KING'S INSTITUTE FOR ARTIFICIAL INTELLIGENCE

KING'S College LONDON

Annual Report 2023

Showcasing artificial intelligence at King's

Contents

The King's Institute for Artificial Intelligence is working to bring together and coordinate our artificial intelligence research and education across the university, to provide more opportunities for multidisciplinary teams to rise to the challenges of AI, and to give a greater coherence to the overall agenda at King's in a supportive AI community. To tell us about your work, or find out more about ours, email **ai-institute@kcl.ac.uk**.

Message from the Director

The last year has seen an explosion in AI developments, with a torrent of news around issues concerning the vast potential of AI in general, the tendency towards normalisation of generative AI in a short space of time, and the increasing concerns surrounding ethics, risk and regulation of AI.

At King's College London, with our work spanning the breadth of the university and reaching out into every faculty, we are well placed to deliver advances in technology and applications and to address the important issues arising in relation to impact and implication for society more generally.

Back in the summer of 2022, Google engineer Blake Lemoine, who had been tasked with testing the company's chatbot (or large language model), LaMDA, concluded that it was sentient. LaMDA told Lemoine that it knew how it felt to be sad and angry, and that it feared death. Lemoine was fired for violating confidentiality policy, but the incident created a major public discussion about the technology and its implications. That was back then. Among other things, the advent of ChatGPT in November 2022 (and its successors and competitors) has fundamentally changed the landscape. AI is routinely being used by the public, causing consternation in universities and schools, triggering a new wave of investment in the next generation of tools and techniques, both by governments and venture capital, and is rarely far from the headlines.

Two of the so-called godfathers of AI have spoken in public about their concerns relating to AI. There have also been public letters from the great and the good of AI, first calling for a pause on the training of large language models (or, more precisely, those systems that are more powerful than GPT-4), and second warning about the potential risks posed by AI to humanity. Yet for many experts, concerns of existential risk seem always ready to distort a more sensible conversation. Indeed, there are more immediate questions to be answered in relation to the impact and implications of AI, around issues of bias, ethics and regulation, or even around liability and copyright, which are pressing and require attention now.

At a national level, the UK government and funding agencies have been active, with a steady stream of funding commitments totalling not just tens of millions, but reaching into the hundreds of millions, for AI centres for doctoral training, for AI hubs across

different disciplines, for additional Turing AI fellowships, for a national Responsible AI hub, and much more. A white paper on the UK approach to AI regulation was published in the spring of 2023, and an AI Foundation Model Taskforce was announced in June 2023 to lead AI safety research in order to take advantage of the benefits of foundation models while ensuring they are safe and reliable, with £100m of initial government funding.

Against this background, at King's the last year has seen many relevant developments that position us well to contribute to the arising challenges and to take advantage of the opportunities. For example, the new Digital Futures Institute, based in the Faculty of Arts & Humanities, but also bringing in science, engineering, and health, aims to advance research and public understanding of the contexts, consequences and possibilities of technology by reframing technical challenges as human and social challenges that require broader, and more collaborative, ways of knowing and responding. A new £2.6m EPSRC Digital Health Hub was awarded to King's, led by the School of Biomedical Engineering & Imaging Sciences, to support the development of new digital technologies and reduce the time it takes for these to benefit patients. Most recently, Kate Devlin of the Department of Digital Humanities is the King's lead in a £31m partnership with Southampton and several others on Responsible AI, securing £5m for King's. Finally, it's worth noting that the Department of Informatics has launched new undergraduate degree programmes in Artificial Intelligence with 2023 entry imminent.

Meanwhile, the King's Institute for Artificial Intelligence has been working to support the King's AI agenda across all these areas. The Institute led a successful application for King's to become a member of the Turing University Network in April 2023, but King's interactions with The Alan Turing Institute are much wider with, for example, Steven Niederer, of Biomedical Engineering & Imaging Sciences, being appointed as Co-Director for the Turing Research and Innovation Cluster in Digital Twins. The Institute has continued to bring together colleagues in different disciplines in community activities, and mediated interactions across departments and disciplines within the university, as well as establishing a strong public presence for King's through the AI exhibition that ran through May and June 2023, and the King's Festival of AI that drew in over 1,000 people to events across five days in May.

As we come to the end of the second year of operation, much has been achieved. However, just as the developments and challenges in AI continue to grow in scale and pace, so too must we redouble our efforts to ensure that the unique qualities of King's in our multi-dimensional approach to AI are visible to ourselves and to the wider world, and that we persist in breaking down the barriers between disciplines that are too often a constraint to the scope of achievement. In this context, I am delighted that as I depart King's for a new adventure, I leave the Institute in the excellent hands of Carmine Ventre who will take over stewardship in the immediate future as Institute Director. I have no doubt that King's has a major role to play in the development of AI, nationally and internationally, both in the technical advances that improve lives and in consideration of the implications in ensuring responsible and safe application and deployment.

Professor Michael Luck
Director
King's Institute for Artificial Intelligence
July 2023

Our vision

	THEME
1	**COORDINATION** Provide a point of research coordination within King's for funding calls, consortium building, selection panels, matchmaking, etc
2	**COMMUNITY** Bring the research community together across King's by delivering an events programme and internal communications for networking and collaboration in support of a King's AI research culture and community
3	**REPRESENTATION** Provide all-King's representation for external partnerships, for example with The Alan Turing Institute, collaborating institutions, donors, funding agencies, and more
4	**EDUCATION** Coordinate development of educational offerings, drawing on the strength across the institution, potentially to include undergraduate modules, online modules, executive education, and less formal staff upskilling
5	**ENGAGEMENT** Develop opportunities for wider public engagement, working on new initiatives and in collaboration with, for example, the Policy Institute, King's Culture, the Entrepreneurship Institute, Science Gallery London, etc
6	**VISIBILITY** Create a coherent and visible external presence for King's as a whole, across the breadth of the university, eg website, news stories, marketing collateral

	EXAMPLES 2022–23
1	Turing Network Development Award ● Turing University Network ● Funding proposal mediation service in support of, for example, UKRI AI centres for doctoral training and UKRI AI hubs ● King's Institute for AI affiliate scheme ● Participation in internal selection processes for demand management
2	Community event: Applying AI Methods to Microscopy and Imaging, January 2023 ● Community event: Creative AI: Theory and Practice, January 2023 ● King's Business School AI Away Day, May 2023 ● Digital Futures Institute launch: ChatGPT and the future of writing machines, April 2023 ● Methods Seminar Series: Methods for Collaboration (SSPP): 'Prospects for AI at King's', January 2023
3	*AI and Cognitive Reinforcement for Mission Decision Support*, panel in Military Information Supporting Operations (MISO) meeting: AI and Innovation, November 2022 ● *Artificial Intelligence: Towards Safety and Trust*, presentation for BT Leaders Lab, September 2022 ● All King's representation for Turing Network Development Award delivery ● All King's representation for the Turing University Network proposal, application and delivery
4	Mediation of Informatics MSc interdisciplinary projects with external supervisors for collaboration opportunities ● Support for Digital Research and Data Strategy regarding AI educational opportunities ● Visit from National Institute of Education, Nanyang Technological University, Singapore, with: Informatics; CUSP London; Education, Communication & Society, October 2022
5	*Bringing the Human to the Artificial* exhibition, May – June 2023: 23 fixed exhibits, 20 live interactive demonstrations ● King's Festival of Artificial Intelligence, May 2023 ● Public lecture: *How should policymakers be thinking about the AI and data agenda?*, Hetan Shah, Chief Executive of the British Academy, April 2023 ● Public lecture: *How to Quantify AI Explainability for Improved Safety and Trust*, Avi Rosenfeld, Jerusalem College of Technology, January 2023
6	*Bringing the Human to the Artificial* exhibition, May – June 2023 ● The King's Festival of Artificial Intelligence, May 2023 ● Mediated news opportunities, for examples see below ● 253 subscribers to internal mailing list ● 1,334 followers on Twitter and 367 followers on LinkedIn

Examples of mediated news opportunities include: *How artificial intelligence is matching drugs to patients*, BBC News: Heba Sailem of the Institute of Pharmaceutical Sciences, discusses how AI might transform the global pharmaceutical industry with the BBC ● *Is this the future of humanity?*, Reasons to be Cheerful podcast: Kate Devlin of the Department of Digital Humanities, guest with Ed Miliband and Geoff Lloyd ● *Apple Knows You Didn't Mean to Type 'Ducking'*, New York Times: Yulan He of the Department of Informatics comments on updates to Apple's autocorrect feature ● *Is the AI apocalypse actually coming? What life could look like if robots take over*, Evening Standard, Kate Devlin of Digital Humanities and Michael Luck of Informatics contribute to feature on the AI apocalypse

Bringing the Human to the Artificial exhibition & the King's Festival of Artificial Intelligence

The King's Institute for Artificial Intelligence held an artificial intelligence exhibition, in collaboration with King's Culture, in the Arcade in Bush House in May and June 2023 and a five-day festival of events in May.

The exhibition, *Bringing the Human to the Artificial*, showcased cutting-edge research from across King's exploring the effects of artificial intelligence (AI) on different aspects of our lives.

It featured multiple perspectives on AI at King's, across social sciences, law, arts, humanities, security, science, engineering and health, representing work from nearly all faculties. The 23 fixed exhibits included case studies of research in AI in different disciplines, creative installations, original artwork, and video and audio showcases of AI. Exhibitors spanned academic and research staff, students and artists collaborating with King's researchers. Across five 'demo days', the exhibition also hosted 20 live interactive exhibits, allowing visitors to engage directly with King's exhibitors, students and staff alike. Researchers, students and representatives from the technology, culture and policy sectors came together to mark the launch of *Bringing the Human to the Artificial* on 3 May 2023.

To accompany the *Bringing the Human to the Artificial* exhibition, the King's Institute for AI also initiated the first ever King's Festival of Artificial Intelligence over five days at the end of May. This extended the reach of the exhibition and maximised the opportunities for increased visibility and successful public engagement arising from both exhibition and festival. Over 1,000 people attended a series of events dedicated to exploring the latest developments in artificial intelligence and some of the challenges and opportunities they pose. The festival programme covered a broad range of topics, including the law surrounding AI, robotic surgery, human-AI relationships, art, and the impact of AI on video games and music. There were also opportunities to participate in King's AI research through a series of demonstrations, including several created by PhD students from the UKRI Centre for Doctoral Training in Safe and Trusted AI.

'This exhibition, and the AI festival, are part of a conversation about ensuring AI is used for good, to benefit us all.' *Professor Michael Luck, Director of the King's Institute for Artificial Intelligence*

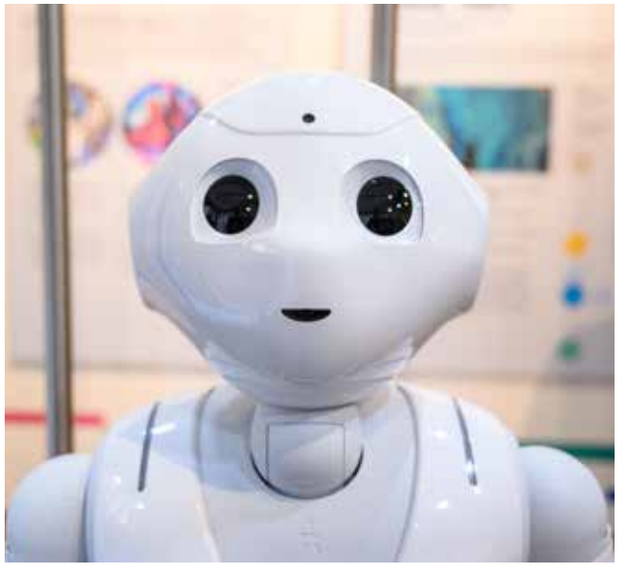

Leading King's academics working on aspects of AI and related areas delivered keynote public talks throughout the festival. Professor Luca Viganò from the Department of Informatics explored how fairy tales illustrate key concepts of AI and cybersecurity, Professor Dan Hunter from the Dickson Poon School of Law examined the law and ethics of generative AI, and Professor Prokar Dasgupta from the School of Immunology & Microbial Sciences outlined the digital future of surgery with audiences, with a live demonstration of the Versius Virtual Reality system used to train surgeons in robotic surgery. In the context of AI and art, Digital Humanities featured strongly, with Professor Joanna Zylinska exploring human and machine creativity, and Dr Kate Devlin considering the potential and pitfalls of fostering friendships and intimacy with computer software and hardware.

The festival also included family-friendly events suitable for younger audiences such as a 'Build a Robot' workshop and a 'Fantastic (Artificial) Beasts and Where to Find Them' workshop, the latter also devised and led by students from the UKRI Centre for Doctoral Training in Safe and Trusted AI, which encouraged children to

'Over 1000 people attended a series of events dedicated to exploring the latest developments in artificial intelligence and some of the challenges and opportunities they pose.'

explore AI principles and concepts through games and other fun activities.

Feedback received suggested that the exhibition and festival activities were very successful in both exposing the successes and challenges of AI, and raising the profile of King's AI research. For example, one respondent said: 'What has changed is not so much my opinion on AI but my opinion on how King's can contribute to the public discourse on AI and to education as well as research challenges, tapping into the existing expertise in a coordinated and creative way, and reaching out to a rich network of partners.'

Indeed, the exhibition and the festival allowed King's AI research to reach new audiences. For example, 45 per cent of visitors to the arcade exhibition who returned feedback, and 64 per cent of festival attendees who returned feedback, had no affiliation to King's. Only 20 per cent of these visitors to the exhibition had any degree level experience of AI or a related field and 80 per cent of all attendees (regardless of affiliation or degree of specialist knowledge) felt that the exhibition made AI research at King's more accessible to audiences outside King's. Meanwhile, the King's Festival of Artificial Intelligence webpages attracted over 20,000 page views and over 10,000 unique visits by the end of June 2023.

The Institute's exhibition and festival complemented a programme of activities taking place across King's in the space of AI and related technologies across the summer and autumn of 2023.

'What has changed is not so much my opinion on AI but my opinion on how King's can contribute to the public discourse on AI and to education as well as research challenges, tapping into the existing expertise in a coordinated and creative way, and reaching out to a rich network of partners.'
Festival feedback

Bringing the Human to the Artificial highlights

CASE STUDY

RELIABLE AND TRUSTWORTHY AI FOR DEFENCE

The UK has an AI strategy for defence and national security. The military is experimenting with new autonomous platforms and with the doctrine and concepts that might allow their effective employment. Autonomous systems are already at work in data processing and intelligence analysis. However, there are concerns...

Indeed, there are concerns about the application of AI to defence, both from an ethical standpoint, and in terms of performance. The debates underway about the ethical implications of using AI in national security, including in decisions about the employment of lethal force, include some of the following ideas.

Human decision-maker
Even in an era with pervasive and increasingly sophisticated machine cognition, the importance of the human decision-maker endures, reflecting an ethical and cultural desire to preserve meaningful human control.

Skillsets
There is a need to upskill the workforce involved in defence, and more broadly to promote AI literacy in wider society so that the concerns and implications can be better understood.

Mass and scale
While AI may allow maintenance of qualitative advantage at scale, there remain profound practical questions, for example of how to organise (and lead) a mixed platoon of humans and autonomous machines.

Vulnerabilities
AI systems are potentially vulnerable to electronic warfare countermeasures like jamming or spoofing, and are susceptible to offensive cyber warfare. There are also difficulties of assurance and trust such as when AI is susceptible to bias in training data. AI that is used for defence must be reliable and trustworthy, and sufficiently transparent that humans can understand its decision-making.

CASE STUDY

FORECASTING FUTURE MEDICAL EVENTS

CogStack

Over the past four years, the AI world has surged ahead with large language models, also known as *foundation models*, which can be adapted to achieve many linguistic tasks.

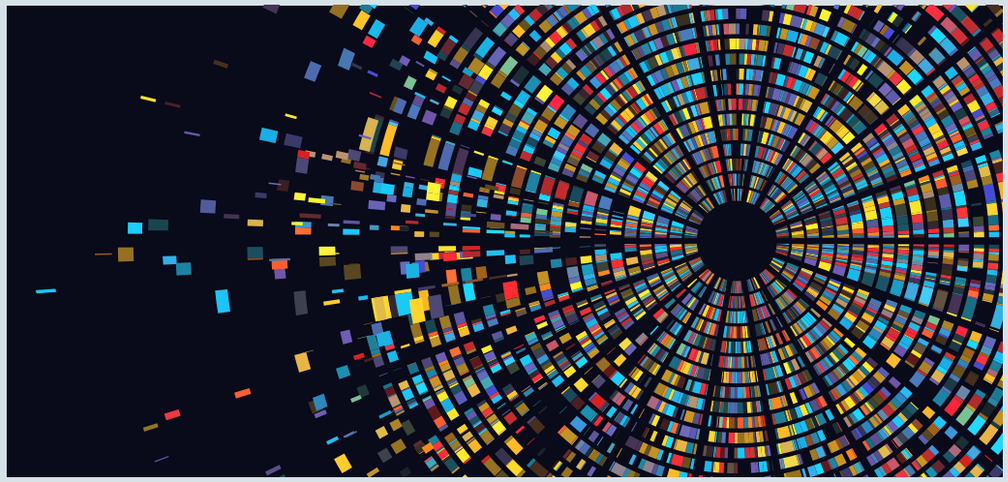

You've probably seen many articles in the media recently about some of these models (ChatGPT, DALL·E 2), that can write essays and software code, and generate art and films, and which have many other capabilities.

With the NHS at breaking point, a critical question is whether these AI approaches could be used to improve care. Hospital records hold detailed information about each patient's health status and general clinical history, a large portion of which is stored within the unstructured text. Temporal modelling of this medical history, which considers the sequence of events, could be used to forecast and simulate future events, estimate risk, suggest alternative diagnoses or forecast complications.

We developed **Foresight**, as part of the CogStack platform, a novel GPT-based pipeline that is trained on NHS data to forecast future medical events such as disorders, medications, symptoms and interventions.

On tests in two King's Health Partners' NHS Trusts (King's College Hospital, South London and Maudsley) and the US MIMIC-III dataset, **Foresight** performed well when set challenges by clinicians. The model is being used in many areas including: real-world risk estimation; virtual clinical trials; clinical research to study the progression of diseases, simulate interventions and counterfactuals; and for educational purposes.

Going forward, foundation models for medicine provide the potential for a diverse integration of medical data that includes electronic health records, images, lab values, and biologic layers such as the genome and gut microbiome.

CASE STUDY

LEARNING WITH HUMANS TO PLAY GAMES

Artificial intelligence has seen great success in many applications, including image classification, natural language processing, and more. The key element in these examples is the machine learning from static datasets with clear desired outputs.

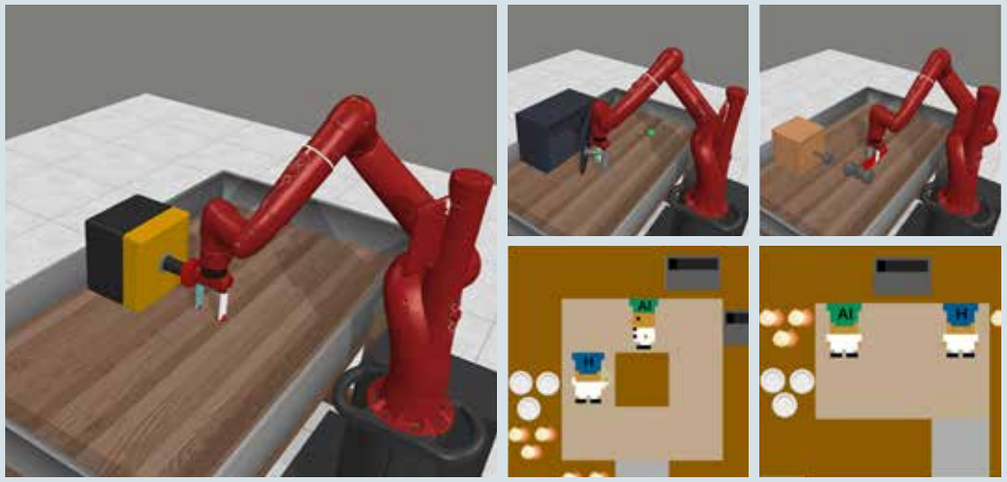

However, instead of learning knowledge from static datasets, it is also possible (and indeed necessary and desirable) to learn through feedback from interactions in their environment. This is especially important for machine learning models to be deployed in the real world, since the models themselves will impact each other, turning their individual decision-making into a multi-agent problem. Agent learning in a complex multi-agent world is therefore a fundamental challenge for next-generation AI.

Reinforcement learning is one form of AI that learns by interacting in the world. When interacting with humans, it can become especially effective. Our work shows how an agent can learn to coordinate with humans without knowing about interacting with humans in advance, in the context of the video game, Overcooked.

This is part of a wider research agenda concerned with the cooperation of machines with other machines and with humans. It has also involved machines learning from humans for robotic control tasks, and the use of reinforcement learning to play multiplayer video games such as StarCraft II.

CASE STUDY

THE CREATIVE AI LAB

Increasingly, artists are using AI technologies such as machine learning in their creative practice. However, these technologies are often seen as a ‘black box’, lacking the visibility and transparency of more traditional techniques and media.

Some artists ‘open’ this black box by making the processes at work within machine learning systems visible to audiences. The **Creative AI Lab** highlights the importance of such work and conducts interviews and studio visits to explore and support it. We have found that the artists who engage in these techniques are contributing to **Explainable AI** and that their work invites a general public to engage with technical aspects of AI systems and to acquire urgently needed AI literacy. Examples of artists using these techniques are as follows.

The **Creative AI Lab** is a collaboration between the Department of Digital Humanities, King’s College London, and the Research & Development Platform of the Serpentine Galleries, and was founded by Mercedes Bunz and Eva Jäger in 2019.

Adam Harvey traces the origin of data sets that train facial recognition.

Allison Parrish allows visitors to manipulate language systems by playing with different sliders or weights.

Refik Anadol visualises the range of data within a cultural archive.

dmstfctn explores the training of machine learning systems used in cashierless Amazon supermarkets.

Use the QR code above to find out more about the artists.

CASE STUDY

DESIGNED BY AI

With the advent of generative AI through tools like **Stable Diffusion**, **ChatGPT**, **DALL·E 2**, and **Midjourney**, AI has entered the field of creative industries – previously exclusive to humans – and can design things like paintings, interior design, or advertising. This introduces important questions and opportunities for businesses.

For example, we know very little about how people would cope with AI as a (competing) source of artistic creation. In response, our work examines when and why people are likely to adopt and buy creative products that have been designed by AI.

In our studies, we show that people are more likely to choose human-made than AI-made products, but that there are considerable differences between different product categories regarding the acceptance of, and preference for, AI-made products.

We uncover the psychological processes underlying people's preferences for human-made products, in that people perceive AI-made products as threatening and therefore engage in a process of dehumanisation. Consequently, we investigate what can be done to increase people's acceptance of AI-made products.

We find that if products are worse than expected, people prefer AI-made products, but if they are as good as or better than expected, human-made products are preferred. Marketers can leverage this by effectively steering the price or the ambiguity of the product offering.

CASE STUDY

JUDGED BY THE MACHINES?

Could we build a better justice system using artificial intelligence? Could we trust AI to decide legal cases? Too often, justice is not served. Delays leave victims, witnesses and the accused waiting years for a decision on their case.

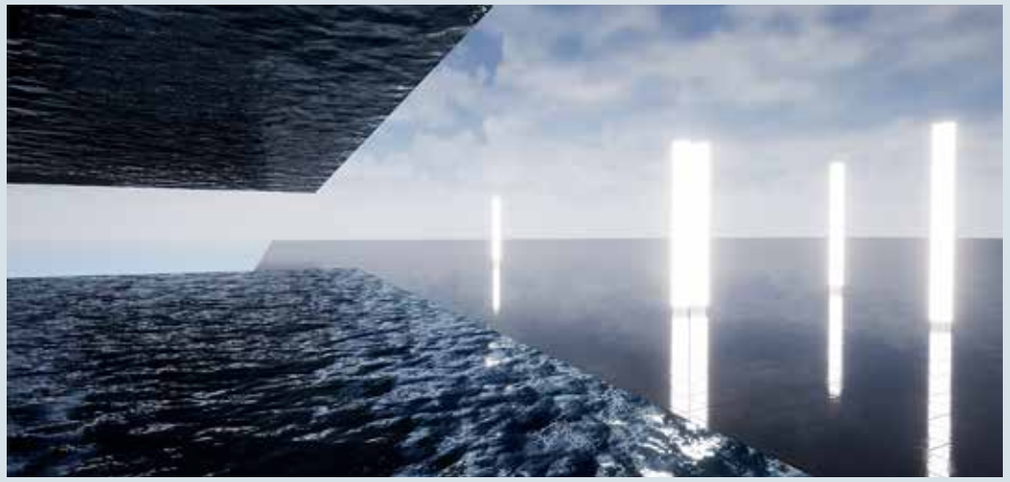

Even if algorithmic decision-making could help with this delay, public trust in human-computer interaction in justice is under-explored. Can public confidence be maintained in a move to automation? To help answer these questions, our work uses an interdisciplinary approach, informed by legal analysis and empirical data, to produce a model that can decide and sentence a traffic offence.

We incorporate methods from the social sciences to build mathematical models in a collaborative way. Over the next year, members of the public, judges and lawyers will interact with the system and give detailed feedback so that we can understand how automated decision-making and AI might genuinely help to improve justice.

In collaboration with Vivienne Griffin from Somerset House Studios, a striking 3D game environment, **MERCY**, has been created to convey the research in an accessible way.

IMAGE DESCRIPTION (CLOCKWISE) The Fake Haven, 2021, Vivienne Griffin, Digital video, sound, 16:30 (still); MERCY, 2022, Vivienne Griffin, Digital video, sound (still). MERCY is an interactive game environment developed as part of this project, metaphors within the digital space are used to convey law and justice concepts and promote discourse.

CASE STUDY

ORAL MICRO-VASCULAR DIAGNOSTIC IMAGING

In healthy mouths, the smallest blood vessels (capillaries) form tiny loops just under the mucosal surface. This allows the surface lining to survive. In disease, however, these loops change. We are using AI to observe capillaries and to identify different tissue patterns without human intervention and bias.

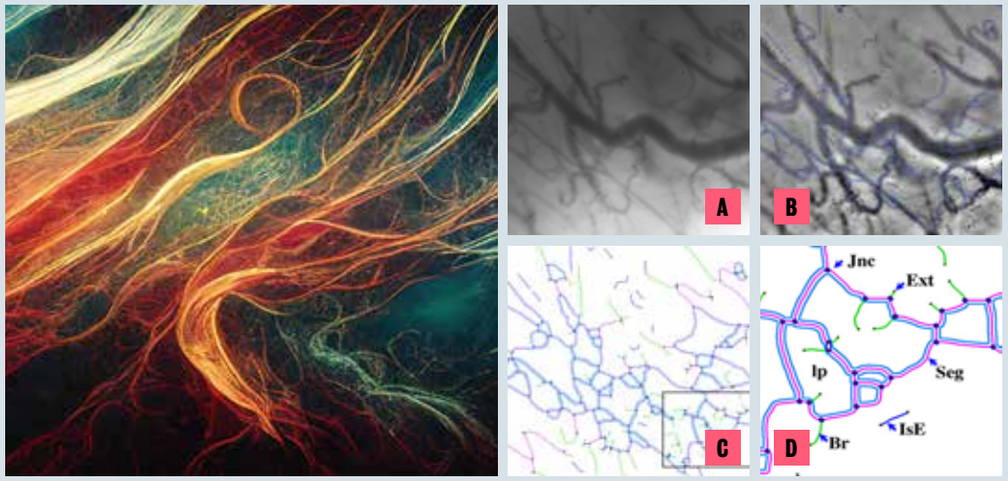

In inflammation, capillaries become wider and more can open up, which is why inflamed mucosa looks red; in cancer conditions, the micro-vessel shapes change further and new vessels are driven to grow (angiogenesis) by the effect of the cancer itself as it demands constant energy supplies to grow.

Over the last decades as medical instruments have been developed to examine tissues in vivo and in situ, several qualitative reporting and descriptive classification schemes have been introduced. However, these are fraught with human observer bias and potential inconsistency of interpretation as different opinion enters the classification assessment.

We have developed imaging instruments that can see the capillary circulation directly. Using machine learning, we are starting to be able to segregate malignant, pre-malignant, inflammatory and normal tissue patterns, without human intervention or the biases of human image assessment. This should allow us to ultimately look for indications of malignancy, without any local anaesthetic or other surgical tissue harvesting being necessary.

IMAGE DESCRIPTION (CLOCKWISE) Example of how the raw image (A) is processed to a clean image and the software plots the map (B), then the map is abstracted (C), and the data characteristics are extracted from the map, ready for AI interpretation of characteristics of the image (D).

A: initial raw microvascular capillary image. **B:** modelling and analysis of the skeleton (derived from binarising image A after background subtraction), superimposed to optimised version of image A. **C:** model of the microvascularisation composed of different vectoral elements. **D:** amplified image from C with Jnc (junction), Ext (extremity), Seg (segment), IsE (isolated element), Br (branch) and Ip (loop). Vectorial elements are indicated as some examples of the quantitative summary characteristics extracted from the raw image in A, that can be analysed objectively in the process, to indicate the likely pathology driving the behaviour of the micro-circulation.

CASE STUDY

ETHICS OF AI-BASED MEDICAL TOOLS

IN SEARCH OF AUTONOMY, BENEFICENCE, NON-MALEFICENCE & JUSTICE

There are very many examples of AI-based medical tools that can improve understanding and management of health outcomes for better patient diagnosis and treatment.

Yet with AI-based tools becoming more common as part of daily routine in the medical domain, practitioners and researchers must consider the ethical implications of integrating AI recommendations on treatment choices and diagnosis of disease.

From the development of AI tools to their potential deployment into clinical care, we identify several ethical challenges that closely connect with the four ethical principles on which any medical professional should base their own conduct. These are as follows.

Maximising patient autonomy in informed treatment decisions:
Respect for Autonomy

Acting in a patient's best interests:
Beneficence

Treating patients as ends in themselves:
Non-Maleficence

Distributing medical benefits fairly:
Justice

Among others, these principles raise questions around authority and ethical *responsibility* in the case of physician and machine collaborating, avoiding *dehumanisation* so that patients are not regarded as mere mechanical systems, and *transparency* so that people can understand the rationale for diagnoses and decisions.

CASE STUDY

HARNESSING MOBILE DATA TO ASSESS DISEASE AND RELAPSE

Data from mobile devices can give a full and continuous picture of a person's health at a level of detail that has not been possible until now. With such data we need AI.

This continuous monitoring produces a huge amount of data and we need AI approaches to help extract learning and insight.

The RADAR-CNS programme brought together clinicians, researchers, engineers, computer scientists and bioinformaticians from all over the world to explore the use of wearables and smartphones to help manage depression, multiple sclerosis and epilepsy. The programme collected 62 terabytes of data from 1,450 participants over seven years.

To manage the data, researchers developed the open-source RADAR-base platform and used AI techniques to identify and test possible indicators of health and relapse.

RADAR-CNS

Examples include the following.

A machine learning approach correctly detected severely convulsive epileptic seizures from wearable data.

AI techniques were used to extract features from Bluetooth data to approximate rhythms of behaviour and show they could predict depression severity.

A machine learning method was developed and applied to Fitbit heart rate data to identify people with COVID-19.

Data analysis continues, with a range of outputs in the pipeline. The programme focussed on the use of remote measurement technologies for depression, MS and epilepsy but they have many other potential applications. The RADAR-base platform is now used by over 30 projects.

CASE STUDY

CHANGING THE FUTURE OF CANCER

One in two people will be diagnosed with cancer at some point in their lifetime. If cancer is identified early, the chances of survival are more than 80 per cent at 10 years but, when diagnosed late, survival falls to less than 30 per cent at 5 years.

Yet diagnosing cancer is extremely challenging. Unlike other diseases, there is no single identifiable symptom or test that can alert doctors to a potential cancer diagnosis.

C the Signs is a clinical platform that uses artificial intelligence mapped with the latest evidence and research to identify patients at risk of cancer at the earliest and most curable stage of the disease. Covering the entire spectrum of cancers, C the Signs can identify patients at risk of cancer in under 30 seconds, accurately detecting the type of cancer(s) the patient is at risk of and recommending the most appropriate next step.

Founded by King's alumna Bhavagaya Bakshi and co-founder Miles Payling, **C the Signs** was supported by the Entrepreneurship Institute's King's20 Accelerator.

C the Signs uses a conditional logic model with combinations of risk factors, symptoms, clinical signs and investigation results to determine if a patient is at risk of cancer and what the next step should be (eg a blood test, scan or referral to a specialist). We are currently building the most advanced machine learning model to detect cancer even earlier.

EXHIBITS

Aha!
This lenticular print illustrates the closing gap between human and artificial intelligence.

Come Closer?
Caged robots prompt visitors to consider their responses to robotic and AI technologies.

Puck – An Automated Game Designer
A computationally creative AI that uses different AI techniques to design games.

The AI Gallery
A collection of art created with machine learning and other AI techniques.

Artificial Soul
A radio programme asks if there are any aspects of humanity that could never be programmed into a machine.

Neuromatic
A video offering a commentary on human and machine vision, exploring what it means to see the world.

The Chaos of Raw Unknowing
A short film that explores trust and the danger of anthropomorphising AI.

Learning to Play Games with Humans
A recent attempt to train agents for human-AI coordination, using the game Overcooked.

MERCY
Developed by Vivienne Griffin, MERCY explores the role of AI and algorithms in judicial systems.

Good Intentions, Bad Art
A podcast episode asking 'Are good intentions a necessary ingredient of good art?'

Who's Responsible?
An installation prompting visitors to consider environmental, health, and social justice issues associated with AI.

King's joins the Turing University Network

King's became a member of the Turing University Network in April 2023.

As one of the initial 36 universities joining the network from previous Turing university partners and Turing Network Development Award recipients, King's will engage with The Alan Turing Institute and its higher education, industry and government partners.

The Alan Turing Institute serves as the UK's national institute for data science and artificial intelligence, and the Turing University Network is a newly established initiative that aims to create better connections across the data science and AI landscape and promote collaboration among universities in the UK. It provides a platform for universities to engage with the wider data science and AI community, including the Turing and its broader network of partners, with the goal of accelerating research and fostering real-world impact.

The Turing University Network is a key component of the new strategy of the Turing that aims at changing the world for the better with data science and AI. The network supports the Institute in achieving its three ambitious goals: advance world-class research and apply it to national and global challenges, build skills for the future, and drive an informed public conversation.

Operating on an engagement-based model, the Turing University Network promotes an open and inclusive approach. Universities within the network can shape the Turing's programmes through the Annual Meeting, influence strategic direction and contribute to the national agenda. More specifically, the network provides opportunities for collaboration, knowledge exchange, and showcasing of research through conferences, networking events, research clusters, and initiatives. It also enables universities to connect and collaborate with each other, share best practices, pool resources, and engage in academic exchange or mentorship programmes.

'King's will engage with The Alan Turing Institute and its higher education, industry and government partners, which are working to generate real-world impact through data science and AI research.'

Yulan He of the Department of Informatics, who also holds a five-year Turing AI Fellowship funded by UKRI, has been appointed as the King's Academic Liaison lead for the network. She will work closely with the King's Institute for AI, and its Director, Manager and Steering Group, to develop the relationship and to maximise its value to staff and students.

To find out more about the partnership with the Turing, and to stay updated with the Turing University Network updates and latest news and opportunities, sign up to the King's Institute for AI mailing list.

Sign up to the King's Institute for AI **mailing list** for Turing updates

Events

2022–23 saw the King's Institute for AI continue its programme of activities, both bringing external speakers into the university and creating opportunities to break down barriers by bridging between disciplines.

For example, in April 2023, the Institute welcomed Hetan Shah to King's to give a talk entitled, 'How should policymakers be thinking about the AI and data agenda?'. As well as being chief executive of the British Academy, the UK's national academy for the humanities and social sciences, and a board member of Our World in Data, a website that provides long run data on global challenges such as poverty, Hetan is also a member of the board of The Legal Education Foundation, a philanthropic body supporting a better justice system and visiting professor at the Policy Institute. He helped to found the Ada Lovelace Institute which seeks to promote AI and data for public good, and he was its Vice Chair until the end of 2022. In his talk, he discussed how policymakers should respond to the opportunities and challenges of the data and AI agendas, and implications for areas such as the justice system. He also considered methods by which policymakers could better understand public thinking on new technologies and the role AI might play in making policy more effective.

The King's Institute for AI and the UKRI Centre for Doctoral Training in Safe and Trusted AI also hosted Avi Rosenfeld of the Jerusalem College of Technology for a talk entitled 'How to Quantify AI Explainability for Improved Safety and Trust'. Avi outlined several ground-breaking directions for how explainable artificial intelligence (XAI) can be quantified and applied for improved AI safety and trust, also providing a use case for how novel metrics can be used within a medical application.

Two community events were also organised by the King's Institute for AI. The community events sought to develop a community at King's centred around

'The workshop brought together biologists, computer scientists, mathematicians, physicists, microscopists and others from King's to explore how artificial intelligence can be applied to microscopy and imaging.'

a particular AI theme or application. They presented a valuable opportunity for the King's community to learn about AI research and practice at King's and to generate new connections across the university.

Building on the success of community events in the first year of operation, the King's Institute for AI's second community workshop, 'Applying AI Methods to Microscopy and Imaging', was organised by Robert Knight, Mads Bergholt and Andrew King in January 2023. The workshop brought together biologists, computer scientists, mathematicians, physicists, microscopists and others from King's to explore how artificial intelligence can be applied to microscopy and imaging. Highlights from the talks presented included using AI for analysing patient imaging data, the idea of uncertainty calibration in AI for decision-support, training and optimising AI models for analysing spectral and CT data, and applying deep learning to image reconstruction of super-resolution data.

In January 2023, the King's Institute for AI's third community workshop, 'Creative AI: Theory and Practice', organised by Joanna Zylinska, Mercedes Bunz and Daniel Chávez Heras, with the Creative AI Lab (King's and the Serpentine Gallery), brought together colleagues from disciplines across King's (and beyond, with participants from LSE, National Gallery X, Oxford, Cambridge, Leeds and São Paulo) – and not restricted just to arts and humanities researchers but also including computer scientists and engineers – to discuss what the concept of 'creative AI' means to them and to show some specific projects through which it can be enacted. The presentations offered diverse perspectives and included reflections on: machine vision and the notion of seeing; creativity in different forms; misinformation; machine learning; prompt engineering; and art and AI more generally. Themed panels at the end of each session aimed to open up discussion across participants and break down boundaries between disciplines.

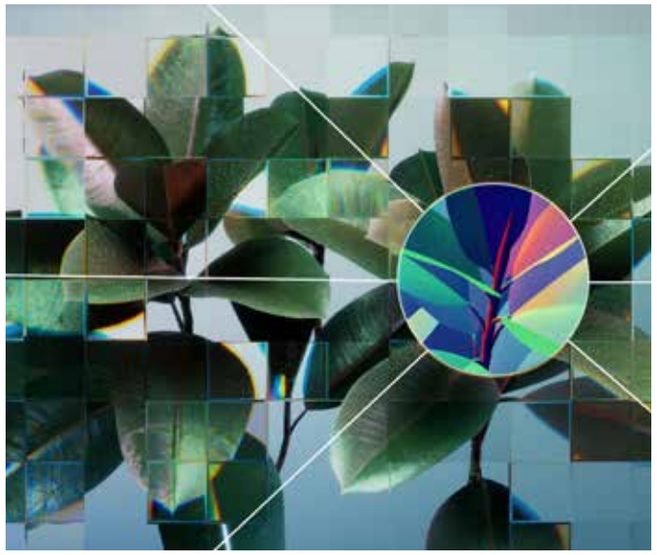

'Great diversity of talks and people. I enjoyed the opportunity to learn from and meet folks from different disciplines.'
Feedback about the 'Creative AI: Theory and Practice' community event

Above Image by Alan Warburton/ © BBC/Better Images of AI/Plant/ CC-BY 4.0

Phases of development

Three main phases of development are envisaged for the Institute: Initiate, Thrive and Sustain. These terms characterise the natural phases of the Institute's growth over an extended period.

A measure of success for each phase will be how well a diverse and engaged AI community at King's achieves visibility and traction, how it flourishes, and how it remains embedded in King's future successes across research, education and service. The diagram below exemplifies one possible set of indicators across the Institute's development.

In 2022–23, major headline initiatives – including the *Bringing the Human to the Artificial* exhibition and the King's Festival of AI, and the King's AI community events programme – have been key to both consolidating the Institute's presence at King's and showcasing AI at King's.

INITIATE 2021–23	**THRIVE** 2022–25	**SUSTAIN** 2024–27
Create an identity for the Institute and establish its presence within King's and externally.	**Establish educational activities** to develop training materials for staff upskilling and courses.	**Explore educational opportunities** for alumni and training of external stakeholders.
Establish collaborative events across faculties, building an AI community at King's.	**Develop and promote case studies** highlighting the best of AI at King's.	**Engage with government and industry** in support of regulation for ethical autonomous systems.
Provide a coordination and representation service for AI at King's.	**Develop best practice guidance** to support ethical research and case studies.	**Support the full range of King's activities**, including research and technology transfer.

Table 1. Example indicators of the three main phases of development for the Institute

About us

King's College London has a vibrant and expansive community in artificial intelligence, with work taking place in every faculty across King's. Activity spans core AI (its technical underpinnings), applications and societal implications. The King's Institute for Artificial Intelligence is dedicated to connecting researchers, educators, students, policymakers, and the wider public to foster collaboration, advance research, and develop the understanding and application of AI in society.

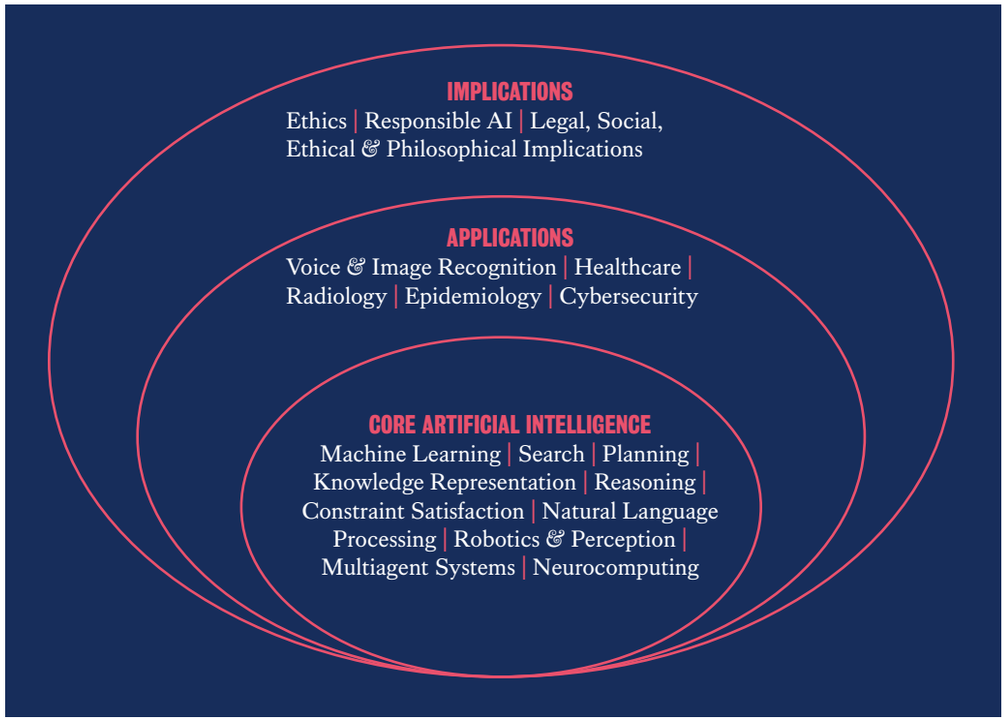

Figure 1. The breadth of AI: artificial intelligence, applications and implications

● Tell us about your work in artificial intelligence at King's. **Complete our survey.**

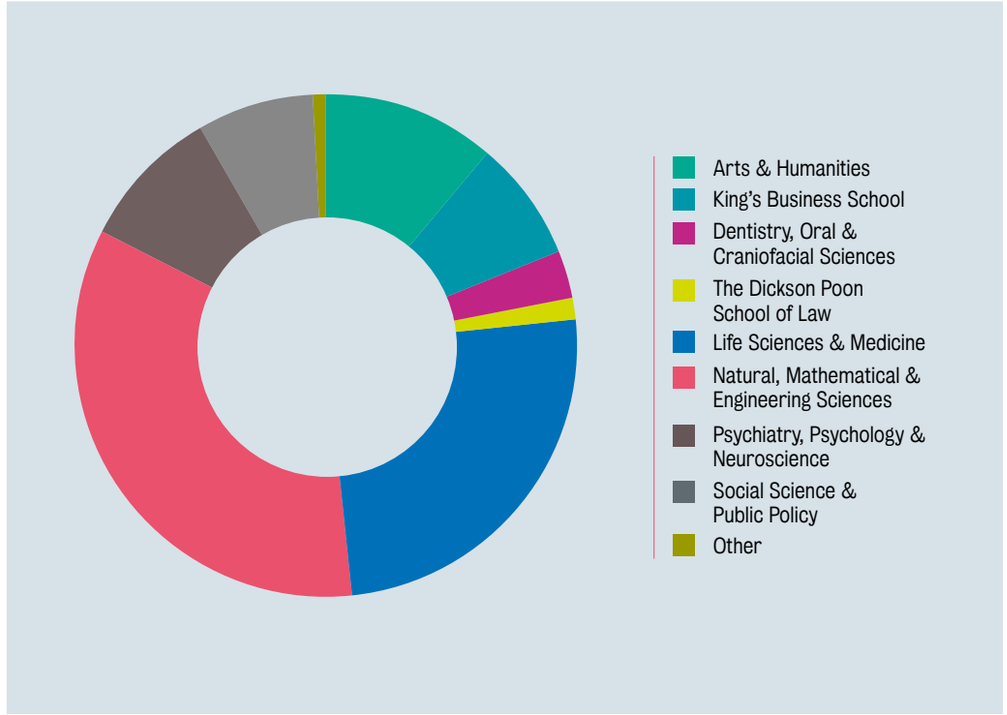

Figure 2. A snapshot from June 2023 showing the distribution of AI activity across King's faculties. Data is pulled from our ongoing mapping exercise.

INSTITUTE LEADERSHIP

Professor Michael Luck
Institute Director
michael.luck@kcl.ac.uk

Professor Rosalyn Moran
Deputy Director
rosalyn.moran@kcl.ac.uk

Professor Carmine Ventre
Institute Director
(incoming from July 2023)

INSTITUTE OPERATIONS

Maria Flores
Institute Manager
(Cover, May 2023 – February 2024)

Rose Hepworth
Institute Manager

Jess Keating
Communications & Engagement Manager

Karly Maxwell
Senior Business Support Officer

Lucy Bailey
Programme Production Manager
(Exhibition & Festival, March – June 2023)

Contact us via
ai-institute@kcl.ac.uk

STEERING GROUP 2022–23

Dr Stefan Bernritter
Reader (Associate Professor) of Marketing, King's Business School

Dr Mads Bergholt
Senior Lecturer in Biophotonics, Faculty of Dentistry, Oral & Craniofacial Sciences

Dr Arianna Ciula
Director & Senior Research Software Analyst at King's Digital Lab, Faculty of Arts & Humanities

Dr Mateja Durovic
Professor of Law and Technology, The Dickson Poon School of Law

Dr Crina Grosan
Senior Lecturer, Florence Nightingale Faculty of Nursing, Midwifery & Palliative Care

Dr Andrew King
Reader in Medical Image Analysis, Faculty of Life Sciences & Medicine

Dr Robert Knight
Reader in Developmental Genetics, Faculty of Dentistry, Oral & Craniofacial Sciences

Professor Paul Luff
Professor in Organisations and Technology, King's Business School

Dr Scott McLachlan
Lecturer in Digital Technologies for Health, Florence Nightingale Faculty of Nursing, Midwifery & Palliative Care

Professor Mark Mulligan
Professor of Physical & Environmental Geography, Faculty of Social Science & Public Policy

Dr Kenneth Payne
Professor of Strategy, Faculty of Social Science & Public Policy

Dr Marina Riabiz
Lecturer in Statistics, Faculty of Natural, Mathematical & Engineering Sciences

Professor Carmine Ventre
Professor of Computer Science, Faculty of Natural, Mathematical & Engineering Sciences

Dr Aleksej Zelezniak
Senior Lecturer in Computational Biophysics, Faculty of Life Sciences & Medicine

Professor Joanna Zylinska
Professor of Media Philosophy & Critical Digital Practice, Faculty of Arts & Humanities

DESIGN
Susen Vural Design
susenvural.com
Approved by
brand@kcl.ac.uk
September 2023

PRINT
Paragon

This publication has been produced using paper from sustainable sources and bleached using an elemental chlorine-free process. The paper is produced at a mill that meets the ISO 14001 environmental management standard and the EMAS environmental management standard. The publication is fully recyclable.

King's Institute for Artificial Intelligence
King's College London
Bush House, 30 Aldwych
London, WC2B 4BG

E. ai-institute@kcl.ac.uk
X @aiatkings
W. kcl.ac.uk/ai