Chinese University

BULLETIN

No.1, 2021

Chinese University
BULLETIN

Contents

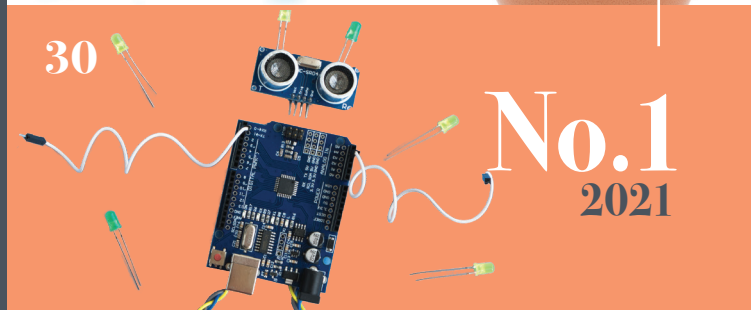

The *Chinese University Bulletin* is published biannually by the Information Services Office

Address all correspondence to
Information Services Office, CUHK,
Sha Tin, N.T., Hong Kong SAR,
The People's Republic of China

E-mail
iso@cuhk.edu.hk

Website
www.iso.cuhk.edu.hk/english/pub/bulletin

The New Gospel According to A.I.

I must have been six when I first heard about AI. Unsurprisingly, that knowledge came from **Steven Spielberg**'s rather too obviously named treatment of the subject, *A.I.* I did not remember much of the film, save for the part where the protagonist, a robot boy named David, goes to the bottom of a submerged Manhattan and prays to a mossy statue of the Blue Fairy from *Pinocchio* for the next 2,000 years to be made human and thus eligible for the love of its owner. Like the anti-tech mob that it wins over with its endearing realism earlier in the story, I could not help but feel for it, seeing perhaps also something of my child self in it.

But I knew as soon as the film ended that I had taken it too seriously. When the DVD player, being already the pinnacle of human technology in the reality of the early 2000s, sluggishly spat out the disc, I knew it was just a fantasy. The invention of sentient robots, the ethical short circuits they create, a full-on machine takeover—all these were too remote a prospect at the time and, as I now suspect, are probably allegories of some of humanity's more familiar problems.

Twenty years on, these scenarios remain in the realm of science fiction. This is not to say, however, that things are as they were 20 years ago. Turn on your TV, and you are guaranteed to hear at least one mention of AI every 15 minutes, be it in a commercial for the latest express loan services or an ad for some electrical appliance that you would never have dreamt of having use for a brain. The world of AI as painted by Spielberg is still in the imagination, but this imagination is now seen as coming closer to materializing and is gaining a foothold in the public discourse, alongside much hype and unease.

Many might have forgotten we were in a similar position with the coming of IT, having an equally frantic conversation, exactly 20 years ago. Not long before Spielberg released his film, the *Chinese University Bulletin* ran an issue titled 'IT: The Name of the Game'. With another shift in our way of life at hand, the *Bulletin* spoke to over a dozen AI researchers across all fields at CUHK in that same spirit of cutting through the noise and foregrounding what is truly at stake. This is that story of the advancements in science, medicine and business we have made with AI, a story of what it means to society, the arts, law and ethics, and education—a story, ultimately, of what it now means to be human.

BY
jasonyuen@cuhkcontents

ADDITIONAL REPORTING BY
ronaldluk@cuhkcontents

COVER AND
LAYOUT DESIGN BY
amytam@cuhkimages

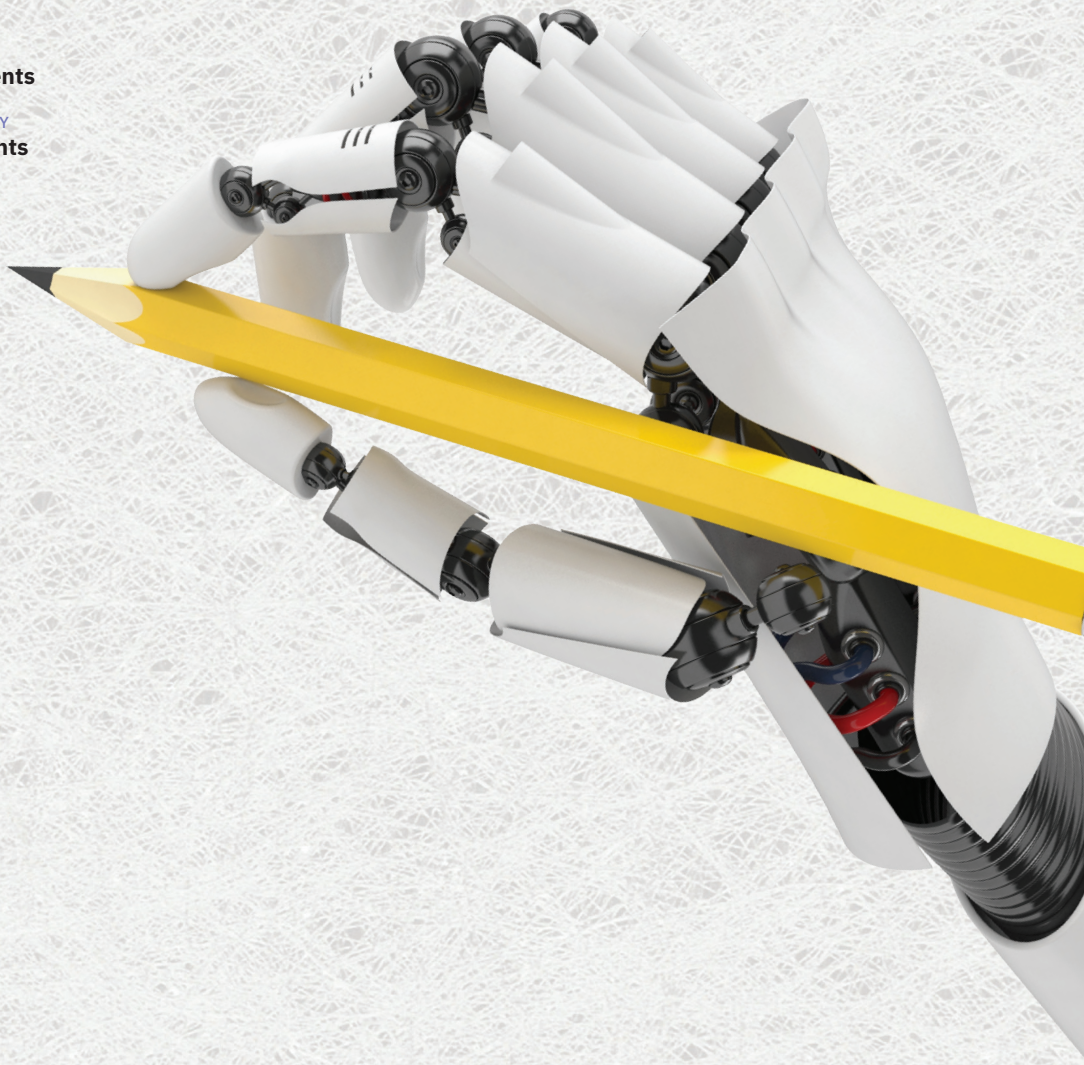

Open Sesame

With AI comes the promise of unlocking the treasure that is big data for social good. There are a few catches, though.


AI: Society

8:30 a.m., Monday. For the third time in a week, the train is held up in the middle of the tunnel. Try as you might to be on time, there seems always to be a way to undo the head start you have made. Trapped like the rest of the hundreds of commuters on board, heavy-eyed, you think there is nothing you can do except, maybe, tweet furiously about it:

'Well don, 3rd delay this week. You ppl at the railways had ONE job.'

Besides being a way to let off some steam, this tweet may seem inconsequential. But for **Luo Shuli**, ramblings on social media like this can make a difference and lead to better transport—given the right tool to process and understand them.

A PHD CANDIDATE at the Department of Geography and Resource Management, Shuli has been using social media data to better understand the perception of a city's public transport system. For her dissertation research, which won her a Best Paper award at the Second International Conference on Urban Informatics, she focused on Shenzhen and collected tens of thousands of public Weibo posts about the city's metro system. What she was interested in were the sentiments they convey and such metadata as when and where they were published, which might point to areas of the system that needed improvement.

Prof. Sylvia He
Department of Geography and Resource Management

Luo Shuli
PhD candidate, Department of Geography and Resource Management

In the US, social media data has been used to detect incidents in public transport systems. AI can help analyze and verify incident reports on social media, based on which operators can make responses and alert passengers

'A recent trend in urban planning is to complement conventional surveys with big data,' said Prof. **Sylvia He**, Shuli's advisor and co-author of the prize-winning paper. While allowing for more comprehensive social demographics to be collected and for customization, surveys are costly and can only be conducted infrequently. Meanwhile, a city-wide Internet of Things (IoT)—made possible by the explosion of smart phones, GPS-enabled vehicles and, of course, social media users in recent years—provides a steady, voluminous stream of data capable of revealing a great deal about the population's travel behaviour.

'A hallmark of smart cities is the widespread use of energy-efficient, zero-emissions electric vehicles, and to promote e-mobility, we need more charging stations. But where should we build them? This is where big data can fill us in,' Professor He explained, giving another example of how big data is used in urban planning. Whereas normally researchers would have to depend on surveys, Professor He's team is now exploring a method of using raw data from existing stations to find out where, when and for how long electric vehicles are more likely to be serviced, thereby ensuring a more reasonable distribution of charging facilities across the city.

As many insights as it may contain, though, there is no way we mere mortals can sift through such a massive amount of data. We need a program that, having been taught the rules, automates the task. Better yet, we can have an AI program, which can figure out the rules without needing us to write them out for it. Indeed, when it comes to social media posts, which are often rife with typos, shorthands and other irregularities like the tweet we have seen at the outset, AI is the clear winner in that it spares us from having to teach the program to recognize the numerous special

cases there are. All it needs is a good amount of training, through which it can learn the rules and exceptions from samples we provide it with.

'Efficiency and personalization are what usually motivate the use of AI in public administration,' said Prof. **Wilson Wong**, director of the Data Science and Policy Studies programme. Aside from unlocking the wealth of data around us, an intelligent automated system can help respond to the many different needs of citizens around the clock. In Japan, for example, chatbots have been employed to give more individualized and accurate information on government services. From e-government to e-governance, public goods provision to policymaking—there is much potential for AI to do social good as many have called for lately.

Prof. Wilson Wong
Director of the Data Science and Policy Studies programme

BUT AI, TOO, HAS ITS LIMITATIONS. For one part of her dissertation, Shuli ended up ditching the AI model and went with the classical statistical approach, having compared how they performed in discerning the sentiments in the Weibo stories. It could be that the model needed more data for training, Professor He suspects, but there might be no way of knowing what went wrong. Indeed, many AI models are what computer scientists call black boxes, which is to say they have such an opaque decision-making mechanism that it is virtually impossible to diagnose the errors they make. At any rate, AI has not had much of an edge over traditional methods to begin with in terms of understanding emotions.

'The model can categorize a sentiment as positive or negative pretty decently, but it doesn't tell you *how* positive or negative it is,' said Shuli of her experience performing sentiment analysis using AI. 'And if you ask the model to be more specific and return anything more descriptive than a label that says "positive", "neutral" or "negative", you'll probably get something wildly inaccurate.'

Things get even muddier when you are dealing with circumlocutions, like the sarcasm in our opening tweet. Solutions have been proposed to give AI models an awareness of contexts, as Professor He noted, but for now, machines are often still dependent on human calibration when it runs into this kind of problem. Beyond the understanding of social media parlance, this lack of tacit knowledge is a major reason why AI is not playing a more decisive role in public administration.

To be data-literate, ultimately, is to have the knowledge to use data in a way that improves your life while not being enslaved by technology.

'There are many misconceptions of what AI can do,' said Professor Wong, who has been part of an Association of Pacific Rim Universities (APRU) project exploring AI's capacity for social betterment. 'With less controversial stuff like renewing driver's licenses and handing out consumption vouchers, which are really just matters of verifying the applicant's eligibility, surely AI can be of help. But how much further can it go?' Consider university admissions. On top of academic results, the board will look for certain personal qualities: being principled, willingness to communicate, honesty, and so on. They are not exactly subjective, but they are hard to define, understood only through socialization. If we are to replace human admission officers with machines, the challenge will be for them to understand these qualities in mathematical and logical terms. How are we to create an algorithm for that?

'The same goes for court trials. It's hard to imagine a formula by which a machine can determine if the defendant is remorseful, however fair it might be to have a robot to be the judge.'

And here AI hits another roadblock: it is rarely even impartial. We have seen that the rules by which an AI model makes judgments stem from the samples we chose to train it on. If the samples are biased—as they

often are by their nature of representing only part of the whole truth and by the simple fact that they were selected by humans—so must the model be. With Shenzhen having a predominantly younger population, it is not too problematic to consult an AI model feeding solely on social media data, as Professor He pointed out. In cases like hiring government employees, though, it is probably a bad idea to rely on AI. Using the current workforce as the template, an AI recruiter would miss out on bright minds that do not fit 'the norm', miss the opportunity to shake things up and, more dangerously, inherit whatever discriminatory practices that characterize the organization in its present state.

'With all its unreasonable judgements, AI does, after all, serve to expose everything that's wrong with its teachers, us. Rather than thinking about letting it run our lives, we should take this opportunity to reflect on the human prejudices that have made it the way it is,' said Professor Wong.

AS ENCOURAGING AS IT IS, the fact that there have been calls for the use of AI for social good, including even a movement that got itself the catchy abbreviation 'AI4SG', is a reminder that things can move—and most certainly have—in the wrong direction. It can be a treasure trove that AI is unlocking, but it can also be a Pandora's box. We have been talking about AI in tandem with big data, and we have seen how they enable each other, albeit imperfectly. This symbiosis comes at a price, one that we might have given up caring about: privacy.

'The relationship between data and privacy is forever a contentious one. At one end of the spectrum, you have a society that withholds all its data for privacy's sake and gives up all the benefits we've talked about; at the other end, you have a society that surrenders all its data to the point where even the faintest of facial expressions could, with the right technology, be monitored,' said Professor Wong. With most people going for the middle ground, the idea of data governance has gained momentum over the past few years.

'It's all about creating a mechanism where data users, including the government, can be held accountable,' Professor Wong explained. Ideally, it will be a legal framework regulating the whats, whens, whos and hows of data collection and use.

'By making the use of our data transparent and keeping ourselves informed of what's happening behind the scene, we might find ourselves closer to a symmetry of information and, therefore, power.'

One particular issue with data use in public administration is its scope. While most people are willing to sacrifice some of their data for whatever benefits they are promised, there will always be those that firmly object to any erosion of their privacy. Though it is getting more and more difficult, they remain free not to use social media to keep the hands of Big Tech away from their information. That is, however, not an option in face of an intrusive public policy, which by definition applies to everyone in the community, given the ubiquity of data-driven technologies.

'It's impossible to go completely off the grid, to be quite frank. What we can think about is how we can minimize the impact for these people,' said Professor Wong. 'Some people are really uncomfortable with the ideas of smart cities and IoT, at the thought of a smart refrigerator looking at your snack stash and intervening in your eating habits in the name of health. What we can do is allow opting out as far as possible. With new policies, we can run pilot schemes with those that are more enthusiastic and let the hesitant wait and see.'

At the end of day, though, no institution is perfect. What is perhaps most needed is an understanding of AI and big data at an individual level, a data literacy.

'As I often tell my students, data deleted is not deleted. There are many ways in which data can be recovered, so it's best to think twice before creating it. This is the sort of alertness you get with data literacy,' said Professor Wong. 'To be data-literate, ultimately, is to have the knowledge to use data in a way that improves your life while not being enslaved by technology.'

WITH ITS BLUNDERS AND FLAWS, AI may have a hard time making critical decisions for us; and given the invasion of privacy it enables, it needs more scrutiny, indeed, than it is getting. But no contribution it makes is too small, whether it be adding to traditional statistics when it comes to thinking about a city's transport, or assisting policymakers in distributing public goods; and with a proper regulatory regime and a keen awareness of the power of data among citizens, it can, after all, do good.

'With all that we've said about smart cities, I believe AI can also encourage civic engagement by making the data the individual citizen generates valuable,' Professor He added, reminding us how in the world of AI even a throwaway tweet from a disgruntled commuter can contribute to the smooth functioning of a city.

'I'd like to think it's here to make lives better.'

Orchestrating Airfield Operations with AI

Developed by
Centre of Cyber Logistics, CUHK and ubiZense Ltd.

Commissioned by
Airport Authority Hong Kong

A prototype Internet of Things (IoT)-augmented airfield service system, AS2, has been developed to provide ground handlers with a bird's eye view of aircraft turnarounds at the Hong Kong International Airport (HKIA) in real time. The project was given the Smart Mobility Grand Award at the Hong Kong ICT Awards 2019. Powered by AI, its latest production version promises detailed descriptive and predictive analytics based on data collected from a network of thousands of equipment on the apron.

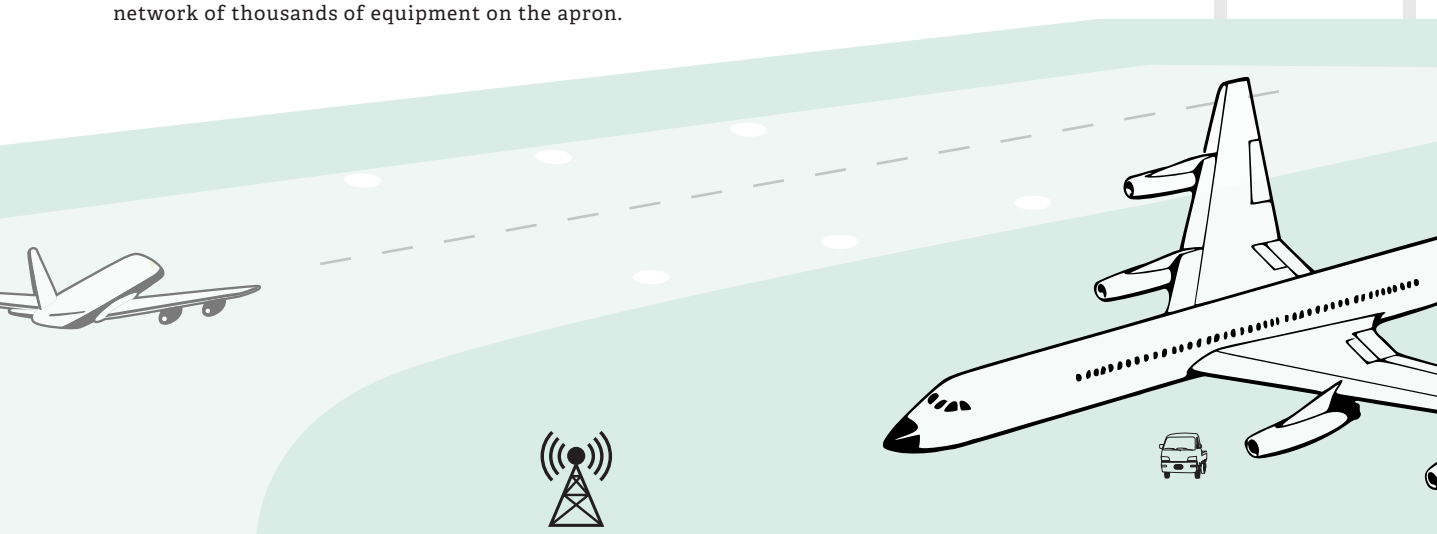

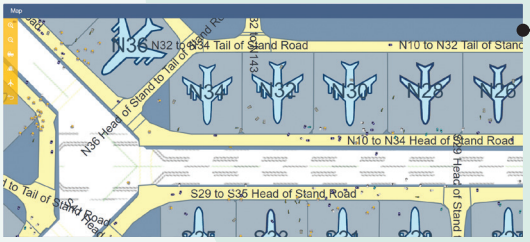

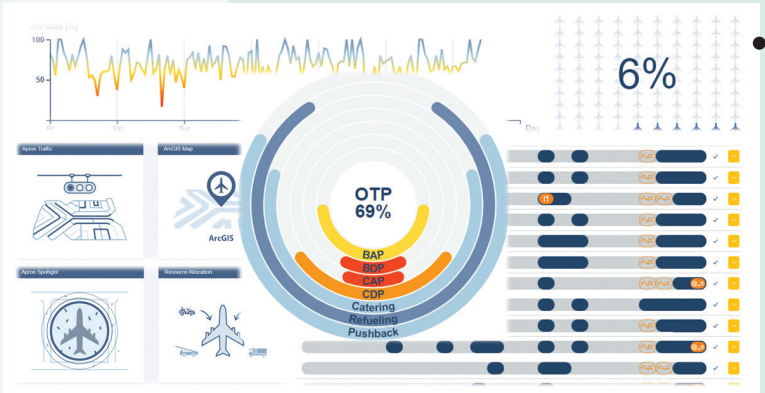

Real-time analytics

Using IoT data from over 6,000 GPS-equipped and/or Bluetooth-ready service vehicles and non-motorized equipment on the apron, flight movement data, and video footage of parking stands and other strategic locations, the AI-powered AS2 can provide in-depth analysis of turn-around processes currently taking place. Warnings are pushed to ramp handlers when such events as delays in baggage delivery are detected.

Predictive analytics

AS2 can also make predictions using the data and alert operators to potential disruptions. This allows for rapid and proactive actions.

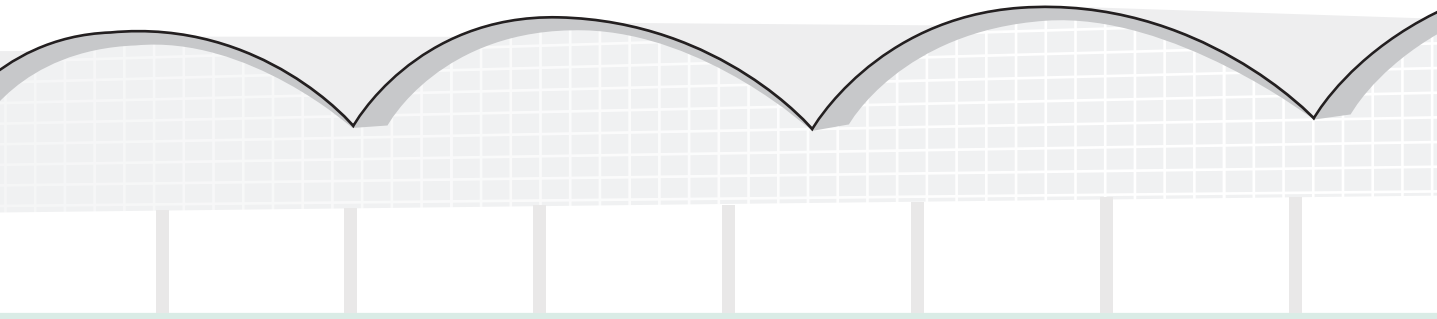

Between landing and take-off
A lot happens after an aircraft lands and before it leaves: baggage and cargo handling, refuelling, catering, and so on. Airfield operators have to juggle with the many ground handling tasks to ensure the smooth alignment of airside and landside operations.

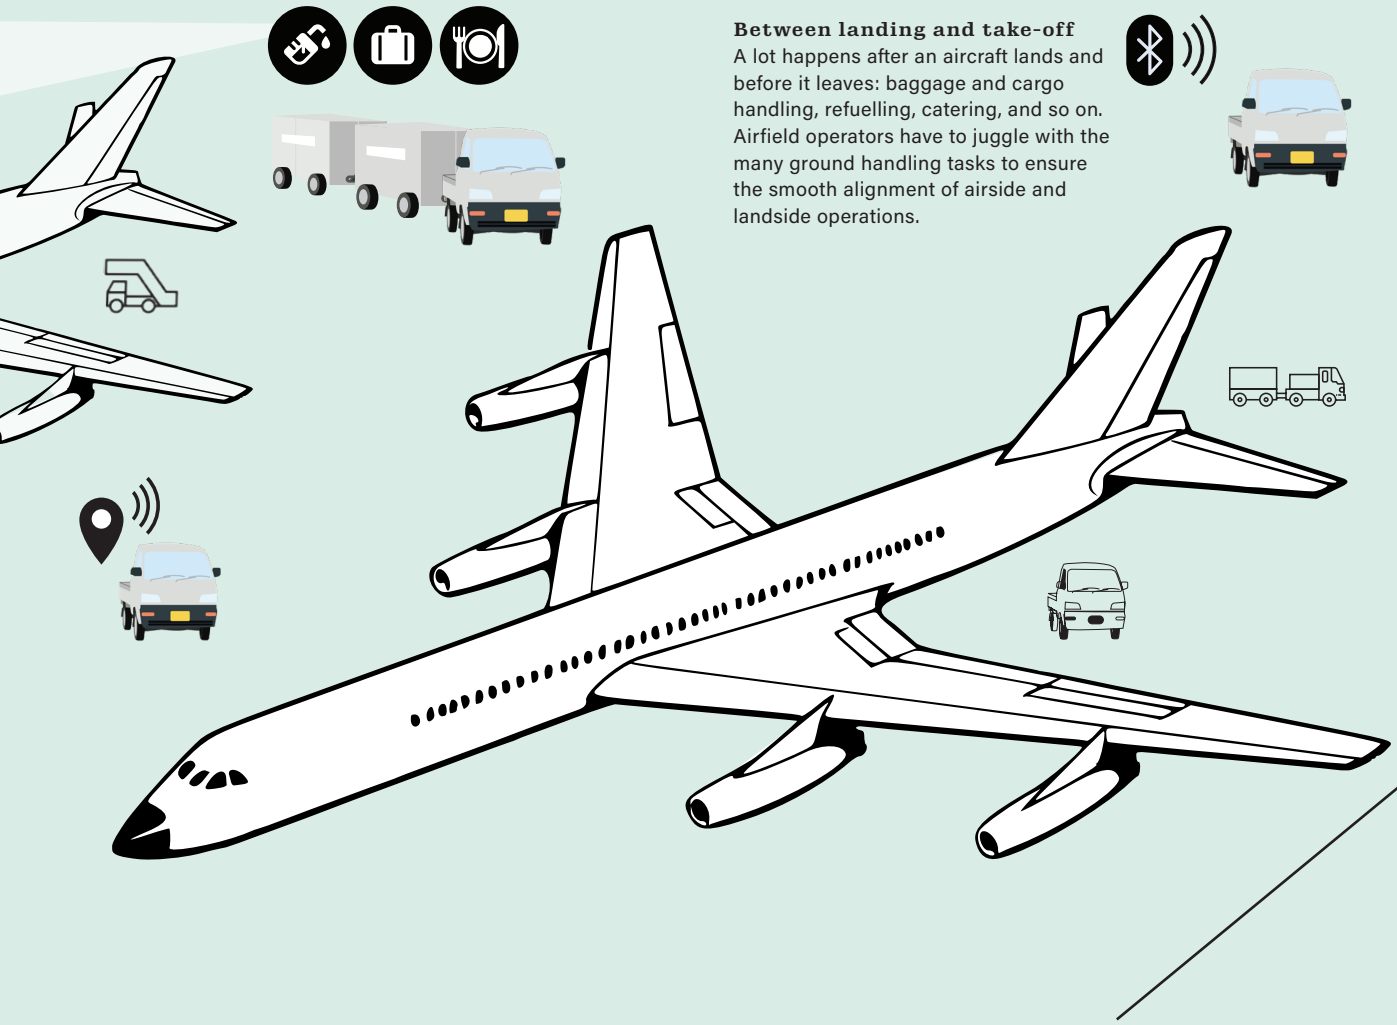

IN 2019, THE HKIA HANDLED OVER

400,000
AIRCRAFTS

70M
PASSENGERS

4.8M
TONNES OF CARGO

Go On, Little Nightingale

The growing use of AI in the arts, that seeming last stronghold of human excellence in the age of machines, is to some as disquieting as it is astounding. But how much ground can AI art really make?

AI: Art

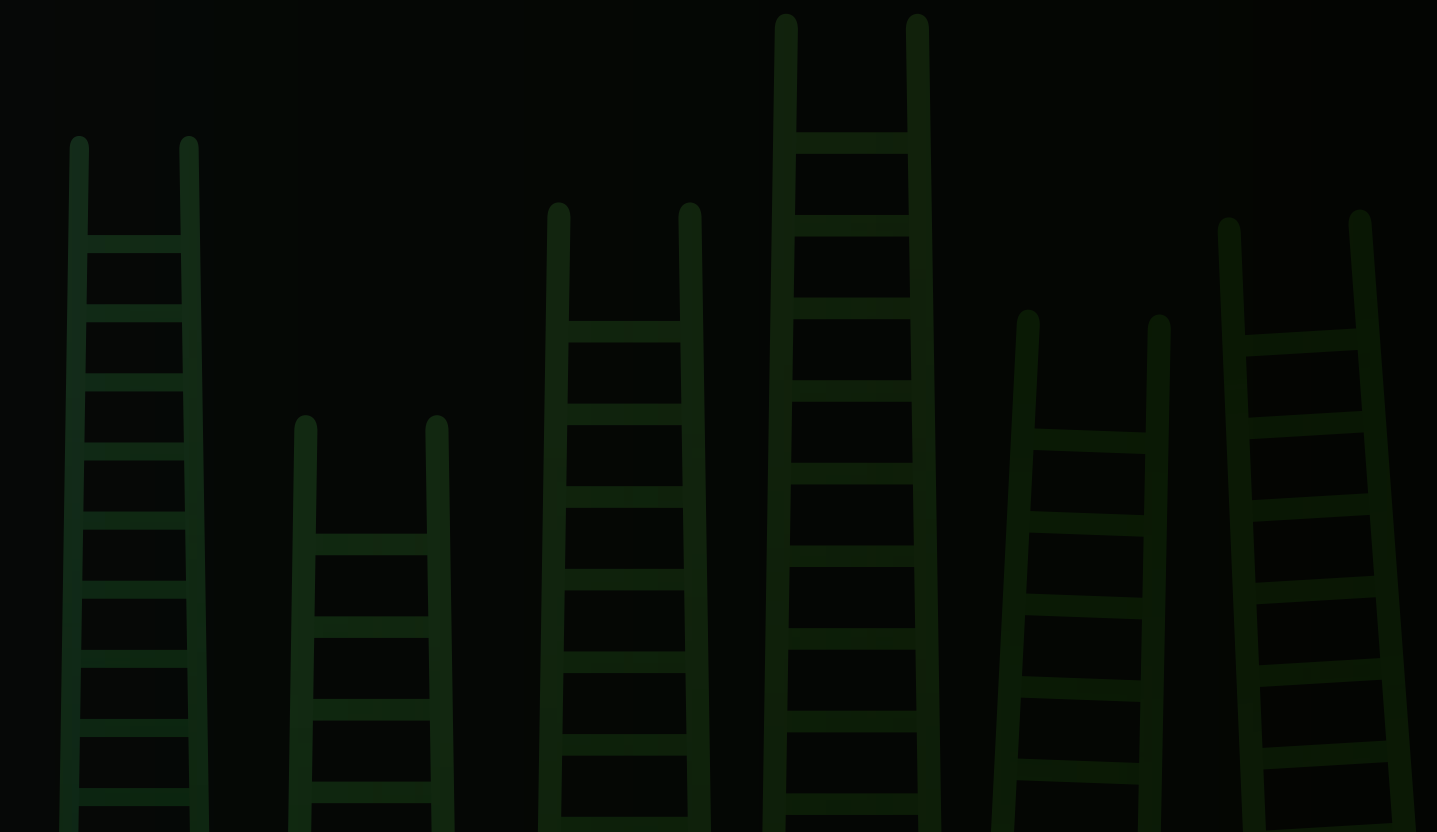

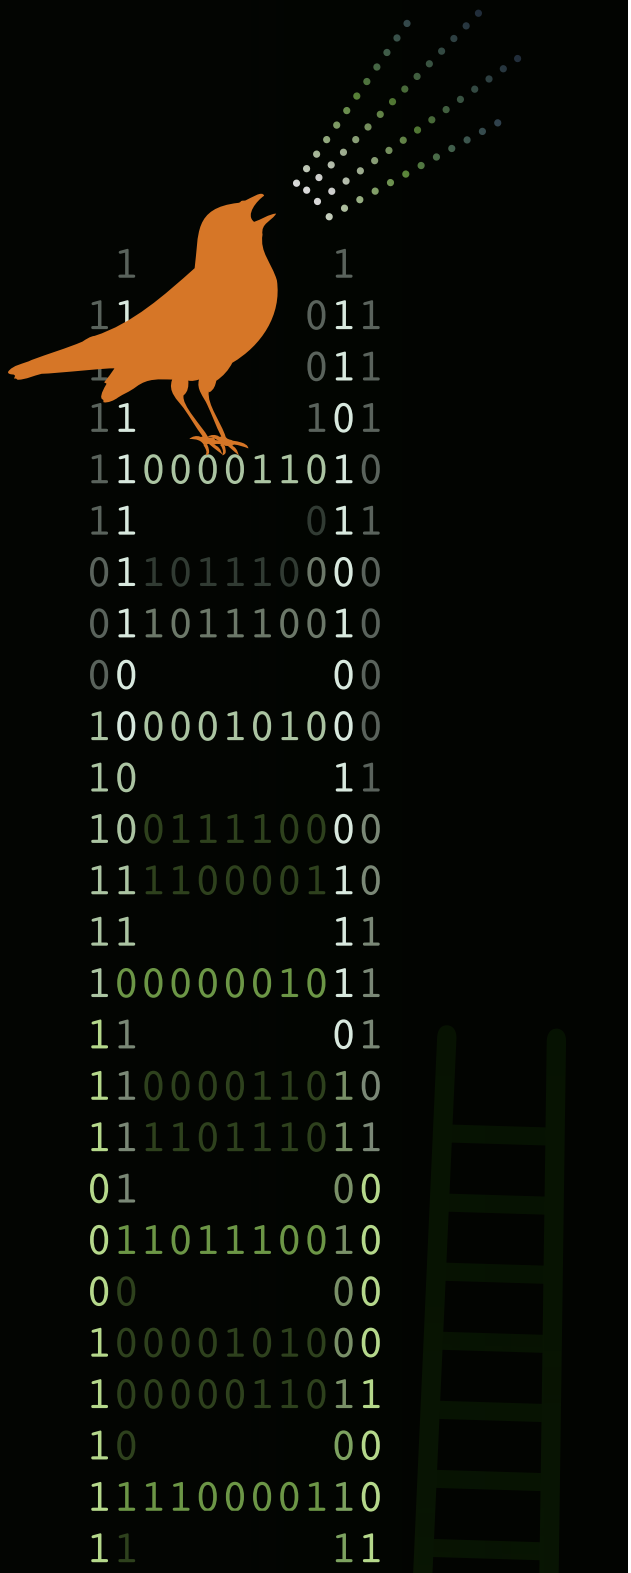

In 1988, **Katsuhiro Otomo** released the first US edition of his manga masterpiece, *Akira*. It had always been problematic to get the traditionally black-and-white Japanese comics into the Western world, which preferred colours. With this culture gap in mind, as the American artist **Steve Oliff** recalls, Otomo decided to publish his work in colour for American readers, and Oliff was given the job to colourize it, panel by panel. As one can imagine, it was an exceedingly arduous task. For this reason, *Akira* has been one of the very few manga works that come in both a black-and-white and a colour version.

This is where Prof. **Wong Tien-tsin** has found yet another use for AI.

PROFESSOR WONG IS A LONG-TIME RESEARCHER of computer graphics at the Department of Computer Science and Engineering. While computer graphics have been around for half a century and used with much success in films and video games, the focus had for decades been on photorealistic rendering, that is the creation of 3D objects, texture and light as they would look in nature. In the 2000s, though, the field saw a major breakthrough—thanks to our fellow gamers. Ever-rising expectations of gaming experience meant better and better graphics processing units (GPUs), which scientists and engineers were quick to exploit. One result was the dawning of computer vision and, consequently, rendering that goes beyond imitations of the real world.

'When we say a computer "sees" or "understands" an image, it means it has extracted information from the image, like its stylistic features, which it should then be able to work into another image,' Professor Wong explained. If we show the computer a painting, then, it should be able to generate something akin to human art—something expressive. An attempt at this is AI Gahaku, a web application that took the Internet by storm last year with its promise of turning any photo into a classical painting in the style of the user's choice. What it does is exactly to study existing paintings from different schools and periods and rework the photo based on what it has

Prof. Wong Tien-tsin
Department of Computer Science and Engineering

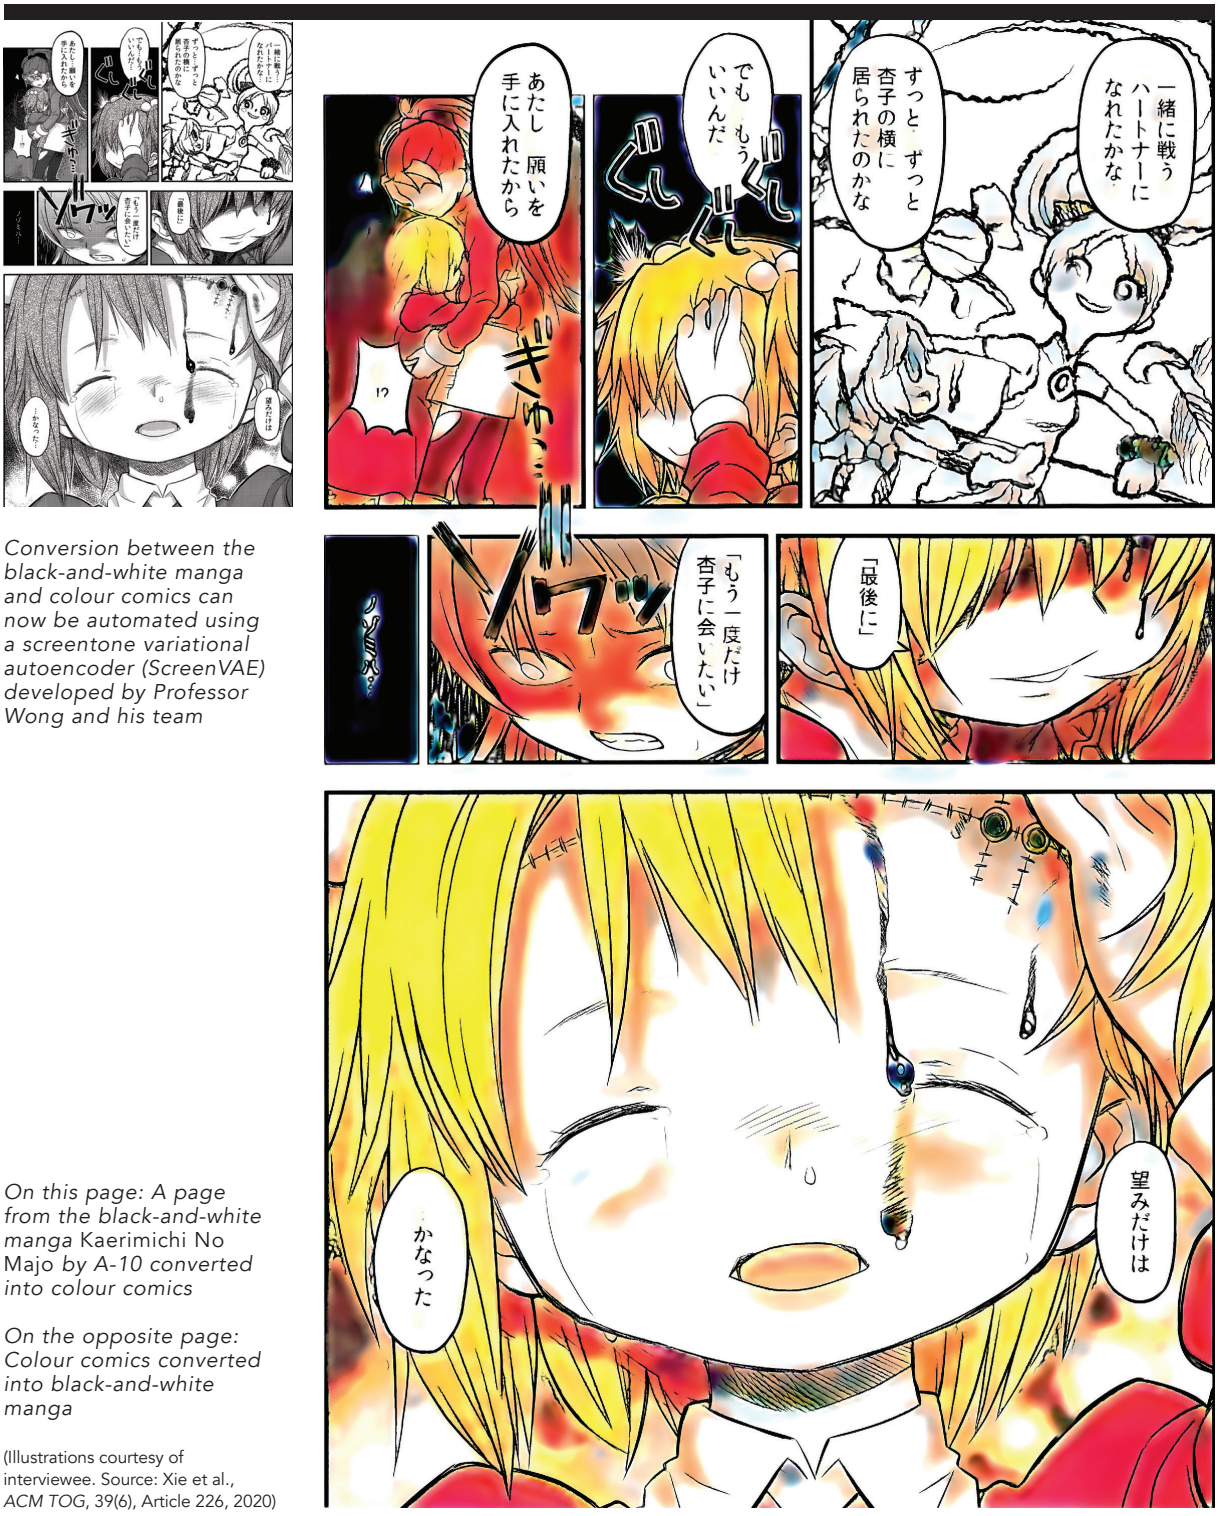

Conversion between the black-and-white manga and colour comics can now be automated using a screentone variational autoencoder (ScreenVAE) developed by Professor Wong and his team

On this page: A page from the black-and-white manga Kaerimichi No Majo *by A-10 converted into colour comics*

On the opposite page: Colour comics converted into black-and-white manga

(Illustrations courtesy of interviewee. Source: Xie et al., *ACM TOG*, 39(6), Article 226, 2020)

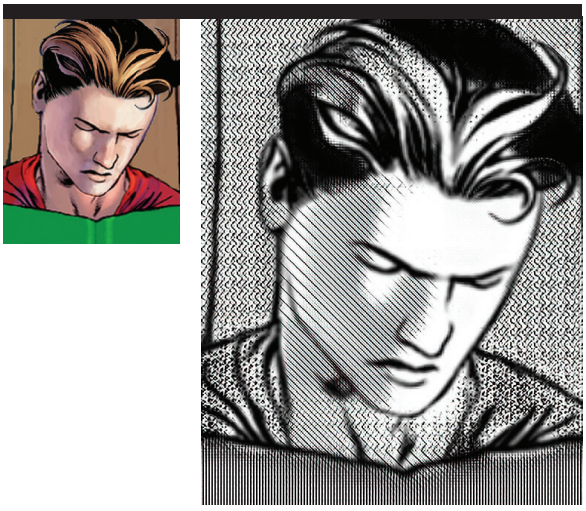

learnt about the chosen style. And as its name suggests, this feat is made possible by AI, which could only have run on the powerful GPUs that we have been blessed with over the past two decades.

Having both the intelligence and the hardware, computers are ready for a bigger role in making art, especially manga. 'It's such a labour-intensive craft,' said Professor Wong of manga production. 'Everything from the writing to the visuals will have to be taken care of by the artist and, if they are lucky, a few apprentices, oftentimes by hand.' Automation is theoretically possible without AI, but then for each task there will have to be a tailor-made algorithm, and each algorithm will involve a number of parameters that must be manually tuned. With AI, which can learn its own parameters and adapt to different problems of a similar nature, parts of the labour can finally be delegated to machines, one example being precisely colourizing black-and-white comics.

'There was a semiautomatic approach, where the artist would throw in dashes of colour and the computer would build on that. Then came our AI-based solution,' said Professor Wong. Their model involves two stages in its initial form. At first, the computer is trained to recognize the black-and-white texture and remove them, leaving only the outlines; at the second stage, it learns what colours and shades are normally used in different situations, say for human skin and hair, and fill the outlines accordingly. In a later iteration, the model performs the task in one go. The artist can provide the model with hints for better results, but even without human guidance, its performance is still quite acceptable.

IN MUSIC, THE USE OF AI HAS ALSO FLOURISHED. As it in visual art, the dream of automation in music has a long history.

'Attempts at having computers make sounds began as early as the invention of computers,' said Dr. **Chau Chuck-jee** of the Department of Computer Science and Engineering, who teaches the first undergraduate course on computer music at CUHK, and himself studied music here alongside computer engineering. In 1951, Alan Turing's Ferranti Mark 1 played a snippet of *God Save the King* for a BBC recording crew, making it the first piece of computer-performed music on tape. In the decades that followed, researchers moved on to exploring ways of making computers write music. Earlier proposals include knowledge-based systems, where experts essentially taught the computer music theory. But as we have seen, writing out the rules are as taxing as it is for us to compose music entirely by hand, if not more. Also, these programs lack general applicability. Thus came the more robust solution of machine learning.

Dr. Chau Chuck-jee
Department of Computer Science and Engineering

'With better hardware after the AI winter at the end of the last century, there was renewed interest in machine-learning approaches to computer music,' said Dr. Chau. Many of these approaches involve what is known as a neural network, a system modelled after the human brain. Given a large enough database—one that covers, say, the entire Western classical music canon—the network can work out using statistics the norms of music composition or the recurring features in the works of a particular composer. Using this knowledge, it can make new music or imitate a certain composer. Programs employing this technique include AIVA, a Luxembourg-developed virtual composer registered at SACEM, the French association of music writers, and granted copyrights. Another example is DeepBach, a model that writes chorales in the style of Bach given the soprano part.

'The researchers brought in a number of people with different levels of musical knowledge and played a series of compositions by DeepBach. On average, around half of them thought that they were really by Bach,' said Dr. **Szeto Wai-man** of the General Education Foundation Programme, who also specializes in computer music and has been involved in promoting public understanding of AI. As Dr. Szeto reported, the model does occasionally deviate

from the norms of music and Bach's style, and experts can easily differentiate one of its works from an actual Bach. For casual listeners, though, it does the trick.

'They're like the student who usually follows the rules and gets the answers right,' said Dr. Chau of the plethora of AI composers debuting recently. 'But is that *good* music?'

Dr. Szeto Wai-man
General Education
Foundation Programme

ON THE MOST BASIC LEVEL, ART CAN BE QUITE SIMPLE. It is the right colour in the right place, the right note at the right time, and these are things that machines have done reasonably well, being now somewhat able to figure out for themselves what the norm says about being right. But art is also about breaking the norm, and that breakaway must be for a good reason. This transformative aspect of art has been achievable only by living creatures, chiefly humans. With AI, can machines catch up with us in this respect?

'I won't rule out that possibility, but it's not going to be easy,' said Dr. Szeto. Making it new, to borrow a phrase from the avant-garde American poet Ezra Pound, is not in itself difficult. Back in the 18th century, composers would mix and match pieces of music in sequences determined by a die in a game called the *Musikalisches Würfelspiel*—and thus was born a new composition. In theory, computers can do the same with randomness added into their otherwise norm-conforming models. But while innovations are made this way, they are all driven by chance, not by some aesthetic motivation. This is where AI comes to a dead end.

'Computers do not have motivations. That's the current state of things,' said Dr. Chau. For them to get closer to being able to make aesthetic decisions, as Professor Wong noted, we might conduct surveys where humans would rate their works and let them learn what it is that appeals to us. But then of course, it will not be the computer's own judgment but ours, and what one portion of humanity enjoys is not necessarily enjoyable to another. Furthermore, we will remember that beauty often transcends its time.

'Stravinsky's *Rite of Spring* was so poorly received at its premiere that there was literally a riot. Now it's a classic,' said Dr. Szeto. 'It's unimaginable for a computer to have the same insight as Stravinsky, to know a work is valuable despite how unpalatable it is to the current taste.'

AUTOMATING MUSIC

Various attempts have been made since the mid-20th century to get computers to write music as **George Papadopoulos** and **Geriant Wiggins** reported. The explorations of the use of AI in composition slowed down during the AI winter of the 1970s and 80s, when the hardware was not powerful enough to implement many of the proposed solutions. At the turn of the century, the dream of automating music production was rekindled with the advent of modern GPUs, which are not limited to processing visual data.

AI composers are like the student who usually follows the rules and gets the answers right. But is that *good* music?

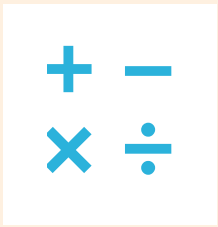

Mathematical and statistical approaches

Examples include *Formalized Music*, a 1963 treatise by the Greek composer-engineer Iannis Xenakis. In his book, Xenakis proposes writing music using set theory and stochastic processes.

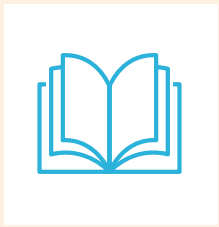

Knowledge-based approaches

These approaches involve experts providing the computer with the rules for writing music. Examples include CHORAL, developed by the American engineer **Kemal Ebcioğlu** in 1988.

Language models

These models treat music as a language and apply natural language processing (NLP) techniques to composition. Examples include the 1987 *Experiments in Musical Intelligence* (*EMI*) by the American composer-scientist **David Cope**.

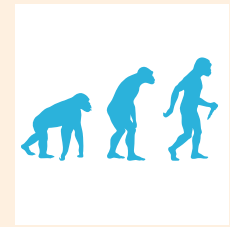

Evolutionary methods

These methods involve the computer generating a certain number of melodies, which are then rated to decide which of them will be retained. They get their name for approximating to evolution in nature, where the fittest survive.

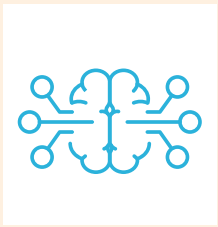

Machine-learning approaches

These approaches work on neural networks, which learn what constitutes music from large collections of samples and use those insights to write music. Examples include DeepBach, AIVA and OpenAI Jukebox, which makes new songs given the lyrics and specifications of the genre and the artist.

Ultimately, for machines to be able to make aesthetic judgements the ways humans do—beyond learning what humans find appealing—is a question of whether they can feel or consciously experience the emotions aroused by an artwork. That will require, of course, consciousness, and it is unlikely the machines will come with one for as long as we, their architects, lack an understanding of how our own consciousness came to be, let alone how it might be recreated for computers.

'AI can have knowledge about emotions the way people who never like heavy metal still know it's exciting, but knowing an emotion is not the same as feeling it,' Dr. Szeto explained. We all *know* heavy metal is exciting by virtue of its loudness, tempo and some of its other objectively definable features, but to *feel* the excitement is to also register that indescribable, visceral rush. This is what AI misses, and this limited grasp of emotions is why computers still need human intervention when it comes to, say, colourizing manga.

'Our model will need a guide if there's an uncommon colour that the human artist wants to use to express a certain emotion,' said Professor Wong. For instance, if the artist wants the normally blue sky to be painted red for a sense of menace, they will have to intervene and give the model a palette of different shades of red. 'There are ongoing efforts to make machines extract emotions from a drawing. If they do end up being able to learn what the emotions of a drawing are while knowing what colours normally express those emotions, they might be able to do without a guide. But it's tough to say how successful it will be, given how hard it is to describe emotions mathematically with all their subtleties.'

'A THING ABOUT PUBLIC UNDERSTANDING of AI is how extreme it tends to get. People like thinking of AI as either inept or godlike,' said Dr. Szeto. Two years ago, the Chinese tech conglomerate Huawei caused a stir on the

Internet when one of their AI programs had completed the melodies of the entire missing third and fourth movements of Schubert's unfinished Symphony No. 8. But commentators soon pointed out that the attempt is a far cry from Schubert's style, and as vast a project as it might be, it was still just the melodies, which could not have worked without being arranged into a full orchestra score and performed by humans. Beyond sensationalism and a blanket dismissal, how can we make sense of AI's place in the world of art?

'It's more than just an intellectual exercise. We're really trying to provide the industry with tools they can utilize,' said Professor Wong of his research, which has also led to a way of automatically converting colour photographs into manga drawings, much to the artists' convenience. He remembers being approached by a comics publisher trying to remove all the speech bubbles in a work they were digitizing and fill in the gaps. That was 10 years ago, and an automated solution was not available then. The publisher ended up outsourcing the task to Vietnam, spending a whopping 40% of their budget on it alone. Had they come to him 10 years later, the professor joked, the story would have been a lot different, now that his team does have a solution, driven by none other than AI.

'They may not work a hundred per cent of the time, but with a bit of tweaking by the human user, they do work quite neatly and help cut costs significantly.'

It is the same story with music. With all the virtual composers out there, video producers who are looking for just a reasonably fair piece of music to go with their content can save a good few dollars and time, now that they do not have to hire a human composer. Similarly, game developers can now bring AI composers into their works, providing personalized, non-repeating ambient music for as long as the gameplay might last. Even to the music producers themselves, AI can be immensely helpful, given how much the industry has compartmentalized. Whereas professional arrangers depended on the composers to give them something to work with, they can now hone their skills with melodies produced by machines.

'AI can empower non-experts by allowing them to quickly achieve things beyond their remits and focus on their specialties,' said Dr. Chau.

We have seen that for a computer to write music, it will first need the ability to understand the rules and norms of music. This opens up possibilities for AI to be used in music education, where it can advise students on what makes a composition or performance right and acceptable.

Letz make it happen, *Op. 23, a choral symphony written by AIVA, was performed at Luxembourg's National Day celebrations in 2017* (SIP / Charles Caratini)

It is even possible for AI to help music researchers with the insights it has gathered through studying countless compositions and performances.

'With its ability to analyze music and, indeed, other forms of art en masse, AI can help scholars gain a bird's-eye view of a particular style and better understand how humans make art in the first place,' said Dr. Szeto, gesturing towards possible new directions for the emerging field of digital humanities.

WITH ALL THAT HAS BEEN SAID ABOUT AI AND THE ARTS, perhaps no one has said it better than Hans Christian Andersen in his 1843 fairy tale 'The Nightingale', proving again that art is often ahead of its time. The story begins with a Chinese emperor bringing in a nightingale to sing for him, only to then desert it for a mechanical songbird that knows no rest. When this new favourite of his, singing nothing but a monotonous waltz day after day, goes unresponsive because no one is there to wind it up, the emperor regrets having abandoned the nightingale and valued the sounds of a machine above those of a living being. He is right to finally recognize nature's superiority in art as we should too. But as the nightingale says in defence of its surrogate, the machine has done well for what it was confined to doing.

'On their own, machines may not be able to make great art. They're ultimately tools that serve at the pleasure of the human artist or, more accurately, curator,' said Dr. Chau, affirming the need for humans in art in the age of AI to provide the direction. 'Within those limits, though, machines do work miracles.'

Keeping a Sharp Eye on Retreating Ice

Prof. Liu Lin
Earth System Science Programme

An AI model can now automate the delicate, costly task of delineating the fronts of retreating glaciers from satellite images, helping scientists to monitor changes in the cryosphere under climate change. Trained and tested on satellite photos of three glaciers in Greenland taken between 2002 and 2019, the deep learning model delivers promising results with deviations from manual delineations being as small as three pixels.

(Background photo courtesy of Prof. Nicolaj Krog Larsen of the University of Copenhagen)

calving front

ice mélange

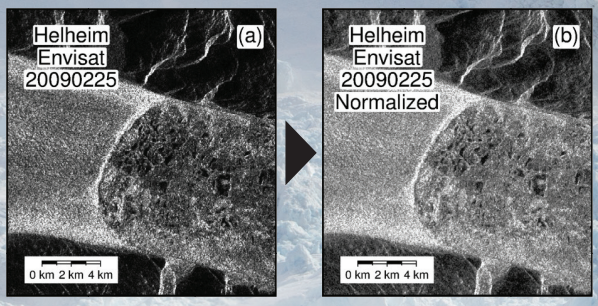

1. Cold data
Greenlandic glaciers have been retreating much faster over the past 20 years. To understand how they react to different climatological and glaciological changes, scientists examine their satellite images and meticulously trace and track their calving fronts, the edges along which they collapse.

2. Preparations
For the AI system to take up the task, the satellite images must first be cropped and augmented. Since they were captured by different satellites, which, for one, work on different colour bands, they also had to be standardized with some having to be denoised.

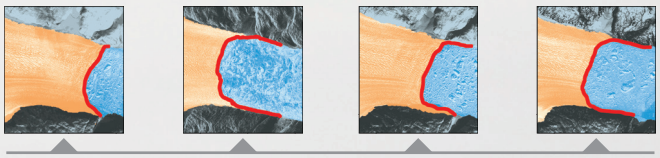

Helheim Glacier

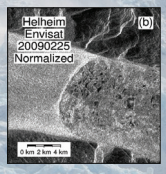

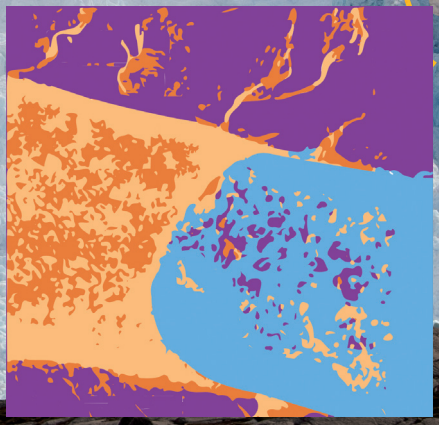

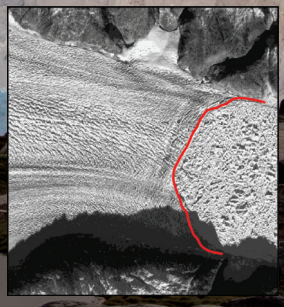

3. Training
The normalized images were fed through a multi-layered algorithm called DeepLabv3+, an example of a common deep learning process known as a convolutional neural network (CNN). Using this algorithm, the computer worked out what constitutes a glacier from the images.

4. In operation
Having figured out what it is that makes something a glacier, the system can now automatically distinguish the glacier from its surroundings and delineate the calving front with any satellite image.

The Perilous Gift of Life

Advances in AI have led to ideas of affording machines the rights and obligations of a human person, but what we really should be talking about is our own responsibilities as developers and users of AI.

AI: Law and Ethics

'I've met Sophia. She's amazing,' said Prof. **Eliza Mik**, quickly realizing how she, too, could not avoid anthropomorphizing the robot, which was famously made a Saudi Arabian citizen in 2017. 'But she's just a piece of rubber.'

To learn, to grow, to catch up with us—such is the rhetoric we like to use to talk about machines. Indeed, it looks like they are becoming one of us with their constantly increasing capabilities to process the world and respond to it in ways that have real impact on our lives, thanks to AI. The reasonable next step, it seems, would be to subject them to the same ethics that govern us. But while it may be an interesting philosophical question, said the technology law specialist at the Faculty of Law, it is really a moot point on a pragmatic, legal level.

'It's completely misconceived.'

THERE IS INDEED ROOM FOR DEBATE as to whether machines can be a person in ethical terms. Depending on the conditions one imposes on moral personhood, according to Prof. **Alexandre Erler** of the Department of Philosophy, it can be argued that a machine with AI is a moral agent.

'There's the more demanding notion that to be a moral agent, you need to have characteristics like having a mind and such mental states as desires, beliefs and intentions. The extreme would be adding things like consciousness to the list,' says the scholar of the ethics of new technologies. If these are the conditions for being a moral agent, he continued, then no existing machine will qualify. However, there is a less demanding school of thought.

'The idea is that as long as the machine has a certain flexibility within a pre-programmed framework and the

Prof. Eliza Mik
Faculty of Law

Prof. Alexandre Erler
Department of Philosophy

Made by the Hong Kong-based Hanson Robotics, Sophia has been criticized as a hype by prominent AI researchers like Yann LeCunn

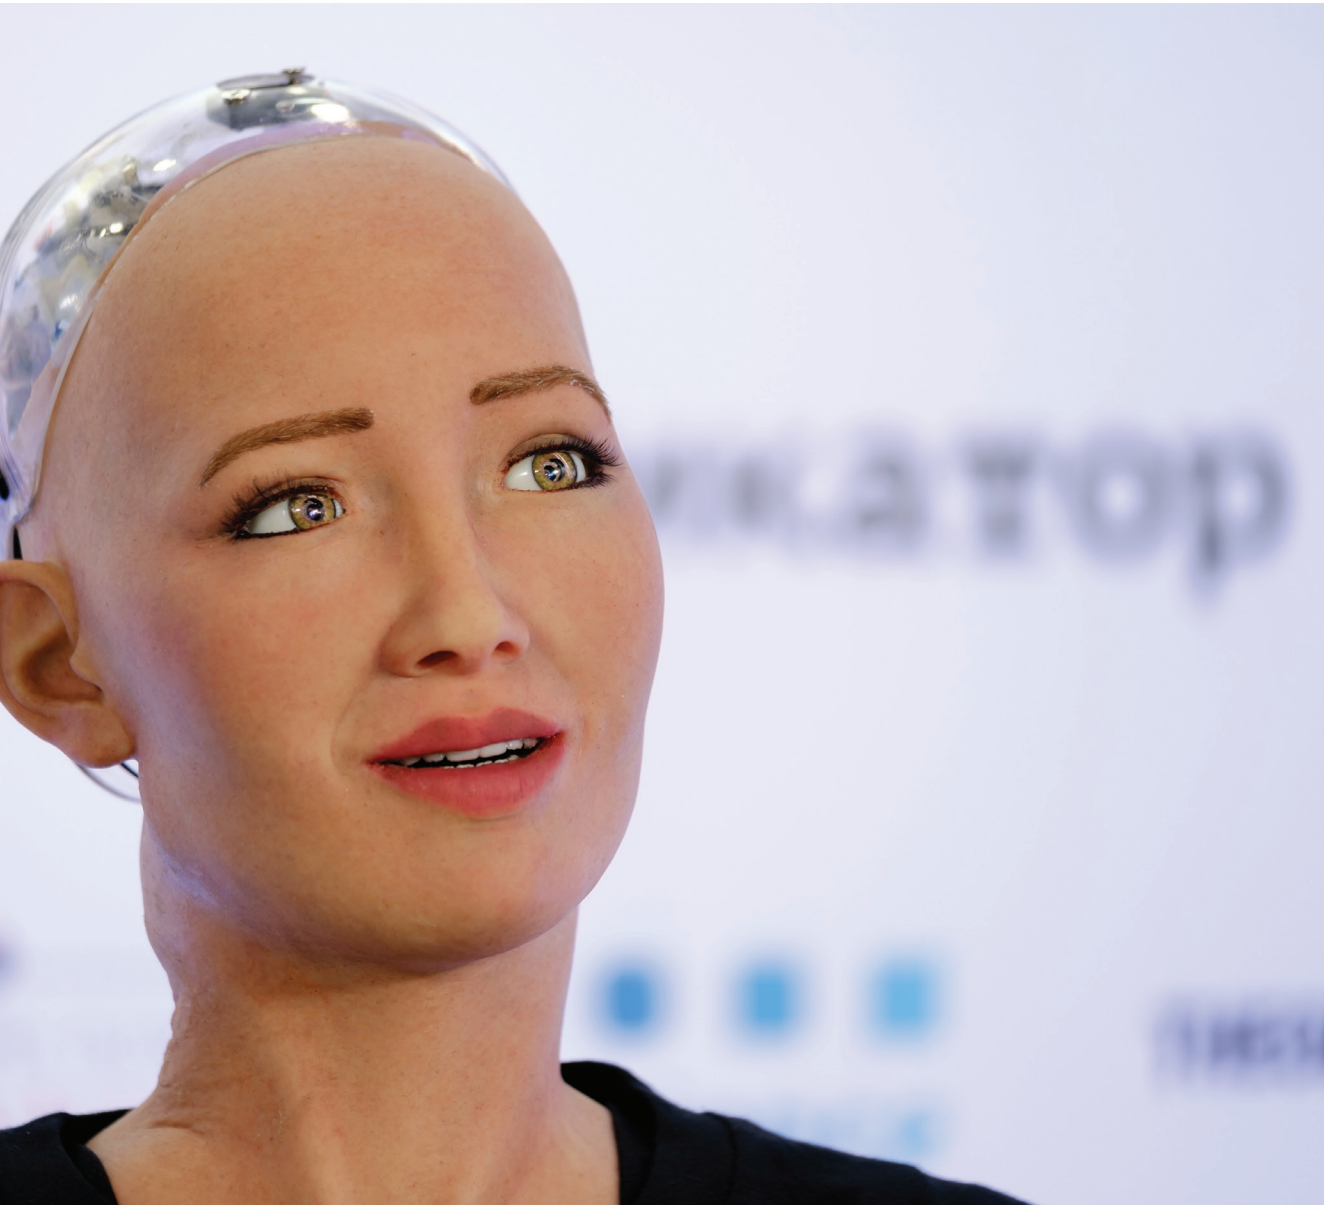

choices it makes have moral significance, like a decision to save a life, it is a moral agent.' One might argue whatever the machine does, it is still driven by the goal set by humans. However, it may also be argued that humans themselves are ultimately motivated by external factors like genetics in each of their actions. If this less stringent view is correct, then we are already seeing or could very soon be seeing machines with moral agency, examples being self-driving cars and autonomous weapons.

But even if a machine can be somewhat of a moral agent on a theoretical level, it is another story in the real world. In 2017, the European Parliament contemplated a proposal for AI machines to be made legal persons considering how they can cause human injuries and deaths as machines like industrial robots have. As Professor Mik pointed out, though, the proposal was plainly impractical.

'Just use your common sense. Let's say this machine that assembles cars is granted legal personhood. What next?

It doesn't have money, it doesn't feel sorry if it kills a human, and you can't put it in prison—it's just a machine,' said Professor Mik, noting how personhood goes hand in hand with liability. While non-humans like companies have been recognized as legal persons, it is always predicated on the fact that the entity can compensate for its actions, something machines are simply incapable of doing. And in any case, it is unnecessary to hold the machine accountable.

'The robot has an owner. If the robot misbehaves and causes any form of damage, the owner can pay. And instead of granting personhood to the robot, just ask the owner to take extra insurance. Problem solved,' said Professor Mik.

It is similarly counterproductive to grant machines rights. Earlier in our story, we took a look at the Luxembourg-developed virtual composer AIVA, which has been given copyrights. In theory, it may be justified to make a creative AI program a copyright holder provided the developers and all the humans needed to fine-tune its work are also duly credited as Professor Mik reminded us. Although its output would inevitably be an amalgamation of pre-existing human works as we have seen, the same can be said of our own artistic productions. As long as its work departs considerably from its predecessors by the standard we judge a human work, it could be copyrighted. But again, it is unhelpful in practice to give machines copyrights or any rights for that matter when, ultimately, it will have to be the human user or developer that receives the actual benefits.

'In a nutshell, there's no advantage in recognizing AI as a legal person,' Professor Mik said. 'Sophia's just a really nice publicity stunt after all.'

It's possible for us to feel less responsible as we delegate more and more tasks to machines. Whether it's justified to feel that way will depend on the state of the technology.

WHILE SEEING HOW MACHINES MAY QUALIFY AS MORAL AGENTS IN THEORY, Professor Erler also stressed the realistic problem of machines being unable to pay for their actions with the current level of AI. But perhaps the real danger of all the talk of giving machines personhood is how it obscures the fact that however free they might be, as both Professor Erler and Professor Mik noted, machines are confined within a larger *modus operandi*, determined by none other than humans. Whether we are dictating to them what principles to follow or, in the case of AI, let them figure them out by observing how we deal with certain situations, humans are at the heart of every action they take. And above all, it is always us who initiate a task and set their goals. Rather than trying to elevate their status and toning down our role, we should focus on how we can use them more responsibly, starting by knowing when we must answer for what a machine causes.

'It's possible for us to feel less responsible as we delegate more and more tasks to machines. Whether it's justified to feel that way will depend on the state of the technology,' said Professor Erler, using the 2009 crash of Air France Flight 447 as an example. One explanation of the tragedy is that the aircraft's automated system stopped working, which goes to show that as much as the operation of the modern aircraft has been automated, the system can still fail. Where it cannot be reasonably believed that the machine works all the time as in this case, the human should stay vigilant and be ready to take over. In the crash, however, the pilots got confused, being unprepared for the system's failure. This is one of those situations where the human user must be held accountable for what the machine does.

'If the machines keep getting better and they become way more reliable than we are at certain tasks, then the users might be justified not to feel responsible for the machines' judgements. It would actually be more reasonable not to step in,' said Professor Erler.

But for now, we will have to take care not to put all our trust in it, not least because many AI systems are black boxes as we have seen. The fact that we do not always know the rationale behind a machine's decisions leads to an important ethical question: is it at all responsible to use something we do not fully understand? Professor Erler said it may be justified if the machine consistently delivers good results and the outcome is the only thing that matters like in a game of chess, where all we care is probably to have a worthy opponent. In cases where procedural justice is crucial, though, such as using AI to predict recidivism, we will have to do better than blindly following the machine's advice.

'What you'd ideally expect is a list of reasons from the machine. That doesn't mean you have to know all the technical details behind it, and you might not be able to. As long as it gives you a justification, you can go on to evaluate it and decide if it's any good,' said Professor Erler, hinting again at humans' irreplaceable role in the age of AI.

Another weakness of AI we must bear in mind is the bias it inherits from the data we use to train it on as discussed earlier in our story. We have seen that a machine feeding on incomplete knowledge can lead to real, lived societal harm aside from producing uninspired art. Worst of all, the bias may go unnoticed with machines exuding a veneer of neutrality. Speaking of the ethics of autonomous vehicles, Professor Erler brought up a particularly chilling example of such harm.

'There've been surveys on how the moral principles guiding a self-driving car may differ around the world. Some societies seem to think that people of higher social status are more important morally and deserve more protection. Do we want our cars to act upon these sorts of beliefs?'

For as long as machines continue to work under our influence, we will need to address their bias as responsible developers and users. A big and diverse data set for them to feed on will of course be imperative, but people other than those involved in building and training them—ordinary citizens like ourselves—also have a role.

'One thing we could do is to report instances of bias when we think we've encountered them, and that can contribute to a conversation,' said Professor Erler, citing the automated recruiting system that carried forward the gender disparity at Amazon and rejected applications from women. 'In some cases, there may be no bias after all and some groups do fit better in certain areas, but in others the bias is real. If there's a discussion and an awareness of the problem, there will be an incentive for it to be rectified like in Amazon's case, where they stopped using the system after becoming aware of the issue.'

'I LOVE SCIENCE FICTION,' said Professor Mik, who is certainly not unfamiliar with the trope of robots becoming human and superhuman in all the films she has watched and rewatched. But when asked about the prospect of sentient machines being created in reality and how that might change our view on AI personhood, she was quick to brush the idea off.

'When we're that far, we're going to have bigger problems than AI personhood. We'll probably not be around. In any case, you'll know what level of technological progress we've reached when you actually read the literature written by those involved in developing the technology.' And as other AI researchers have pointed out, there is no reason why we would want to invest in making sentient machines when the very purpose of having machines in the first place is to make them serve us as we please.

'What does it give you to create a robot that feels? It only gives you trouble. You can't talk back to your Alexa anymore,' said Professor Mik.

Professor Erler agrees it is a remote prospect, but he suggested how some of these more speculative scenarios might be worth thinking about on a philosophical level as his former colleague at Oxford Dr. **Toby Ord** does in his book *The Precipice*.

'The argument is that if we wait until AI reaches that level of development, it would probably be too late. We would no longer be able to place constraints on its design and prevent catastrophes.'

It is always interesting to get ahead of ourselves. In fact, it is important that we do—how else can we know beyond this narrow slice of existence we call the present? But the present has its own problems, pressing ones indeed. When it comes to our current day-to-day negotiations with AI, there is a broader truth to what Professor Mik said as a fan of sci-fi:

'Keep sci-fi away from law.'

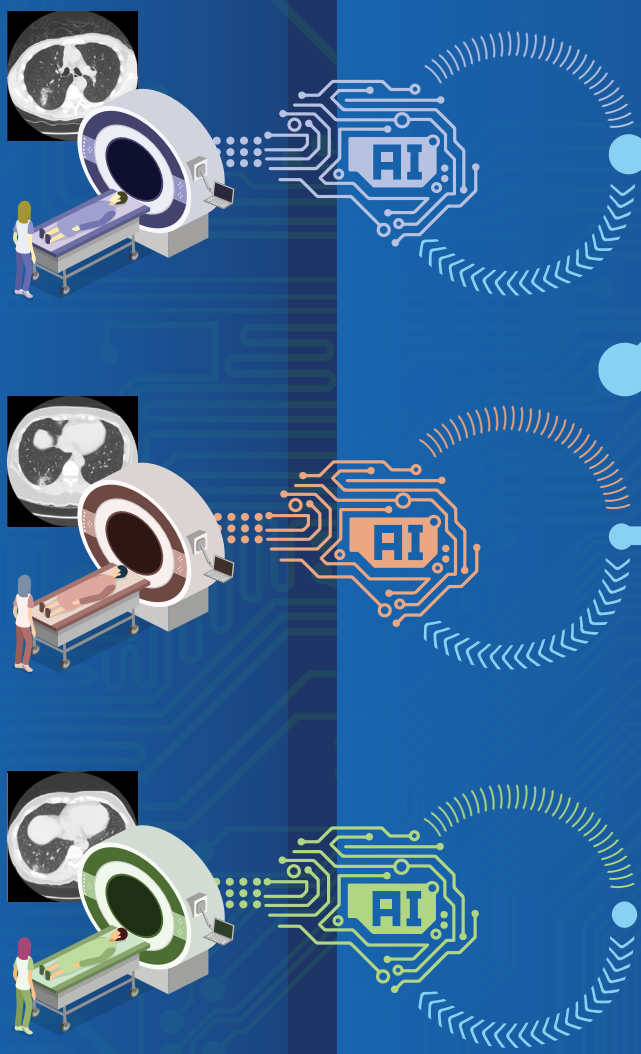

Each hospital trains its own lesion detection models using its own stock of CT scans. The insights but not the images themselves are sent to the main server.

TRAINING STAGE 1

The main server improves the central model in light of the insights. An updated version of the central model is then distributed to each hospital, where it is further trained for more insights and further improvement. This cycle goes on until the central model achieves a certain accuracy.

TRAINING STAGE 2

RAW CT IMAGE

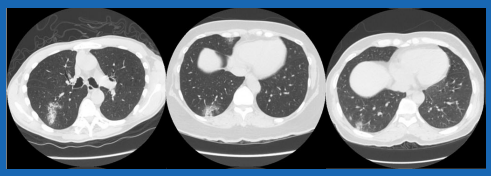

A doctor will need five to 10 minutes to study a chest CT scan. Assessments are made in qualitative terms.

AI-PROCESSED CT IMAGE

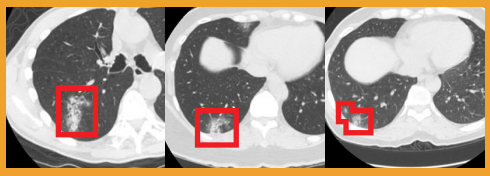

In just around 40 milliseconds, the model, working on a convolutional neural network (CNN), delivers its diagnosis with indications of the lesions' locations and severity in quantitative terms.

COVID Diagnosis Made Better, Faster and Safer

Prof. Dou Qi and **Prof. Heng Pheng-ann**
Department of Computer Science and Engineering

Dr. Tiffany So and **Prof. Simon Yu**
Department of Imaging and Interventional Radiology

An AI system can now rapidly and accurately detect COVID-19 infections in chest CT images using the federated learning technique. The model is trained to recognize COVID-19 lesions independently at multiple hospitals with their own stock of CT scans. Their insights are then gathered and contribute to a central model, which returns correct diagnosis 95% of the time. Throughout the process, the training images remain with individual hospitals without having to go into the same pool and potentially leading to privacy breaches.

Come Away, O Human Child

As AI takes root in every aspect of our life, our single most important mission will be to equip future generations with the technical literacy and moral awareness this new age calls for.

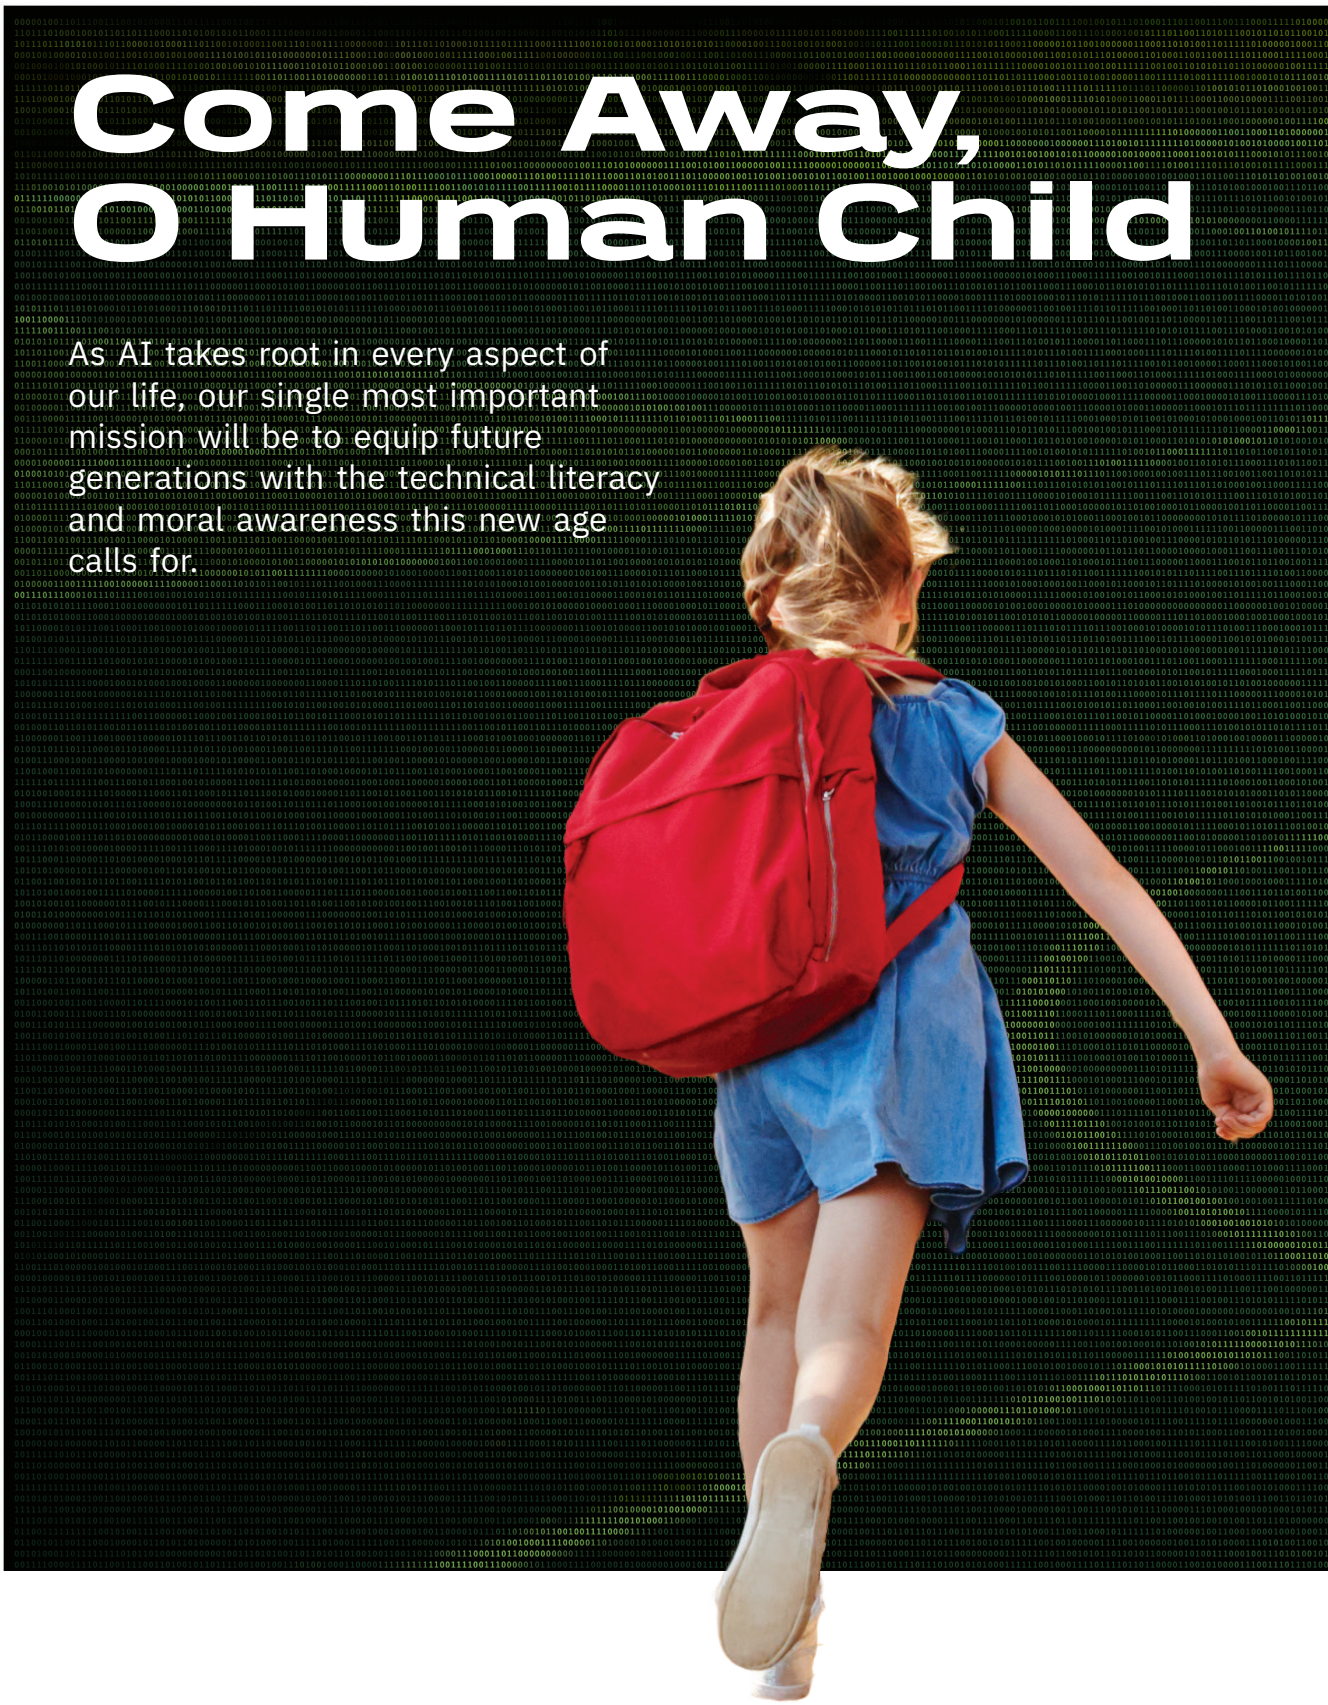

AI: Education

If there is one thing I have learnt from covering AI, it is the importance of education.

Speaking to our expert sources, each having their own focus, one really begins to see how AI presents unique opportunities and challenges in different fields. And yet with every one of our experts, the issue of education cropped up at some point in their response. Indeed, with all the different revelations AI may have for each area of our life, they face the same problems of being misunderstood or not getting enough attention to begin with. I myself constantly struggled to ask the experts the right questions, having never been formally taught about AI.

How should we move forward from here?

IN THE PAST FEW YEARS, AI education has quickly become a key area of development for universities around the world, CUHK being no exception. In 2019, the Department of Computer Science and Engineering (CSE) launched the Artificial Intelligence: Systems & Technologies (AIST) programme, the first of its kind in Hong Kong. Two years into its inception, Prof. **Irwin King**, chairman of CSE, counted the programme a success in terms of getting strong students from around the world and with students achieving stellar academic results. And over the past two years, changes have been made to enhance the programme, an example being the addition of a first-year survey course and an elective on applied deep learning to the curriculum. However, Professor King also noted a challenge they are facing.

'In some ways, it's a small-scale way of training students. We only have around 40 of them in the programme.'

Prof. Irwin King
Chairman of the Department of Computer Science and Engineering

Surely engineering is not the only field that sees an ever-greater need to teach AI. For instance, the undergraduate Urban Studies Programme now covers an area of concentration on smart sustainable cities, which will involve training students to make use of big data. In a similar vein, the Faculty of Social Science launched the Data Science and Policy Studies (DSPS) programme in 2019. Looking back on the last two years, Prof. Wilson Wong, the programme

AI FOR KIDS

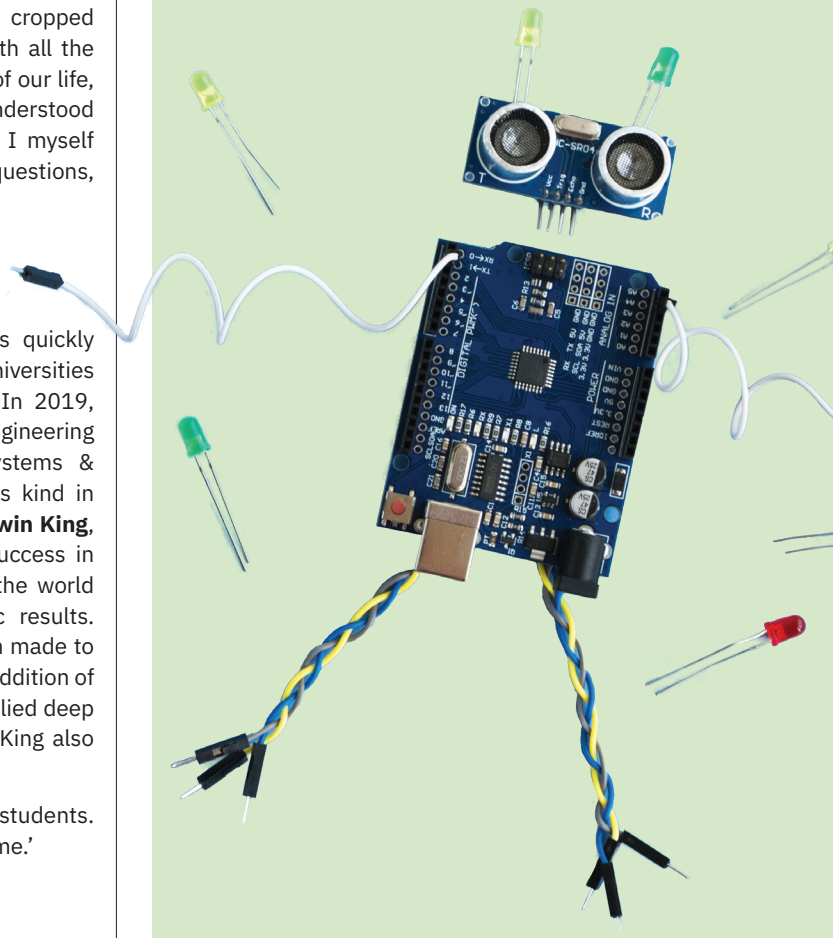

CUHK's AI for the Future Project, funded by the Hong Kong Jockey Club Charities Trust, has proposed a pre-tertiary AI curriculum. The curriculum comes in 12 chapters, each further organized into five modules. Piloted in 2019, it is expected to be incorporated into relevant existing school subjects like information technology, mathematics and social education in the future.

The objectives of each module are as follows:

Awareness

To be aware of the history, background and development of the AI technology in question

Ethics and impact

To explore the use of AI for social good with reference to real-life cases, to appreciate AI's transformative effects on work and to consider the ethics of AI

Knowledge

To identify the key concepts and the impact of AI with reference to some of its more interesting applications, especially those that have local relevance

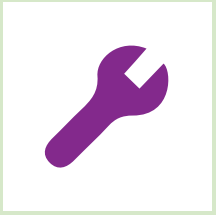

Interactions

To experiment with AI technology in a laboratory setting

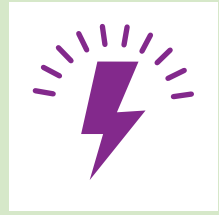

Empowerment

To acquire the ability to develop end-to-end systems using AI technology

director, is glad that it has got AI through to all its students despite their different academic backgrounds.

'As much as we've tried to tear down the wall between the arts and the sciences in secondary education, old habits die hard, and yet AI is exactly one of those things that take both worlds to do what it should. This was where our first intake of students ran into some problems,' said Professor Wong, recalling how those with backgrounds in arts subjects battled with programming and those from science backgrounds had trouble doing humanities studies.

'It eventually worked out, though, and now all our students can handle both the technical and the human aspects.'

BESIDES PICKING UP SPEED IN HIGHER EDUCATION, what more can we do to respond to AI's rapidly increasing pervasiveness? For one thing, not everyone will attend university, much less specialize in AI. Having in mind an AI education for all, the Faculty of Engineering and the Faculty of Education have partnered to create an AI

curriculum for junior secondary students under the CUHK Jockey Club AI for the Future Project. Currently in its pilot phase, the curriculum contains 12 chapters, which cover topics like the principles of AI and how machines perceive and interact with the world. A designer of the curriculum, Prof. **Thomas Chiu** of the Department of Curriculum and Instruction said it is far more than just about the textbook knowledge and passing exams.

Prof. Thomas Chiu
Department of Curriculum and Instruction

'While we definitely want students to learn something about AI, it's also about stimulating interest in the subject such that they will want to find out more about it later in their life,' said Professor Chiu. 'What we're hoping to cultivate is a readiness for AI, which will allow them to live with AI without the anxiety we're seeing a lot of, even if they're not pursuing the subject in an academic setting.'

One of the things the curriculum will alert students to is the disruptions AI will bring to the job market. It is inevitable that AI will eliminate certain occupations as both Professor Chiu and Professor Wong pointed out, but it will also free us from mechanical work and allow us to focus on creative tasks. Meanwhile, work in general will require some knowledge of AI the way it requires IT knowledge now. The curriculum, according to Professor Chiu, is precisely designed to prepare students for this future of work.

'With all the virtual composers, for example, students who want to make music will have to be prepared to take on a more creative role,' said Professor Chiu, echoing what other researchers have said about humans' place in art in the age of AI. As Professor Wong reminded us, this will be a great opportunity to finally give attention to developing personal qualities like innovativeness and humanistic thinking, which have been neglected for far too long.

One might question if a junior secondary student is ready to learn about AI given the complexity of the maths and science it involves, but Professor Chiu has an even more ambitious vision.

'I believe there will eventually be an AI curriculum for primary schools. In some countries, they're teaching AI at kindergartens, specifically the kind of thinking that goes into developing and using AI,' he said. 'We're well aware this is foundational education, and there's no way students can learn everything at this level. We can, however, teach them concepts that are appropriate to the things they're already learning in related subjects, like computer literacy, maths and social studies.'

WHAT ELSE MUST WE DO? Early on in our story, we discussed the importance of digital literacy, an awareness of the ethical implications of digital technology as well as the technology itself. Indeed, no one will be truly ready for AI without an education of its ethics, a serious reflection on all the moral concerns we have touched on and many more. In the AIST programme, students are required to take a course addressing some of the ethical issues, and at the Department of Philosophy, a course covering AI ethics is being offered. Similarly, DSPS students are always asked to consider the moral significance of the technological trends they are learning about.

'Technology should not only be a matter of speed and performance,' said Professor Wong. 'What are the pros and cons of a smart city? When does it become a surveillance city? These are the sorts of questions our students get with a chance to look at real-life cases.'

And as grim as some of these problems are, it is never too early for our children to confront them. Aiming for what Professor Chiu calls 'a local understanding from a global perspective', the curriculum they are proposing will encourage students to consider topics like autonomous vehicles and machine bias with reference to examples around them while recognizing the extents to which these problems are universal.

But as anyone who have studied in Hong Kong will know, ethics is rarely given much attention in the traditional curriculum, either on its own or in the context of another subject. After all, it is rarely covered in assessments. How are we to ensure that ethical discussions do not become mere footnotes in the new curriculum? While Professor Chiu said the curriculum is not designed to go with assessments, the team did make the effort to put ethics centre stage.

'In addition to a separate chapter for AI ethics in general towards the end, each of the curriculum's other chapters get time for ethical discussions specific to its theme,' he said, explaining how it should not get left out along the way. 'We're really giving quite a lot of emphasis to it in our design.'

SO WHAT DO OUR EDUCATORS HAVE IN MIND for the future of AI education? For both the AIST and the DSPS

What we're hoping to cultivate is a readiness for AI, which will allow them to live with AI without the anxiety we're seeing a lot of.

programmes, a way forward is to encourage students to have a focus.

'At the end of the day, it's impossible to learn everything about the field. We would encourage students in their senior years to figure out whether they're more interested in the tech side of things or policy work and concentrate,' said Professor Wong. 'Either way, they will have developed a reasonable capacity for both understanding the social implications of data and working comfortably with data to formulate a policy.'

As for the future of pre-tertiary AI education, Professor Chiu believes it will be necessary for all teachers to have some familiarity with AI, which will probably be incorporated into every school subject at some point. For instance, language education in the future might involve the use of natural language processing (NLP): teachers might find it helpful to teach grammar by letting machines demonstrate how they breakdown and interpret a text. But no matter what we do with AI in classrooms, said Professor Chiu, it must be shown to be at work in the everyday world, which will require cooperation between the school and the community.

'After all, AI has everything to do with our daily life.'

Indeed it does.

Getting to the Root of Mangrove Methane Emission

Prof. Derrick Lai
Department of Geography and Resource Management

Mangroves are known to cool the climate by taking up carbon, but they also give off the powerful greenhouse gas methane, which can offset their climate benefit by more than 50% over a period of 20 years. With an AI algorithm known as the *random forest*, soil temperatures and salinity have now been identified as the main drivers of methane emission from a mangrove at the Mai Po Nature Reserve.

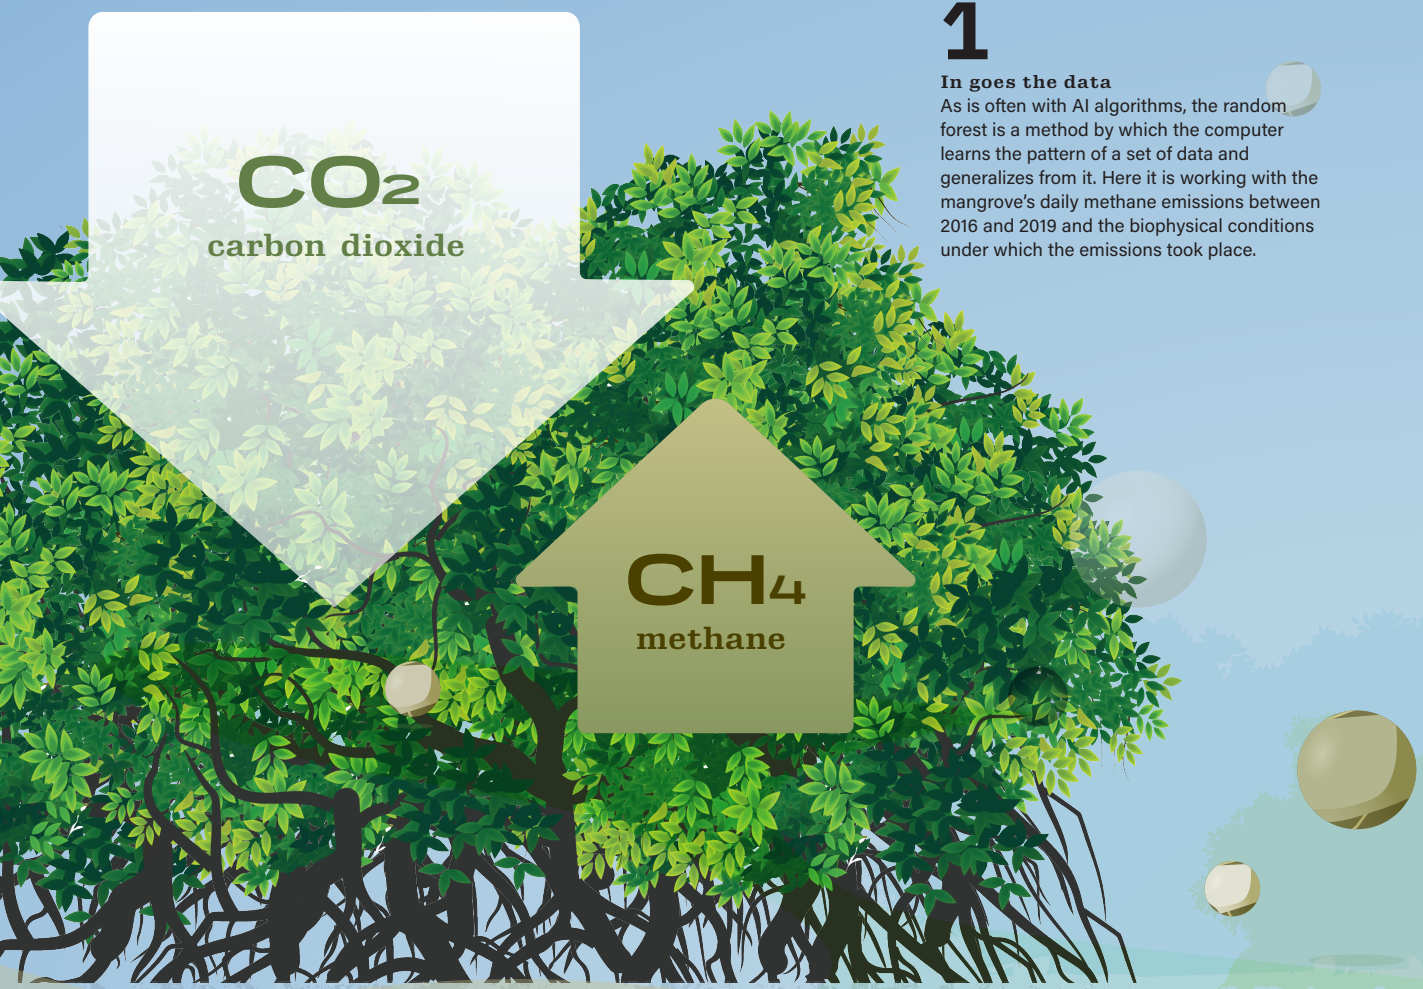

1

In goes the data
As is often with AI algorithms, the random forest is a method by which the computer learns the pattern of a set of data and generalizes from it. Here it is working with the mangrove's daily methane emissions between 2016 and 2019 and the biophysical conditions under which the emissions took place.

2

The sorting tree
The computer takes random samples from the emission data and feeds them through what is called a *decision tree*, where they are sorted into two groups by, say, the air temperatures at which the emissions occurred. The sorting is repeated with another biophysical variable to further organize the emission data.

3

Planting a forest
The computer takes another batch of random samples from the emission data and makes another decision tree. This is repeated until there are many of these trees, hence a random 'forest'. A single decision tree looks at just one iteration of the data and derives a pattern specific to that iteration, which may be too specific. In AI lingo, this is called *overfitting*. With an ensemble of trees, this overfitting averages out.

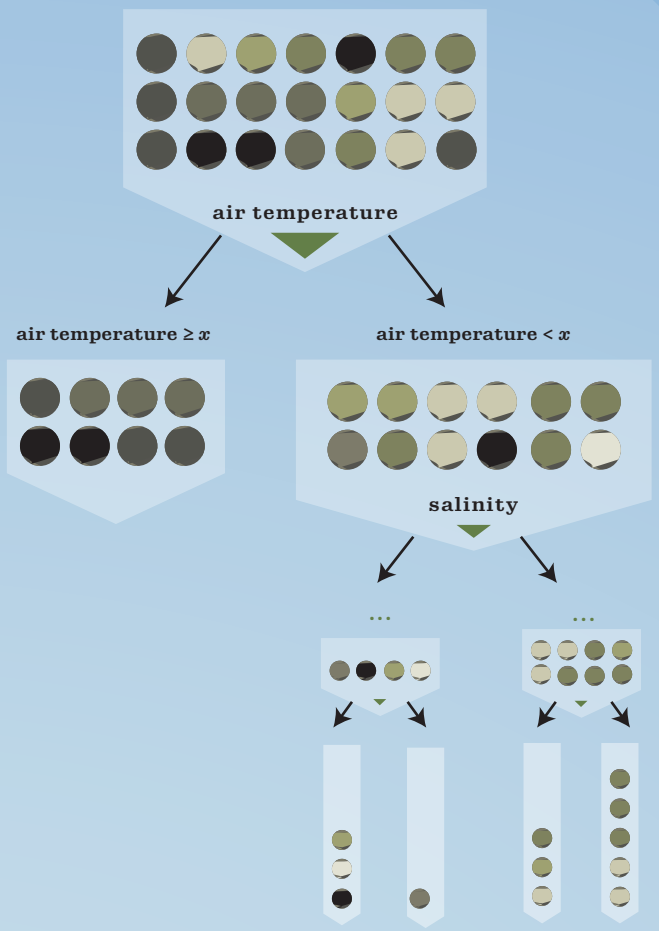

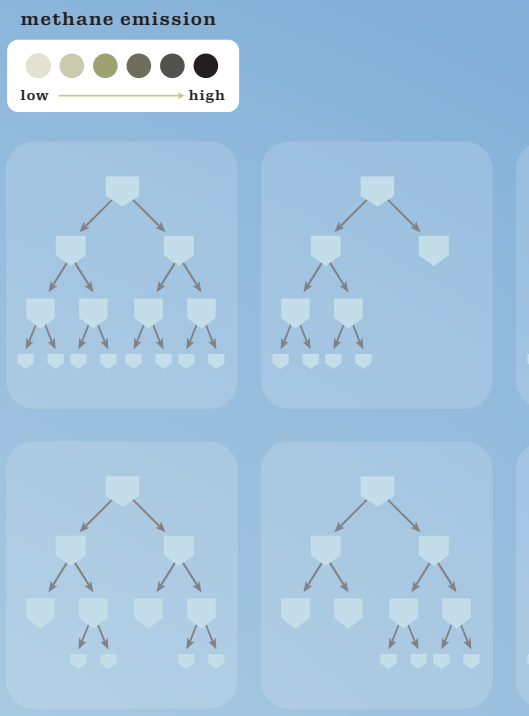

4

Finding the culprit
Now that a pattern emerges from the emission data, we can find out how much each biophysical variable has contributed to putting the data in order and bringing out the pattern using a metric known as the *mean decrease in impurity*. The variable that contributes the most or, in other words, has the greatest effect in shaping the data is precisely the most significant driver of the emissions.

A whiter box
Unlike many AI models, which have gained infamy for being black boxes, a model based on the random forest algorithm has a reasonably straightforward and transparent decision-making process. This is what allows us to assess relatively easily the extents to which individual variables determine the emission pattern.

A Traveller's Guide to the World of AI

From pipe dream to reality, AI has come a long way, going through countless trials and errors before reaching its current state. The *Chinese University Bulletin* has constructed a timeline chronicling some of the milestones in the global history of AI and its development here at CUHK. We have also compiled 10 burning questions about AI, which we have invited Prof. Irwin King, chairman of the Department of Computer Science and Engineering, to shed some light on.

BY **ronaldluk@cuhkcontents**

01 What exactly is AI? Is there a common definition for it?

The term 'AI', which stands for 'artificial intelligence', was coined in 1956 by Prof. John McCarthy, a computer scientist and mathematics professor at Dartmouth College at the time. It originated from a summer research project led by Professor McCarthy to explore the notion that every aspect of human intelligence can be precisely simulated by machines.

Simply put, AI is about developing machines that can function in ways that are highly similar to humans, mimicking human intelligence in such areas as visual perception, language use and decision making.

02 What are some key AI concepts that we should all know?

A concept central to AI is *machine learning*. It is the process by which AI learns to make informed decisions from the data fed to it using certain algorithms.

Another important concept is the *neural network*, a class of computing systems that makes it possible for machines to learn in the first place. A neural network usually consists of mathematical models known as artificial neurons, which are designed to function the way biological neurons do.

One of the hot topics in AI research is *natural language processing*, more commonly known as NLP. It is concerned with how machines can study human languages and learn their patterns, thereby acquiring the ability to perform such tasks involving language as sentiment analysis, text classification, and question and answer generation.

03 Which countries play significant roles in the development of AI?

As mentioned, the US is the birthplace of AI and has been the pioneer of AI developments. European countries have also made significant contributions like AlphaGo, a British AI program famously known for beating professional *go* players. In the past decade, China has been catching up. It's now a leader in AI research publications and patents and poised to spearhead AI-powered businesses.

04 What are the most common uses of AI in daily life?

There's literally a myriad of uses. Applications of AI are usually classified into weak ones and strong ones. Weak AI refers to applications that perform simple automated tasks, examples being virtual assistants like Siri and online chatbots that answer questions in real time. Strong AI refers to applications capable of performing much more complex tasks, which include autonomous vehicles and service robots.

05 What are the benefits of AI?

Undoubtedly, AI can improve human life. Indeed, people from all walks of life can enjoy the benefits of AI in various ways. For instance, autonomous vehicles can increase traffic efficiency and offer more travel options. Meanwhile, AI-enabled robots can also take up tasks that are dangerous or difficult for humans, be it neutralizing bombs, exploring the ocean's depths or mining natural resources like coal and oil.

06 What challenges come with increased use of AI?

One significant issue with AI is how it disrupts the job market. Many jobs can or will be taken up by robots or machines. Outcompeted in both speed and accuracy in many situations, human workers may be replaced.

Another challenge AI poses is how difficult it is to ensure that it consistently makes the 'right' decision, notably in the case of autonomous vehicles. Imagine a self-driving car is now going full speed toward an old person and a child. Suppose the brakes have failed and the only option is to steer away from one of them and sacrifice the other. Who should the AI save? Should it prioritize the lives of the passengers instead? Is there only one right answer? Does a right answer exist at all?

07 Is it true that AI is stealing our jobs?

Yes and no.

Indeed, AI can take on many repetitive tasks, but that doesn't mean everyone will lose their jobs. Take the accounting industry as an example. Granted, many accounting software applications are now equipped with AI and can handle various basic tasks. However, many aspects of the profession still require such human input as the instinct of an experienced accountant who can spot a problem before it emerges in the numbers. Humans are also needed in face-to-face conversations and day-to-day interactions to garner clients' trust.

So, no. Although repetitive tasks may be automated, humans' involvement is still crucial in many settings and cannot be replaced in any way. As a matter of fact, more jobs will be available for those with qualification in AI-related fields as machines become more widely used. With this in mind, though, it'll be important for us to upskill.

08 Will a full-on AI takeover happen some day?

The short answer is no. There are fundamental obstacles that prevent AI from becoming human.

AI is limited to performing a particular kind of task that it is programmed to deal with and is unable to do anything beyond that. Moreover, they operate only in terms of logic, principles and theories. Meanwhile, humans are more than just a program or an algorithm—they have the free will to think, to understand and to feel. And while humans often defy rules and make mistakes, it's what makes life interesting and unique, something that's out of AI's reach.

09 How does the future look for AI? What does the future hold for us in a world of AI?

I think the future is a bright one. AI is already all around us, so we can expect more and more sectors to take advantage of the technology.

It's understandable that people might worry about the coming of AI, but remember there's a silver lining to every cloud. Surely AI has taken away some jobs, but it has also created new opportunities. Computer scientists and data analysts, for example, are getting more chances to shine.

10 Do you need to be a genius to learn about AI?

The answer is a flat no.

AI is a field overlapping with many areas of human endeavour, including computer programming, engineering, mathematics, statistics, language, and so on. Naturally, those who can readily think critically and logically will have an edge, but there's nothing stopping you if you have the heart and passion for it. The world of AI is as welcoming as it's exciting.

A brief history of AI

c.700 BCE

The ancient Greek poet Hesiod wrote of the mythological intelligent robot Talos

1912

Leonardo Torres y Quevedo built the world's first chess computer, El Ajedrecista

1948-9

William Grey Walter built the world's first autonomous robots, Elmer and Elsie, which could follow light and navigate around obstacles

1950

Alan Turing proposed the Turing test to measure a machine's intelligence

1956

John McCarthy coined the term 'artificial intelligence' at Dartmouth College

Artificial intelligence (n.)

1960

Ray Solomonoff developed the idea of algorithmic probability, one of the mathematical foundations of modern AI

1966

Joseph Weizenbaum created the NLP program Eliza, the world's first chatbot

1969

The idea of backpropagation, a core part of AI that allows it to learn from mistakes, was conceived

1972

Shakey, the first mobile robot to reason through its actions, was created

1974-80

1987-93

Reduced funding led to two AI winters

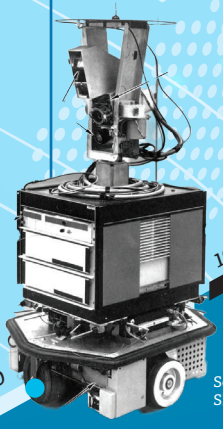

Source: SRI international

1950 1960 1970 1974-80 1987-93

Source: juergvollmer, Alan Levine

1997

IBM's Deep Blue defeated world chess champion Garry Kasparov

1999

Sony created AIBO the robot dog, the first robot designed for home entertainment

2000

2001

The CUHK Multimedia Laboratory, a pioneer in deep learning and computer vision, was established

2010

2011

Apple released Siri ; IBM's Watson, which understands tricky questions using NLP techniques, defeated two former champions on the US TV quiz show *Jeopardy!*

2012

CUHK launched an optional intelligence stream under the undergraduate Computer Science programme

2014

NVIDIA established Hong Kong's first CUDA Research Center at CUHK to encourage graphics processing unit (GPU) research

2015

Four CUHK engineering students came first in the 'Huawei Cup 2015' Intelligent Design Competition for College Students

2016

The CUHK Multimedia Lab was named one of the world's pioneers in AI research at the GPU Technology Conference, being the only Asian institution to receive this recognition

2018

CUHK co-launched the Global AI Academic Alliance

2019

CUHK launched the Artificial Intelligence: Systems and Technologies (AIST) programme, Hong Kong's first Bachelor of Engineering programme in AI

2020

2020

Eleven CUHK engineering researchers made the AI 2000 Most Influential Scholar Annual List; PhD student Xu Hao came first at the global college AI training camp DeeCamp

2021

PhD student Gao Yifan was named one of the Top 100 Chinese New Stars in Artificial Intelligence by Baidu Scholar

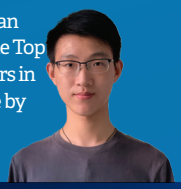

AI at CUHK

Appointments

Council Members

Prof. Wing Yun-kwok

Mr. Kelvin Yeung Yu-ming

Dr. Norman Chan Tak-lam

		Name	Period
New	Member	Prof. Wing Yun-kwok	16.4.2021—20.1.2023
		Mr. Kelvin Yeung Yu-ming	24.4.2021—31.8.2023
		Dr. Norman Chan Tak-lam	1.8.2021—10.2.2025
Re-appointed	Vice-Chairman	Dr. Chien Lee	2.3.2021—1.3.2023
	Member	Mr. Thomas C.B. Liang	15.4.2021—14.4.2024

University Officers and Senior Staff

Prof. Wong Suk-ying

Prof. Anthony T.C. Chan

Prof. Nick Rawlins

		Name	Period
New	Master of CW Chu College	Prof. Wong Suk-ying	1.7.2021—30.6.2025
	Pro-Vice-Chancellor	Prof. Anthony T.C. Chan	1.8.2021—31.7.2024
	Pro-Vice-Chancellor	Prof. Nick Rawlins	1.8.2021—31.7.2024
Re-appointed	Pro-Vice-Chancellor	Prof. Chan Wai-yee	1.8.2021—31.7.2024
	Master of Wu Yee Sun College	Prof. Anthony T.C. Chan	1.8.2021—31.7.2024

Emeritus Professors

4.1.2021

Prof. Zhang Junsen
Department of Economics

15.1.2021

Prof. Christine Huang Yi-hui
School of Journalism and Communication

1.3.2021

Prof. Joseph J. Y. Sung
Department of Medicine and Therapeutics

1.7.2021

Prof. Andrew Chan Chi-fai
Department of Marketing

1.8.2021

Prof. Cai Xiaoqiang
Department of Systems Engineering and Engineering Management

1.8.2021

Prof. Chow Hak-fun
Department of Chemistry

1.8.2021

Prof. David Charles Donald
Faculty of Law

Prof. Fok Tai-fai
Department of Paediatrics

Prof. Fong Wing-ping
School of Life Sciences

Prof. Charles David Gomersall
Department of Anaesthesia and Intensive Care

Prof. Hau Kit-tai
Department of Educational Psychology

Prof. Esther Ho Sui-chu
Department of Educational Administration and Policy

Emeritus Professors

1.8.2021

Prof. Hua Wei
Department of Chinese Language and Literature

Prof. Lau Kwok-ying
Department of Philosophy

Prof. Kenneth Lee Ka-ho
School of Biomedical Sciences

Prof. Li Hon-lam
Department of Philosophy

Prof. Gordon Clark Mathews
Department of Anthropology

Prof. Wong Chi-sum
Department of Management

1.8.2021

Prof. Lawrence Wong Wang-chi
Department of Translation

1.8.2021

Prof. Jimmy Yu Chai-mei
Department of Chemistry

1.9.2021

Prof. Shige Makino
Department of Management

2.9.2021

Prof. Gabriel Lau Ngar-cheung
Department of Geography and Resource Management

香港中文大學
The Chinese University of Hong Kong

We all like the feel of paper. But this brochure will increase your carbon footprint. So share a copy with friends or read it online at your own leisure (*www.iso.cuhk.edu.hk/english/pub/bulletin*).
Thank you for supporting the environment.