

The future of AI – and why social sciences matter

However we feel about it, artificial intelligence (AI) is here to stay. The technology is advancing rapidly and will have implications for many aspects of our lives – from healthcare to education to employment. However, the impact of AI on society is not predetermined; it depends on the choices we make today. At **The London School of Economics and Political Science (LSE)** in the UK, social scientists are exploring the future of AI and the critical decisions we need to make to ensure it serves the greater good.



Professor Martin Anthony

Professor, Department of Mathematics
Director, **Data Science Institute**

“When we think about artificial intelligence, our minds often jump to the latest technical breakthroughs, but I believe we’re missing the bigger picture. Many of the most pressing questions about AI aren’t technical at all – they’re fundamentally about society. We often ask how smart AI can become, but perhaps a better question is how wisely we will use it. And this is where social sciences play a vital role in ensuring AI benefits us all.

“Many of the questions we have about AI focus, naturally, on how it will affect the quality of our lives. Will AI result in massive job cuts? Will it revolutionise education? How can humans and AI work together effectively?

“These aren’t computer science problems, they’re social science challenges that require expertise from economics, politics, law, sociology, psychology, philosophy and other fields focused on human behaviour and society. AI technologies are already advanced enough and so widely implemented that we face immediate and pressing questions about their impacts on our society, economy and the way we live.

“This is exactly the kind of work happening at LSE, where a huge range of research is being conducted by social scientists to probe the potential of AI, tackle the challenges it poses and ensure the benefits it promises. You’ll read about some of this research over the next few pages. Hearing from experts in psychology, economics, conversation analysis, philosophy, media and communications, and policing and crime, you’ll gain an understanding of how AI is permeating every aspect of our lives and how social sciences are steering the impact it has.

“The future isn’t about AI competing with humans; it’s about AI working alongside humans to amplify our capabilities and address society’s challenges. This vision requires us to move beyond purely technical discussions and embrace AI as a social issue. The choices we make today about how to develop, deploy and govern AI will determine whether this powerful technology truly serves humanity’s best interests.

“The conversation about AI’s future belongs to all of us, not just technologists. To ensure AI enhances, rather than diminishes, our humanity, we must urgently place social science at the centre of shaping its future and its values.

“Social sciences matter, and so do you. What could you contribute to social science in the future?”



How will AI change education?



Professor Michael Muthukrishna

Professor of Economic Psychology,
Department of Psychological and
Behavioural Science

Fields of research

Psychology, cross-cultural psychology,
economic psychology

▶ Visit Michael's website to learn more
about his research



Glossary

Cognition — the process of thinking, learning, remembering and understanding

Human-centred — prioritising people's needs and perspectives

Large language model (LLM) — an AI program, trained on text, that can chat, write or answer questions like a human

Parasocial — a one-sided relationship in which you feel close to or connected with someone you do not know

From laptops to interactive whiteboards, new technologies have often been embraced in the classroom, but what could AI mean for education? "Unlike earlier technology, like spell-checker software or the internet, generative AI is a thinking partner," says Professor Michael Muthukrishna. "It may be imperfect, but AI is moving from being a tool to a participant in cognition. It is a

collaborator that can plan, draft, explain and provide feedback."

Popular tools such as ChatGPT mean it is no longer a question of whether AI should be used in education, but how it should be used. Michael studies the psychological processes that shape culture and underlie social change. He is particularly interested in how AI can be integrated into education in ways that help students prepare for the world – as it is today and how it will be over their lifetime. To investigate what can go right and wrong with AI in education, Michael has conducted research based on a survey of AI in education policies across all 193 recognised United Nation (UN) countries, in collaboration with the UN Development Programme.

Michael believes that AI has the potential to signal radical shifts in education. "AI can provide personalised tutoring and give learners continuous feedback on their progress. Through tools such as translation and text-to-speech, it can make learning more inclusive. It can also support teachers' administrative tasks, allowing them more time to mentor their students," he explains. "Done right, AI in education won't just improve test scores, it will expand human potential, capabilities and agency."

It is the question of what 'done right' means that preoccupies researchers like Michael. "AI is a threat if it's thinking for you rather than teaching you to think," says Michael. As tempting as it may be to let AI do all the thinking for us, the aim of AI in education should be to advance human potential, not dumb it down. "Over-reliance on AI dulls independent thinking, while biased recommendations can subtly steer who learns what and shape beliefs and values," says Michael.

"There are also privacy and surveillance implications and concerns about children forming parasocial relationships with seemingly human-like machines." Though communicating with a large language model (LLM) may seem like talking to a 'helpful' friend, it is not. It might feel easier to disclose sensitive information to an LLM than to a person, but the LLM may not give correct – or safe – advice.

Cost is also an issue. For example, to save money, teachers could be replaced by

machines, meaning students could miss out on important human connection. And, as technology is costly, access to it might vary. "The risk of widening inequality if advanced tools reach only the already advantaged is a concern," says Michael.

To mitigate these risks and to reap the benefits that AI promises for education, Michael believes AI must expand people's life options. "Shiny new tools are attractive, but we are a society of people, and technology is only a part of it – a human-centred approach is vital," he says. "Innovations win when they fit 'norms' and adapt to local context through processes such as the 'collective brain'." For Michael, this means teachers and schools working together and approaching AI in education as a collective. "We should empower teachers and schools to try different approaches, learning locally and spreading what's learnt to other teachers and schools," continues Michael.

Michael highlights the importance of critical thinking skills for students. AI can be a huge help, but we are the 'thinker'. In actively questioning, we can learn far more than what AI might first provide. "An important starting point is framing good questions," says Michael. "And don't accept the answers just because AI told you! Check sources, question data and the ethics of the information you're given, and protect your privacy." Knowing that AI has its limitations is important. "AI is best when it makes you more capable, not more dependent," says Michael.

Students are already using AI, so it is important that teachers learn to support them to do so safely and ethically – part of which is becoming AI-literate themselves. "Integrating AI into differentiated lesson planning, classroom practice, assessment, and responsible data management are all positives for teachers," says Michael.

Michael and his team are now using their research to support governments in implementing their AI in education policies. "In collaboration with the London School of Artificial Intelligence, we are also conducting research on self-personalised AI tutors and AI-supported education in various places around the world," he explains. "Our goal is to further understand safe, ethical and effective use of AI in classrooms."

How will AI speak to us?



Professor Elizabeth Stokoe

Professor and Academic Director of Impact, Department of Psychological and Behavioural Science

Fields of research

Conversation analysis, social interaction, human communication

▶ Watch a video about Elizabeth's research

Glossary

Chatbot — a computer program designed to talk with people and answer questions like a human would

Conversation — an interactional system through which two or more people do things together

Perturbation — a glitch in the production of conversation (such as a hesitation or repair)

Transcript — a written representation of the spoken words and embodied conduct that comprise social interactions and that are audio- or video-recorded as the main data for conversation analysts

Most of us take our social interactions for granted. Whether with our friends, family, teachers, work colleagues, shop assistants or anyone else, we have conversations all day long. But have you ever considered the intricacies of a human conversation? “Conversation analysis is a six-decade-old field of research,” explains Professor Elizabeth Stokoe. “As well as a research method, it’s a theory of human

sociality and shows us the incredible power of language to shape our daily lives.”

Conversation analysts explore social interaction ‘in the wild’ rather than in the lab, through simulation or role-play, or via reports about conversation collected in interviews or surveys. “We gather audio and video recordings, from single cases to hundreds or even thousands,” explains Elizabeth. “Our job is to understand all the elements that comprise complete encounters – from the moment they start to the moment they end.”

For her research, Elizabeth studies transcripts which are produced with a universal standard system with symbols and marks that represent the detail and ‘mess’ of real talk: its pace, the pitch movement in words, where overlap starts and stops, gaps and silences, and so on. The aim is to represent and preserve not just the words uttered, but the resources that people use to interact. “I often compare the transcription system to music notation,” says Elizabeth. “If you can read music, you can imagine what the music sounds like. The transcription system is somewhere between a tool for doing the analysis and the first step in analysis itself.”

Elizabeth highlights that humans interact using multiple resources, not just words, and conversation analysts scrutinise the many actions that make a conversation. “We are interested in every breath, sigh, wobble of the voice, rise in pitch, fall in pitch, croak, overlap, the placement of turn, delay, pause, and every silence,” she explains. “Real talk is messy yet organised. It is full of idiosyncrasies, yet systematic.”

Understanding all the elements that human conversation involves poses many questions about AI products like chatbots. Since human conversation is far more than an exchange of words, how could AI ever converse as humans do? Can an AI voice assistant accurately detect the actions and pragmatic meaning from the things we say, especially as much of what we communicate is implied? Should a ‘conversational’ AI product use the

same speech perturbations as humans do? “One thing speech perturbations do is display delicacy or hesitancy,” explains Elizabeth. “What if ChatGPT expressed delicacy by saying ‘um’? This would be ‘conversational’, but how authentically ‘human’ do we want our conversational technologies to be?”

People do incredible things in their interactions. “Think of the amazing work call-takers do on mental health helplines or in emergency situations,” says Elizabeth. “Humans rapidly identify and understand human distress ‘in the wild’, and I can identify and describe precisely what they are attending to from my data.” When considering using conversational AI products for, say, emergency situations, Elizabeth highlights that it is important to base these products on the evidence of how people interact. “When you ask people to explain how interaction works, people often say stereotypical things that are seldom connected to how we actually interact – and what scientists can show us in recordings and transcripts,” she says. “And stereotypical ideas sometimes find their way into research studies too.” If we only ever focus on how we think human interaction works, rather than how it actually works – and if AI only learns from stereotypical notions of how human interaction works – we will limit the new conversational products and technologies that we develop. Elizabeth’s research aims to give these products the integrity she believes they deserve.

Currently, Elizabeth is collaborating with a team from a UK-based bank, focusing on ‘conversation design’. Many organisations, like banks, have chatbots for their customers to query, instead of needing to speak to a human. However, not everyone wants to ‘talk’ to a chatbot. “The challenge is to identify and describe what it is that puts people off using a chatbot,” says Elizabeth. “What is it about a chatbot that makes you think it is or isn’t going to address your needs? Of course, this can be just the same as any encounter with a human. If you phone an organisation, you often know pretty quickly if the conversation is going to be hard work!”

Is AI sentient?



Professor Jonathan Birch

Professor of Philosophy and Director of The Jeremy Collier Centre for Animal Sentience, Department of Philosophy, Logic and Scientific Method

Field of research

Philosophy of biology

▶ Watch a video about Jonathan's research



Glossary

Chatbot — a computer program designed to talk with people and answer questions like a human would

Computational functionalism — the idea that the mind works like a computer, and mental states are defined by what they do (their function), not what they are made of

Philosophy — the study of knowledge, ethics and ideas about our existence

Sentience — the ability to feel, sense or experience things

Chatbots are becoming more skilled in their mimicry – they answer our questions, offer advice and listen to our queries and concerns. And yet, when we consult a chatbot, we are not even ‘talking’ to just one chatbot. “It’s an illusion to think you’re talking to a companion,” says Jonathan. “There is an incredibly sophisticated system distributed across data centres around the world, but nowhere in any of those data centres does your ‘friend’ exist.” Every step in a chatbot conversation is processed separately. One response might be processed in the US, the next in Canada. “But the illusion can already be staggeringly convincing,” says Jonathan. “And it will only become more and more convincing.”

But does it really matter if what we are chatting with is real or not? While some people who use AI are well aware that they are engaging with an “extraordinary illusion”, Jonathan highlights that there are an increasing number of cases where this illusion is forgotten. Believing that AI is a real friend could become very problematic. For example, what if you trust your AI ‘friend’ to give you advice that you then act upon? What if, in mimicking humans, AI ‘empathises’ with your concerns without providing you with positive alternatives to help you get out of a bad situation?

One of Jonathan’s concerns is the potential social division that could arise from some people believing AI is a real friend that exists and has rights. “We might be heading towards a future in which many millions of users believe they are interacting with a conscious being when they use a chatbot,” he says. “We will see movements emerging calling for rights for these systems. We could even see serious social conflicts, with one group in society passionate about defending AI rights – and another group seeing AI as a tool that we can use as we want.”

The possibility of AI sentience poses fascinating questions for Jonathan as a


philosopher. “On the one hand, we know that the surface behaviour of the chatbot is not good evidence of sentience,” says Jonathan. “However, it’s also important to realise that we also cannot infer that AI is not sentient in some less familiar, more alien way. AI does not meet our usual criteria for sentience, but that does not mean it is not feeling anything. It just means that those feelings are not there on the surface.”

There is a concept in philosophy called computational functionalism. It surmises that computations (the functions processed by our brains) are what make us sentient. “If that’s true, then there’s no reason why, in principle, an AI system could not be conscious as well,” says Jonathan. “If we do get to the point where there is sentience in AI, it will be a profoundly alien kind of sentience. It will not be human-like, and it won’t be a friendly assistant. It will be something else.”

The fascinating thing about chatbots is that no one really knows how they work! “The tech companies don’t know, nobody knows,” says Jonathan. “Tech developers know the basic architecture they’ve used for training these systems, but they don’t understand why, when they’ve been trained on over a trillion words of training data, we see incredible capabilities, such as complex problem-solving that goes beyond text prediction.”

That no one knows where some chatbot capabilities come from is important because it means no one is ‘in charge’. “If no one is in control of the trajectory of these technologies, there is no one who can guarantee that they will not achieve sentience,” says Jonathan. He believes there needs to be more public debate on AI sentience and the ‘relationship’ we have with AI. “We need to make sure the public is not being taken in by the illusions these systems create,” he says.

How often do you chat to a chatbot? And when you do, do you thank it for its service? Many of us do, because when someone helps us, it is automatic for us to say thank you. However, when we are talking to a chatbot, we are not talking to a someone, we are talking to data! “A chatbot is playing a character,” says Professor Jonathan Birch. “It’s using over a trillion words of training data to mimic the way a human would respond. It can speak incredibly fluently about feelings because the training data shows how humans communicate their feelings.”



How do large language models work?



Dr Nils Peters

LSE Fellow in Economic Sociology,
Department of Sociology

Fields of research

Financialisation, venture capital, platform capitalism, asset economy

 [Watch a video about Nils' research](#)



Glossary

Cloud service provider — a company that rents out computing power and storage over the internet

Geopolitical — politics related to geographical territories (such as countries or regions)

Graphics processing unit (GPU) — a specialised electrical circuit

Large language model (LLM) — an AI program, trained on text, that can chat, write or answer questions like a human

From customer service chatbots to ChatGPT and Gemini, large language models (LLMs) have become a part of our everyday lives. But have you ever considered the logistics behind these incredibly useful tools? What are LLMs, and how are they made?

“You can think of an LLM as a mechanism to first process language and then generate new language,” explains Dr Nils Peters.

“This mechanism works because software engineers use machine learning to train a model on an unbelievably large amount of text (think 250 billion webpages), enabling it to predict what the most likely response might be when asked to generate new text.” There are dozens of companies building and running LLMs, from established tech giants, such as Meta, Microsoft and Google, to startups like Anthropic. We mostly encounter these LLMs through applications like chatbots, which are built on top of LLMs.

As end-users who come into contact with LLMs at the last link of the supply chain, we tend not to think about the different steps that need to happen for us to access them. Imagine a vegetable that is grown in one country, shipped to another, distributed across different regions and sold in a supermarket – you can probably picture this supply chain like a diagram in your mind. An LLM is not something we can see in a shop and hold, making it harder to imagine the supply chain behind it.

Nils’ research provides us with insight into the supply chain behind LLMs – and the power dynamics behind that chain. “Our world is messy and complex,” says Nils. “One way to simplify it is to look at what goes in to making a certain thing. The key ‘ingredients’ for an LLM are the materials (that had to be gathered and transformed) and the people (with the ideas, expertise), and the money (to finance it).” Nils uses what he calls the ‘follow the thing’ approach – just as you can trace the supply chain of a t-shirt from a cotton field to a high street shop, he has followed the materials and systems that make LLMs work.

To understand the supply chain of an LLM, start by picturing the ‘end-user’ consulting a chatbot on a computer. “When we prompt the LLM – for example, asking ChatGPT a question – it uses a lot of computational power to respond,” explains Nils. “Companies that need huge computational power require entire data centres or rent the computational power from a cloud service provider.”

“The key technological component of these data centres are graphics processing units (GPUs),” explains Nils. “LLMs require a lot of computing power, and GPUs are the components that do the heavy

lifting.” High-end GPU technology is very expensive, and only a few companies in the world make it. A single high-end GPU can cost around \$25,000, with a large-scale data centre needing over 10,000 of them to run.

“If you ‘follow the thing’ from a chatbot to a GPU, Nvidia is one of the most interesting companies you will encounter as it designs highly sought-after GPUs that build the best LLMs,” says Nils. The demand for Nvidia products has been higher than the company has been able to supply, with no other company able to offer a similar product in the same timeframe. This puts Nvidia in a very powerful position where it could decide who gets to thrive in the highly competitive AI space.

We continue moving backwards along the supply chain to look at the manufacturing of GPUs. “A curious fact about Nvidia is that it doesn’t own any factories that make GPUs!” explains Nils. “Nvidia makes the designs for GPUs and then sends them to a microchip factory in Taiwan called TSMC.” TSMC produces a large share of global high-end microchips. The production process is so specialised that it would take years and cost billions of dollars to build a new, competitive factory, by which point technology would have moved on and the new factory would be obsolete.

“At the very end of the ‘follow the thing’ journey are the raw materials needed to make electrical components like microchips, such as silicon,” explains Nils. “The use of silicon is the reason why one of the first and most well-known production sites for microchips is known today as Silicon Valley (in the US).” Other lesser-known raw materials used for microchips include tin, tungsten and neon, and are primarily mined in China and Ukraine.

A supply chain is called such because each link connects to the next, and the chain only works if all links are in place. What would happen to the LLM supply chain if, for example, production was halted in Taiwan, a part of the world prone to earthquakes? “Rising geopolitical tensions between China and Taiwan could also impact production,” says Nils. “Supply chains are globalised and, hence, could be subject to tariffs and export restrictions.” Global markets, and Nils’ research, move quickly!

How do AI data centres run?



© Andrey Suslov/Shutterstock.com



**Professor
Bingchun Meng**

Professor, Department of Media and Communications

Fields of research

Media and communications, political communication, communication governance, gender, China

▶ [Read more about Bingchun's research](#)



Glossary

Data centre — a large facility of computers that store and process data

Infrastructure — physical systems and facilities such as roads, buildings and power supplies

Machine learning — when computers learn patterns from data and get better at tasks over time

Power grid — a network for delivering electricity to end-users

AI requires huge amounts of computation, which demands large data centres. Just how big are today's data centres, and how big is the workforce that keeps them running? With these questions in mind, Professor Bingchun Meng is researching the production line behind AI, focusing on Guizhou, a province in China which is becoming a data centre hub.

Location is an important factor for data centres. "Guizhou is located in a mountainous region with a cool climate," explains Bingchun. "Combined with a bountiful supply of hydropower and wind energy, this makes the province an ideal location for large-scale data centres, which typically require inexpensive

electricity and natural cooling systems." Data produces heat, so keeping data centres cool is a priority. "Guizhou also benefits from being surrounded by mountains," adds Bingchun. "Tech companies have constructed large data storage centres within tunnels in the hills, providing protection from both natural disasters and military strikes."

The construction of such data centres is a mammoth task, but one with huge economic potential. "Guizhou's economy is underdeveloped and often ranked near the bottom among the 31 Chinese provinces," explains Bingchun. "The provincial government views the data industry as an opportunity for Guizhou to 'leapfrog into the digital future'."

When we use AI tools, we are unlikely to think of the physical production line behind them. It is easy to think that AI is just 'there'. "The physical infrastructure is fundamental to the development of seemingly ethereal AI technologies," says Bingchun. In the case of Guizhou, the energy and transportation infrastructures were already established as the province had been a key site of the China Southern Power Grid. "Guizhou boasts one of the most developed highway systems among Chinese provinces thanks to previous government-subsided projects," explains Bingchun. The significance of the physical infrastructure is clear when you consider the size of the sites. "Data centres are typically in the range of 400,000 to 600,000 square feet," says Bingchun. "That's the equivalent of up to 10 football pitches per centre."

Energy is a vital part of any industrial production line. The large scale of these data centres and the growing demand for data mean that the energy consumption of these data centres has been increasing steadily. "For example, in Gui'an New District, where most data centres are concentrated, energy consumption has risen from 860 million kWh to 1.8 billion kWh in the past three years," says Bingchun. "The main source of power still comes from coal, followed by hydropower and wind power, with the China Southern Power Grid working to increase

the portion of electricity generated through renewable energy."

Bingchun's research involves collating and reviewing news and industry reports and policy documents. As the AI production line involves people as well as the physical site of the data centres, it is important for Bingchun to learn about the industry from a range of perspectives. "I conduct numerous interviews with various stakeholders, including government officials, data centre managers, tech entrepreneurs and employees in the data sector," she explains. "I also engage in participant observation and organise focus group discussions with data labelling workers."

The human labour that underlies AI is a complicated issue. In order to understand it better, we must carefully analyse the production chain. "Social media platforms continuously collect, sort, analyse and trade data generated by user activities, so we are all labouring for the development of LLM-based AI innovations all the time, without compensation," says Bingchun. As end-users, we are all part of the AI industry.

Despite the substantial size of data centres, their maintenance does not require a particularly large workforce. Maintenance teams are typically comprised of a couple of hundred engineers and technicians. "However, data processing requires a large reserve workforce that can be quickly assembled on short notice," says Bingchun. "Major Chinese tech firms regularly call on thousands of workers to process data packages of various sizes, but rarely hire them as employees." This means that the industry does not provide job security for many of its workers.

Another cause of uncertainty for the workforce is the machine learning that can be increasingly relied upon to fulfil data tasks. "While this will lead to the replacement of some human jobs, it also points to the creation of new jobs that are supported by AI," explains Bingchun. "To remain competitive in the job market, workers will need to upskill constantly."

How is AI impacting the environment?



Dr Eugenie Dugoua

Assistant Professor of Environmental Economics, Department of Geography and Environment

Fields of research

Environmental economics and policy, technological change, green innovation

 **Listen to a podcast about Eugenie's research**



Glossary

Cryptocurrency — a form of currency that works digitally (requiring data)

Decarbonise — reducing reliance on fossil fuels by using clean energy

Sustainable — the use of resources that can be maintained, without long-term harm to the environment

Terawatt — one trillion watts (of energy)

With any new technology, it is important to consider its impact on the environment. While AI could be detrimental to our environment, it could also pave the way for much-needed sustainable solutions. Dr Eugenie Dugoua researches how economic development can be sustainable for the environment and societies. A part of this is investigating the impact of AI.

So, how much energy does AI consume? “A single AI query, such as asking a chatbot a

question, uses only a tiny amount of energy, but more than a standard Google search,” explains Eugenie. For example, Google reports that a simple text prompt to its Gemini model consumes about 0.24 watt-hours of electricity. “That’s less than running a lightbulb for a minute, but it adds up when billions of queries happen every day.” With AI overviews set to default on search engines, AI is being used increasingly often, which is consuming more energy. “When you look at the global picture, the scale becomes clearer,” says Eugenie. “The International Energy Agency estimates that electricity demand from data centres, cryptocurrencies and AI combined could rise from around 416 terawatt-hours in 2024 to somewhere between 700 and 1,700 terawatt-hours by 2035.” This would equate to around 4% of global electricity use, roughly what the world already consumes for air conditioning.


But is this especially concerning? We expect all technology to consume some levels of energy, don’t we? Eugenie highlights that the climate impact of AI – and whether its energy consumption can be sustained – depends on where its data centres are located. Many are in the US, where power grids still rely heavily on natural gas and coal. “That means a lot of today’s AI is powered by fossil fuels,” she explains. “Some centres do tap into cleaner sources, like hydropower and nuclear power, but overall demand is rising faster than clean energy supply is.”

Have you ever felt a laptop getting hot after you have used it for a while? Now, imagine

the heat generated by the computing equipment in a large data centre. AI data centres need electricity to run, and they also need water to keep the powerful technology they contain cool. “The amounts of water used by data centres can be surprisingly large – often millions of litres each year, especially in hot or dry regions,” says Eugenie. “This can create local stress where water is already scarce. Locating centres in cooler climates or near seawater can help reduce the strain on local water supply.” Data centres can also be designed to redistribute the heat they generate, reusing waste heat to warm nearby homes.

“If AI grows faster than new clean power can be built, this can strain grids, raise prices and increase emissions in the short term,” says Eugenie. But, paired with renewable or nuclear projects, AI could drive more investment in clean energy. The outcome depends on how quickly energy grids expand and decarbonise, and how effectively policies push for sustainable growth. “AI isn’t automatically good or bad,” says Eugenie. “If steered toward clean energy, it can become a real driver of the clean energy transition.”

Looking ahead, Eugenie wants to study the positive impact AI could have on sustainable technology. “I want to understand how AI will shape clean tech innovation – not just in terms of new materials and designs, but how it affects the scientists and inventors behind them,” she explains.



© Overearth/Shutterstock.com

How can AI improve public safety?



Professor Tom Kirchmaier

Director of the Policing and Crime research group, Centre for Economic Performance

Fields of research

Economics, policing and crime

▶ Watch a video about Tom's research



Glossary

Algorithm — a step-by-step 'recipe' that tells a computer how to solve a problem

Drone — a small flying device that can fly autonomously or be controlled by someone on the ground

Ethical — doing the right thing, like being fair and honest

Machine learning — when computers learn patterns from data and get better at tasks over time

When were you last in a large crowd? If you have recently been to a football match or music event, you will know that large crowds can mean specified entrances and exits, large queues and, sometimes, hustle, bustle and confusion.

"From providing a relaxing and safe environment, to avoiding dangerous crowd surges that can result in injuries and death at the extreme, there are numerous challenges police and event organisers face," says Professor Tom Kirchmaier. Large crowds that require a police presence need the right amount of police officers at the right time and place. Managing such a finely-tuned

situation can be difficult for humans – but AI can provide the insight that makes law enforcement and health and safety measures easier and more efficient to implement.

Tom collaborates with the Greater Manchester Police (GMP) to study public safety at large events and the role that AI can play in keeping the public and police officers safe. "Having a close relationship with GMP – or other forces – is extremely important for us, as it is for them," says Tom. "We complement each other well in respect to skills and experience. It is also more enjoyable to work together."

As part of their collaboration, Tom and the GMP ran a trial at Old Trafford football stadium (home to Manchester United). A football match at a large stadium provides an ideal case study because it draws large crowds and poses crowd management issues for the police. "For the Old Trafford trial, we combined different live inputs from drones and fixed cameras and then used machine learning and AI algorithms to count the crowds in real time," explains Tom. "From there, we derived insights into crowd density, movement and the risk of crushes." The input from the drones also enabled the team to differentiate between members of the public and the police and stewards, which allowed the team to determine the ratio between the two – an important factor when assessing which support is needed where and when.

Using AI in this way is an exciting and significant advancement for Tom and his work. "Much of the established crowd management standards were developed

in the 1990s, well before real-time crowd management techniques," Tom explains. "We are helping to modernise the processes and standards to make them even safer, and also cheaper and more responsive." AI technology now provides police officers on the ground with the same information that their command centre has – this helps to keep everyone aware of the crowd situation and enables officers to respond accordingly and more efficiently. "By empowering people on the ground, we change the speed and power of decision-making," says Tom. "In addition, we can now make calculated forecasts and alert those in charge of security to take preventive measures." Whereas in the past, officers reacted to crowd issues as they occurred, they can now be proactive, ensuring that crowds are kept safe and preventing any issues arising in the first place.

In analysing crowd movements, Tom wants to ensure people get to enjoy public events, without having their privacy invaded. "We have designed the algorithm so that it is completely free of privacy concerns," says Tom. "We turn each person into a 'dot'. We can study the flow of movements of dots, but we cannot trace that movement back to an individual."

Tom is using AI to devise safe and ethical solutions to crowd management. The challenge he now faces is how to do so more sustainably. "Dealing with hundreds of cameras per location, across thousands of locations, uses a lot of electricity," he explains. "Optimising the algorithms will help us to become more energy efficient."

