

MANCHESTER  
1824

The University of Manchester

# INSTITUTE FOR DATA SCIENCE AND ARTIFICIAL INTELLIGENCE

---

A DIGITAL FUTURES ACTIVITY

## ADVANCES IN DATA SCIENCE AND AI CONFERENCE

MANCHESTER | 20-21 JUNE 2022



**The  
Alan Turing  
Institute**

The 2022 Advances in Data Science and AI Conference is funded by  
The University of Manchester and The Alan Turing Institute.

[www.manchester.ac.uk](http://www.manchester.ac.uk)

[www.turing.ac.uk](http://www.turing.ac.uk)

# Welcome



Professor Magnus Rattray, Director of the Institute for Data Science and AI, The University of Manchester

The University of Manchester’s Data Science Institute was founded in 2016 and our first annual conference took place in 2017.

In 2018 we joined the Alan Turing Institute and rebranded to the Institute for Data Science and Artificial Intelligence (IDSAI).

This year we look forward welcoming our participants, virtually and in-person, for an exciting programme that includes the latest developments in explainable and causal approaches to AI, ethical and responsible AI, AI for health and natural language processing.

- [3](#) Welcome
- [4](#) Programme
- [6](#) Speakers
- [12](#) Abstracts
- [18](#) Contributors
- [20](#) Features
- [32](#) Campus Map

Conference organisers:

Professor Magnus Rattray  
Professor Sophia Ananiadou  
Dr Mauricio Álvarez López

Follow and contribute to live conference discussion [#ADSAI2022](#)

Follow the Institute for Data Science and AI [@idsai\\_uom](#)

[www.manchester.ac.uk/ADSAI22](http://www.manchester.ac.uk/ADSAI22)

# Programme

## Day One

Time	Theme	Speaker/Event	Title
1000	Introduction		
1015	Explainable and causal approaches to AI (part 1)	Virginia Aglietti, Deepmind and The Alan Turing Institute	Causal decision-making in static and dynamic settings
1050		Arthur Leroy, The University of Sheffield	Multi-Means Gaussian Processes: A novel probabilistic framework for multi-correlated longitudinal data
1100	Break		
1115	Explainable and causal approaches to AI (part 2)	Sijia Li, University of Leeds	Scalable Bigraphical Lasso: Two-way Sparse Network Inference for Count Data
1125		Umang Bhatt, University of Cambridge, Turing Institute	Challenges and Frontiers in Deploying Transparent Machine Learning
1200	Lunch Poster Session		
1330	Ethical and responsible AI (part 1)	Sylvie Delacroix, University of Birmingham	Data trusts as bottom-up data empowerment infrastructure
1405		Youngmin Kim, The University of Manchester	Benchmarking Safe Optimization Algorithms: Methodology, Software and Validation
1415		David Leslie, Turing Institute	From principles to practice and back again: Building a responsible AI ecosystem from the ground up
1450	Break		
1500	Ethical and responsible AI (part 2)	Darminder Ghataoura, Fujitsu	What are the barriers for closer Human-Machine Teaming?
1535		Panel session	
1635	Close		

All sessions will take place in Engineering Building A, in Lecture Theatre A and the main event foyer. See back page for location details.

# Programme

## Day Two

Time	Theme	Speaker/Event	Title
0910	Introduction		
0915	Explainable AI and NLP (part 1)	Andre Martins, Instituto Superior Tecnico, Lisbon	Towards Explainable and Uncertainty-Aware NLP
0950		Tom Quilter, Pinpoint Learning	Large-Scale Predictive Models in Education - New Data and Improved Accuracy
1000		Louis Filstroff, Aalto University	Approximate Bayesian Computation with Domain Expert in the Loop
1010		Naoaki Okazaki, Tokyo Institute of Technology	Towards controllable, faithful, and explainable text generation
1045	Break		
1100	Explainable AI and NLP (part 2)	Nikolaos Aletras, The University of Sheffield	How can we improve explanation faithfulness in NLP
1135		Christopher McGreavy, Aalto University	Towards Virtual Laboratories: AI Assistance in Experimentation
1145		Federico Tavella, The University of Manchester	American Sign Language acquisition via Imitation Learning
1155		Christina Lioma, University of Copenhagen	Pitfalls with ablation in neural network architectures
1230	Lunch		
1320	AI for Health	Alfredo Vellido, Polytechnic University of Catalonia	Explain yourself: XAI as social responsibility
1355		Benediktas Valys, University of Sheffield	HMM-based clustering of multimorbidity trajectories
1405		Tianyu Cui, Aalto University	Improving Neural Networks for Genotype-Phenotype Prediction Using Published Summary Statistics
1415		Matthew Lyon, The University of Manchester	Angular Super-Resolution in Diffusion MRI with a 3D Recurrent Convolutional Autoencoder
1425		Melanie Fernandez Pradier, Microsoft Research	From the research to the clinic: aligning ML systems with clinicians for improved mental care
1500	Close		

# Speakers

## Virginia Aglietti

Virginia Aglietti is a Research Scientist at DeepMind within the Data Efficient and Bayesian Learning Team.

Virginia's research focuses on causal inference with an emphasis on methods that combine causality and machine learning algorithms for decision making.

Prior to joining DeepMind, Virginia received a PhD in statistics from the University of Warwick where she was part of the OxWaSP program, working on Gaussian processes methods for structured inference and sequential decision making under the supervision of Professor Theo Damoulas.




## Nikolaos Aletras

Nikolaos Aletras is a Senior Lecturer in Natural Language Processing at the Computer Science Department of the University of Sheffield, co-affiliated with

the Machine Learning (ML) group. Previously, he was a Lecturer in Data Science at the Information School, University of Sheffield. Nikolaos has gained industrial experience working as a scientist at Amazon (Amazon ML and Alexa) and was a research associate at UCL, Department of Computer Science. Nikolaos completed his PhD in Natural Language Processing at the University of Sheffield, Department of Computer Science.

Nikolaos' research interests are in NLP, Machine Learning and Data Science. Specifically, He is interested in developing computational methods for social media analysis and the law, ML for NLP and information retrieval methods for improving access to large document collections. His research has been funded by the EPSRC, ESRC, Leverhulme Trust, EU, Amazon and Google.



## Umang Bhatt



Umang Bhatt is a PhD candidate in the Machine Learning Group at the University of Cambridge.

His research lies in trustworthy machine learning.

Specifically, he focuses on algorithmic transparency and its effects on stakeholder decision-making. His work has been supported by a JP Morgan AI PhD Fellowship, a Mozilla Fellowship, and a Partnership on AI Research Fellowship. He is currently an Enrichment Student at The Alan Turing Institute, a Student Fellow at the Leverhulme Centre for the Future of Intelligence, and an Advisor to the Responsible AI Institute. Previously, he was a Fellow at the Mozilla Foundation and a Research Fellow at the Partnership on AI. Umang received a B.S. and M.S. in Electrical and Computer Engineering from Carnegie Mellon University.

## Sylvie Delacroix



Professor Delacroix's research focuses on the intersection between law and ethics, with a particular interest in habits and the infrastructure that moulds our habits

(data-reliant tools are an increasingly big part of that infrastructure).

She is considering the potential inherent in bottom-up Data Trusts as a way of reversing the current top-down, fire-brigade approach to data governance. She co-chairs the Data Trust Initiative, which is funded by the McGovern Foundation: see <https://datatrusts.uk>.

Professor Delacroix has served on the Public Policy Commission on the use of algorithms in the justice system (Law Society of England and Wales) and the Data Trusts Policy group (under the auspices of the UK AI Council). She is also a Fellow of the Alan Turing Institute. Professor Delacroix's work has been funded by the Wellcome Trust, the NHS and the Leverhulme Trust, from whom she received the Leverhulme Prize. Her latest book -*Habitual Ethics?*- is forthcoming with Bloomsbury / Hart Publishing in July 2022.

@SylvieDelacroix  
<https://delacroix.uk>

# Speakers

## Melanie Fernandez Pradier

Melanie Fernandez Pradier is a senior researcher at Microsoft Research working on probabilistic models for healthcare and the live sciences.



Her research interests include generative models for Immunomics, Bayesian modelling, interpretable ML, and out-of-distribution generalization. Melanie received her PhD on Bayesian Nonparametrics from University Carlos III in 2017. She was a postdoctoral fellow at the Harvard Data Science Initiative until 2020, working on interpretable ML and deep Bayesian models with Finale Doshi-Velez. Melanie is a co-founder of the ICBINB Initiative, and an editor of the MDPI special issue on 'Foundations and Challenges of Interpretable ML'.

## Darminster Ghataoura

Dr. Darminster Ghataoura is Head of AI and leads Fujitsu's offerings and capabilities in AI and Data Science within the Defence and National Security space.



Darminster manages the strategic and technical AI relationships with partners, academic institutions and UK government and is a Strategic Advisory Network member for the UKRI Trustworthy Autonomous Systems (TAS) Hub.

Darminster is also a Simon Industrial fellow at the University of Manchester, Decision and Cognitive Sciences Department, with his main focus in the area of 'Human machine-teaming'.

Darminster has over 15 years' experience in the design and development of AI systems and services across the UK Public and Defence sectors as well as UK and international commercial businesses.

He was awarded with the Fujitsu Distinguished Engineer recognition in 2020 and holds an Engineering Doctorate (EngD) in Autonomous Military Sensor Networks for Surveillance Applications, from University College London (UCL).



## David Leslie



David Leslie is the Director of Ethics and Responsible Innovation Research at The Alan Turing Institute. So far in his career, David has held positions

at Princeton, Yale and Harvard.

David has served as an elected member of the Bureau of the Council of Europe's Ad Hoc Committee on Artificial Intelligence (CAHAI), and is the author of the UK Government's official guidance on the responsible design and implementation of AI systems in the public sector, understanding artificial intelligence ethics and safety (2019) and a principal co-author of Explaining decisions made with AI (2020)

He is also Principal Investigator of a UKRI-funded project called PATH-AI: Mapping an Intercultural Path to Privacy, Agency and Trust in Human-AI Ecosystems. Most recently, he has received a series of grants from the Global Partnership on AI, the Engineering and Physical Sciences Research Council, and BEIS to lead a project titled, Advancing Data Justice Research and Practice.

## Christina Lioma



Christina Lioma is a professor in machine learning at the Department of Computer Science, University of Copenhagen.

Her research focuses on applied machine learning, information retrieval and web search technologies, web data mining and analytics, recommendation systems and natural language processing.

She has a track-record of research collaboration with Danish and international industry, and an alumni of >20 PhD students and postdocs. Since 2012, Christina Lioma has attracted more than 50 million USD in external funding.

More detailed bios can be found at [manchester.ac.uk/ADSAI22/speakers](https://manchester.ac.uk/ADSAI22/speakers)

# Speakers

## André Martins

André Martins (PhD 2012, Carnegie Mellon University and University of Lisbon) is an Associate Professor at Instituto Superior Técnico, University of Lisbon, researcher at Instituto de Telecomunicações, and the VP of AI Research at Unbabel.



His research, funded by a ERC Starting Grant (DeepSPIN) and other grants (P2020 project Unbabel4EU and CMU-Portugal project MAIA) include machine translation, quality estimation, structure and interpretability in deep learning systems for NLP. His work has received best paper awards at ACL 2009 (long paper) and ACL 2019 (system demonstration paper). He co-founded and co-organizes the Lisbon Machine Learning School (LxMLS), and he is a Fellow of the ELLIS society.

## Naoaki Okazaki

Naoaki Okazaki is Professor in School of Computing, Tokyo Institute of Technology, Japan.



Prior to this faculty position, he worked as a post-doctoral researcher at the University of Tokyo (in 2007-2011), and as an associate professor at Tohoku University (2011-2017).

He is also a visiting research scholar of the Artificial Intelligence Research Center (AIRC), National Institute of Advanced Industrial Science and Technology (AIST). His research areas include Natural Language Processing (NLP), Artificial Intelligence (AI), and Machine Learning.

## Alfredo Vellido



Alfredo Vellido is Associate Professor in Computer Science at the Polytechnic University of Catalonia.

Alfredo completed a B.Sc. in Physics at the Universidad del País Vasco (UPV-EHU), Spain in 1996. He earned his Ph.D in Neural Computation at Liverpool John Moores University (LJMU). Alfredo then became Ramón y Cajal Research Fellow at Universitat Politècnica de Catalunya (UPC), Barcelona, Spain (2003-2008) and has been Associate Professor at UPC since then.

His research focuses on ML applications in healthcare and medicine and, lately, on their societal impact.

# Abstracts

## Causal decision-making in static and dynamic settings

**Virginia Aglietti**

In this talk I will consider the problem of understanding how to intervene in a causal system so as to optimize an outcome of interest. I will present two methodologies that, by linking causal inference, experimental design, and Gaussian process modeling, allow one to efficiently learn the causal effects and identify an optimal intervention to perform, both in static and dynamic settings. In this first part of the talk I'll focus on static settings and discuss how finding an optimal intervention to perform requires solving a new optimization problem which we call Causal Global Optimization.

I'll then introduce Causal Bayesian Optimization, an algorithm that allows solving these problems by incorporating the knowledge of the causal graph in Bayesian Optimization thus decreasing the optimization cost and avoiding suboptimal solutions.

In the second part of the talk, I will then show how the approach developed for static settings can be extended to select actions at different time steps.

Longer abstracts can be found at [manchester.ac.uk/ADSAI22/speakers](http://manchester.ac.uk/ADSAI22/speakers)

## Challenges and Frontiers in Deploying Transparent Machine Learning

**Umang Bhatt**

Explainable machine learning offers the potential to provide stakeholders with insights into model behavior, yet there is little understanding of how organizations use these methods in practice. In this talk, we will discuss recent research exploring how organizations view and use explainability. We find that most deployments are not for end-users but rather for machine learning engineers, who use explainability to debug the model. There is thus a gap between explainability in practice and the goal of external transparency since explanations are primarily serving internal stakeholders. Providing useful external explanations requires careful consideration of the needs of stakeholders, including end-users, regulators, and domain experts. Despite this need, little work has been done to facilitate inter-stakeholder conversation around explainable machine learning. To help address this gap, we report findings from a closed-door, day-long workshop between academics, industry experts, legal scholars, and policymakers to develop a shared language around explainability and to understand the current shortcomings of and potential solutions for deploying explainable machine learning in the service of external transparency goals.

## Data trusts as bottom-up data empowerment infrastructure

**Sylvie Delacroix**

It proceeds from an analysis of the particular type of vulnerability concomitant with our 'leaking' data on a daily basis, to argue that data ownership is both unlikely and inadequate as an answer to the problems at stake.

There are three key problems that bottom-up data trusts seek to address:

1. Lack of mechanisms to empower groups, not just individuals
2. Can we do better than current 'make belief' consent?
3. Can we challenge the assumed tradeoff between promoting data-reliant common goods on one hand and addressing vulnerabilities that stem from data sharing?

## From principles to practice and back again: Building a responsible AI ecosystem from the ground up

**David Leslie**

The recent history of AI ethics and governance has been characterised by increasingly vocal calls for a shift from principles to practice. In this talk, I will explore some of shortcomings of those who have championed this shift. I will argue that many staunch supporters of the "principles to practice" movement have, in fact, taken too much of a documentation- and audit-centred point of view, thereby neglecting the social, cultural, and cognitive preconditions of the responsible innovation practices they aspire to advance. Beyond off-the-shelf tools and documentation-centred governance instruments, closing the gap between principles and practice requires a transformation of organisational cultures, technical approaches, and individual attitudes from inside the processes and practices of design, development, and deployment themselves. Achieving this requires researchers, technologists, and innovators to establish and maintain end-to-end habits of critical reflection and deliberation across every stage of a research or innovation project's lifecycle.

# Abstracts

## Open challenges for closer Human-Machine Teaming

**Darminder Ghataoura**

Interactions with artificial intelligence (AI) is now commonplace. We increasingly rely on intelligent systems to extend our human capabilities, from chat-bots that provide technical support to virtual assistants like Siri and Alexa. However, today's intelligent machines are essentially tools, not true collaborative partners.

If the vision for Human-Machine teaming is to augment human processes and improve productivity, intelligent machines will need to be flexible and adaptive to the states of the human, as well as the environment. This poses interesting challenges such as understanding human capabilities, intentions, the ability of the machine to generalise to new situations and the all important trust dimension. In this talk we will look to highlight some of these challenges and to understand where we are currently on this journey.

## Towards Explainable and Uncertainty-Aware NLP

**André Martins**

Natural language processing systems are becoming increasingly more accurate and powerful. However, in order to take full advantage of these advances, new capabilities are necessary for humans to understand model predictions and when to question or to bypass them.

In the first part of this talk, I will describe how sparse modeling techniques can be extended and adapted for facilitating sparse communication in neural models.

In the second part, I will present an uncertainty-aware approach to machine translation evaluation.

This is joint work with Taya Glushkova, Nuno Guerreiro, Vlad Niculae, Ben Peters, Marcos Treviso, and Chryssa Zerva in the scope of the DeepSPIN and MAIA projects.

Longer abstracts can be found at [manchester.ac.uk/ADSAI22/speakers](http://manchester.ac.uk/ADSAI22/speakers)

Towards controllable,  
faithful and explainable  
text generation

**Naoaki Okazaki**

Deep neural networks have made a breakthrough in various NLP tasks in the recent ten years. Benchmark scores of neural machine translation were nearly doubled from those of the statistical machine translation with the advancements in sequence-to-sequence models, attention mechanism, and the Transformer architecture. This success also broadened natural language generation (NLG) applications into abstractive summarization and image caption generation. However, relying on a large amount of supervision data, the current NLG models are not flexible to adapt to a slightly different task, reduce unreliable outputs, and explain the reason behind outputs. In this talk, I will present our recent studies on controllable and faithful abstractive summarization and explainable grammatical error correction.

How can we improve  
explanation faithfulness in  
LLM

**Nikolaos Aletras**

Large pre-trained language models (LLMs) currently dominate performance across language understanding tasks. However, their complexity and opaque nature have opened up new challenges on how to extract faithful explanations (or rationales), which accurately represent the true reasons behind a model's prediction when adapted to downstream tasks. In this talk, I will present recent work from my group on how we can improve faithfulness of LLM predictions and a study of explanation faithfulness in out-of-domain settings.

# Abstracts

Pitfalls with ablation in neural network architectures

**Christina Lioma**

Ablation tests are frequently used for analysing machine learning performance. Generally, ablation refers to the removal of a component in order to understand its contribution to the overall decision-making process. For instance, it is expected that the ablation of a feature during classification will affect the output somehow analogously to the importance of that feature for the classification task at hand. Ablation is therefore routinely used to attribute feature importance, as well as to explain machine learning output in a partial, approximate, and model-agnostic way.

This talk will point out a core problem when using ablation with neural network architectures. The problem stems from the tendency of neural network architectures to ignore complex predictive features in the presence of few simple predictive features, even when the complex features have significantly greater predictive power than the simple features. This talk will provide evidence demonstrating the existence of this tendency, even in small neural network architectures, and show how this may entirely invalidate the standard interpretation of ablation tests. A discussion about why this is important and ways of moving forward will be provided.

This is joint work with Qiuchi Li (University of Copenhagen).

Explain yourself: XAI as social responsibility

**Alfredo Vellido**

Machine-learning-based systems are now part of a wide array of real-world applications seamlessly embedded in the social realm. In the wake of this realization, strict legal regulations for these systems are currently being developed, addressing some of the risks they may pose. This is the coming of age of the concepts of interpretability and explainability in machine-learning-based data analysis, which can no longer be seen just as an academic research problem.



From the research to the clinic: aligning ML systems with clinicians for improved mental care

**Melanie Fernandez Pradier**

Current Machine Learning (ML) approaches can achieve high levels of performance on clinical data that matches or even exceeds human clinicians, e.g., by leveraging the power of deep neural networks and large data repositories. However, success stories of deploying ML systems in healthcare domains remain limited. Lack of interpretability and misalignment with clinical expertise are some frequent caveats of such systems. In this talk, I will share the insights that my colleagues and I learned when trying to bridge the gap from the research to the clinic in the context of antidepressant prescriptions.

I will start with a word of caution based on a within-subject factorial user study: ML recommendations and explanations may impact clinicians decisions negatively. I will then present preferential MoE, a novel human-ML mixture-of-experts model that rely on human expertise as much as possible, demonstrating its potential in the management of Major Depressive Disorder.

# Contributors: short talks

## **Multi-Means Gaussian Processes: A novel probabilistic framework for multi-correlated longitudinal data**

Arthur Leroy, Research Associate in Computer Science, The University of Sheffield  
Co-authors: Mauricio A Álvarez, The University of Manchester

## **Scalable Bigraphical Lasso: Two-way Sparse Network Inference for Count Data**

Sijia Li, PhD candidate, University of Leeds  
Co-authors: Martín López-García, Neil Lawrence and Luisa Cutillo

## **Benchmarking Safe Optimization Algorithms: Methodology, Software and Validation**

Youngmin Kim, PhD Candidate, The University of Manchester  
Co-authors: Richard Allmendinger and Manuel López-Ibáñez

## **Large-Scale Predictive Models in Education - New Data and Improved Accuracy**

Tom Quilter, Senior Engineer, Pinpoint Learning  
Co-authors: Richard Turner, Anastasija Ilic and Karen Poon

## **Approximate Bayesian Computation with Domain Expert in the Loop**

Louis Filstroff, Postdoctoral Researcher, Aalto University  
Co-authors: Ayush Bharti and Samuel Kaski

## **Towards Virtual Laboratories: AI Assistance in Experimentation**

Christopher McGreavy, Postdoctoral Researcher, Aalto University  
Co-authors: Carlos Sevilla-Salcedo, Arto Klami and Samuel Kaski

## **American Sign Language acquisition via Imitation Learning**

Federico Tavella, PhD Candidate, The University of Manchester  
Co-authors: Aphrodite Galata and Angelo Cangelosi

## **HMM-based clustering of multimorbidity trajectories**

Benediktas Valys, PhD Candidate, University of Sheffield  
Co-authors: Ilaria Bellantuono, Amaia Calderon Larrañaga, Alessandra Marengoni, Davide Liborio Vetrano, Mauricio Álvarez and Mari-Cruz Villa-Uriol

## **Improving Neural Networks for Genotype-Phenotype Prediction Using Published Summary Statistics**

Tianyu Cui, PhD Candidate, Aalto University  
Co-authors: Khaoula El Mekkaoui, Aki Havulinna, Pekka Marttinen and Samuel Kaski

## **Angular Super-Resolution in Diffusion MRI with a 3D Recurrent Convolutional Autoencoder**

Matthew Lyon, PhD Candidate, The University of Manchester  
Co-authors: Paul Armitage and Mauricio Alvarez

# Contributors: posters

## **Genetic Algorithms for the Design of Manufacturing Systems**

April Bryan, Lecturer in Advanced Manufacturing, The University of Manchester

## **Detecting Fraud in Financial Payments**

Tamsin Crossland, Senior Architect, Icon Solutions

## **Signed Graph Convolutional Networks for node classification**

Thu Trang Dinh, PhD Candidate, The University of Manchester

## **Robustifying mixture models using integral probability metric**

Adam Farooq and Yordan Raykov

## **Structure Learning of Bayesian Networks: Challenges and Opportunities**

Zhigao Guo, Postdoctoral Research Associate, The University of Manchester

## **AI in the Boardroom: Let the Law be in the Driving Seat**

Joseph Lee, Reader in Corporate and Financial Law, The University of Manchester

Co-author: Peter Underwood

## **On NFT Art and its Price**

Min-Bin Lin, PhD Candidate, Humboldt University, Berlin

Co-authors: Vanessa Emanuela Guarino and Cathy Yi-Hsuan Chen

## **Realised Volatility Forecasting: Machine Learning via Financial Word Embedding**

Eghbal Rahimikia, Lecturer in Finance and PhD Candidate, The University of Manchester

Co-authors: Stefan Zohreen and Ser-Huang Poon

## **Topological Data Analysis for Economics and Finance**

Simon Rudkin, Senior Lecturer, The University of Manchester

Co-authors: Pawel Dlotko and Wanling Qiu

## **Exploring Rawlsian Fairness for k-means Clustering**

Stanley Simoes, Marie Curie Early Stage Researcher, Queen's University, Belfast

Co-authors: Deepak P and Muiris MacCarthaigh

## **Dialogue Explanation With Reasoning For AI**

Yifan Xu, PhD Candidate, The University of Manchester

MANCHESTER  
1824

The University of Manchester

# DIGITAL FUTURES

---

## TRANSFORMING OUR WORLD

Digital Futures is a highly interdisciplinary network that operates across the whole range of The University of Manchester's digital research.

.....  
MORE THAN 1700 RESEARCHERS ACROSS 30 DISCIPLINES  
.....

UNDERSTANDING AND DRIVING DIGITAL TRANSFORMATION  
.....

ENGAGING WITH CITIZENS, BUSINESS AND GOVERNMENT  
.....

PROVIDING THOUGHT LEADERSHIP, SHAPING THE FUTURE  
.....

[www.digitalfutures.manchester.ac.uk](http://www.digitalfutures.manchester.ac.uk)

# Digital Futures at The University of Manchester

Digital Futures is The University of Manchester's strategic response to the challenges and opportunities presented by the digital revolution.

It brings together a multidisciplinary community of over 1700 researchers from across the University's three faculties with the aim of:

- providing an integrated view of our digital research, creating a 'front door' for potential partners
- bringing together cognate research communities, and exploiting synergies
- developing a coherent strategy for addressing major societal and technical challenges

“ By creating a coherent framework of societal challenges and cross-cutting capabilities, we've been able to create an integrated and accessible view of our digital research for external stakeholders whilst exploiting synergies and building critical mass internally.

Professor Chris Taylor

”


The programme is built around challenges and crosscutting capabilities.


**Societal Challenges** focus on real-world activities that are economically and socially important, building on multidisciplinary strengths within the University.

**Institutional Challenges** focus on the potential for digital technology to transform what we do as a University and how we do it: how we conduct research, what and how we teach, and how we function as an organisation.

**Cross-cutting Capabilities** are digital frameworks, technologies and methods that are important areas of research in their own right and provide the underpinning for addressing the challenges.


Find out more:  
[digitalfutures.manchester.ac.uk](http://digitalfutures.manchester.ac.uk)





To hear about more events like this, use the QR code to access our newsletter sign-up page.

You can also follow us on Twitter at [@DigitalUoM](https://twitter.com/DigitalUoM)



# Natural language processing and text mining at The University of Manchester

The National Centre for Text Mining (NaCTeM), based in the Department of Computer Science in the School of Engineering at the University of Manchester, is the first publicly-funded text mining centre in the world. It was established to provide support, advice, and information on text mining (TM) technologies and to disseminate information from the larger TM community, whilst also providing tailored services and tools in response to the requirements of the academic community.

NaCTeM researchers have excelled in community shared tasks and challenges, notably in BioCreAtIvE III, IV and V, in BioNLP 2011 and 2013 (for the most complex task of event extraction) and most recently obtained two first places in tasks of the 5th CL-SciSumm Shared Task 2019. Moreover, NaCTeM's participation in DARPA's \$45m Big Cancer Mechanism initiative, in a consortium led by the University of Chicago, saw it produce in 2015 the top performing system for extracting information to support cancer pathway modelling. NaCTeM's academic and industrial research projects range over many domains from biology and biomedicine to biodiversity, toxicology, neuroscience, materials, history, social sciences, insurance, and health and safety in the construction industry, with funding coming from EPSRC, ESRC,

MRC, AHRC, Wellcome Trust, NIH, Pacific Life Re, Lloyd's Register Foundation, AstraZeneca, DARPA, EC Horizon 2020, JST and the cosmetics and extracts industry, among others.

Applications arising from such research include Thalia, a semantic search engine over more than 20m biomedical abstracts; Facta+, to find unsuspected associations in the biomedical literature; HoM, allowing semantic search of historical medical and public health archives; and RobotAnalyst, supporting the hitherto laborious screening stage of systematic reviewing through active learning techniques. NaCTeM also collaborates closely with the Artificial Intelligence Research Center, National Institute of Advanced Industrial Science and Technology, Japan.

Find out more:  
[nactem.ac.uk](http://nactem.ac.uk)



# Alan Turing

Turing's name is synonymous with Data Science and Artificial Intelligence at The University of Manchester, from his work here in the 1940s and 50s to our current partnership with the Alan Turing Institute.


Alan Turing joined The University of Manchester in 1948, as Reader in the Mathematics department. In 1949, he took up the post of Deputy Director of the Computing Laboratory. This coincided with Manchester's work to develop the world's first modern computer: the 'Manchester Baby'.

In 1950 Turing published a seminal paper entitled 'Computing Machinery and Intelligence', in which he first addressed the issue of what was to be labelled artificial intelligence (AI).

In his paper, Turing developed a method to determine whether a machine can be recognised as 'intelligent' by demonstrating human-like thinking - this challenge was called the 'Imitation Game', and is now known as the 'Turing Test'. Computing Machinery and Intelligence would have a significant influence on AI, a research area that continues exponentially today and where Manchester is still a global centre of excellence.

Turing contributed to the development of the Manchester Mark 1 and the Ferranti Mark 1, the world's first commercially-available digital computer, designed at The University of Manchester by Freddie Williams and Tom Kilburn.

Turing, Kilburn and Williams all now have Manchester buildings bearing their name.



The University's Alan Turing Building

**// Turing's pioneering work in mathematics, computing and artificial intelligence helped to distinguish and enhance our reputation in these academic areas, something that continues to this day.**

**Here at The University of Manchester, Turing's legacy lives on as future generations of mathematicians and physicists study in a building that bears his name. This latest recognition is richly deserved and a fitting tribute to one of the greatest scientists of the 20th Century. //**

Professor Dame Nancy Rothwell





The Alan Turing Memorial,  
Sackville Gardens, Manchester

The Alan Turing Institute was created as the national institute for data science in 2015, adding artificial intelligence to their remit in 2017. The institute is a collaborative hub, with roots in universities and centres of research excellence across the UK, and strong links to a growing network of industry, public sector, and third sector partners.

Since 2018 there have been thirteen University partners: Birmingham, Bristol, Cambridge, Edinburgh, Exeter, Leeds, Manchester, Newcastle, Oxford, Queen Mary University of London, Southampton, UCL and Warwick.

The Alan Turing Institute also collaborates with businesses and public and third sector organisations to apply this research to real-world problems, with lasting effects for science, the economy, and the world we live in.

Our researchers collaborate across disciplines to generate impact, both

through theoretical development and application to real-world problems. They are fuelled by the desire to innovate and add value.

The institute has three ambitious goals:

- Advance world-class research
- Train the leaders of the future
- Lead the public conversation


Being a national institute enables the institute to deliver benefits that a single university could not deliver alone. Breaking down disciplinary boundaries, computer scientists, engineers, statisticians, mathematicians, and scientists work together under one shared goal.

Crucially, the Turing Institute is a convening power, bringing together the best talent in the data science and AI community to speak to industry, policy-makers, and the public.

In October 2021, The University of Manchester announced 33 Manchester-based Turing Fellows across all three of our Faculties.

Find out more: [turing.ac.uk](https://turing.ac.uk)

Use the QR code to find out more about Manchester's Turing Fellows.



# Samuel Kaski

Professor of Artificial Intelligence  
Turing AI World-Leading Research Fellow

The University of Manchester is developing unique research teams to help cure humanity's increasingly complex future health and societal problems by partnering researchers with Artificial Intelligence (AI).

The University of Manchester's Professor Sami Kaski is among the first Turing Artificial Intelligence (AI) World-Leading Research Fellows. The fellowships, named after AI pioneer Alan Turing, are part of the UK's commitment to further strengthen its position as a global leader in the field.

Through his fellowship, Professor Kaski aims to overcome a fundamental limitation of current AI systems; that they require a detailed specification of the goal.

Artificial intelligence is still limited by the fact that human intervention is needed to set appropriate objectives and rewards to tell AI systems which outcomes are desired. This is difficult when we only partially know the goal, as is the case at the beginning of scientific research.

As part of this AI driven approach, The University of Manchester has also received a share of £4.4 million research funding from UKRI, in addition to contributions from the partners and the university totalling over £10 million.


**“ Professor Kaski has a world-leading reputation in AI and our collaboration with Aalto University strengthens our global reach to deliver new interdisciplinary research, innovation and impact for the benefit of the university, city and region.**

Professor Martin Schröder

With these investments the university is further strengthening its fundamental research in AI and will be launching the Manchester Centre for AI Fundamentals in 2022.

The full story is available via the QR code below.

To find out more about the Manchester Centre for AI Fundamentals, visit [manchester.ac.uk/fun-ai](http://manchester.ac.uk/fun-ai)



# Manchester Centre for AI Fundamentals


The University of Manchester's new Centre for AI Fundamentals is a key component of a number of significant recent investments The University of Manchester has made into AI education, innovation and industrial collaboration.

We boldly focus on fundamental AI research, which includes probabilistic modelling, deep learning, reinforcement learning, causal modelling, human-in-the-loop ML, explainable AI, ethics, privacy and security.

This centre brings together renowned academic expertise in AI with the latest research taking place across our growing institutes including The Institute for Data Science and AI (IDSAI), The Christabel Pankhurst Institute and our partnerships with The Alan Turing Institute and the European Laboratory for Learning and Intelligent Systems (ELLIS).

The University has already recruited a number of key staff to the new centre for AI Fundamentals, including lecturers and researchers, and there are a number of further vacancies currently being advertised. For further information on these, please use the QR code below.

To find out more about the Manchester Centre for AI Fundamentals, visit [manchester.ac.uk/fun-ai](https://manchester.ac.uk/fun-ai)



**“ We need new kinds of AI assistants which can learn to work well with humans and complement their skills. That requires new fundamental AI research, and Manchester has recognized this opportunity and is considerably strengthening its AI research. Manchester is a top-notch place to build and apply new AI which matters and has impact. ”**

Professor Samuel Kaski

# The Christabel Pankhurst Institute

For health technology research and innovation

In 2021, a consortium led by The University of Manchester launched the Christabel Pankhurst Institute for Health Technology Research and Innovation. This new multimillion pound institute is building on Manchester's academic strengths in digital health and advanced materials to discover innovative health and care solutions.

This institute is part of an ambitious plan set out in the Greater Manchester (GM) Local Industrial Strategy to boost the city-region's provision in this area.

The initiative will build on investments from the University, Manchester Science Partnerships (MSP), the Engineering and Physical Sciences Research Council (EPSRC), and The Alan Turing Institute, creating a total budget of more than £25m.


The institute will move into a flagship building at the centre of the University's campus on the Oxford Road Corridor, as well as having bespoke, state-of-the-art research and business development spaces at MSP's Citylabs campus. This location and partnership will provide support for business growth by facilitating better collaboration between the NHS, researchers and industry through MSP, MFT, Health Innovation Manchester and the University.

**Health and scientific innovation is needed now more than ever, so it is terrific news that we can support the launch of The Christabel Pankhurst Institute in Greater Manchester.**

The launch will see The University of Manchester continue to be a pioneer in digital health. It will come as a boost to the business sector by creating employment opportunities, at the same time as delivering further long-term health benefits to our city-region.

Andy Burnham,  
Mayor of Greater Manchester

To find out more, visit  
[pankhurst.manchester.ac.uk](https://pankhurst.manchester.ac.uk)



# Robotics and AI

Manchester's new centre dedicated to state-of-the-art AI technologies

The University's Centre for Robotics and AI pulls together experts and projects from across the academic disciplines who share the challenge of working on the front line of applied robotic technologies.

For example, Manchester researchers are looking to develop robotic systems that are able to explore in the most extreme environments, such as those found in the nuclear industry, power generation or agriculture. Other expertise includes designing robots to support digital manufacture or work in the field of medicine and health.

**“ Robotics is now an important field that can be found in research areas across the University's academic portfolio – which is not surprising, as robotic and autonomous systems are being applied in all parts of our lives.**

With the launch of this Manchester centre of excellence in robotics and AI we are providing a new focus to our multidisciplinary, world-class work in this field.


Professor Richard Curry

”

While driving developments in cutting-edge robotic systems, the centre will also have a commitment to ensure autonomous systems are compatible with the values and expectations of society. Some of the breakthrough Manchester-led research work will include:

- designing control systems with a focus on bio-inspired solutions to mechatronics
- designing trustworthy autonomous systems through the development of new software engineering and AI methodologies for verification
- researching human-robot interaction, with a pioneering focus on the use of brain-inspired approaches to robot control, learning and interaction
- research in ethics and human-centred robotics issues, for the understanding of the impact of robots and autonomous systems with individuals and society.

To find out more, visit:  
[robotics.manchester.ac.uk](https://robotics.manchester.ac.uk)



# European Laboratory for Learning and Intelligent Systems

The European Laboratory for Learning and Intelligent Systems (ELLIS) recently added The University of Manchester as a partner of its global members who strive towards a meaningful contribution to securing Europe's sovereignty and leadership in the research field of modern artificial intelligence (AI).

Four new international units have been announced including; Manchester, Jena and Stuttgart in Germany and Milan in Italy. The new units join a network of world-class institutions across 14 European countries and Israel.

The University of Manchester has recently strengthened its position as a centre for research into AI fundamentals and impactful applications of AI to improve health, security and sustainability. Last year the University appointed AI Chairs in each of its faculties followed by several excellent machine learning faculty appointments in the department of Computer Science.

The University of Manchester has been a partner of the Alan Turing Institute since 2018 and is home to 33 Turing Fellows.

It has a thriving community of data science and AI researchers, with over 900 researchers affiliated to its Institute for Data Science and AI (IDSAI). Manchester's ELLIS unit brings together experts in AI fundamentals

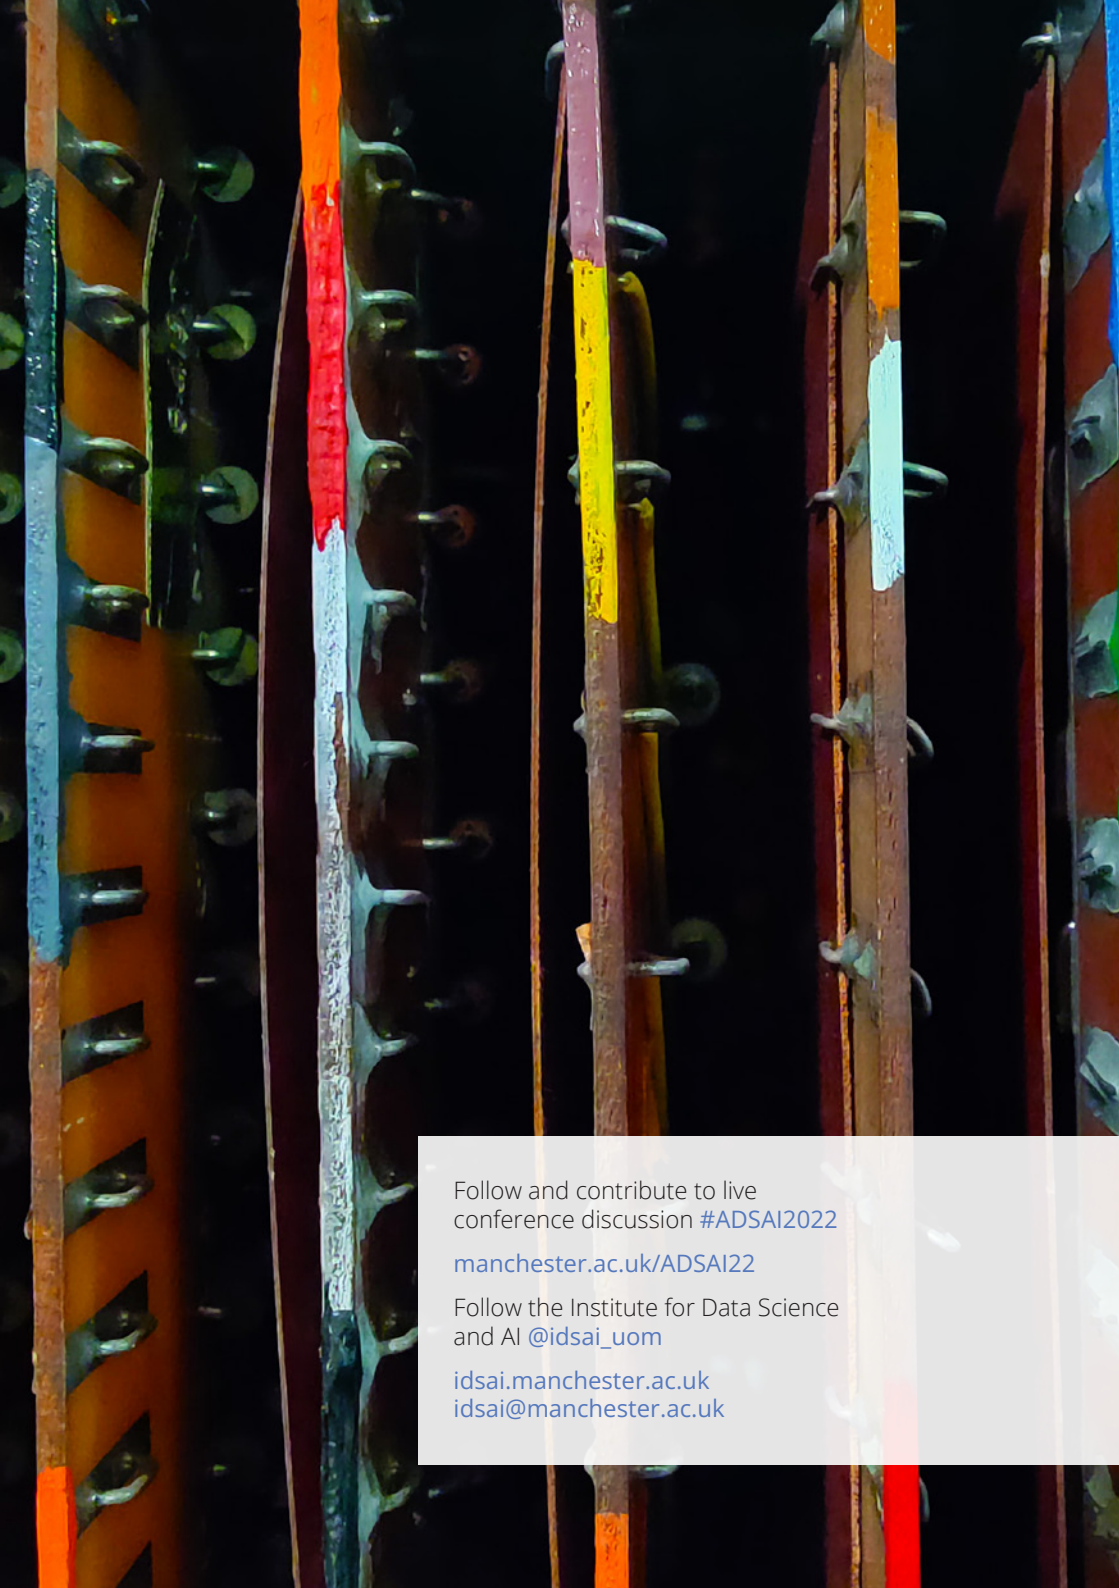
**“ The University of Manchester continues to grow as a centre of excellence for AI research and the new ELLIS unit will further strengthen this activity. Through the new ELLIS unit Manchester will be able to better link machine learning researchers across Europe with impactful applications across many disciplines.**

Professor Magnus Rattray

with experts in the application of AI in other fields, with particular strengths in health, and will connect with other leading experts in the ELLIS network across Europe.

The University of Manchester has established a strategic partnership in a shared professorship with the director of the ELLIS unit Helsinki, Samuel Kaski from Aalto University, Finland. This Northern link will be used in the future to set up the ELLIS units in Manchester and Helsinki as a twin unit, with tight collaboration already under way through research collaboration and exchange.

Find out more: [idsai.manchester.ac.uk/connect/partnerships/ellis](https://idsai.manchester.ac.uk/connect/partnerships/ellis)



Follow and contribute to live  
conference discussion [#ADSAI2022](#)

[manchester.ac.uk/ADSAI22](http://manchester.ac.uk/ADSAI22)


Follow the Institute for Data Science  
and AI [@idsai\\_uom](#)

[idsai.manchester.ac.uk](http://idsai.manchester.ac.uk)  
[idsai@manchester.ac.uk](mailto:idsai@manchester.ac.uk)

# Campus Map

- 1 Kilburn Building
- 2 Alan Turing Building
- 3 Christabel Pankhurst Institute
- 4 Henry Royce Institute
- 5 National Graphene Institute
- 6 Graphene Engineering Innovation Centre
- 7 Manchester Institute of Biotechnology

Find out more:  
[manchester.ac.uk/discover/maps](http://manchester.ac.uk/discover/maps)



Engineering Building A

