COMMENT

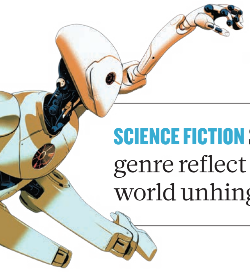

SCIENCE FICTION Six giants of the genre reflect on their role in a world unhinged **p.329**

POLLUTION A call for Pakistan and India to craft air-quality policy together **p.334**

TECHNOLOGY AI developers need a checklist for human-rights violations **p.334**

LAB LIFE Here's to coffee-room chat, catalyst of collaborations **p.334**

COURTESY OF MICROSOFT

Microsoft, in collaboration with others, is using algorithms to convert satellite images into information about categories of land cover, such as forests.

AI for Earth

Create an artificial-intelligence platform for the planet, urges **Lucas N. Joppa**.

Earlier this year, I became Microsoft's first chief environment scientist. I've been tasked with deploying the company's deep investments in artificial intelligence (AI) research and technology to help people around the world monitor, model and ultimately manage Earth's natural systems.

Most people I meet are surprised that one of the world's leading technology companies has a role such as mine. Yet I believe that in the next few years, every major tech firm will be working on applying AI to sustainability.

It is the ethical thing to do. It is good for business. And the time is right: applications of the type that I and others have long been developing can now work at scale. That is why, in 2017, Microsoft put US$50 million into a 5-year programme called AI for Earth (see 'Microsoft for Earth' and www.microsoft.com/aiforearth).

FLYING BLIND

Today, we know more than ever about human activity. More than one-quarter of the 7.6 billion people on Earth post detailed information about their lives on Facebook at least once a month. Nearly one-fifth do so daily (see go.nature.com/2bwmejp). Those data are fed to increasingly powerful algorithms that link people to others, products or information. Yet we are flying blind when it comes to understanding the natural world.

Scientists still struggle to predict the effects of climate change at the resolution of cities or regions, or over timeframes of months or weeks — largely because they don't have the kinds of data needed to make such predictions, or because they ▶

lack the algorithms to convert data into useful information. In the United States, the best available data sets on land cover, at a resolution of 30 metres, were last updated nearly 7 years ago. Globally, the picture is much less complete. Yet without accurate information, housing developers, foresters or other land planners can't make evidence-based decisions about which parcels of land to use for which purposes, and how much to leave untouched.

Meanwhile, almost 95% of oceans — which cover more than two-thirds of Earth's surface — remain unexplored. And scientists have described only around 1.5 million of the estimated 10 million species on Earth, and know little more than the names and collection locations for most of those[1].

AI systems could help in all of these domains. In fact, after seven years of working at the intersection of environmental and computer science, I'm convinced that the technology is now mature enough and the global environmental crisis acute enough to justify the creation of an AI platform for the planet.

What I'm envisioning is a portfolio of AI-infused 'Earth applications' available to people in diverse domains, from forestry to fisheries. These would be analogous to the application programming interfaces (APIs), such as those for searching or mapping, that have enabled people to build software services using components already made by engineers at technology companies.

EARTH APPLICATIONS

AI and environmental-science researchers are now applying algorithms to topics as varied as pollution modelling, agricultural-yield optimization, animal-migration tracking and Earth-system modelling. (Many of them will come together next month in Austin, Texas, at the 17th Annual Conference on Artificial and Computational Intelligence and its Applications to the Environmental Sciences.)

Numerous advances in these areas are coming from AI breakthroughs in non-environmental ones. For instance, the increasing demand for low-cost camera systems for smartphones and other devices has necessitated cheaper algorithmic (instead of hardware) approaches to improving image resolution. These 'super resolution' AI techniques can be used to improve Earth-system models by statistically 'downscaling' low-resolution projections of around 100 square kilometres to high-resolution ones of around 12 square km that are more relevant to local land-use planners[2].

In short, AI systems can now be trained to classify raw data from sensors on the ground, in the sky or in space, using categories that both humans and computers understand, and at appropriate spatial and temporal resolution. With enough data on which to train, and with human feedback, these systems can learn to tag photos, acoustic recordings and genetic information with species names; or to convert satellite imagery into information on water availability at a landscape scale.

Various organizations are already making impressive advances in Earth applications. iNaturalist and eBird, for instance, are identifying species using communities of citizen scientists. So far, iNaturalist's 575,000 members have recorded nearly 7 million observations of more than 128,000 species (see 'Assisted identification').

iNaturalist relies on experts to validate users' recordings, but deep convolutional neural networks are reducing the amount of expert input required — one of the biggest bottlenecks in the organization's growth. Currently, for more than 5,000 species of plant and animal, its algorithms are able to make an accurate identification four out of every five times, on average. And when the algorithms' top five predictions are considered for each sighting, the correct species is included 95% of the time[3].

"The global environmental crisis is acute enough to justify the creation of an AI platform for the planet."

The same types of deep-learning algorithm now used by iNaturalist's vast community of users are also helping ecologists to classify millions of underwater snapshots of corals. Other non-profit organizations, such as WildBook, are scouring the photographs of a particular species posted on social-media platforms such as Flickr. By identifying each individual animal, together with where and how often it's photographed, these algorithms are providing new ways of producing global population estimates for endangered species[4].

ASSISTED IDENTIFICATION

Experts validate species identifications made by users of the iNaturalist app, which launched when the organization was founded in 2008. Artificial intelligence is reducing the amount of expert input required.

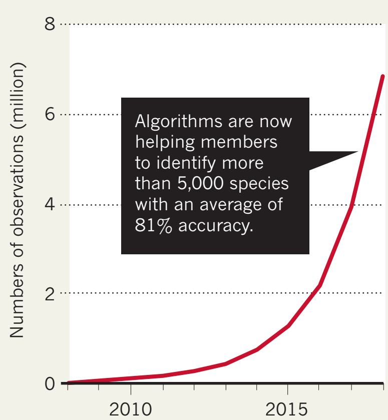

SOURCE: INATURALIST

Environmental applications of AI are already attracting the attention of investment firms, particularly in the agriculture sector. Using data from high-resolution satellites orbiting Earth, the tech company Descartes Labs in Santa Fe, New Mexico, is monitoring crop production globally. Through the application of neural-network algorithms to more than three petabytes of satellite and weather data (1 petabyte is 10^{15} bytes), users can obtain information about projected yields that outperform all other available forecasts in resolution and accuracy. Likewise, Blue River Technology, a California start-up, is using AI algorithms combined with high-resolution cameras attached to tractors and other field equipment to produce automatic weed-detection and weed-removal systems. In my view, the recent acquisition of Blue River by the US agricultural company John Deere for $300 million represents just the beginning of the agricultural industry's AI transformation.

Others are using AI to inform land-use management decisions — such as how to establish wildlife corridors for species such as lynxes and wolverines across the US Rocky Mountains as effectively and efficiently as possible. This is a difficult computational problem because there are so many possible solutions. But AI advances from operations research, developed for instance to work out how best to route traffic along a network of roads with the least cost and delay, are offering guidance[5]. Likewise, game-theory researchers are using AI to help law-enforcement officers to efficiently monitor the vast protected areas they are typically assigned to cover[6]. (AI algorithms identify what monitoring strategy will maximize the probability of patrols detecting illegal activity in a way that minimizes the probability of criminals working out what that strategy is.)

But these inspirational examples need to be the norm, not the exception. Too often, researchers publish exciting results about the application of AI to an environmental problem and those results are never translated into applications. Or an AI system is handed to a non-profit organization or government agency that lacks the resources and expertise to take advantage of it. Worse, traditional AI innovators in industry and academia rarely consider working on environmental applications.

TWO QUESTIONS

I believe that for every environmental problem, governments, non-profits, academia and the technology industry need to ask two

NICOLAS ASFOURI/AFP/GETTY

Artificial-intelligence systems can be trained to classify data, such as that from this motion-activated camera, into categories that humans and computers can understand.

questions: 'how can AI help solve this?' and 'how can we facilitate the application of AI?'

Governments. The public sector must ensure that environmental-data collection platforms continue to be produced, and that the data are made broadly and freely available in formats easily ingestible by AI algorithms. Various projects are already making key data broadly available. One is the European Space Agency's (ESA's) Copernicus programme, the world's largest single Earth-observation programme, and its associated Sentinel missions for land, ocean and atmospheric monitoring. Another is a dedicated satellite system for tracking wildlife, supported by ESA along with the German and Russian governments (see go.nature.com/2jwanje).

Governments, which tightly regulate the use of wireless communication channels, can also provide the funding for projects that are focused on the efficient use of microwave (cellular), radio and other spectra. Efforts such as the US Defense Advanced Research Project Agency's Spectrum Collaboration Challenge, the world's first "machine intelligence competition to overcome spectrum scarcity", could expand our ability to collect environmental information from remote locations (see go.nature.com/2bbhuo8). Ultimately, 'intelligent' sensors that can identify which spectrum to tap into at any moment, depending on availability and cost, could transform the efficiency of data-collection efforts.

Finally, government-led projects can incentivize investigators from diverse domains to work on AI approaches to environmental monitoring. For instance, the creation of a national, sustained ecosystem assessment in the United States — as originally called for in 2011 by the President's Council of Advisors on Science and Technology[7] — would motivate government agencies, academic researchers, non-profits and the tech industry to explore scalable solutions for monitoring the nation's natural resources[8].

Non-governmental organizations. The implementation of environmental sustainability projects often falls to resource-constrained non-governmental organizations (NGOs). Thus, the recognition within NGOs of the power of AI to aid their work — and of the importance of investing in technological solutions — is critical.

A set of core AI technologies, such as species-recognition algorithms, could enable multiple organizations to work more cheaply and efficiently. Currently, what's emerging largely depends on the priorities and interests of the tech sector. Domain-specific NGO consortiums (including those focused on water, say, or on biodiversity) could help to identify which problems funders, researchers and the private sector should prioritize. Such consortiums could also provide guidance on how to develop general infrastructure that multiple organizations could build to create their own specific applications.

Academia. Several changes would encourage more AI researchers in academia to focus on environmental sustainability, and to translate their findings into applications that others can use.

Computer scientists generally publish their work in conference proceedings. (Conference organizers use a peer-review process similar to that used by editors at traditional academic journals.) Awarding prizes at leading conferences for the best solution, rather than for the best paper, could motivate students and faculty members to invest more efforts in engineering.

Currently, AI researchers tend to test their algorithms on a few standard data sets. For instance, image-recognition software is generally tested on ImageNet, a database of around 14 million photographs. (Subjects include people, scenes and objects, as well as plants and animals.) Earlier this year, iNaturalist made its data set of 5,000 photographs of birds, mammals, amphibians and other taxonomic groups available for attendees of the Computer Vision and Pattern Recognition Conference in Honolulu, Hawaii. More environmental scientists should be prioritizing efforts to collate key data sets and make them available. Computer scientists can also help by communicating what kinds of data they need.

There is growing interest in funding AI research for environmental applications. Since 2008, the US National Science Foundation has funded the Computational Sustainability Network, a collection of AI researchers working on environmental sustainability. Submission tracks on sustainability are now included at leading AI conferences. And environmental journals are dedicating more space to the topic.

Fully engaging the academic AI community in environmental issues, however, will require the creation of dedicated academic centres focused on translating research to applications that could be used at scale. These could be attached to leading AI research institutes, such as the University of Southern California Center for Artificial Intelligence in Society, and involve institutions around the world.

Technology sector. The full participation of the technology sector is needed in efforts to provide key data in appropriate formats, as well as the algorithms, the infrastructure to train those algorithms on the data, and the means of making the end services available to as many people as possible. But why should companies whose primary business is not environmental sustainability engage?

At Microsoft, striving to democratize access to our technology is part of our culture. The initial mission to 'put a computer on every desk and in every home' has become a mission to 'empower every person and organization on the planet to achieve more'.

More broadly, market trends indicate that demonstrating social responsibility is good for business. Investors and owners of a collective $100 trillion in assets are now requesting that companies provide information on climate change, water and forest issues through the reporting platform run by the charity CDP and through the United Nations Principles of Responsible Investment. Also, for all major corporations, any type of uncertainty — financial, political or environmental — is generally bad for business. Conversely, in an increasingly resource-constrained world, sustainability is key to resilience; reducing the likelihood of environmental disasters safeguards supply chains, for instance.

The opportunity to work on societally important topics can be a huge pull for potential employees in a competitive industry. Since Microsoft launched AI for Earth, employees have approached me from all corners of the company, especially data scientists and engineers, wanting to apply their skills to issues that they care about.

Finally, there are major market opportunities in this area. High-resolution maps of natural resources are crucial to those working on environmental issues. They are also useful to the military, urban planners and providers of insurance, among others.

Microsoft is one of several tech companies working to provide technology to help solve environmental challenges. Google, Amazon, the software company Esri in Redlands, California, and the data provider Planet Labs in San Francisco, California, are among those striving to make their software, services and data available for environmental applications.

As companies ramp up their efforts, however, they must address the fact that today's AI technologies often come at a price and require computational expertise that puts them out of reach of many. Schemes to address access and education are crucial. Good examples of these include Microsoft's multi-day instructor-led AI for Earth education courses (see 'Microsoft for Earth') and Google's annual Geo for Good workshop, designed to help non-profits and other organizations make better use of Google's map data.

BEST PRACTICE

Researchers and developers must ensure that Earth applications are trustworthy, transparent and fair. The Partnership on AI — a consortium of technology companies and others founded in 2016 that came together this year to decide on best practices — should agree on a set of standards for producing algorithms for Earth observations. The partnership could draw on work from the 2015 Transformations conference in Stockholm. This produced a preliminary Biosphere Code Manifesto, with the aim of writing ten guiding principles for the application of algorithms for environmental sustainability.

Some will argue — rightly — that it was largely the technology of the first and second industrial revolutions that caused the environmental issues of today, and that we already know what we need to do: reduce the emission of greenhouse gases, the destruction of forests and the overfishing of oceans. But decisions about what actions to take will be easier to make — and less vulnerable to politicization — if we know what is happening on Earth, when and where. AI can help to provide that information.

"Artificial intelligence is not a panacea for environmental problems."

Time is too short, and Earth's resources too important, for companies such as Microsoft to ignore what is likely to be humanity's biggest challenge yet: mitigating and adapting to changing climates, ensuring the resilience of water supplies and sustainably feeding a rapidly growing human population, all while stemming an ongoing and catastrophic loss of biodiversity.

AI is not a panacea for environmental problems. But history will judge the success of the Information Age by our ability to deploy its resulting technology in stewardship of the planet. It's a big challenge, but an even bigger opportunity. ■

MICROSOFT FOR EARTH

A US$50-million programme for the environment

Microsoft's AI for Earth programme, announced this month, is dedicated to finding solutions to challenges relating to climate, agriculture, water and biodiversity — in which users of the resulting services and applications could grow to tens or even hundreds of millions of people. (Such services may be produced by Microsoft or by other organizations with Microsoft's support.)

It involves pairing the company's technical expertise on artificial intelligence (AI) projects for environmental monitoring and modelling with that of other organizations. For example, six months ago, we partnered with the California firm Esri, an international supplier of geographic information system software, and the Chesapeake Conservancy in Annapolis, Maryland. In this collaboration, we aim to apply deep-learning techniques to convert freely available high-resolution imagery from the US Department of Agriculture's National Agricultural Imaging Program into land-cover categories (forests, fields, water and impervious surfaces) at 1 metre resolution — all in a system that can be easily updated when new images are acquired.

When complete, this system should allow organizations such as the Chesapeake Conservancy to gain detailed insights into how the lands they care about are changing, and so achieve a type of precision conservation that is desperately needed. Knowing exactly where agricultural run-off is entering the Chesapeake Bay, for instance, could indicate which square metres of land should be replanted.

In another collaboration, with researchers at Johns Hopkins University in Baltimore, Maryland, and the University of Pittsburgh in Pennsylvania, among others, we are working on ways to use organisms such as mosquitoes as self-powered data-collection devices. The aim is to create an AI-powered metagenomics pipeline, and to gain insights about an ecosystem from information about the animals that mosquitoes feed on.

We have made our AI technologies available to more than 35 organizations in more than 10 countries in the past few months, for applications ranging from species-abundance modelling to live poacher detection in drone imagery. We have also hosted multi-day courses to ensure that environmental scientists and managers are able to take full advantage of the tools. **L.N.J.**

Lucas N. Joppa *is chief environment scientist at Microsoft, Redmond, Washington, USA.*
e-mail: lujoppa@microsoft.com

1. Mora, C., Tittensor, D. P., Adl, S., Simpson, A. G. B. & Worm, B. *PLOS Biol.* **9,** e1001127 (2011).
2. Vandal, T. Preprint at https://arxiv.org/abs/1703.03126 (2017).
3. Van Horn, G. *et al.* Preprint at https://arxiv.org/abs/1707.06642 (2017).
4. Berger-Wolf, T. Y. *et al.* Preprint at https://arxiv.org/abs/1710.08880 (2017).
5. Le Bras, R. *et al.* In *Proceedings of the 27th AAAI Conference on Artificial Intelligence* 1305–1312 (2013); available at http://go.nature.com/2nxgycl
6. Nguyen, T. H. *et al.* In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems* 767–775 (2016); available at http://go.nature.com/2atkzyi
7. PCAST. *Report to the President: Sustaining Environmental Capital: Protecting Society and the Economy* (White House, 2011); available at http://go.nature.com/2btene2
8. Jackson, S. T. *et al. Science* **354,** 838–839 (2016).