cherishde

Swansea University Prifysgol Abertawe

Festival of Ideas

1st - 2nd May 2019

Swansea University Prifysgol Abertawe

Computational Foundry
Ffowndri Gyfrifiadol

UNDEB EWROPEAIDD
EUROPEAN UNION

Llywodraeth Cymru
Welsh Government

Cronfa Datblygu Rhanbarthol Ewrop
European Regional Development Fund

Festival of Ideas

1st - 2nd May 2019

Swansea University Prifysgol Abertawe | Computational Foundry Ffowndri Gyfrifiadol

Two days of provocative, stimulating and motivating interactions addressing critical computational topics

Curated by globally leading scientists – Professor Ben Shneiderman, Professor Jennifer Preece and Professor Alan Dix – we'll explore Algorithmic Accountability and Citizen Science. The event will be held in the new Computational Foundry - www.swansea.ac.uk/science/computationalfoundry/

During the event we'll also be celebrating the UK's Digital Economy strengths with demos, posters and opportunities to talk with researchers across the UK.

Professor Alan Dix, Director of the Computational Foundry, shares his thoughts on the topics discussed as part of the Festival:

Algorithmic Accountability: "From computers that play Go, to cars that drive themselves and websites that seem to know what you want almost before you do; our world is filled with computer algorithms that are becoming ever more complex and ever more 'intelligent'. Yet, there is also another side: self-driving car crashes, racially biased search results, and leaked social media data used to influence elections. Will digital technology save the world or destroy it? Join us as world experts in computer science and artificial intelligence discuss these key issues of our time."

Ctizen Science: New Agendas, Broader Impacts: In 1950s B-movies white-coated scientists poured frothing mixtures from test-tubes behind doors that said "No Entry". Today science has broken free of the lab: during the annual BBC Springwatch tens of thousands of volunteers contribute to projects counting bees and butterflies, ladybirds and hedgehogs. In 2015 the Open University iSpotNature application hit half a million observations. This 'Citizen Science' is real science, creating new knowledge that is helping us understand the world during a time of rapid environmental change. We have gathered together a panel of international experts, but this day is for everyone: how we can all contribute and all learn."

Festival of Ideas
Day 1 & Day 2

Swansea University Prifysgol Abertawe | Computational Foundry Ffowndri Gyfrifiadol

Day 1: Algorithmic Accountability

Professor Shneiderman is holding a day's workshop to think deeply around the subject of algorithmic accountability on 1st May, together with a panel of experts comprised of:

- Professor Christopher Hankin, Co-Director of Institute for Security Science and Technology at Imperial College London
- Professor Richard Harper, Co-Director of the Institute for Social Futures at Lancaster University
- Dr Sandra Wachter, Research Fellow at Oxford Internet Institute
- Dr Brent Mittelstadt, Research Fellow and British Academy Postdoctoral Fellow at Oxford Internet Institute
- Professor Andrew Crabtree, Professor of Computer Science, University of Nottingham
- Ryan Carrier, Executive Director at ForHumanity

Day 2: Citizen Science: New Agendas, Broader Impacts

Professor Preece is holding a day's workshop to think deeply around the subject of Citizen Science: New Agendas, Broader Impacts on 2nd May, together with a panel of experts comprised of:

- Professor Muki Haklay, Professor of GI Science at University College London
- Professor Julia Parrish, Associate Dean College of the Environment at the University of Washington
- Dr Helen Spiers, University of Oxford, Biomedical Research Lead at the Zooniverse
- Prof Eileen Scanlon, Associate Director of Research and Innovation in the Institute of Educational Technology, Open University
- Professor Geoff Proffitt, Director of Knowledge Economy (Biosciences) at Swansea University

Professor Ben Shneiderman

Dept of Computer Science & Human-Computer Interaction Lab, University of Maryland

Ben Shneiderman (http://www.cs.umd.edu/~ben) is a Distinguished University Professor in the Department of Computer Science, Founding Director (1983-2000) of the Human-Computer Interaction Laboratory (http://hcil.umd.edu), and a Member of the UM Institute for Advanced Computer Studies (UMIACS) at the University of Maryland. He is a Fellow of the AAAS, ACM, IEEE, and NAI, and a Member of the National Academy of Engineering, in recognition of his pioneering contributions to human-computer interaction and information visualization. His widely-used contributions include the clickable highlighted web-links, high-precision touchscreen keyboards for mobile devices, and tagging for photos. Shneiderman's information visualization innovations include dynamic query sliders for Spotfire, development of treemaps for viewing hierarchical data, novel network visualizations for NodeXL, and event sequence analysis for electronic health records. Ben is the co-author with Catherine Plaisant of Designing the User Interface: Strategies for Effective Human-Computer Interaction (6th ed., 2016). His book Leonardo's Laptop (MIT Press) won the IEEE book award for Distinguished Literary Contribution. His book on The New ABCs of Research: Achieving Breakthrough Collaborations (Oxford, 2016) promotes an integration of applied and basic research using the methods of science, engineering, and design.

Professor Richard Harper

Institute for Social Futures, Lancaster University

Richard Harper was originally trained as a sociologist but is now a computer scientist, concerned with how new technologies shape us and how we in turn shape our technologies - in the space that is often known as Human Computer Interaction or HCI. He has written 13 books, including the award winning Myth of the Paperless Office and Texture - human expression in the age of communications overload. Topics in his 200 or so journal articles include the social impact and design of mobile phones, the future of search engines, the latest incarnation of artificial intelligence, the role of video-mediated connection in family life and even on the use of Wittgenstein in technology design. He is currently Co-Director of the Institute of Social Futures (ISF) at the University of Lancaster and Professor of Computer Science at the same institution. Amongst other endeavours, he is Principal Investigator on a Leverhulme Trust doctoral training centre on Material Social Futures. This entails looking at the relationship between new material forms of computing and energy production and the social arrangements they enable. Key to this research is seeking a balance between material possibility and environmental impact.

Dr. Sandra Wachter
Oxford Internet Institute

Sandra Wachter is a lawyer and Research Fellow in Data Ethics, AI, robotics and Internet Regulation/cyber-security at the Oxford Internet Institute at the University of Oxford. Sandra is specialising in technology-, IP-, and data protection law as well as European-, International-, human rights and medical law. Her current research focuses on the legal and ethical implications of Big Data, AI, and robotics as well as governmental surveillance, predictive policing, and human rights online. Sandra also serves as a policy advisor for governments, companies, and NGO's around the world on regulatory and ethical questions concerning emerging technologies."

Dr. Brent Mittelstadt
Research Fellow, Oxford Internet Institute, University of Oxford

Brent Mittelstadt is a Research Fellow and British Academy Postdoctoral Fellow in data ethics at the Oxford Internet Institute, a Turing Fellow at the Alan Turing Institute, and a member of the UK National Statistician's Data Ethics Advisory Committee. He is a philosopher specialising in digital ethics, bioethics, and ethical governance of algorithmic systems, with a particular focus on the EU General Data Protection Regulation and fairness, accountability and transparency in automated decision-making systems.

Professor Andy Crabtree
University of Nottingham

Andy Crabtree is Professor of Computer Science at the University of Nottingham. He has a background in sociology and interdisciplinary design and has particular expertise in ethnography and design. He is currently involved in a number of research projects that focus on the Internet of Things and the challenges to privacy occasioned by the widespread introduction of connected devices in everyday life, including the Databox project. Andy currently holds a 5-year EPSRC Fellowship that focuses on 'building accountability' into an increasingly intelligent Internet of Things. At stake here is whether or not legal accountability and the explanations enabled by the machine learning community will suffice, or whether other approaches and forms of accountability that respond to the mundane expectations and concerns of society's members are needed too?

Ryan Carrier

Executive Director, ForHumanity

Ryan Carrier founded ForHumanity after a 25 year career in finance. His global business experience, risk management expertise and unique perspective on how to manage the risk lead him to launch the non-profit entity, ForHumanity, personally. Ryan focused on Independent Audit of AI Systems as one means to mitigate the risk associated with artificial intelligence and began to build the business model associated a first-of-its-kind process for auditing corporate AIs, using a globally, open-source, crowd-sourced process to determine "best-practices".

Ryan serves as ForHumanity's Executive Director and Chairman of the Board of Directors, in these roles he is responsible for the day-to-day function of ForHumanity and the overall process of Independent Audit. Prior to founding ForHumanity, Ryan owned and operated Nautical Capital, a quantitative hedge fund, which employed artificial intelligence algorithms. He also was responsible for Macquarie's Investor Products business in the late 2000's. He worked at Standard & Poor's in the Index business and for the International Finance Corporation's Emerging Markets Database. Ryan has conducted business in over 55 countries and was a frequent speaker at industry conferences around the world. He is a graduate from the University of Michigan. Ryan became a Chartered Financial Analyst (CFA) in 2004.

Professor Chris Hankin

Co-Director, Institute for Security Science and Technology, Imperial College London

Professor Hankin joined Imperial College London in 1984 and was promoted to Professor in 1995. He is Co-Director of the Institute for Security Science and Technology. His research is in theoretical computer science, cyber security and data analytics. He leads multidisciplinary projects focussed on developing advanced visual analytics and providing better decision support to defend against cyber attacks for both enterprise systems and industrial control systems.

He is Director of the UK's Research Institute on Trustworthy Inter-connected Cyber-physical Systems (RITICS). He is Chair of the UK's Academic Resilience and Security Community (Academic RiSC) and sits on the ministerial oversight group of the Security and Resilience Growth Partnership. He is Chair of the Association for Computing Machinery (ACM) Europe Council. He is also a member of the ACM Publications Board.

Festival of Ideas

1st May 2019, Events Programme

Swansea University Prifysgol Abertawe | Computational Foundry Ffowndri Gyfrifiadol

9:00 - 10:00am: Launch breakfast (Research Crucible, 1st Floor)

10:00 - 10:20am

Welcome & Introduction:
Prof Matt Jones, Head of College of Science, and **Prof Alan Dix**, Director of Computational Foundry (Research Crucible)

10:20 - 11:20am

Algorithmic Accountability: Designing for Safety
Professor Ben Shneiderman, Dept of Computer Science & Human-Computer Interaction Lab, University of Maryland (Lecture Theatre 2, Ground Floor)

11:20am - Noon

The Future is not AI; the Future is HCI
Professor Richard Harper, Institute for Social Futures, Lancaster University (Lecture Theatre 2, Ground Floor)

Noon - 1:00pm

Lunch (Room 102, 1st Floor)

1:00 - 1:40pm

Algorithmic discrimination and group privacy in European law Dr. Sandra Wachter, Oxford Internet Institute (Lecture Theatre 2, Ground Floor)

1:40 - 2:20pm

Too Principled to Fail: Will 'AI Ethics' deliver algorithmic accountability?
Dr. Brent Mittelstadt, Research Fellow, Oxford Internet Institute, University Oxford (Lecture Theatre 2, Ground Floor)

2:20 - 2:40pm

Break (Research Crucible)

2:40 - 3:20pm

Right to an Explanation Considered Harmful Professor Andy Crabtree, University of Nottingham (Lecture Theatre 2, Ground Floor)

3:20 - 3:40pm

Independent Audit of AI Systems
Ryan Carrier, Executive Director, ForHumanity (Lecture Theatre 2, Ground Floor)

3:40 - 4:30pm

Discussant and Panel
Discussant: **Professor Chris Hankin** (Lecture Theatre 2, Ground Floor)

Ben Shneiderman, Dept of Computer Science & Human-Computer Interaction Lab, University of Maryland

Algorithmic Accountability: Designing for Safety

Vital services, such as communications, financial trading, healthcare, and transportation depend on sophisticated algorithms, some relying on unpredictable artificial intelligence techniques, such as deep learning, that are increasingly embedded in complex software systems. As high-speed trading, medical devices, and autonomous aircraft become more widely implemented, stronger checks become necessary to prevent failures. Design strategies that promote human-centered systems, which are comprehensible, predictable, and controllable can increase safety and make failure investigations more effective. Social strategies that support human-centered independent oversight during planning, continuous monitoring during operation, and retrospective analyses following failures can play a powerful role in making more reliable and trustworthy systems. This talk proposes a National Algorithms Safety Board to study failures and make recommendations for improvement. I believe that clarifying responsibility for failures stimulates improved design thinking.

Richard Harper, Institute for Social Futures, Lancaster University

The Future is not AI; the Future is HCI

I think there are two main ways of dealing with AI and its consequences for the role of the human in person-machine symbioses. The first entails treating the topic from the perspective of ergonomics. Here the task is to determine the balance between automation and human agency. The starting assumption is that the technology is more able than a human at some tasks but a human needs to retain some kind of role in the overall system so that they can pick up the pieces when the technology fails. This is the riddle that drives research - getting the right balance between human engagement and machine automation.

The second view treats the problem in terms of how computing systems can be designed to support the discretion of the user. This is the view from the field of Human Computer Interaction (HCI). In this perspective, the task of HCI research is to enable users to appropriate computers. This is the riddle that motivates the discipline. Solutions have entailed defining the right levels of abstractions to convey the functions of systems, the design of those systems to interrogate data of relevance to user purposes, and, beyond this, helping shape user ambitions around the new grammars of action that computer processing can enable. Clearly, AI offers new ways of enabling human discretion. Not all AI technologies will do so, of course, but some. Other types of AI maybe better dealt with from the ergonomics view.

However, it seems to me that, whereas with ergonomics there is a strong case that prior research offers deeply relevant insights for the current topic of AI, the same is not being claimed by the HCI community. Indeed, the manner of many HCI researchers implies that they no longer think human discretion with computers requires novel solutions. They do not invoke any of the canonical papers in the discipline to justify a role in the 'age of AI'. One reason for their lack of morale, if that is what it is, might be the distracting claims that the AI community make about 'natural interfaces', such as those to do with voice-based dialogues. These suggest that interaction between persons and machines is simply a question of communication. This is plainly egregious – it is about purposes and competences as well. Another has to do with the idea that AI systems can take on all discretionary tasks. This too is egregious. One only has to remember that the invention of HCI turned around the realisation that desktop editing tools should be designed to enable the user to apply their discretion in the content of documents, not the computing applications. Document layout, printing and such could be automated, but the reasoning purposes of documents remained in the users' hands, so to speak. Likewise with AI, in my view. The question is not whether AI automates, it is what AI technologies (and the systems they are part of) are designed to let users do that matters.

Despite their low morale, I am convinced that it is HCI researchers that must determine what these might be – these new AI enabled systems for user discretion. Other communities cannot do this; not ergonomists, nor AI researchers. Only those preoccupied with the HCI riddle can do so. To date, HCI researchers have not risen to the task – for whatever reason. My talk is intended to encourage them to do so.

Sandra Wachter, Oxford Internet Institute

Algorithmic discrimination and group privacy in European law

Big Data analytics and artificial intelligence (AI) draw non-intuitive and unverifiable inferences and predictions about the behaviours, preferences, and private lives of individuals. These inferences draw on highly diverse and feature-rich data of unpredictable value, and create new opportunities for discriminatory, biased, and invasive decision-making, often based on sensitive attributes of individuals' private lives. European data protection law affords greater protection to processing of sensitive data, or 'special categories', describing characteristics such as health, ethnicity, or political beliefs. However, this talk will show that the law does not sufficiently guard against privacy risks.

Big Data and AI also aim to identify unintuitive small patterns and meaningful connections between individuals and their data. The analytics behind much automated decision-making and profiling is not concerned with singling out or identifying a unique individual, but rather with drawing inferences from large datasets, calculating probabilities, and learning about types or groups of people.

These technologies thus expand the scope of potential victims of discrimination and other potential harms (e.g. privacy, financial, reputational) to include ephemeral groups of individuals perceived to be similar by a third party. European anti-discrimination laws, which are based on historical lessons, will fail to apply to 'ad hoc' groups which are not defined by a historically protected attribute (e.g. ethnicity, religion).

I conclude by arguing that a 'right to reasonable inferences' could provide a remedy against new forms of discrimination and greater protection for group privacy interests.

Brent Mittelstadt, Research Fellow, Oxford Internet Institute, University of Oxford

Too Principled to Fail: Will 'AI Ethics' deliver algorithmic accountability?

AI Ethics is now a global topic of discussion in policy and academic circles. Many initiatives (at least 29 at the time of writing) aim to develop high-level principles and values to guide the development and deployment of AI across the public and private sectors. I refer to 'AI Ethics' with a capital E, meaning public and private, often multi-stakeholder initiatives which purport to develop governance frameworks, codes of conducts, principles or values to guide the development and deployment of AI. I am not addressing the state of academic work on ethical aspects of AI. Comparisons have recently been drawn between these initiatives and a principled approach to medical ethics. This convergence of AI Ethics around principles defined in medical ethics is opportune. 'Principilism' in medical ethics is perhaps the most prominent and successful implementation of a principled approach to applied ethics. However, there are reasons to be concerned about the future impact of the field. Significant differences exist between medicine and AI development and governance, which suggest a principled approach in the latter may not enjoy success comparable to the former. This talk considers historical lessons from the development and application of medical ethics to critically assess the potential impact of a principled approach to ethics in AI development and governance.

Andy Crabtree, University of Nottingham

Right to an Explanation Considered Harmful

Lay and professional reasoning has it that newly introduced data protection regulation in Europe – GDPR – mandates a 'right to an explanation' of automated decision-making. This has been read as requiring that the machine learning (ML) community build explainable machines to enable legal compliance. However, we argue that this reading should be considered harmful as it creates unrealistic expectations for the ML community and society at large, not least because GDPR does not require that intelligent machines provide explanations. We consider a) what is required of an explanation by GDPR; b) the kind of explanation enabled by ML methods of interpretability and their fit with GDPR; c) the inability of either legal or ML explanations to adequately address the societal imperative of accommodating intelligent machines in everyday life; and d) the need to broaden societal representation beyond the ethics-law-technology matrix, to include HCI and other salient disciplines that seek to engage citizens in shaping AI & Society around their accountable expectations, interests and concerns.

Ryan Carrier, Executive Director, ForHumanity

Independent Audit of AI Systems

A discussion about the merits and operational process of Independent Audit of AI Systems. AS AI and Automation expand exponentially throughout industry and deeply into the lives of individuals, what governance, accountability and oversight exists on these systems. How do we ensure the optimal outcome for humanity from the expansion of these systems. This talk will focus on the need for 3rd party oversight on AI and Algorithms covering ethics, bias, privacy, trust and cybersecurity. It will explain the corollaries to FASB/financial accounting and the subsequent benefits to society broadly. The talk will also dwell on the similarities and differences between financial auditing and AI/Algorithmic auditing. Furthermore, we will discuss the process creating the audit rules. Examples will be provided for potential rules and how these efforts can dramatically improve AI Safety. The talk will also explore, pitfalls and challenges with creating an audit process for artificial intelligence. Who will volunteer? Can the process be mandated? Will the process have a single, global standard or localized standards? Who can participate and how?

Swansea University
Prifysgol Abertawe

Computational Foundry
Ffowndri Gyfrifiadol

cherishde

Swansea University
Prifysgol Abertawe