Report on
New Perspectives
on AI Futures

Autumn 2023

Table of Contents

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Table of Contents

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Table of Contents

Overview

'New Perspectives on AI Futures' was a workshop series held in spring 2023, funded by the Alan Turing Institute and led by Dr Atoosa Kasirzadeh, Chancellor's Fellow and Research Lead at the Centre for Technomoral Futures. The workshops were co-designed by the Centre for Technomoral Futures and the Data + Design Lab, both of the Edinburgh Futures Institute at University of Edinburgh. Three half-day, hybrid workshops were organized to surface new approaches to human flourishing with AI, bringing together expertise, insights and provocations from regions, sectors, and stakeholder groups often left out of high-level discussions of AI and our futures. Each workshop focused on a particular theme and set of motivating questions:

Sustainability and AI Futures

What is missing from current efforts to align AI with sustainable futures, understood broadly in terms of the UN sustainable development goals? How can AI be used to fight climate change and promote environmental justice? What will be the role of AI in supporting sustainable development and food security?

Health, Wellbeing and AI Futures

What aspects are overlooked in current-day AI implementations to enhance health and wellbeing and enable sustainable human flourishing? How can AI be leveraged to detect and manage future risks to health and wellbeing? In what ways can AI innovations improve the accessibility and quality of healthcare today, while also securing future human wellbeing?

Work, Democracy and AI Futures

How will AI change our social, political and economic institutions? What is the future of work in the context of AI? What risks does AI pose to our democratic institutions? How can AI be governed democratically?

These questions have all been posed before. However, they are usually posed to already recognised 'experts' in AI ethics, policy and responsible innovation. These experts typically come from relatively homogenous and well-resourced organisations and institutions in the 'Global North', and their views are frequently cited in academic and policy circles as well as mainstream media. We wanted these workshops to bring together a more diverse group of experts, including those whose direct knowledge and experience of the impact of AI and digital technologies – particularly their impact on the most vulnerable communities, groups and regions – has to date been insufficiently recognised.

The workshop series aimed to generate fresh ideas, questions, and collective visions for AI, ultimately producing a set of 'provocations' from these new perspectives. These provocations were presented to AI thought leaders and policymakers and a global audience on March 6, 2023 at

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

a University of Edinburgh livestreamed event: 'Edinburgh Futures Conversations – The Future of Artificial Intelligence: Shaping our AI Futures'. Panellists at that event were Petra Molnar, Kate Crawford, Stuart Russell, Pascale Fung, and William Isaac. The event was introduced by Professor Shannon Vallor and chaired by Professor Michael Rovatsos, both of the University of Edinburgh.

The provocations also formed the basis of a hybrid workshop called 'AI for the Next Generation: Realising an Inclusive Vision for Scottish AI', held at the Scottish AI Summit on March 29, 2023 in Glasgow. This workshop explored how to operationalise a high-level vision for AI that can truly serve a broad and diverse public. Participants in the workshop worked in groups to consider one of the provocations and identify a set of actions in response to the provocation that could help build better AI futures in Scotland and beyond.

Members of the organizing and facilitation team for the workshops were: Dr Atoosa Kasirzadeh (Principal Investigator), Professor Shannon Vallor, Dr Gina Helfrich, Joe Noteboom, Aditya Singh, and Lara Dal Molin.

Watch the workshop here:

Workshop Participants

Each workshop included a mix of academic researchers from the UK and abroad as well as participants from regions, sectors, and stakeholder groups that are often left out of high-level discussions of AI. Our aim was to bring a diverse range of voices together to contribute to a fresh vision of what is desirable and achievable for the next generation of AI. Our final list of participants included 19 individuals from external organisations. The participants came from a broad sample of countries and regions, including: East Africa, Austria, Brasil, Canada, Cambodia, Germany, India, Iran, Malaysia, Nigeria, Scotland, the UK, and the US.

THE UNIVERSITY of EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Overview

Sustainability and AI Futures
- Favour Borokini – University of Nottingham
- Sophia Georgescu & Frances Lingard – Children's Parliament
- Elina Noor – Asia Society Policy Institute
- María Pérez-Ortiz – UCL Computer Science
- José Renato – LaPin
- Michael Rovatsos – University of Edinburgh
- Jennafer Roberts – LXAI

Health, Wellbeing and AI Futures
- Nevicia Case – University of Alberta
- Seliem El-Sayed – University of Vienna
- Glenn Liddall – People Know How
- Shweta Mohandas – Centre for Internet and Society
- Olasupo Oyedepo – African Alliance of Digital Health Networks & University of Buckingham
- Ixone Saénz Paraíso – SignVideo
- Steve Felix-Uduh – National Primary Health Care Development Agency
- Xiao Yang & Stuart Anderson - University of Edinburgh

Work, Democracy and AI Futures
- Payal Arora – FemLab Co.
- Callum Cant – University of Oxford
- Christina Colclough – The WhyNot Lab
- Karen Gregory – University of Edinburgh
- Cansu Safak – Worker Info Exchange
- Willow Wong - Digital Futures Lab/Singapore Management University

Other Participants and Audiences
- The Edinburgh Futures Conversations – The Future of Artificial Intelligence: Shaping our AI Futures event on March 6 attracted ~350 attendees.
- The AI for the Next Generation: Realising an Inclusive Vision for Scottish AI workshop at the 2023 Scottish AI Summit had ~35 participants.

Provocations

Our framing of the workshop outputs as 'provocations' was intentionally modest, as three half-day workshops are not sufficient for the co-construction of novel solutions, interventions or policies for AI Futures. Instead, we wanted to use these workshops to generate questions that could interrupt, challenge or enrich familiar narratives around AI ethics and governance. We wanted to see what provocations would come from different configurations of 'experts' that could be posed to those experts who already enjoy prominence in AI conversations – not only to broaden those conversations, but to prompt critical questioning of what those conversations all too often take for granted.

While the workshops elicited many diverse insights and inspirations for future research and inquiry, and did not force or seek to reach any particular consensus, the facilitators noted five strong themes that cut across all three workshops and were echoed in different ways by multiple participants in each. These five themes became the basis of our provocations:

Design Inclusively.

AI-based technologies are often designed for a specific group of people, rather than with and by that group. For instance, AI-based devices are being designed for use in care homes for the elderly, and AI-based applications are being designed for children to enhance their educational progress. Similarly, 'AI for Good' proposals often target vulnerable communities in the global majority. Yet there are serious questions about whether these tools truly serve these groups well or meet their most urgent needs. Why should AI-based technology be imposed on vulnerable groups, even when they are supposedly the intended beneficiaries, rather than being inclusively designed with them? How can we shift from a **design AI futures-*for*** culture to a **design AI futures-*with-and-by*** culture? Moreover, how can vulnerable groups claim the right to explicitly refuse an AI-based product in its entirety?

Slow Down.

We face a dilemma about time when designing and governing AI systems. The need for quick design and development to remain economically competitive, often framed as a 'race' for AI superiority, conflicts with the need to carefully and thoughtfully govern and regulate socially embedded AI systems, through deliberative and consultative processes that take time. As a result, governance and regulations are either delayed, or hastily and superficially framed. For example, pressures around remaining economically competitive are shaping the UK's AI strategy in the direction of being 'pro-innovation', arguably prioritising friendliness to business over safety and reliability for the public. To escape this trap, do we need something like a 'Slow AI' movement?

Provocations

Go Together.

There is another time-related dilemma about AI Futures. The speed at which AI is progressing threatens to worsen digital divides within and between communities, further marginalising those with limited technological access and expertise. For instance, AI tools are rapidly becoming integrated into all aspects of life in the most economically privileged communities, whereas basic digital needs like internet access are still lacking within many disadvantaged communities. As observed by our workshop participant Steve Felix-Uduh, a digital health innovation leader from Nigeria, making AI work for everyone will require a form of digital implementation ambidexterity, with one hand moving quickly to develop and govern AI for the public benefit while the other hand focuses on bridging and narrowing digital divides. A proverb in Africa cited by a participant in the workshop is: 'If you want to go fast, go alone. If you want to go far, go together.' Do we want to **merely go *fast* with AI?** Or do we want to **go *far*, bringing people and communities along with AI?**

Develop Social and Ethical Literacy.

As AI becomes more integrated into various aspects of our daily lives, it is vital to ensure that people have the necessary knowledge to engage with this technology in an informed and empowered way. But the dominant AI narrative places the burden on impacted communities and stakeholders to develop greater technical literacy, imposing even further costs on vulnerable and excluded groups. For example, many women and members of minority groups have had to learn how AI tools frequently bias search results against them and reinforce discriminatory stereotypes, in order to alert developers and media to the worst examples of this harm and demand remedies. Why do we not instead, or at least equally, stress the **obligation of AI researchers and designers** to develop greater **social and ethical literacy** about the communities and contexts they claim to build AI for, moving the burden for knowing how to foresee and prevent AI harm to communities away from those who currently suffer that harm and onto those who profit from it?

Broaden the Focus.

While it is important to discuss the potential risks and benefits of AI for society, what about the risk that focusing our energies too much on AI may divert attention from more urgent issues? For example, the promise of AI in healthcare may not benefit those in the global majority who still lack basic health infrastructure and resources to access it. Even data-driven harms are arguably less urgent social problems than access to medical care, clean water and other basic goods, not to mention the climate and biodiversity crisis that threatens us all. Our current media environment, following industry-fed hype and speculation about hypothetical AI-fueled 'existential risks', is centring AI. But might we actually need to **de-centre AI** in order to avoid neglecting other more fundamental and pressing issues?

Commentary

While our workshop participants came from a variety of places, backgrounds, and professions, there were some thematic elements that tended to unify their perspectives. Perhaps most importantly, the perspective of the provocations centres indigenous peoples and people of the global majority. For example, the overarching focus on the individual in the West tends to be reflected in the lack of responsibility assumed by institutions as well as an excessive focus on individual action (e.g. in the context of getting consent for one's data), whereas in cultures from the global majority there is often a greater recognition of the role of the collective and the community.

Furthermore, conversations among workshop participants had a habit of returning again and again to the need to shift societal, political, and economic power structures in order to enable conditions that could lead to AI futures worth wanting. That is, rather than focusing on technology and tech development, our participants tended to locate the root of many AI-related problems with larger systems like strands of capitalism that shape how AI is deployed and operates in society. For example, everyone seems to agree on the need for participation and co-production of AI and its governance with and by affected communities, but experience shows that this is costly. Even where there is funding for co-design, it is often time-limited with no provision for ongoing engagement. Who should *pay* for co-production, when the only parties who can afford it are large corporations with the power and incentive to neglect it? How can we create a new system of incentives that will ensure sustainable resources for community participation in AI development and governance?

We found that in the context of trying to engage largely comfortable professional audiences with the provocations, both during the Edinburgh Futures conversation as well as at the Scottish AI Summit workshop, it was not uncommon for the response to evade the crux of a provocation. For example, when asked to address the provocation that AI researchers and designers have an obligation to develop greater social literacy about the communities and contexts they claim to build AI for, initial conversations quickly pivoted to the importance of STEM education efforts and upskilling the general public's knowledge around AI. This move twisted the question of how technologists might begin to learn social and ethical literacy into a question of how to make society more technologically literate, which was precisely the converse of the provocation's intent. The provocations highlight difficult challenges, uncomfortable value conflicts and deep complexities in our relation to AI. Audiences to the provocations often seemed to struggle to grapple directly with these, resorting instead to evading the provocation by changing the subject, or reducing it to a more familiar or comfortable framing.

We hypothesize that one potential reason for this distortion is the difficulty of dealing with provocations that stretch across particular disciplines and methodologies. While there are robust dialogues on AI futures within different fields, it remains rare and difficult to weave together knowledge across disciplinary boundaries and to think in a future-oriented way that encompasses inter- and multidisciplinary perspectives. Moreover, not all of these perspectives, disciplines and forms of expertise hold equal power and status in elite academic, technical and policy circles; and our provocations challenge the 'expert consensus' of those who already enjoy a privileged standing.

THE UNIVERSITY of EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Workshop participants Sophia Georgescu and Frances Lingard note that from the perspective of children's human rights, this idea of the individual expert also centres adulthood. Adults are considered to be able to design, impose and regulate AI systems without involving children, who are among the first generation to grow up with them integrated into everyday life. Children are perceived to be passive consumers and users and thus excluded from decision-making.

A related factor is the fact that expert conversations about AI futures still tend to take place against a background of resolute techno-optimism: the belief that new technologies are always intrinsically aligned with human progress and advancement. Such a belief implies that ethical issues with new technology like AI only require some modest 'tinkering around the edges' of the existing technology, softening its points of friction with society without changing its basic trajectory. The fact that some of our provocations sharply call this into question is another reason why some expert audiences are likely to resist direct engagement with them.

The provocations themselves also seem to point to one another. We may need to 'slow down' in order to 'go together'. And to 'design inclusively' seems to presuppose that technologists must 'develop social literacy'. These connections may imply a deeper underlying flaw or misalignment in our current sociotechnical milieu, something that we collectively struggle to name, and that seems to lie beyond the reach of the familiar tools of AI ethics and governance work: modest policy fixes, research investments and design interventions. It may be necessary for us to think 'beneath' the provocations to identify this deeper, broader sociocultural challenge and a fitting response to it.

Scottish AI Summit Workshop

At the Scottish AI Summit, we asked workshop participants to consider one of the provocations and identify a set of actions in response that could help build better AI futures in Scotland and beyond. Below, we summarise the suggested actions that participants generated in response to each provocation.

How can we shift from a design AI futures-for culture, to a design AI futures with-and-by culture?

Develop multiple efforts to engage communities from the beginning of AI development efforts. For example, governments, NGOs, and communities could create accessible, centralized user groups and leverage social media for wider participation. Trust people: establish a set of standards on user engagement from the start, such as through a governing body. Provide funding for initiatives that engage the community, such as by establishing community hubs at local libraries. Increase the accessibility of skills, empowering end users as designers, similar to how TikTok made regular people into filmmakers. Promote open data access to support transparency with affected communities. In addition, players central to the tech development cycle should consider inclusivity. For example, developers could build responsible AI usage into the basic development cycle through the creation of tools and standards. Funders and investors should push for design principles over profit, encouraging purpose-led investment and ensuring that good practise is carried out.

THE UNIVERSITY of EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Commentary

How do we design and govern AI wisely and inclusively amid the 'race' for AI superiority – do we need something like a 'Slow AI' movement?

Governments have a responsibility to consider the needs of all citizens, using AI as an enabler for people. Decoupling implementation and innovation would allow appropriate time for ethical considerations while still promoting advancement in research and development. Ethical AI certification could be introduced, similar to B-corp status, to ensure responsible practices. Everyone involved in the ecosystem should consider the ways that money continues to incentivise the 'race' for AI superiority, and devise strategies that affect financial outcomes in order to incentivise a slower pace. Finally, we should foster knowledge exchange across disciplines, such as through mentoring or conversation exchanges.

Why do we not stress the obligation of AI researchers and designers to develop greater social literacy about the communities and contexts they claim to build AI for?

Researchers should be educated on ethical issues from the beginning and prioritise understanding the needs of the communities instead of being solely focused on their own research problems. Establish a library of mistakes and learning to improve design and community engagement practices across the industry. Researchers and designers should spend time with diverse groups informally, not just for formal ethics discussions, to gain better social understanding.

Might we need to de-centre AI in the current media environment to avoid neglecting other fundamental and pressing issues?

Instead of starting with AI as the solution, identify the problems that need solving first. Engage in sober conversations about what AI can and cannot do, resisting the hype and fear of the current media environment. Prioritise user-centred design when developing public services, considering accessibility and co-creation throughout the process and assessing who truly benefits from the services provided.

Further Related Questions and Work

We asked the workshop participants for their views on further questions and work prompted by the workshops. Their responses follow below.

Nevicia Case

- The International Council on Alcohol, Drugs and Traffic Safety (which I was part of during my PhD) has an interesting model of addressing the disparities in their research field that may be worth considering. Research labs from the 'Global North' are paired with a 'twin' lab from the 'Global South'. The two labs are expected to share in regular knowledge exchange (e.g., the council provides funding for research scholars and select trainees to visit their twin lab for a brief stay) and resources (e.g., sharing of research equipment, datasets, etc. between the labs). Adopting a similar model may help AI researchers in the 'Global North' to become more attuned to considerations for AI in the 'Global South' while also providing AI researchers in the 'Global South' access to more resources.

Christina Colclough

- AI systems must be 'inclusively governed' - this means, amongst others, governed by deployers in cooperation with representatives of those who are subjects of the AI system. In the workplace and with regards to automated management systems, these subjects are the workers, and in unionised workplace the representatives would be the shop steward.
- This governance body shall oversee system performance relative to workers' rights, fundamental rights, discrimination/bias, impact on working conditions, occupational health and safety, and more.
- At all times, it is the deployer who is responsible, accountable and liable for (un)intended harms caused by using these systems.
- Trade-offs between system accuracy and, say, fairness must never be taken by developers alone. Instead, the governance body must be informed and jointly consider whether the trade-off is acceptable.
- Employers introducing automated systems must be bound to 'disruption's obligations' - i.e. the obligation to retrain affected workers, and/or support their career development in other ways. This must be done in cooperation with worker representatives and/or their union.

Sophia Georgescu and Frances Lingard, Children's Parliament

- "If [adults] want to know how AI can be better, [they] should just ask us [children]!" Words of a Member of Children's Parliament, aged 10.
- As rapid AI development increases children's exclusion from decision-making to influence their everyday lives, especially for more vulnerable children, adults are responsible for listening and positively responding to the impacts, whether positive or negative, on their lives.
- This is critical along intersectional lines. We have seen that children with the least access to education on AI and their human rights are the most likely to experience harms from algorithmic decision making, which lacks ethical regulations and data representation for younger ages.
- We were inspired in our workshop by the creative and genuinely challenging work that others are doing globally and urge everyone to incorporate human rights by including children as active users of AI and participants in a fairer system of reflexivity and accountability.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Karen Gregory

- How will workers be drawn into conversations about AI and the future of work? In which sectors might workers' voices be most beneficial to broader debates about regulation or development of AI? Through what channels will workers be able to report issues with automated systems? What lessons from platform workers' current legal challenges can be adopted into discussion of AI development?
- How will trade unions be engaged in these conversations? Which third sector organizations are moving forward here to lead a discussion on work and AI regulation, and how are they engaged in conversations? How, if at all, are organisations that represent workers involved in government discussion of AI development and regulation?
- How are platform companies engaged in these discussions and how, if at all, are companies held to "fair work" standards as they develop and invest in AI?

Elina Noor

- At heart is the question of what ethical design of AI might look like were it reimagined by 'users' in the global majority if (i) global power imbalances were reconfigured and (ii) a relational approach were taken (to each other as human beings and to nature in the life cycle of digital tech, from the extraction of raw materials to the energy consumption it demands as well as the waste it generates every few years). On this basis, I suggest at least two paths for future work:
- A reassessment of the economic – and therefore, power – structures that currently incentivise the entire life cycle of data-driven tech, including AI; and
- Relational approaches to AI, including e.g. alternative data governance frameworks, the notion of community vs individual data

Olasupo Oyedepo

- How do we define & appropriately contextualise AI ethics to fit various communities & cultures?
- Are conversations like this "one-off" or are we working to incorporate such discussions into the very fabric of AI development, deployment & governance?

Seliem El-Sayed

- Identifying Regulatory Focus: To effectively regulate AI, we must identify crucial elements within the AI pipeline that require closer scrutiny. By doing so, we can govern the development and deployment of AI technologies while minimizing potential risks.
- Regulators as Ethical Referees: Regulators and governments should assume the role of impartial referees in the AI race. Their objective is to determine what AI applications are permissible and what requires restrictions. This approach fosters ethical standards and protects society's interests.

Based on our experiences of the Edinburgh Conversations event and the Scottish AI Summit workshop, we believe that more multi-stakeholder discussions are needed to engage effectively with the provocations and illuminate the path that will lead us to futures with AI that can support human flourishing – to futures that are worth wanting.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

New Perspectives on AI Futures Workshop Participant Biographies

Sustainability and AI Futures

Favour Borokini

Favour Borokini is a PhD student at the Horizon Centre for Doctoral Technology at the University of Nottingham where she is researching the ethics of avatars and virtual human creation. In the past, Favour worked with various organisations researching the impact of emerging technologies on social equality.

At Pollicy, an award-winning East African feminist collective, Favour led the African Women in AI which investigated the potential impact of AI on African women as well as the challenges of African women who are AI professionals, as a Data and Digital Rights Researcher.
 Favour enjoys fantasy, romance and detective novels and is a huge fan of anime.

Sophia Georgescu

Sophia Georgescu is a PhD candidate at the University of Stirling specializing in children's participation and political ecology, as well as project worker for Artificial Intelligence, Learning for Sustainability and other climate change projects for Scotland's Children's Parliament. She worked to facilitate the participation of children under the age of 14 in decision making around the climate emergency in Scotland, including as a part of Scotland's Climate Assembly.

Before moving to Scotland, she obtained her BA in Geography from the University of Cambridge. Drawing together spatial relations, the politics of environmental justice, and activism, she is keen to explore the links between the climate crisis and the children's rights crisis to enact the changes that are desperately needed to tackle both. This includes understanding how current structures exclude children from exercising agency and power, including in AI system development as it continues to emerge as a techno-optimist solution to climate and other social crises.

Further information about Children's Parliament's current work on AI and children's human rights can be found here.

Frances Lingard

Frances Lingard is the Digital Media officer and project worker for the Artificial Intelligence project at Children's Parliament. Working with children across Scotland, their work spans projects with an interest in how digital media and engagement can support a rights-based practice. The Artificial Intelligence project, with the Alan Turing Institute and the Scottish AI Alliance, is working with schools from Glasgow to Shetland gathering children's ideas on AI and thinking about what needs to happen for AI to play a role in keeping all children happy, healthy, and safe. The next stage of the project will investigate how children can become involved in AI development and policy.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Frances has a background in programming, creative technology and education. At Children's Parliament they coordinated and produced 'Rightsburgh', an online game which was made by children for adults to explore how actions can impact on children's human dignity. They teach Access to Coding and digital rights courses through the Workers Education Assocation.

Their academic interests lie in data rights education and access in the context of AI governance. They are keen to explore the urgency to make stronger links between the possibilities and limitations of AI in the climate crisis. They want to think about the role a rights-based practice can play in framing these discussions. Further information here.

Elina Noor

Elina Noor is a senior fellow in the Asia Program at Carnegie where she focuses on developments in Southeast Asia, particularly the impact and implications of technology in reshaping power dynamics, governance, and nation-building in the region.

Previously, Elina was director of political-security affairs and deputy director of the Washington, D.C. office at the Asia Society Policy Institute. Prior to that, Elina was an associate professor at the Daniel K. Inouye Asia-Pacific Center for Security Studies in Honolulu. She spent most of her career at the Institute of Strategic and International Studies Malaysia, where she last held the position of director, foreign policy and security studies. Elina was also formerly with the Brookings Institution's Project on U.S. Relations with the Islamic World.

Between 2017 and 2019, Elina was a member of the Global Commission on the Stability of Cyberspace. She currently serves on the ICRC's Global Advisory Board on digital threats during conflict.

Elina read law at Oxford University. She obtained an LL.M (Public International Law) from the London School of Economics and Political Science, University of London, graduating with distinction at the top of her class. A recipient of the Perdana (Malaysian Prime Minister's) Fellowship, she also holds an MA in security studies from Georgetown University, where she was a Women in International Security Scholar.

María Pérez-Ortiz

Maria Perez-Ortiz is Assistant Professor at the Centre for Artificial Intelligence and the Department of Computer Science at University College London (UK). She is co-founder and Director of the first MSc programme on Artificial Intelligence for Sustainable Development, where she leads two courses on the intersection of emerging AI technologies, sustainability and ethics.

Maria describes her research and teaching agendas as "planet-centered AI", focusing on responsible, safe, sustainable and impactful AI technologies for people, the planet and tackling the challenges of the Anthropocene (e.g. climate change mitigation/adaptation, scenario-analysis methods for supporting policy making and science).

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Her research contributions and real-world development/deployment of AI technologies for sustainable development have had to date impacts in areas related to biomedicine (organ transplantation, skin cancer, infectious diseases, visual perception), multiple dimensions of environmental sustainability (climate change, polar conservation, sustainable agriculture and renewable energies) and educational technology.

Her current areas of interest span topics related to responsible tech innovation for ecology, education and policy making, where she explores both the opportunities and risks that these emerging technologies bring.

José Renato

José Renato is a Brazilian researcher and activist working on digital policy focusing on artificial intelligence regulation. He co-founded the non-profit think tank Laboratório de Políticas Públicas e Internet - LAPIN, and is now its representative as a civil society member at the Brazilian Federal Administration's Central Committee on Data Governance. He is an Alexander von Humboldt Stiftung's German Chancellor Fellow alumnus, having conducted research in partnership with the German think tank iRights.Lab and the European Parliament to investigate AI regulation in Brazil and Europe.

Michael Rovatsos

Michael Rovatsos is Professor of Artificial Intelligence and Deputy Vice Principal of Research at the University of Edinburgh, where he also heads up the Bayes Centre, the University's data science and AI innovation hub. He has over 20 years of experience in AI research, where the focus of his work has been on multiagent systems, i.e. intelligent systems where artificial and/or human agents collaborate or compete with each other. In recent years, he has primarily been working on ethical aspects of AI, developing methods for designing fair and diversity-aware AI algorithms that respect the values of human users. Michael has published around 100 scientific articles on various topics in AI, and has been involved in research projects that have received over £18 million of external funding. He received his PhD from the Technical University of Munich in 2004, and has been living in Scotland ever since, after spending most of his early years in Germany away from his native Greece.

Jennafer Roberts

Jen has been an affiliate researcher with Accel AI Institute since 2017 after she graduated with a masters degree in anthropology and social change. Her thesis work was on what it means to care in a world of neoliberalism and individualism. With Accel AI, she has led workshops and talks on mindset training for AI engineers and enthusiasts. In 2021, after moving to Cambodia, she focused her work on researching ethics in AI, and has since been composing research papers, textbook chapters, reports, meta-analyses, and blog posts on various topics in ethical AI including top-down and bottom-up implications of ethics for AI, machine learning bias and mitigating racial bias in medical machine learning, Indigenous Data Sovereignty and neocolonialism in data, and environmental and human rights issues within AI ethics.

Health, Wellbeing and AI Futures

Nevicia Case

Dr Nevicia Case is a postdoctoral fellow in the Faculty of Law at the University of Alberta. Her research maps and analyzes regulatory policies and frameworks for clinical translation and market authorization of regenerative medicine products and therapies. Having co-founded Health Innovation Initiative, Dr Case is experienced in entrepreneurship and investment in health/medical startups. She is certified in private capital investment from Ivey Business School and is presently an Impact Consulting Fellow with Propel Impact.

As a former One Young World Scholar with Johnson & Johnson, Dr Case has led research on impact frameworks for clinical research. She is also a member of the Montreal Hub of the World Economic Forum's Global Shapers Community, where she served as Vice-Curator (interim) in 2022.

Seliem El-Sayed

Seliem El-Sayed is a PhD student at the Department of Political Science at the University of Vienna. Until July 2022, he was a project staff member at the Austria-wide Digitize Project: Computational Social Science in Digital and Social Transformation and has since been part of the "Ethics & Society" team of an AI company in London. He researches ethical and social standards for data collection and -use in computational social science. He also investigates when and how data use creates public value, as well as the role the relationship between law and ethics plays in this. His graduate thesis (SciencesPo Paris/FU Berlin) assessed the extent to which the German legal system can protect citizens from algorithmic discrimination by state entities. His most recent publication, co-authored with Professor Barbara Prainsack, proposes a solidarity-based approach to governance.

Glenn Liddall

Glenn established People Know How in 2013 having worked and volunteered in the third sector for over 25 years. He is values-driven to support and empower people and communities to thrive and has trained in leadership and counselling with a commitment to continuous learning and development. Glenn's role is to lead and develop People Know How, through our strategy for growth and innovation, ensuring leadership and development opportunities for staff and VIPs. He also works to establish, develop and nurture excellent collaborative relationships with all stakeholders to ensure People Know How continues to deliver the most productive and impactful difference to communities aligned to our vision, mission and values.

Shweta Mohandas

Shweta is a Researcher at the Centre for Internet and Society, India. Her areas of work and interest include Artificial Intelligence, Privacy, and Intellectual Property Rights and India's policies surrounding them.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Olasupo Oyedepo

Olasupo Oyedepo or "shoopo" is the co-founder and Director of the African Alliance of digital health networks (African Alliance). He is a technology strategist with over 20 years of experience working across digital transformations. He has led and managed several projects and programmes including the development of Nigeria's inaugural eHealth strategy, the hosting of a regional workshop on digital health strategies, and a ministerial dialogue alongside the World Health Assembly. He believes that the transformational power of technology should have appropriate ownership and governance, if these investments are to be sustainable and impactful. With a Masters in Digital Health Governance, he is currently pursuing a PhD in the Ethical use of Artificial Intelligence, in healthcare at the University of Buckingham.

Ixone Saénz Paraíso

Ixone Saenz is a Data Scientist at SignVideo. Ixone sees the many possible benefits leveraged from AI technologies and is a passionate advocate for empowering disabled communities to impact and shape the developmental path that these technologies may follow.

Steve Felix-Uduh

Steve Felix-Uduh is the Digital Health and Innovation Specialist at the Maternal & Child Health Emergency Operations Centre in Nigeria. He led the Digital Health team within the NPHCDA where he defined and structured the application of technology in strengthening primary healthcare outcomes across Nigeria. He contributed extensively to the development of Nigeria's inaugural Digital Health Strategy and has been a consistent player on the Global Digital Health stage. Steve is currently conducting research at the University of Buckingham School of Computing to investigate how trends in digital technologies influence the lives and lifestyles of people, particularly in the 'Global South'.

Xiao Yang

Xiao Yang is a PhD student based in Centre for Doctor Training (CDT) for Biomedical AI and Institute for the Study of Science, Technology and Innovation (ISSTI) in Edinburgh. She is looking at the arena of diagnostic AI (specifically machine learning based tools in radiological diagnosis), interested in the how the start-ups orient themselves around innovation opportunities and their collaboration strategies with various stakeholders.

Stuart Anderson

Stuart Anderson is Professor of Dependable Systems in the School of Informatics. His current interests are in the governance and regulation of socio-technical systems that incorporate machine learning and how the need for effective governance and regulation can help drive responsible design practice.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Work, Democracy and AI Futures

Payal Arora

Payal Arora is a digital anthropologist and author of several books including the award-winning "The Next Billion Users" with Harvard Press. She is a Professor and Chair in Technology, Values, and Global Media Cultures and an Academic Director in UX and Inclusive Design at the Erasmus Centre for Data Analytics, Erasmus University. She is co-founder of FemLab, a feminist future of work initiative. Her expertise lies in user experience (UX) among resource-constrained and precarious communities, cross-cultural AI ethics, and inclusive design. She comes with two decades of fieldwork in the 'Global South'. She sits on several boards such as Columbia University Earth Institute and World Women Global Council in New York. She currently lives in Amsterdam.

Callum Cant

Dr Callum Cant is a postdoctoral researcher at the Oxford Internet Institute, where he works on the Fairwork AI project. His research focuses on how AI systems are transforming workers' experiences of work.

Christina Colclough

Widely regarded as a thought leader on the futures of work(ers) and the politics of digital technology, Dr Christina J. Colclough is an advocate for the workers' voice and for strong, quality public services. She founded the Why Not Lab with the aim to reshape the current digitalisation trajectory, so human rights, freedoms and autonomy are respected and protected.
Christina's background is in labour market research and in the global labour movement, where she led their future of work policies, advocacy and strategies for a number of years. She was the author of the union movement's first principles on Workers' Data Rights and the Ethics of AI.

Karen Gregory

Karen Gregory digital sociologist, ethnographer, and Programme Co-Director of the MSc in Digital Sociology. She currently co-leads the Digital Social Science Research Cluster at the Center for Data, Culture and Society at the University of Edinburgh. She is also an Associate Editor at the Journal of Cultural Economy.

Cansu Safak

Cansu Safak is a researcher with experience in investigating artificial intelligence and data driven technologies, particularly in the fields of policing, welfare and public sector procurement. She previously worked with the Ada Lovelace Institute, The Bureau of Investigative Journalism and Big Brother Watch.

Willow Wong

Willow Wong is a graduate student in BA English and Philosophy with an interest in the philosophy of technology, existentialism, and Buddhist philosophy. She is a passionate advocate for young adult mental health with strong skills in research, lobbying, relationship management and policy development in the charities/higher education/NGO sector. In her current research role, she re-formulates and clarifies issues of artificial intelligence design and governance (and its big data outputs) by examining its theoretical foundations.

Workshop Facilitators

Atoosa Kasirzadeh

Atoosa Kasirzadeh is a Chancellor's Fellow (tenure-track assistant professor) in the University of Edinburgh philosophy department and Research Lead of the Centre for Technomoral Futures at Edinburgh's Futures Institute. Previously, she was a visiting research scientist at Google DeepMind (London, UK) and a philosophy postdoctoral fellow for the Humanizing Machine Intelligence Grand Challenge at the Australian National University. She holds two Ph.D.s: one in philosophy, specializing in philosophy of science and technology, from the University of Toronto and one in mathematics from the Ecole Polytechnique de Montreal and the Group for Research in Decision Analysis (GERAD).

Shannon Vallor

Shannon Vallor is the Baillie Gifford Professor in the Ethics of Data and Artificial Intelligence in the Edinburgh Futures Institute (EFI) at the University of Edinburgh. She is Director of the Centre for Technomoral Futures at EFI, and co-director of the UKRI research programme Enabling a Responsible AI Ecosystem. She regularly advises government and industry on ethical use and design of AI, and is the author of 'Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting' (2016) and the forthcoming 'The AI Mirror: How to Reclaim our Humanity in an Age of Machine Thinking.'

Gina Helfrich

Gina Helfrich is Baillie Gifford Programme Manager for the Centre for Technomoral Futures at Edinburgh Futures Institute. She leads on operations and strategy, working to deliver the Centre's mission and activities, which are focused on the ethical implications of present and future advances in artificial intelligence (AI), machine learning, and other data-driven technologies. She also serves as Deputy Chair of the University's AI and Data Ethics Advisory Board. She holds a Ph.D. in Philosophy from Emory University with a Graduate Certificate in Women's, Gender and Sexuality Studies. Prior to joining University of Edinburgh, Dr. Helfrich held several programme management and operations roles in the non-profit technology sector.

THE UNIVERSITY *of* EDINBURGH
Edinburgh Futures Institute | **Centre for Technomoral Futures**

Joe Noteboom

Joe Noteboom is a PhD student in Sociology at the University of Edinburgh and part of the Centre for Technomoral Futures and the Centre for Research in Digital Education. His research focuses on the social and ethical implications of data-driven technologies and practices in higher education. Alongside his PhD research, Joe is a member of the university's AI and Data Ethics Advisory board.

Lara Dal Molin

Lara is a PhD student in Science, Technology and Innovation Studies at the University of Edinburgh, part of the joint programme in Social Data Science with the University of Copenhagen. She is interested the intersection between Artificial Intelligence, language and gender, and investigates this through the study of Large Language Models. In her PhD project, she is exploring hybrid qualitative-quantitative methods, inspired by the Design Justice framework, to address gender bias in open-source language models. Lara is also a tutor within the School of Social and Political Sciences, and a PhD assistant at EFI.

Aditya Singh

Aditya Singh is a Baillie Gifford Fellow pursuing a PhD with the Centre for Technomoral Futures and the Global Academy of Agriculture and Food Systems at the University of Edinburgh. His research investigates models of collective data governance in food systems. He is a lawyer by training, with an interest in critical approaches to data and AI ethics. He has previously worked in a range of research, advocacy and consultancy roles on digital rights, platform regulation, and data governance.