



art-ai

UKRI Centre for Doctoral Training in Accountable, Responsible and Transparent AI (ART-AI)

Annual Report 2020-21



UNIVERSITY OF
BATH



UK Research
and Innovation

Eamonn O'Neill, Director:

Welcome to the 2020-21 annual review of the University of Bath's Centre for Doctoral Training (CDT) in Accountable, Responsible and Transparent Artificial Intelligence (ART-AI).

Within a year of our launch, the Covid pandemic forced us to move our events and training online and it was only over this summer (2021) that we have been able to hold a few in-person events. Nonetheless, we've had an exciting and productive year. We look forward over the next academic year to gradually increasing the number of face to face meetings and events, obviously while taking sensible precautions as we do so.

Our third cohort started in September 2021 and at 18 students is the largest cohort we have recruited so far. We received over 250 applications so our new intake represents the best of the best. I am certain that they will prosper and succeed just as their predecessors in our first two cohorts.

Once again, I would like to thank all our partners and staff for their work over the last year supporting our students, and for a job well done in often challenging circumstances.

September 2021

List of contents

Foreword and contents	2
2019 intake	3
2020 intake	4
2021 intake	5
Student activities	6
Student research outputs	7
Review of the year	8
Our programme	9
How to apply	10
Partnerships	11
Academic management and governance	12

Oscar Bryan: *Machine learning for the detection of unexploded ordinance using synthetic aperture sonar*

George Fletcher: *Intelligent 3D character creation and animation*

Catriona Gray: *The adoption and use of Artificial Intelligence Technologies (AITs) by humanitarian organisations*

Akshil Patel: *Autonomous skill acquisition*

Mafalda Ribeiro: *AI-inspired closed loop electrical neuromodulation in multiple domains*

Elena Safrygina: *Translating spatio-temporal imaging data into clinical data using machine learning*

Jack E. Saunders: *Reactive collision avoidance using deep reinforcement learning for the application of UAV delivery*

Jack R. Saunders: *Crossing the Uncanny Valley: Using deep learning for realism in facial animation*

Elsa Zhong: *How AI may influence human judgement and decision making*

Damian Ziubroniewicz: *Fair machine learning*

Student Profile: Elsa Zhong at Zhejiang Lab



[Zhejiang Lab](#) is a research institute in Hangzhou, China, founded jointly by [Zhejiang University](#), the [Zhejiang Provincial Government](#), and [Alibaba Group](#). With a strong focus on research into artificial intelligence, Zhejiang Lab has been a partner of the [University of Bath](#) since 2019.

As part of the academic exchanges in the field of AI ethics and safety between the two institutions cohort 1 student [Elsa Zhong](#) visiting Hangzhou for several months in 2021. During her visit, she delivered a guest lecture on AI and Ethics that discussed issues that have become increasingly prominent, such as algorithm bias and discrimination, the public's lack of trust in new technologies, algorithm manipulation, and data privacy. She highlighted that these issues cannot be avoided in the development of artificial intelligence, and governments internationally face the same challenges to solve them.

There are many similarities in AI ethics principles between the UK and China. However, differences in culture, societal systems and social environment will profoundly affect the interpretation of relevant principles and the formulation of specific industry standards, laws and regulations, and further affect the future formulation of international laws and regulations in the field of AI ethics and safety. Elsa analysed the challenges in the field of AI ethics and safety; introduced the data and AI ethics and safety policies in the UK and China; and discussed the future construction of AI ethics and security governance frameworks with Chinese characteristics and the development direction of global AI governance by comparing the similarities and differences of the two countries.

Ed Clark: *AI tactical decision aid for management of naval sensors and autonomous vehicles*

Thao Do: *AI for identification and support for victims of sexual exploitation in southeast Asia*

Tom Donnelly: *AI for the control of upper limb prosthetics*

Andy Evans: *AI and its consequences for the labour market*

Tory Frame: *Enhancing adult sleep with AI*

Joe Goodier: *Interpretable diagnostic imaging using generative priors*

Finn Hambley: *Bayesian modelling of low-voltage networks and demand-side management schemes*

Matt Hewitt: *Hierarchical reinforcement learning for transparency in AI*

Emma Li: *Improved interpretability methods towards more responsible AI*

Pablo Medina: *Ecological rationality of choices between gambles*

Deborah Morgan: *AI regulation as an institutional change-maker within the European Union and the United States*

Jessica Nicholson: *Robotics processing architecture: a skill acquisition approach to artificial cognitive development*

Alice Parfett: *Using history to predict and prevent negative outcomes of population-wide racial biometric classification systems*

Brier Rigby Dames: *Modelling brain-like intelligence in an evolutionary context for AI applications*

Alex Taylor: *Machine learning in safety critical engineering*

Scott Wellington: *Decoding imagined speech from the brain signal*

Student Profile: Ed Clarke at Sea!



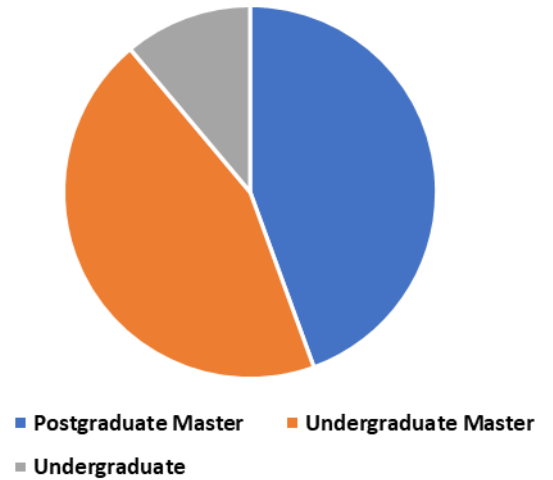
In maritime operations, it is necessary to optimally deploy and position a finite number of sonar sensing vehicles (manned and autonomous) to maximise their combined detection coverage over a geographical area of interest. Performance prediction of the sonar systems is crucial. Currently, performance prediction is complicated and user-intensive. It involves environmental data gathering, acoustic propagation modelling, sonar system modelling, and manual interpretation of the modelled outputs. Environmental data can be highly variable in space and time and is obtained from different sources (direct in-situ measurements, numerical models, and historical databases). The underwater acoustic models predict how sound propagates between sources and receivers, which depends upon the environmental data as well as the source and receiver characteristics (e.g., frequency and depth). The sonar performance models predict detection performance from the acoustic characteristics output from the acoustic models and the sonar characteristics (e.g. aperture size and types of signal processing applied to the acoustic signals). Users interpret the outputs of the performance models and use these to develop plans for the deployment of naval assets.



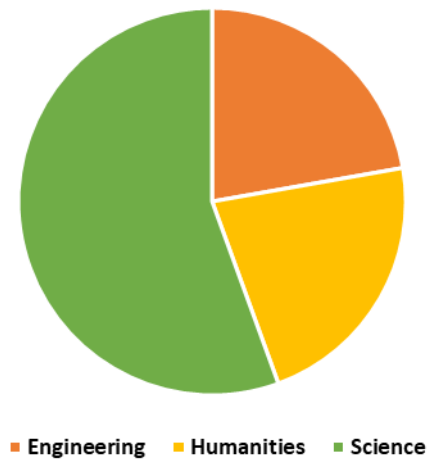
Ed's PhD aims to incorporate AI into the chain to aid and enhance decision-making for optimal deployment. We envisage that with exposure to a sufficiently comprehensive example dataset (i.e., training), optimal deployment could be achieved directly from data earlier in the chain, ultimately directly from key environmental features. Such features might include the spatial/temporal distribution of the oceanic mixed layer depth or the seabed type for shallow water areas, both of which have a significant effect on acoustic propagation and therefore sonar performance. Incorporating AI within a maritime decision chain will require a high degree of transparency. This is crucial for developing trust with operators, for accountability and for responsible use.

- Dan Beechey: *Autonomous development of action hierarchies*
- Tom Cannon: *Task agnostic efficient and adaptive edge devices using hierarchical reinforcement learning, task inference and transfer learning*
- Robert Clarke: *Shared control for brain-actuated robotic arms towards fusion of AI and human intelligence*
- James Elson: *Autonomous option discovery in hierarchical reinforcement learning*
- Harshinee Goordoyal: *Machine learning for predictions of haemodynamics in cardiovascular devices*
- Sophia Jones: *Algorithms for AI inspired by the bounded rationality of humans*
- Ferdie Krammer: *Big data and machine learning for reaction design*
- Toby Lewis-Atwell: *A transparent machine learning approach to chemical safety assessment*
- Jack McKinlay: *Explainable AI in Bayesian machine learning*
- Ben Rogers: *Neural network optimisation of auxetics*
- Tom Ryder: *EEG-based brain signal processing and applications using statistical machine learning*
- George Sandle: *Deep generative energy-based models*
- Shashank Sharma: *Using history to predict and prevent negative outcomes of population-wide racial biometric classification systems*
- Joshua Tenn: *Debiasing AI advice-taking in human-AI collaborative decisions*
- Doug Tilley: *Smart cyber-physical systems for multimodal human-robot collaboration*
- Todd van Steenwyk: *Measuring corporate risk exposure with machine learning*
- Syeda Zhara: *Algorithmic auditing*
- Yue Zhang: *Improve perceived trustworthiness and fairness in automated decision-making: the role of mental models*

Intake by highest qualification



By academic faculty

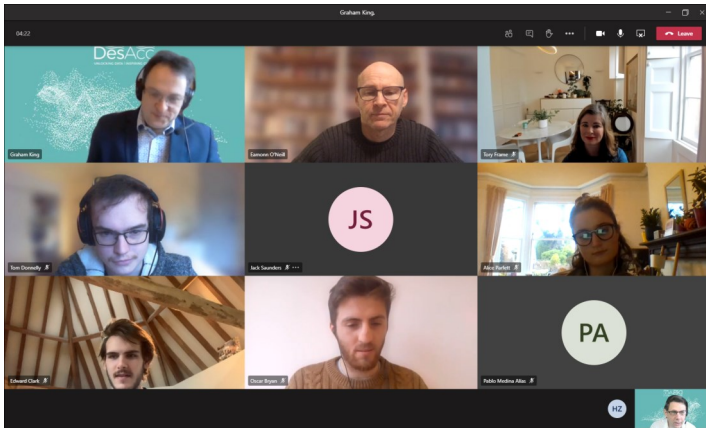


ART-AI's 2021 intake: 15 of our 18 new students were able to meet in-person at the start of term.

AI Challenge Days

This year, due to the pandemic, our AI challenge days on Wednesday 27 and Thursday 28 January 2021 were held online.

ART-AI students came together virtually for this annual event. This year, in their peer groups, students attended 4 sessions spread across both days hosted by ‘table hosts’ chosen from our panel of external partners. The objectives of these sessions were for ART-AI students to spend a thought-provoking hour discussing the host’s topic, discussing a ‘challenge’ and relating it to the context of their research. At the end of each session, the students gave feedback from their discussions.



“I enjoyed the focus provided by the fact that we had partners bring their problems/topics to discuss instead of us trying to come up with issues from such a broad field. It meant that we could go a bit deeper on topics instead of having superficial conversations which we have had before”

“Very enjoyable and informative event.”

“I thought the event was very well-organised with some interesting questions and thoughts raised by both the partners and students”



Diversity solutions to strategic questions

Student Conference

ART-AI student [Jack R Saunders](#) led the organisation of a Student Conference on 16 and 17 September 2021. This was an in person event on campus at the University of Bath.


Day 1 had a talk delivered by Ali Eslami from ART-AI partner [DeepMind](#) and a tutorial on Explainable AI delivered by Vaishak Belle from the [Alan Turing Institute](#). In the afternoon, ART-AI students [Tory](#) and [Deb](#) ran an engaging session on policy, entitled; ‘Why and how organisations engage in policy with examples drawn from the EU proposal for a Regulation on AI’.



Day 2 saw Graham King and Steve Collins from [DesAcc](#), an ART-AI partner, delivering a highly interactive session on ‘Healthcare in AI’. After a mind focusing Tai Chi session delivered by ART-AI staff Marina and Julian, the students then heard from Ben Miles from [Spin Up Science](#) about entrepreneurship.



At the end of Day 2, a 24 hour Hackathon was organised which started at 5pm on the 17 September and finished at 5pm on Saturday 18 September. The students spent 24 hours working on the problem of “how to program a self-driving car that drives around a track using just the camera with deep learning and AI”.



Oscar Bryan, Paul M. Bayle, Christopher E. Blenkinsopp, Alan J. Hunter (2020). Breaking Wave Imaging Using Lidar and Sonar. *IEEE Journal of Oceanic Engineering* **45** (3): 887–897.

<https://doi.org/10.1109/JOE.2019.2900967>

Chris E. Blenkinsopp, Paul M. Bayle, Daniel C. Conley, Gerd Masselink, Emily Gulson, Isabel Kelly, Rafael Almar, Ian L. Turner, Tom E. Baldock, Tomas Beuzen, Robert T. McCall, Huub Rijper, Ad Reniers, Peter Troch, David Gallach-Sanchez, Alan J. Hunter, **Oscar Bryan**, Gwyn Hennessey, Peter Ganderton, Marion Tissier, Matthias Kudella, and Stefan Schimmels (2021) High-resolution, large-scale laboratory measurements of a sandy beach and dynamic cobble berm revetment. *Scientific Data* **8** (1): 22.

<https://doi.org/10.1038/s41597-021-00805-1>

Alice Parfett, Stuart Townley, and Kristofer Allerfeldt (2021) AI-based healthcare: a new dawn or apartheid revisited? *AI & Society* **36**: 983–999.

<https://doi.org/10.1007/s00146-020-01120-w>

Mafalda Ribeiro, Paulo Rocha, and Benjamin Metcalfe (2020) An ultra-sensitive biosensor to investigate Random Telegraph Noise in human breast cancer cells. The 1st International Electronic Conference on Biosensors: 2 Nov 2020 → 17 Nov 2020. *Proceedings 2019* **3**: x.

<https://sciforum.net/manuscripts/7223/manuscript.pdf>

Mafalda Ribeiro, Aya Elghajji, Scott P. Fraser, Zoë D. Burke, David Tosh, Mustafa B. A. Djamgoz and Paulo R. F. Rocha (2020) Human Breast Cancer Cells Demonstrate Electrical Excitability. *Frontiers in Neuroscience* **14**: 404.

<https://doi.org/10.3389/fnins.2020.00404>

Mafalda Ribeiro, Pamela Ali, Benjamin Metcalfe, Despina Moschou, and Paulo Rocha (2021) Microfluidics integration into low noise multi-electrode arrays. *Micromachines* **12** (6): 727.

<https://doi.org/10.3390/mi12060727>



The Global AI Standards Landscape: An Extended Seminar

Our extended seminar on the global AI standards landscape took place online on 14 and 15 December 2020. The objectives of the event were to raise awareness of standards in general, share knowledge of particular standards activities related to AI and foster community-building between standards making and AI interests.

The speakers focused on particular topics that fall within the scope of the UKRI Centre for Doctoral Training in Accountable, Responsible and Transparent Artificial Intelligence (ART-AI) and included activities on emerging standards on certification, algorithmic bias, data management and the design of (semi-)autonomous systems, as well as providing perspectives from national standardisation, consumer and industry bodies. The event finished with a panel discussion which focussed on individual perspectives on the standards-making process.



FinTech: Regulatory Challenges and Sustainability Opportunities

ART-AI and [CGR&IS](#) (The Centre for Governance, Regulation and Industrial Strategy) held a workshop entitled 'FinTech: Regulatory Challenges and Sustainability Opportunities' on 28 April 2021.

The workshop brought together 100 participants and leading experts from the financial industry, regulatory bodies, and academia to discuss the role of Artificial Intelligence (AI) in:

- Enhancing and restricting market processes
- Improving the quality of financial services
- Supporting sustainable investing
- Supporting solutions to societal problems

"I thought this was a very interesting and well-organised event. FinTech is not my field of expertise, so it was fascinating to learn a bit more about it, see what sort of work is being done in this area, and how AI comes into play."

Global Challenges Research Fund (GCRF) MX-UK Integrative Think-Tank (ITT) on Environmental Resilience

Our first [Integrative Think Tank \(ITT\)](#) event was held online from 10 to 14 May 2021. There were 60 participants, made up of students, staff and partners, working across different time zones: ourselves, the [SAMBa CDT](#) at the University of Bath, two academic institutions in Mexico ([CIMAT](#) and [UNAM](#)), [CENAPRED](#), the Mexican National Center for the Prevention of Disasters, [SACMEX](#), the water utility overseers for Mexico City and the state government of Jalisco.

ITTs are mechanisms to develop the formulation of a multidisciplinary problem at a fundamental research level from high level challenges presented by partners outside the basic research area. Participants at the ITT explored challenges including how to classify seismic signals from the Popocatepetl volcano, how to debug a seismic catalogue, how to calculate final users' water consumption, how to optimise water supply and how to map forest fire risk.



ART-AI Spotlight on Equality, Diversity and Inclusivity series

ART-AI launched a new 'ART-AI Spotlight on Equality, Diversity and Inclusivity' seminar series. The aim of these talks is to raise awareness about Equality, Diversity and Inclusivity issues within our sector. The first talk in the series took place on 10 June 2021 and we were delighted that Maria Skoularidou joined us for this talk chaired by ART-AI student [Mafalda Ribeiro](#). Maria is a PhD Student at the University of Cambridge and also [Dis}Ability in AI Chair](#).



We train “specialists with perspectives”: graduates with AI skills as well as a broader understanding of the engineering applications of AI and the ethical, policy and socio-economic consequences of intelligent technologies. We combine deep training in a specialism which builds on each student’s background, capabilities and aims, with integrated training across the breadth of three disciplines: engineering, computer science and social sciences.

Key themes include:

- transparency and intelligibility of AI
- risk and decision making with AI
- safety and trust in human-machine systems
- policy-making with and for AI
- innovation in data-driven and classical AI
- engineering applications of AI
- diversity and bias in AI

The course is designed around a 4-year full-time programme of study, starting in late September/early October each year, where students progress together as a cohort. Applicants are expected to be available for full-time study although a part-time option is available at the discretion of the CDT directors, depending on the needs of the student and the project.

Year 1

In year 1, a Masters-level MRes programme draws together AI and computer science, engineering and technology, and humanities and social sciences.

Students choose specialist courses in their "home discipline" and introductory courses in the other two disciplines, alongside five compulsory units:

- Humans and intelligent machines
- Statistics for data science/mathematics for AI
- AI challenge
- Interdisciplinary research preparation
- Interdisciplinary thesis formulation report

The range of optional units to choose from includes:

- Intelligent agents
- Machine learning
- Robotics software
- Autonomous systems navigation, mapping and communications
- Knowledge, data and our digital social world
- AI as a social and political practice

The taught units are assessed using coursework and

examinations (both oral and written). Students are required to reach a pass mark of 60% overall in order to progress onto the research phase.

Years 2-4

In years 2-4, students focus on their PhD research, while taking occasional short courses on relevant topics, transferable skills, and specific training for their research interests. Progress is monitored by an interdisciplinary supervisory team and the CDT’s Director of Training via regular meetings, six-monthly progress reports and a progression board of examiners in year 3.

Once a year, students will come together with academic staff and CDT partners in an AI challenge co-creation workshop to identify key challenges for AI.

There are masterclasses and annual ART-AI conference, building a wider network of peers and potential employers. Students normally go on a placement or research visit to a project partner.

The final stage of the PhD is submission of a substantial thesis and an oral or viva voce examination, in which students are required to defend the thesis to a Board of Examiners.

As this is an integrated programme, graduates receive both the MRes and the PhD certificates together at the end of the course.



We are looking for highly motivated students with a demonstrable interest in their research topic and in working in an interdisciplinary environment. We are inclusive of all types of people, and are actively searching for diversity to build a more comprehensive basis from which to generate positive innovation in AI. We are looking for a range of individuals with different experiences, perspectives and backgrounds who are willing to share their understanding and learn collaboratively to collectively become “specialists with perspectives”.

Entry requirements and funding eligibility

You should have a first class or strong second class bachelor's honours degree or international equivalent in computer science, engineering, social sciences, policy research, psychology, or a related subject.

Normally, you will also need to have taken a mathematics course or a quantitative methods course at university or have at least grade B in A level maths or international equivalent.

UKRI-funded ART-AI CDT studentships are available on a competition basis for up to 4 years. Funding will cover tuition fees and maintenance at the national doctoral stipend rate. We will offer at least ten studentships for 2022 entry, up to three of which can be awarded to international students.

English language requirements:

- IELTS: 6.5 overall with no less than 6.0 in all components
- The Pearson Test of English Academic (PTE Academic): 62 with no less than 59 in any element
- TOEFL IBT: 90 overall with a minimum 21 in all 4 components

If your first language is not English but within the last 2 years you completed your degree in the UK, you may be exempt from our English language requirements.

If you need to improve your English language skills before starting your studies, you may be able to take a [pre-session course](#) to reach the required level.

Application forms should be submitted online (course code RDUCM-FP02) and be accompanied by a research proposal.

[On the application form](#) you will need to provide:

- personal details, including your educational background
- passport details if you need a visa
- personal statement, with details of your motivation for applying and your research interests. We also wish to know what you personally bring as a contribution to equality, diversity and inclusion and how your research will contribute to societal goals.
- names of your intended supervisors

- names and contact details of two referees, including at least one from academia if you have been studying recently

You will need to upload:

- a research proposal
- scans of your degree certificate(s) and transcript(s)
- an up-to-date CV

We may ask for additional documents.

If you are an international student you should also provide:

- an authorised translation of your degree certificate and transcript if they are not in English
- your English language assessment certificate (if available).

Research proposals need not be the research you will end up pursuing in the CDT, but gives the selection panel the opportunity to assess your suitability for the programme and your research potential. Research proposals should be about 1,000 words in length and include:

- **Aim and objectives:** What are the central aims and research questions that will guide your research? Explain what questions you are trying to answer or which hypotheses you want to test. Prioritise one or two central questions from which you can derive secondary ones.
- **Rationale:** Contextualise your aims and objectives in a broader field of study, identifying the main literature or authors that you are addressing.
- **Methodology:** Explain how you are going to conduct your research; what information you would need, how you would collect it and how you are going to analyse it. This only needs to be indicative at the moment.
- **Timeline:** Provide a summary of what you are planning to do and when.
- **Bibliography:** Provide a short bibliography of key publications related to your research question.

Feel free to:

Propose your own research project in an area suitable for ART-AI. If you are unsure about fit, [contact Marina De Vos, Director of Training](#).

Get in touch with an individual academic to discuss doing a PhD under their supervision. You can search our academics' research interests at our [research portal](#). If you're not sure which individual to get in touch with, [contact Marina De Vos, Director of Training](#).

Occasionally, we advertise specific projects [on FindAPhD.com](#). If you are applying for a specific project, please make this clear in your application and proposal.

We work with a wide range of partners, including companies, the public sector, and international universities, to encourage networking and the co-creation of research projects on issues that matter to them and us.

Our students also go on placements and student exchanges to experience different research environments and work on real problems in partner organisations.

Partners can fund or co-fund studentships, act as hosts for internships or academic visits, send visiting students to work with us, host masterclasses or give tutorials and seminars, and exchange expertise, access to facilities, or external supervision.

If you are interested in working with us please contact us at info@cdt-art-ai.ac.uk.

'The development of AI is an extremely exciting growth area and Rolls-Royce view it as one of the key technologies to enable our digital strategy. It's crucial we understand the accountable, responsible and transparent elements of Artificial Intelligence, in order to ensure we can maximise the vast potential in the Defence and Aerospace sector.' - **Gareth Hetheridge, Rolls Royce Senior IT Business Partner**

'We see huge implications for the future of industry resulting from developments in AI, all of which are built on the premise of trust. The CDT will build a community who'll have the skills to develop the AI solutions of the future whilst ensuring these systems are ethically sound, and capable of making decisions reliably and without bias.' - **Professor Ian Risk, Chief Technology Officer, CFMS**

'Like most exciting new developments in science and technology, AI has immense capacity for good -- but raises many unanswered questions about how it will affect people's relationships, their communities and the distribution of power within society. We're delighted to partner with the CDT and look forward to working with its partners and the doctoral students it will recruit.' - **Revd Dr**

Our existing network includes:





Eamonn O'Neill
Director



Marina De Vos
Director of Training



Emma Carmel
AI and public
policy



Alan Hunter
Partnerships



Hugh Lauder
Impact of AI on
society



Julian Padget
Innovations in AI
technologies



Özgür Şimşek
Machine
learning



Peter Wilson
Autonomous systems
and robotics

Strategy and Partnerships Board

Independent Advisory Board

Lord Hollick — Chair (until May 2021)

Husna Siddiqi — Chair (from May 2021)

Commodore Barry Brooks

Prof Mark Coeckelbergh (until
November 2020)

Dr Fiona Colquhoun

Dr Rachel Free

Matt Griffith

Gail Kent

Dr Oli Lanning

Dr Clara Neppel

Chris Rees

Len Rinaldi

Antony Walker

Phil Williams

Ian Risk (CFMS) — Chair

Alistair Nottle (Airbus)

Mishka Nemes (Alan Turing
Institute)

Phil Johnston (AutoNaut)

Esin Turkbeyler (BMT)

Malcolm Brown (Church of England)

Liz O'Driscoll (Civica)

Mae Soto (Dalhousie University)

Obum Ekeke (DeepMind)

Hugh Lyshkow (DesAcc)

Simon Fanshawe OBE (Diversity by
Design)

Tom Begley (Faculty AI)

Karen Crosson (FCA)

Mandy Chessell (IBM)

Carles Sierra (IIIA-CSIC, Spain)

Kenji Takeda (Microsoft)

Rob Harper (Moogsoft)

Samantha Dugelay (NATO CMRE)

Ken Satoh (NII, Japan)

Leah Chapman (NPL)

Duncan Russell (Ocado)

Alison Adams (ONS)

Euan Cameron (PwC)

Gareth Hetheridge (Rolls Royce)

Jon Hargroves (SEA Ltd)

Mark Burnett (Seiche)

Shi-Min Hu (Tsinghua University)

Virginia Dignum (Umeå University)

George Buchanan (University of
Melbourne)

Andreas Austeng (University of
Oslo)

Torsten Schaub (University of
Potsdam)

André de Carvalho (University of
São Paulo)

Jaime Sichman (University of São
Paulo)

Atte Oksanen (Tampere University)

Maneesh Singh (Verisk Analytics)

Wu Fei (Zhejiang University)

UKRI Centre for Doctoral Training in Accountable, Responsible and Transparent AI
(ART-AI)

1 West

University of Bath

Claverton Down

Bath BA2 7AY

 @ARTAIBath

 <https://www.linkedin.com/company/ukri-cdt-in-accountable-responsible-and-transparent-ai/>

<https://cdt-art-ai.ac.uk/>



art-ai