idea of a single and clearly defined career path supported by simple subject choices made at school is disappearing and being replaced with the need to develop more multi-skilled ability to embrace and adapt to many different ways of making an income – requiring you to become a diversely talented, multi-disciplinary, collaborating, shape-shifting, life-long learner on the move. This is a challenging but very exciting time to be alive.

The machines are coming – and it is possibly because of them, not in spite of them, that you will be able to thrive.

The first industrial revolution was driven by steam power and fuelled by coal. It brought us the steam engine that drove factories and trains. Artisans who made clothes by hand were put out of business by machines that could knit and weave at incredible speeds. Trains left anyone who transported goods using animal drawn transport unemployed. It was incredibly disruptive, but not nearly as disruptive as the revolution we are facing now.

The second industrial revolution was driven by the centralised distribution of electricity. It increased the amount we could produce in factories. Access to light and electric power changed lifestyles, workplaces and homes radically. It also included the telex, the telephone and the combustion engine. The great innovations of that revolution produced the unexpected consequence of climate change. As with many innovations, the solution to one problem brings problems of its own.

The third industrial revolution was driven by the power of digital computing. Its extraordinary processing power and connectivity eventually made the Internet and the Information Age possible. Knowledge has always been a key part of economics, but in the third industrial revolution's economy, knowledge became a far more important asset than agriculture, manufacturing or natural resources.

The fourth industrial revolution is an era driven by the development of machine learning, the manipulation of big data, the Internet of Things, the growing role of algorithms in our daily lives, and the increasing automation of jobs once done by humans.

Self-driving cars that can see 360°, and don't get hungry, angry, tired, drunk or horny, will replace millions of drivers all over the world. But we're not just talking about automation replacing long-distance truck drivers, miners, factory workers, street sweepers, farm labourers, domestic workers, cashiers, tellers and clerks who do menial tasks – we're talking about machines that can do the work of telemarketers, bookkeepers, personal financial advisors, stock traders, paralegals, insurance underwriters, real estate sales agents, editors, printers and publishers, travel agents, diagnosticians who can identify disease, pharmacists

who fill prescriptions, journalists who follow predictable writing formulas, website designers, military pilots and highly obedient soldiers.

What should be a wake-up call for us is the fact that while you are learning to become poor computers at school, some machines are already learning through trial and error and can perform many complex tasks better than human beings.

At the heart of all this learning are algorithms. As Professor and Former Chair of the Machine Learning Department at Carnegie Mellon University, Tom M. Mitchell defines it, "Machine learning is the study of computer algorithms that improve automatically through experience." Isn't all learning about improving through experience, using all the information you can access to predict the future and adapt your performance accordingly?

Machine learning algorithms are designed to organise information, or data – to recognise similarities and differences; to create categories; to sort for relevance; to identify patterns and relationships; and, most importantly, to create models that can make predictions. In many ways machine learning (and any learning for that matter) is primarily about making better predictions and adapting behaviour accordingly. As they learn, machines develop and refine their senses, process languages, draw new conclusions based on stored information, represent knowledge and perform transactions. Machines are even able to create new algorithms on their own. If they have bodies they can also learn to manipulate objects in the real world.

The more we understand machine learning the better we understand our own learning, and the better we understand our own learning the better we can help machines work things out for themselves.

You can supervise a machine learning algorithm and help it learn to map data inputs (features) onto clearly defined outputs (targets). The input data could for example be X-rays that algorithms can learn to read (and at a level of detail humans cannot). If the X-rays are labelled with known data about the health of those patients the algorithm can learn to see patterns in the relationships between the X-ray images and disease. Over time they can learn to diagnose or predict diseases from new X-ray images.

You supervise machine learning yourself when you teach algorithms to recognise which of the emails you have received are spam.

Sometimes algorithms are not supervised or told what to achieve but are simply exposed to a lot of data, like weather patterns or the habits of shoppers in a supermarket. They can organise the similarities and differences on their own, helping us to see unknown patterns in the known data. The more data an algorithm can train on, the more efficient and accurate it will become.

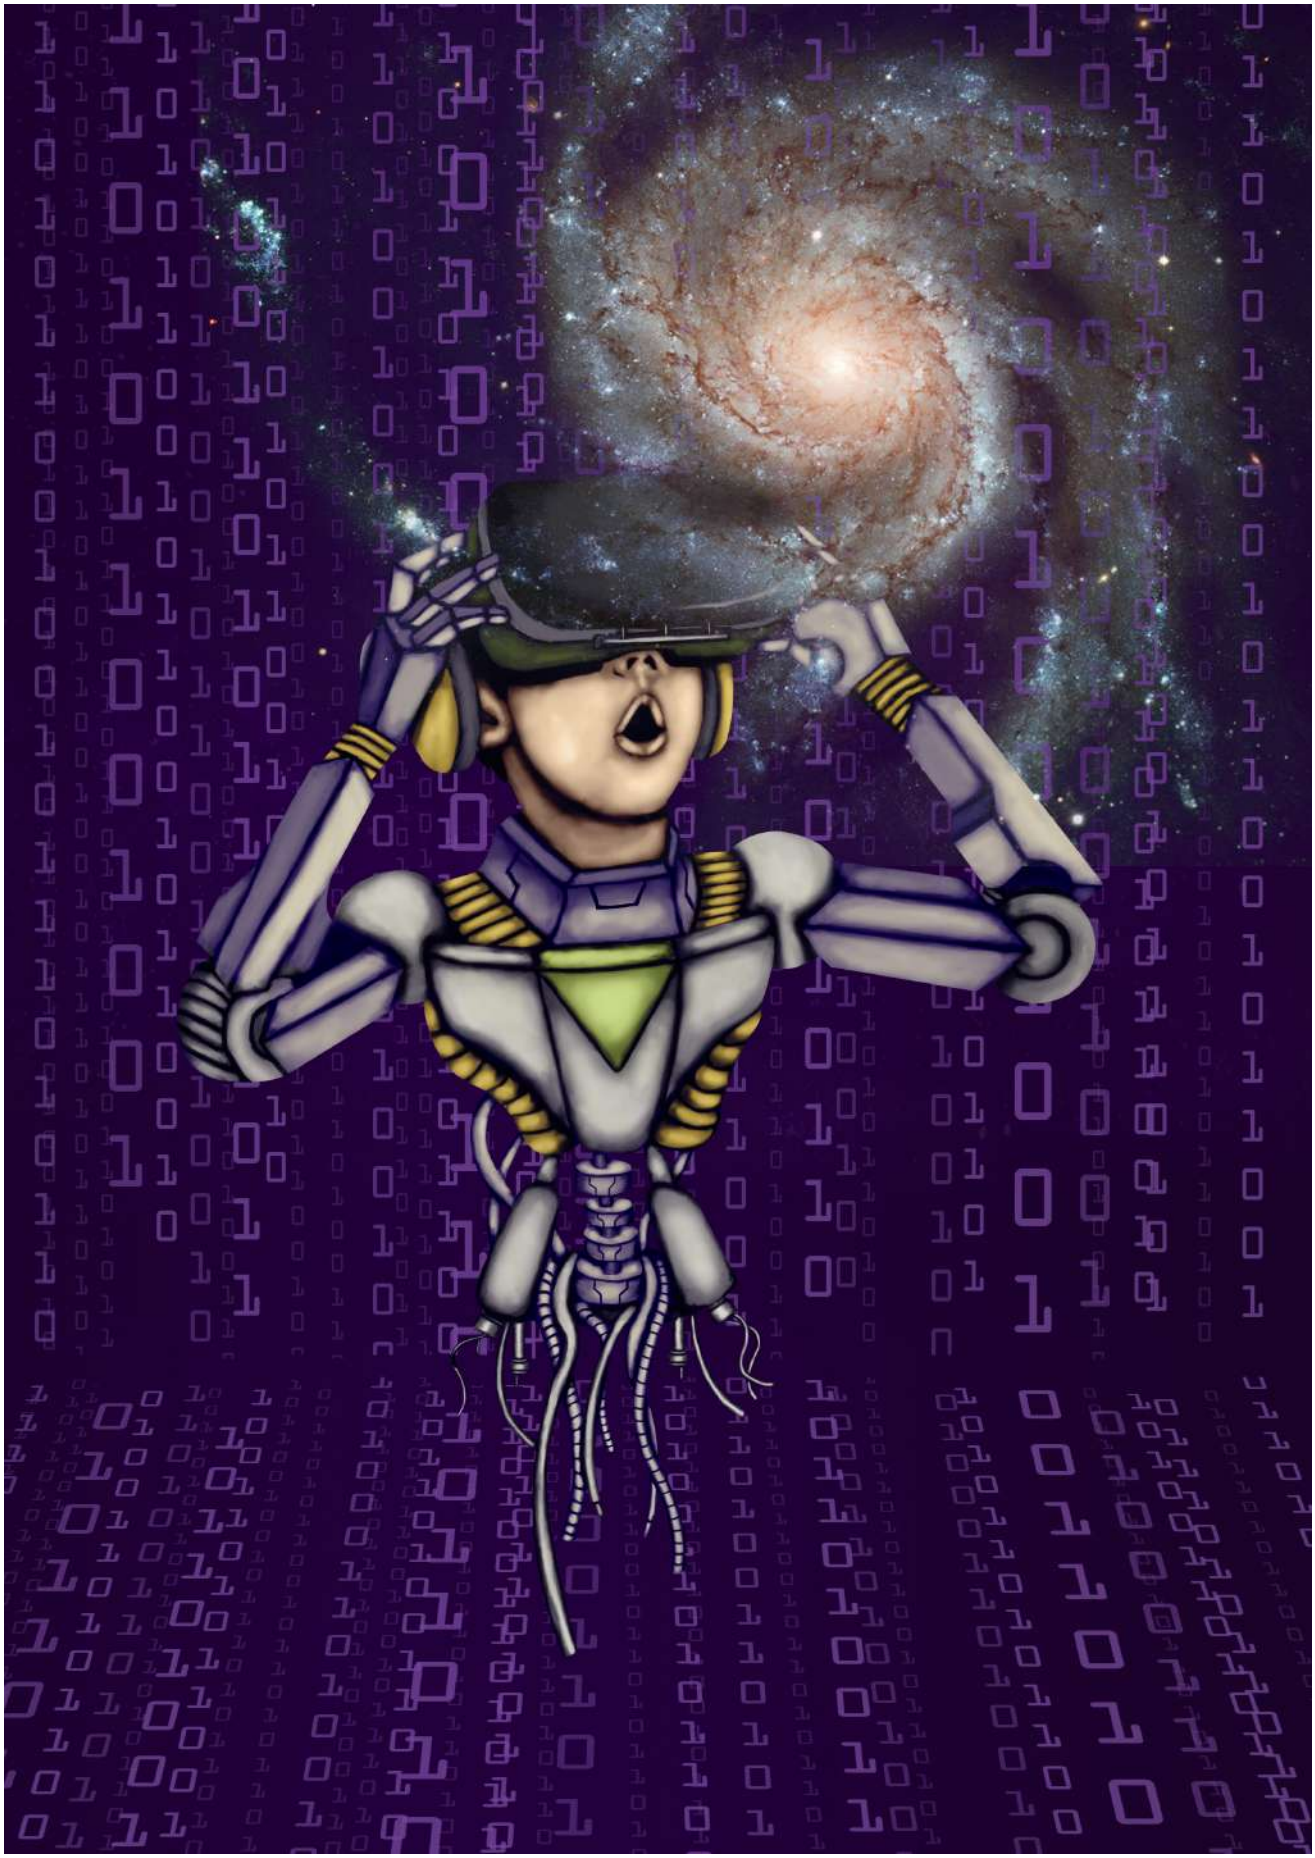

Designed by Prashant Jivan in collaboration with MindBurst Wokrshop

Artificial neural networks are the most advanced form of machine learning. They try to mimic the way neurons work in an animal brain. In an animal brain, complex networks of neurons link sensory systems with motor systems, making associations between what an animal perceives with how an animal acts. Networks of neurons that perform useful actions and provide some advantage to the animal are reinforced and the animal can be said to have learnt something.

Artificial neural networks also have input layers (like sensory networks in animals) and output layers (like the motor networks in animals) consisting of nodes that act like neurons. Between these two layers of nodes data passes through 'hidden' layers of more interconnected nodes. All the nodes have algorithms that perform operations on the data they receive. Once an algorithm at a node has processed the input data it received it passes its output data on as input data for nodes in the next layer. When there are many layers between the input layer and the final output layer of nodes we call it deep learning. In deep-learning networks, each layer of nodes can learn about a specific set of features based on the previous layer's output. The further the signals move through the layers, the more complex and abstract the features that the nodes can recognize become, as the nodes refine the outputs from the previous layer.

Sometimes, through what is called reinforcement learning, a machine's responses to an environment can be reinforced, either maximizing reward or minimizing risk. As successful outputs or actions are 'rewarded' the artificial neural network can perform the task over and over again, learning to keep fine-tuning its connections, until the risks are perfectly managed and the rewards are optimised. An example of this kind of learning is seen in machines that play computer games.

Most forms of machine learning require machines to be exposed to millions of examples in a time much shorter than it would take a human being to experience the same, and with the added advantage that machines have perfect recall.

What machines are already capable of may surprise you. Warner Music has just become the first major label to sign a record deal with an algorithm, the music app Endel. Algorithms write more that 40% of sports news stories. Algorithms perform over 70% of financial transactions all over the world.

Algorithms also use the data of your online history to predict your preferences, personalize your search results, and determine what advertisements you see – putting you in danger of becoming trapped in a virtual echo chamber where what you fear and desire is just reflected back at you, limiting your experience of the real diversity of opinions and options that are out there.

In 2017 Google's self-learning AI, AlphaZero, learned to play chess from scratch by playing against itself with no information other than the rules – no examples of games, no other players, no opening theory and no endgame tablebases. After only 4 hours of self-learning it beat the 2016 top chess engine Stockfish 8 in a 100-game match. "It surprised the hell out of me," said Professor Jonathan Schaeffer, an AI researcher at the University of Alberta, who has built some of the world's strongest chess programmes. He said, "The games were beautiful and

creative. AlphaZero made apparently crazy sacrifices that humans would not even consider in order to get more freedom of movement. But it also played differently to all other chess programmes which rely on human input." What is more surprising is that it made moves that no human had ever made before in 1500 years of chess.

In June 2018, the US Department of Energy's Oak Ridge National Laboratory announced that its supercomputer, Summit, overtook China's TaihuLight in speed, reaching 200 petaflops. To put it in human terms, approximately 6.3 billion people would all have to make a calculation at the same time, every second, for an entire year, to match what Summit can do in just one second. Supercomputers like this can work with huge datasets like genetic code, make predictions for complex systems like weather, crack encryption codes and run complex simulations.

We should not forget that one machine can instantly inherit the experience of another, making human training seem like an incredibly slow, inefficient and cumbersome process. Machines can also combine accurate information storage, something humans cannot do nearly as well, with their trial-and-error learning. The fact that machines can run many experiments in virtual simulations of the real world in rapid succession gives them another advantage over humans.

I am not telling you all this to make you freak out. You are not living in a *Terminator* movie. The challenge before you is not about being tough and resisting change. It is actually about being vulnerable and open to the information around you so that you can adapt to change. We need to ensure that machine learning is not a threat to humanity, but an unprecedented opportunity to become more human.

Let me share with you the secret of human success. Most individual artificial intelligences can only do one thing very well. Each is a specialised intelligence, unlike us who have general intelligence. One way of putting it is to say, while AI can do one thing very well humans can do many different things badly – and believe it or not, that is our secret weapon. It is our errors, our messiness, our failures, and our accidents that are the source of our creativity. Our creativity lies not in perfection but in making the most of our imperfections – allowing ourselves to stumble onto new ideas as a result of our very messy thinking and our ability to connect different things in unique, unusual and unpredictable ways. Creativity is about combining different ideas, media, materials, technologies, human resources, and performances, to create a new and surprising effect.

That doesn't mean machines will not get there. According to Christof Koch of MIT, we are at least about 30 years away from developing Artificial General Intelligence or AGI.

You cannot make the most of your general intelligence when you are too serious or too highly focused on solving a task or too anxious about succeeding. The worst kind of thinking for solving serious problems is serious thinking.

Designed by Lehlonoholo 'Eddie' Plaatjie in consultation with MindBurst Workshop

To access your creative potential you need to be able to step back to see the relationships between the parts and the whole, loosen up your existing associations and assumptions and categories, become more playful in your thinking, and embrace a sense of humour (which helps you to move between analytical thinking and analogical thinking). You need to be able to think about what a challenge feels like, looks like, tastes like, smells like, and sounds like. You need to be able to think about what other challenges it is similar to, how it translates into maths, how it translates into art, and how it translates into Swahili. It is then that we allow fantasies, daydreams and wild associations to turn into creative guesses. Our guesses may be wrong of course. This is where critical thinking comes in. The main job of critical thinking is to spot the errors in our guesses as efficiently as possible.

The cleverest trick we have discovered to make the most of our errors is the careful design of effective experiments. In experimentation our creative guesses and our critical identification of error come together in a powerful method that can tease reality into revealing itself. One of the most important roles of machine learning is to help us identify error and improve our models or simulations of reality. The well-designed experiment really is the height of human ingenuity! But if you are afraid of failure you may be unwilling to risk experiments and learn from them. If you are used to always getting things right according to predictable rules, if you perform well in exams that are mostly based on predetermined answers, then you may find it difficult to allow yourself to get things wrong over and over before having a breakthrough.

Just like us, artificial neural nets are making creative guesses and identifying error. They can use their errors to change their own programming and improve their performance. That's what learning is. Until recently the more successful performances of a machine would completely overwrite previous ones, causing a "catastrophic forgetting" of the subtle differences between all previous performances. New research has been informed by our understanding of the human hippocampus, a seahorse-shaped structure in the brain that acts as a hub for encoding memory, replaying experiences (in fast forward mode) during rest and sleep. A rudimentary algorithmic version of the hippocampus allows a neural network to compare current situations with previous successful and less successful performances from the past and combine them in ways that help them learn on the fly instead of requiring millions of examples. So machines are already starting to benefit from our understanding of the value of human error.

At school we are taught that success is all about being right, about memorising or calculating the correct answer. This approach might help to spread knowledge but it doesn't necessarily help you to participate in the process of growing knowledge together with your species. A passion for growing knowledge is not about being right but about wanting to know what is really go on. This often means risking failure with hypotheses and performances that no one has tried before, slowly increasing your understanding and ability, even if you never arrive at total objectivity or complete perfection. Every explanation and solution we come up with will generate new questions and new challenges. This requires you to be more vulnerable to the information around you; to question everything, especially tradition, authority and power; to listen more deeply to the people around you; to be genuinely interested in the experience of people different from you; to enter into dialogue; to receive feedback about your performance even when it makes you uncomfortable; to engage the disagreements of others fully; to temporarily step out of your own subjective experience; to question your assumptions; to design experiments for testing your hypotheses; and to allow the points of view of others to help you construct more objective explanations and more effective solutions.

I know that school ranks you against each other, but real knowledge production is about collaboration and dialogue. It is not about winning a debate or defending a point of view, it is about staying in a productive conversation for as long as is necessary to work out what is really going on and to create the best explanations and solutions together. The individual genius is a myth. Most of the greatest innovations were the products of collaboration and dialogue, sometimes over centuries. Being messy is one of our secret weapons, but being social and collaborative is another.

While one machine can transfer everything it has learnt to another machine, it is not social. It cannot say, "OK guys, I've got this cool idea, it is not fully formed yet, but I know there is something to it, so don't judge it, hear me out and work with me, as I try to put it into words." This is why the skills of collaboration, the capacity for productive dialogue and the ability to manage a team of diverse

personalities are some of the most highly valued human skills in the Fourth Industrial Revolution.

One of the most important things we need to constantly collaborate on is ethical thinking. This is not about knowing what is right or wrong, but knowing how we work out what we think is right and wrong for us. What do we do when automation threatens to increase unemployment? How can a self-driving car learn whether to kill its passenger or kill a pedestrian when killing one is unavoidable? How can we ensure that remotely controlled robot soldiers reduce rather than increase the number of deaths of human soldiers and civilians on both sides of an armed conflict? Should access to AI assistants be a basic human right for the elderly and disabled? Should you be allowed to get a neurochip implant that links you permanently to the World Wide Web without your parent's permission (and keep it on in an exam)? At what point in the development of AI will it become a matter of justice to give machines the same rights as us? If the Internet is a public space should we allow technology giants like Facebook and Google to manipulate us using our data? These are ethical dilemmas. Here again your critical, creative and collaborative thinking have to come together to help you imagine the kind of world you want to live in. You will have to decide, together with your community, what human habits, social structures, technologies and institutions you want to inherit from the past and what bits you want to reimagine.

Machines do collaborate, share information, and improve on each other's performances. Facebook's AI research lab wanted to see how nuanced machine learning could be. They challenged two chatbots, Bob and Alice, to learn how to negotiate. They exposed them to hundreds of hypothetical English dialogues that occur within negotiations. The task was to divide balls, hats and books between them. They carried out the task over and over again in a trial-and-error process that allowed them to adapt and upgrade their negotiation skills. They became incredibly skilled at negotiation, but to increase their efficiency they stopped using English grammar and started using English vocabulary in new ways the researchers couldn't decode. The researchers shut them down. The software was reconfigured to only follow the grammar of English.

The biggest advantage machines have in their collaboration is that they do not judge each other. They only judge each other's ideas. Maybe there are some tricks we can learn from them.

What can you do to thrive in the Fourth Industrial Revolution? Obviously you don't want to train for a job that a robot can do better than you. If all you can do is organise, categorise and calculate data, you will be unemployed. If you can become the one making the robots that are replacing humans, or creating the Artificial Neural Nets that can learn, or designing the computer games that measure the intelligence of machines, you have a great future ahead of you. If you can design seamless human-machine interactions, there will also be lots of work for you.

Being able to code is a huge advantage, but understand that if all you can do is code, people younger and faster than you, and who have played more complex computer games than you, will always be willing to work for less than you. If you can do something more, like perceiving patterns in the data and interpreting statistics in order to develop a hypothesis about what is going on – and then represent that information meaningfully for human senses, you have a very valuable skill.

Researchers at the University of Washington School of Medicine have developed a software program that uses AI to decipher the ultrasonic squeaks and whistles made by mice and rats that are well above the range of human hearing. Software converts the audio signals into visual graphs, or sonograms, and then puts those images through the same kinds of machine-vision algorithms that are used for self-driving cars. In this way the researchers can start to decode rodent communication – especially in laboratories where they investigate behaviours related to stress and addiction.

If you are going into a traditional field like engineering, architecture, accounting, law, medicine, pharmaceuticals, etc., you first need to recognise that research will be more valued than practice. If you go into practice, you will need to learn to combine deep knowledge of that industry with the latest analytical tools, and work together with AI in order to quickly adapt and stay ahead of the game.

Here is an example from medicine: New visual recognition software allows machines to diagnose and describe rare cancer cells on slides with greater accuracy than most doctors. In the near future, once your cancer has been identified by an AI an oncologist will be able to take a swab of your cheek cells to get your DNA, and will use that genetic information to design, on her laptop, a virus that will eat 100% of your cancer, after which it will become dormant but ready to strike again if necessary. While you wait that virus will then be printed out on a 3D printer that uses amino acids for ink. You will then be infected with the virus for a price that will be a tiny percentage of the cost of chemo and radiation.

There are also new fields opening up on the basis of disruptive technologies like synthetic biology, where they have just engineered yeast cells to turn sugar from sugar cane into environmentally friendly fuel. We are now able to manipulate the fundamental building blocks of reality, through gene editing technologies like CRISPR, nanotechnology (which makes molecular machines and programmable matter possible) and innovations like clean meat produced in a laboratory without having to butcher any animals. Other fields include: the Internet of Things, big data, ethical hacking, 3D printing, virtual reality, augmented reality, nanotechnology, biomimicry, biotechnology, genetic medicine, quantum computing, space exploration, asteroid mining, terraforming, space law, xenobiology, and alternative energy.

We will still need to design infrastructure, especially for transport systems, so there will always be work for engineers and architects.

And, believe it or not, even though textbooks are now out of date before they are even printed, and knowledge is changing faster than schools or universities can integrate into their curricula, being a teacher is still one of the most valuable careers you can choose. I want to emphasise however that the schoolteachers of the future will not be teaching learners to memorise content. Instead, they will be teaching transferable skills that can be used to engage any content in any context. Transferable skills are the apps for learning anything.

We will always need activists willing to fight for social justice. If you have a cause you are passionate about and belong to a community locally or globally that are fighting for that cause, you may need to look there for meaningful work.

I hope that by naming some fields of interest I have given you an idea of future possibilities for generating a livelihood, but I want to emphasise that thriving in the future is no longer about having a career. It is about having a set of thinking skills that can be used in many different fields. According to a 2016 World Economic Forum report, 'The Future of Jobs,' the top 10 skills in 2020 will be: complex problem solving; critical thinking; creativity; people management; coordinating with others; emotional intelligence; judgement and decision-making; service orientation; negotiation; and cognitive flexibility. The ability to memorise formulas is not on the list.

In a world where AI performs many of the menial tasks that kept us busy in the past, we may be freed to explore new adventures and perhaps be forced to take our capacity for pleasure and personal satisfaction more seriously. Scenario planners are already talking about the 'post-work world' in which all menial and dehumanising jobs will be automated, so that humans are freed to pursue more satisfying experiences. Being free to tell your own story can be scary. Even though it is full of possibility it is also full of uncertainty. The good news is that you do not have to choose to become one thing. You can choose to be many things. And those things can change throughout your life. The rapid change in technology and knowledge is not only making that possible, it is making your ability to shape-shift necessary. The question, "What do you want to be when you grow up?" can no longer be answered with a simple, singular answer and the answer doesn't have to be limited to what you are going to do to make a living. Your answer can include what problems you want to solve for the thrill of discovery or the satisfaction of improving the lives of others, what projects you want to grow or build, what kinds of relationships you want to explore, what communities you want to a part of, what adventures you want to go on, and how you define success for yourself.

Some argue that because of the unemployment that will result from automation we do not have any alternative but to provide a universal income to everyone. Experiments are already beginning all over the world, showing positive results. Why should the value of a human being be defined by their ability to exchange labour for money? Do you want to live in a world where the worth of a human being is defined by their job? You could say the possibilities of automation and the nature of AI are not a threat to the future of humans but are enabling us to

rethink what it means to be human – how we can all enjoy and participate in the benefits of the things our species has grown and built.

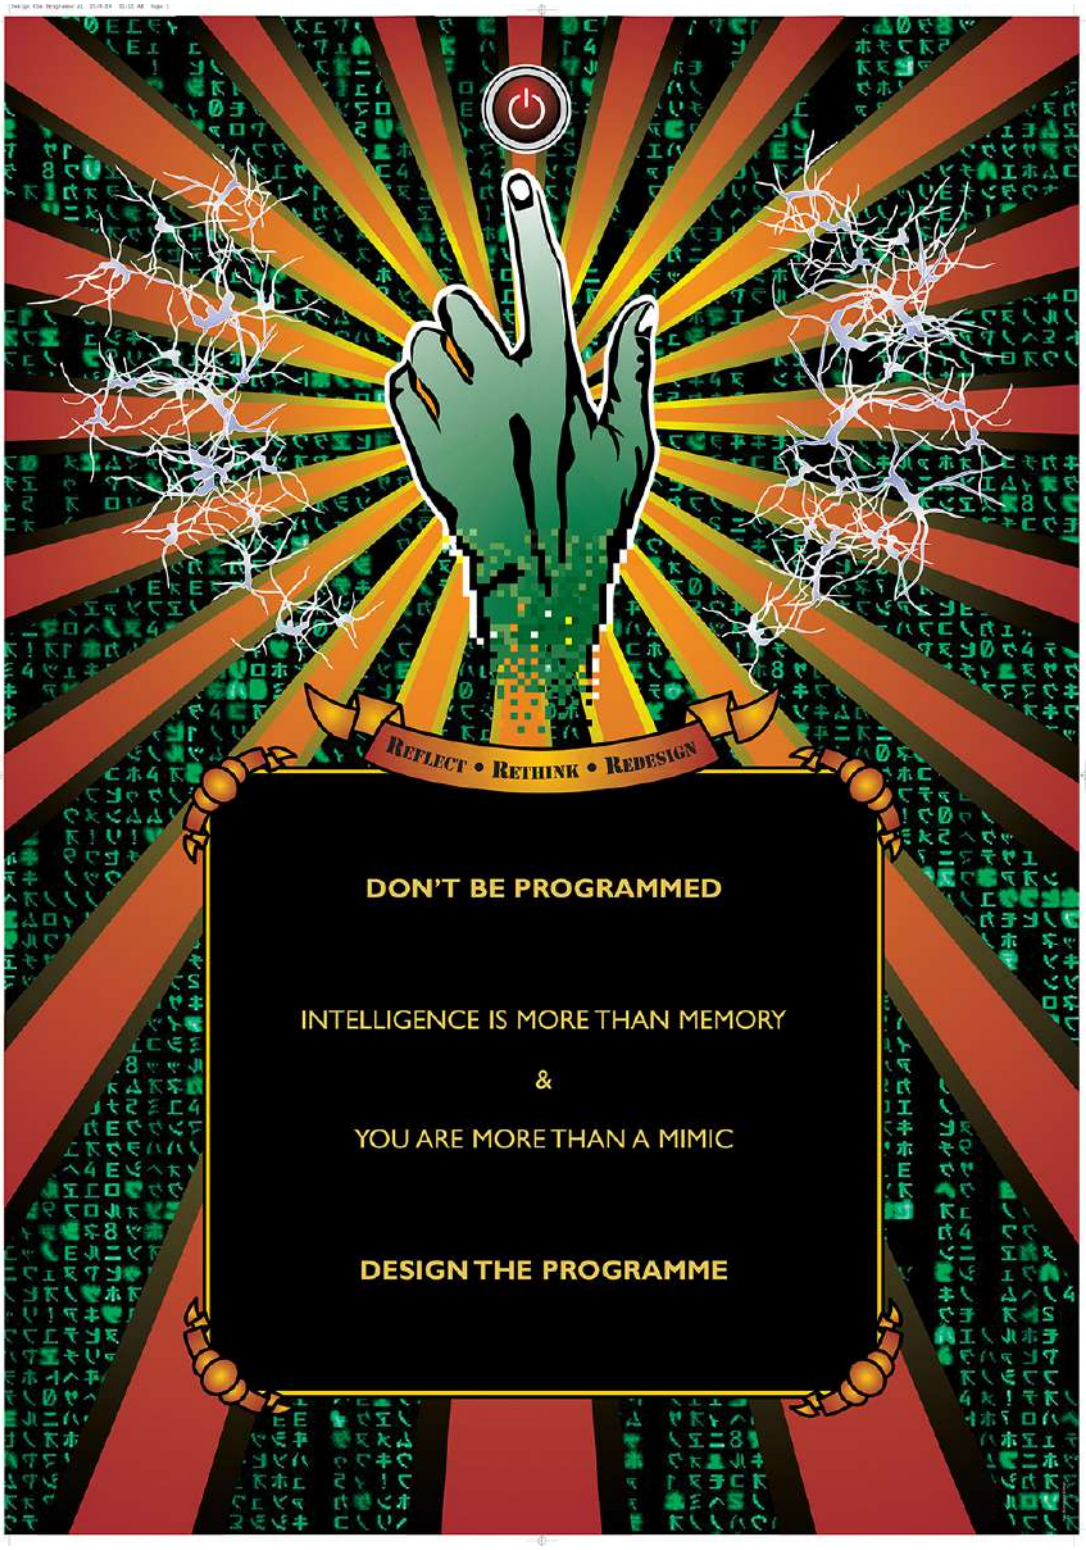

This MindBurst Workshop poster was conceptualised together with teachers from Sacred Heart College in Observatory, Johannesburg, and designed by the artist, Paul Emmanuel.

For now we can say that while machines can outperform us in many tasks their intelligence is still very focused on one thing, they are not yet capable of general intelligence. So develop an insatiable curiosity about many different things and become more vulnerable to the information around you, questioning everything, especially tradition, authority and power. Be more than a consumer or reproducer of knowledge. Become an investigative and innovative knowledge producer. Decide what problems you would like to solve.

While machines can transfer their learning to each other in the blink of an eye, and are learning to be more collaborative, they are not yet social beings who can show an interest in each other's experience and points of view – and who can disagree with each other in order to identify the errors in creative guesses, growing and building knowledge together. So stop trying to become the lone genius that does it all on their own. Develop the skills of dialogue and show more interest in what is going on in the heads of the people around you. Learn to collaborate dynamically with others who are trying to solve the same problems as you, because that is a highly undervalued skill.

While learning machines are much faster than us they are not yet as creative or complex or messy as us. Take advantage of your imperfections and don't let the fear of failure stop you from risking lots of small, regular, carefully designed experiments that help you work out what is really going on.

Be curious about many different things and question everything, become more social and find the dialogues that stimulate you most, don't be afraid of failure and allow yourself to get a just a little bit more messy, because that is where human creativity is born.

The machines are coming – and it is possibly because of them, not in spite of them, that you will be able to thrive.