The Alan Turing Institute

2018 Doctoral Students **Yearbook**

Foreword

Since its foundation in 2015, one of the institutional goals of The Alan Turing Institute has been training new generations of data science and AI leaders. This yearbook is therefore not only a celebration of the individual success of our 2018 doctoral cohort, but also of their contribution to the success of the Turing as a whole.

The Turing has grown massively since 2018 and this cohort helped shape that growth, in particular by providing our team and the wider organisation with valuable input and feedback on how we support students and train the leaders of the future.

The breadth of research areas, vibrancy of ideas and sense of community forged amongst our 2018 doctoral students, and between them and the wider research community, also reflect the Turing's ethos of collaboration.

The challenges brought by the pandemic resulted in a very different PhD experience for this cohort. Whilst face to face collaboration and travel to conferences and workshops ceased for nearly 2 years, our 2018 doctoral students continued to collaborate and inspire with their research output. They produced over 50 research papers and developed work with industry and public bodies through internships at organisations including DeepMind, Snapchat, Spotify and the DECOVID project.

Since graduation members of the cohort have gone on to work as researchers and data scientists at leading companies, taken the next steps in their academic careers as post-docs and founded their own companies, using their skills to solve real world problems.

We have fond memories of supporting our 2018 doctoral students over the past four years; from meeting them at their interviews to competing in a quiz during student induction and keeping in touch through ongoing dialogue. The end of lockdown enabled us all to reconnect in person and we've loved seeing members of the cohort at the office.

We are proud of both your individual achievements and collective success. Thank you for being a delight to work with and great ambassadors for the Turing.

Georgia Koumara (Academic Services Officer) and Sam Selvarajah (Academic Services and Wellbeing Manager)

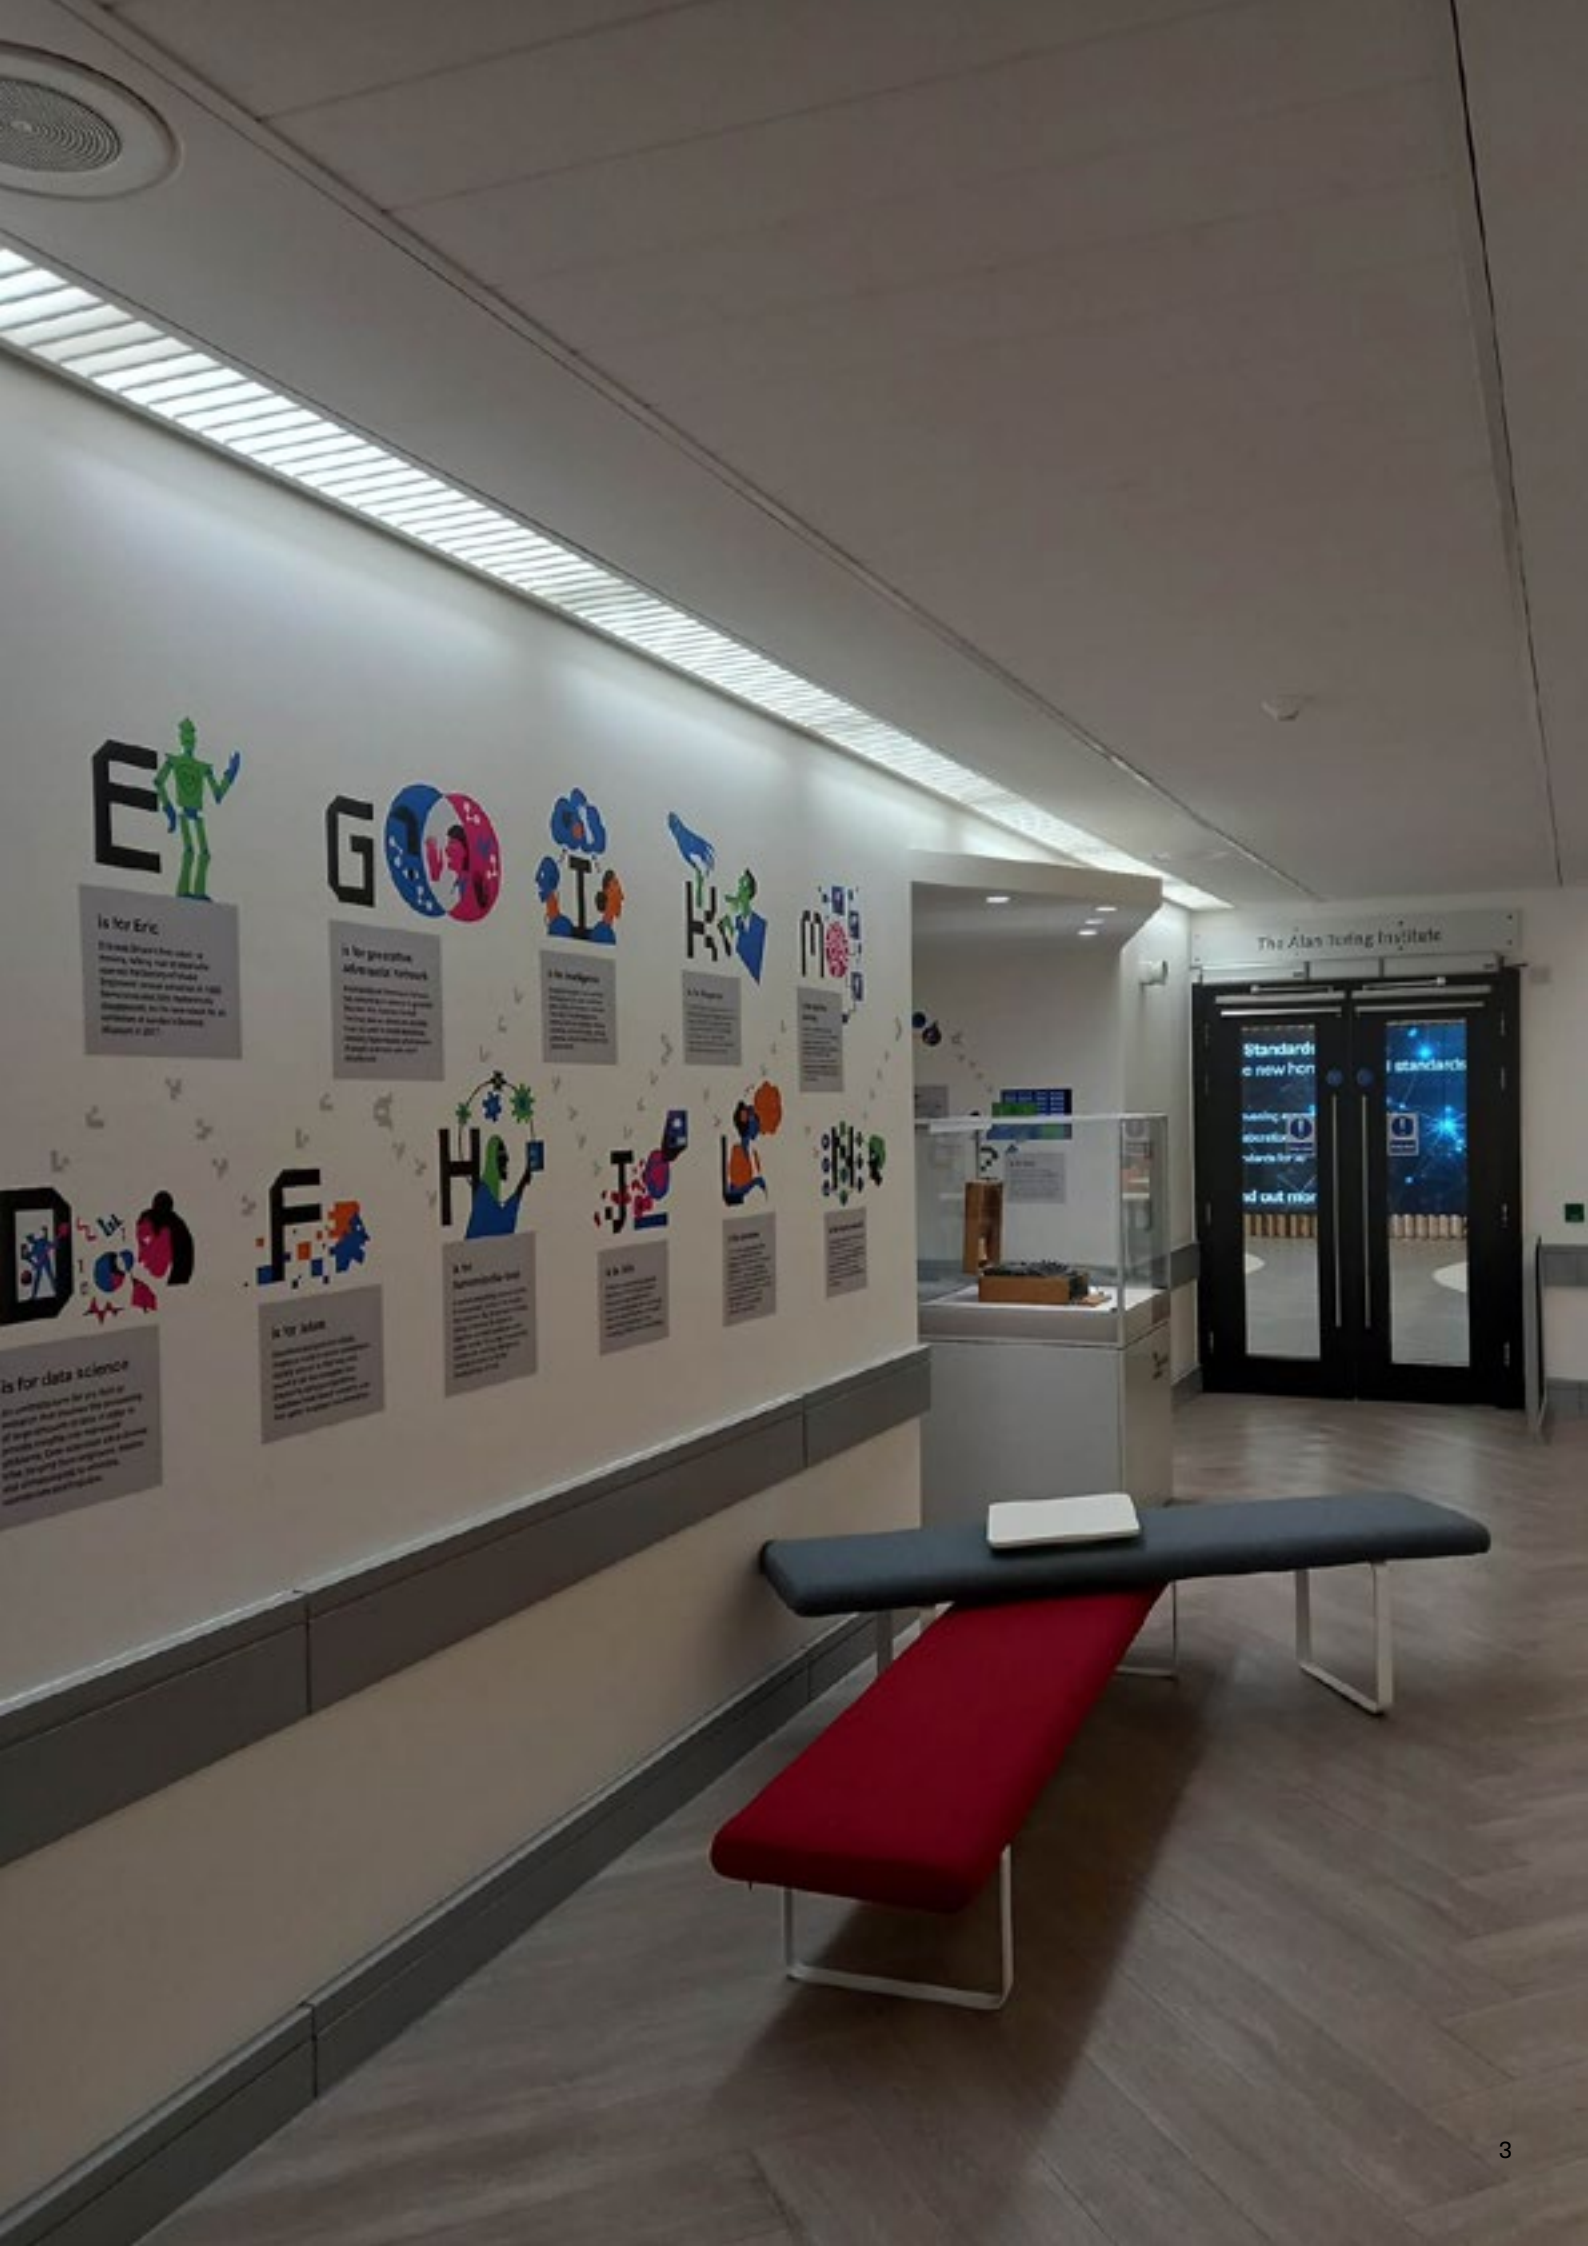
is for Eric
is for data science
The Alan Turing Institute

Students

The Alan Turing Institute is a truly unique place to complete a PhD. We had the opportunity to work with people from a broad range of academic backgrounds, ranging from neuroscience to computer science and even philosophy. This is so important given the interdisciplinary nature of data science and AI.

The networking benefits of the Turing extend far beyond academia. We also met colleagues from software companies, big and small, alongside policymakers and venture capitalists. They are all important stakeholders in the AI ecosystem, and the Turing did a great job of bringing them together to enrich our PhD experience.

The highlight of the program was definitely the cohort itself. Despite our diverse backgrounds, we bonded early over regular lunches and evening socials. These socials weren't limited to just pub trips and meals out - we also went rock climbing and even paid a visit to Bletchley Park. We are grateful to the Academic Services team given how involved they were to enrich our Turing experience.

The pandemic was a particularly difficult time, and sadly made it harder to maintain these bonds. We missed the vibrancy of the Turing office space which outclassed any of our previous faculties. Nevertheless, we finish the doctoral program extremely grateful for the opportunities the Turing gave us and the doors opened for our future.

Anujan, Ryan and Sam (2018 Doctoral Student Representatives)

Note: Due to a later submission date, the following students from the 2018 cohort will be included in a future yearbook: Ilan Price, Sedar Olmez, Christian Blex.

Andrew Manderson

Research areas

Probabilistic programming	Monte Carlo methods
High dimensional inference	Uncertainity quantification

Selected papers & workshops

Papers:

Manderson, A. A. and Goudie, R. J. B. (2022) A numerically stable algorithm for integrating Bayesian models using Markov melding. Statistics and Computing, 32, 24. DOI: 10.1007/s11222-022-10086-2.

Manderson, A. A. and Goudie, R. J. B. (2021) Combining chains of Bayesian models with Markov melding. Forthcoming in Bayesian Analysis. https://arxiv.org/abs/2111.11566

Workshops:

"Joining chains of models with Markov melding" -- presented at the RSS international conference (2021), Joint Statistical Meetings (2021), and at several symposiums/seminars around Cambridge.

Thesis title

Bayesian methodology for integrating multiple data sources and specifying priors from predictive information

Supervisors

Robert Goudie

Main outcomes of my research

In this thesis I developed methodology for combining multiple sources of information in Bayesian models, and specifying priors for such models given information about the observable quantity. Information sources can include multiple data sets, expert opinions, or Bayesian models. These disparate sources of data are combined with Markov melding, in order to fully quantify uncertainties in each information source. This is done by specifying separate models for each information source, which also alleviates some of the difficult in model specification.A crucial step in Bayesian model specification is that of choosing an appropriate prior. For complex models, it is difficult to select a prior encoding available knowledge of the phenomena of interest, whilst still permitting some variation from current expectations in future experiments. Domain experts often possess a good understanding of the scale and distribution of the data they are likely to observe. I developed methodology to translate such information into suitable priors for complex Bayesian models.

Impact of my work

The methodology developed in my thesis assists practitioners of Bayesian inference in combining and specifying appropriate Bayesian models. Complex Bayesian statistical models are the norm in many research areas. By building freely available software implementing my methodological developments, and demonstrating its use through cases studies in many domains, I hope to enable users of Bayesian statistics to build better, more understandable models.

Being a Turing Doctoral Student

I met a huge number of fascinating people in many fields, both close to my own and completely different. I worked in an incredible space inside the British Library. Generally I was exposed to the enormous array of applied data science projects being undertaken across the UK. The time and funding also enabled me to learn several important and universal skills, including an understanding of the effort, processes, and strategies required to produce important and high quality research; how to conduct research independently, and the challenges/benefits thereof; and how best to manage myself over a long time period (though three and a half years has passed unbelievably quickly).

Plans following my PhD completion

Looking for jobs in industry in and around London.

Anujan Poologaindran

Research areas

- Time series
- Neural networks
- Neuroscience
- Graph theory

Selected papers

Section I: Introduction to Connectomics

1. Poologaindran A, Lowe SR, Sughrue ME. The cortical organization of language: distilling human connectome insights for supratentorial neurosurgery. J Neurosurg. 2020;134(6):1959-1966.

2. Poologaindran A, Suckling J, Sughrue ME. Letter: elucidating the principles of brain network organization through neurosurgery. Neurosurgery. 2020;87(1):E80-E81.

3. Briggs RG, Allan PG, Poologaindran A, et al. The Frontal Aslant Tract and Supplementary Motor Area Syndrome: Towards a Connectomic Initiation Axis. Cancers. 2021.

4. Doyen S, Nicholas P, Poologaindran A, et al. Connectivity-based parcellation of normal and anatomically distorted human cerebral cortex. Hum Brain Mapp. 2022;43(4):1358-1369.

Section II: Structural Connectomics

5. Poologaindran A, et al. Longitudinal Connectome Analyses following Low-Grade Glioma Neurosurgery: Implications for Cognitive Rehabilitation. British Neuro-Oncology Society. July 2021. (online due to COVID19)

Thesis title

Elucidating the principles of brain network organisation through neurosurgery

Supervisor

John Suckling

Main outcomes of my research

I demonstrated how brain tumours spatially colocalize with brain networks governing higher-order cognition. Using network science and graph theory, I also demonstrated how we can predict the long-term cognitive outcomes in neurosurgical patients before they have even had surgery. Together, this work ushers in a new paradigm to advance pre-surgical brain mapping to avoid long-term cognitive deficits while offering new rehabilitation approaches to brain tumour patients in the immediate post-operative period.

Impact of my work

Each year over 5000 people die in the UK from cancerous brain tumours. While neurosurgery improves survival, patients have unpredictable post-operative trajectories including impairments to speech, memory, and attention that diminish quality of life. In order to mitigate these issues, I use brain mapping and network science to track how the brain reorganises after tumour removal. Through this process, I discovered that early rehabilitation of cognitive processes is critical for full recovery and that we can distinguish individuals requiring more intensive rehabilitation early in their post-operative care. To advance this discovery with industry collaboration, we are now piloting brain stimulation technology to accelerate how the brain reorganises following surgery; this strategy is applicable to many other conditions, including traumatic brain injury and epilepsy.

Being a Turing Doctoral Student

It allowed me to network and learn from a highly interdisciplinary network of experts, and to learn new skills in AI and machine learning that will be useful in my medical career. It also allowed me to build lifelong friendships with Turing staff and students.

Plans following my PhD completion

To finish my medical training in neurosurgery while continuing academic research part-time.

Selected papers (continued)

6. Poologaindran A, et al. Diffuse Infiltrating Gliomas and the Multiple Demand System for Domain-General Cognition: A Systems-Wide Analyses. medRxiv. (2022)

Section III: Functional Connectomics

7. Poologaindran A, et al. Serial Connectome Mapping of Domain-General Higher-Order Cognition in Cortical Neurosurgery. American Association for Neurological Surgeons. July 2021. (online due to COVID19)

8. Poologaindran A, et al. Gradient Mapping Stratifies Cognitive Trajectories in Glioma Neurosurgery. medRxiv. (2022)

9. Poologaindran A, et al. Presurgical Mapping of Mesoscale Cortical Dynamics to Preserve Higher-Order Cognition: Towards 'Functional' Glioma Surgery. medRxiv (2022).

Continued at turing.ac.uk/people/doctoral-students/anujan-poologaindran

Edward Bartrum

Research areas

Machine learning

Deep learning

Computer vision

Selected papers

Lifting AutoEncoders: Unsupervised Learning of a Fully-Disentangled 3D Morphable Model using Deep Non-Rigid Structure from Motion

Mihir Sahasrabudhe* Zhixin Shu* Edward Bartrum Rıza Alp Guler Dimitris Samaras Iasonas Kokkinos

Geometry Meets Deep Learning Workshop, ICCV 2019

1/11/2018 - 1/8/2019

StyleMorph: 3D aware deformation-based image synthesis

Eric Tuan-Le* Edward Bartrum* Iasonas Kokkinos

Currently under review for 3DV 2022

1/9/2021 - 1/7/2022

Thesis title

Implicit surface-based neural rendering

Supervisors

Iasonas Kokkinos

Main outcomes of my research

I have been researching the nascent field of Neural Rendering, replacing traditional 3D rendering pipelines with neural networks. I have worked mainly on the generative synthesis side, incorporating controllability and geometric priors into generative image models.

I have worked on training 3D models using weak supervision to learn 3D representations from 2D image data alone, and incorporating deformation fields into 3D-aware image synthesis (StyleMorph).

In addition to purely unsupervised generative models, I have worked on 3D lifting - training models to lift 3D representations from unstructured 2D data. In this setting I have also worked on disentangled image synthesis, disentangling multiple factors of variation (shading, albedo, pose, facial expression) from image synthesis (Lifting AutoEncoders).

I have also worked on few-shot novel view synthesis, training models to generate novel views of a scene conditioned on a single input view.

Impact of my work

The main applications of my work are in Augmented Reality. On the generative side, I have made incremental progress towards replacing traditional rendering with learned networks, with applications in 3D graphics. 3D aware image synthesis provides a rich source of 3D-annotated synthetic images, which further down the road could be exploited for learning downstream tasks.

The ultimate goal of my research on 3D lifting is “inverse rendering”, reversing the rendering process to obtain 3D scene representations from raw images. Allowing users to generate 3D controllable scenes from image data captured on a smartphone would open the door to new augmented reality experiences, narrowing the gap between the real and virtual worlds.

Being a Turing Doctoral Student

It allowed me to spend 4 years doing research in an area that I find fascinating. I had complete freedom to choose what I would work on, inspired by the amazing work in AI and Data Science that was happening around the Turing.

The Turing brings together a group of AI and Data Science world-class researchers. Despite this, it still manages to maintain a friendly and open culture. I was able to chat with people who work in completely different areas; their perspectives were so valuable as a PHD can be very isolating. I’m grateful that the Turing offered me that opportunity.

Plans following my PhD completion

I am starting an internship at Snapchat in September, where I am looking forward to working on real-world AR applications of my research. Unfortunately, I’m not at liberty to give any specifics about the project I’ll be working on!

I intend to continue working in the field of Neural Rendering, as it’s an incredibly exciting and fast moving field, and I feel that I have barely scratched the surface of it during my PhD.

Edwin Fong

Research areas

- Modelling (Statistical methods & theory)
- Uncertainty quantification
- Non-parametric & semi-parametric methods
- Machine learning

Selected papers

Fong, E., & Holmes, C. C. (2020). On the marginal likelihood and cross-validation. Biometrika, 107(2), 489-496.

Fong, E., Lyddon, S., & Holmes, C. (2019, May). Scalable nonparametric sampling from multimodal posteriors with the posterior bootstrap. In International Conference on Machine Learning (pp. 1952-1962). PMLR.

Fong, E., & Holmes, C. C. (2021). Conformal Bayesian Computation. Advances in Neural Information Processing Systems, 34.

Fong, E., Holmes, C., & Walker, S. G. (2022). Martingale posterior distributions. To appear in the Journal of the Royal Statistical Society: Series B as a Discussion Paper.

Fong, E., & Lehmann, B. (2022). A Predictive Approach to Bayesian Nonparametric Survival Analysis. In 25th International Conference on Artificial Intelligence and Statistics.

Thesis title

The predictive view of Bayesian inference

Supervisors

Chris Holmes

Main outcomes of my research

My thesis develops statistical methodologies which generalise the Bayesian update to tackle the challenges of computation and model misspecification. With a particular focus on the connections between prediction and inference, I introduced a framework for Bayesian inference requiring only the predictive distribution, which is termed the martingale posterior distribution. The method is scalable and nonparametric, and posterior sampling is achieved without the need for Markov chain Monte Carlo. My thesis also investigates the computational scalability of the Bayesian bootstrap and the connections between the marginal likelihood and cross-validation in Bayesian model selection.

Impact of my work

Bayesian inference in an elegant and holistic framework for statistical inference, as it combines prior information with new data in a coherent manner. However, there are stringent assumptions of the model being correctly specified which are unlikely to hold as datasets become more complex. Furthermore, Bayesian inference can struggle with computational scalability. The methodologies introduced in my thesis attempt to address some of these issues by being nonparametric and parallelizable. Furthermore, the connections between Bayesian prediction and inference sheds light on how the Bayesian method fundamentally deals with statistical uncertainty.

Being a Turing Doctoral Student

My time at the Turing allowed me to meet many talented people with diverse research interests, with whom I have had many inspiring conversations. Being part of the Turing doctoral student community also helped me through difficult times during my PhD, and I have made some good friends on the way.

For my research, the access to the Turing's fantastic computational resources has been invaluable, as many of my projects depended critically on parallel computing. The training and conferences that the Turing enabled me to attend also contributed greatly to my growth as a researcher. Finally, the PhD would not have been possible in the first place without the generous financial support of the Turing, for which I am grateful.

Plans following my PhD completion

I'm now a data scientist at Novo Nordisk, investigating the applications of causal inference in healthcare.

Ioana Bica

Research areas

- Time series
- Causality
- Uncertainty
- Neural networks
- Deep learning
- Machine learning

Selected papers

1. Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. "Estimating counterfactual treatment outcomes over time through adversarially balanced representations." in International Conference on Learning Representation (2020)

2. Ioana Bica, Ahmed Alaa, and Mihaela Van Der Schaar. "Time Series Deconfounder: Estimating treatment effects over time in the presence of hidden confounders." in International Conference on Machine Learning (2020)

3. Ioana Bica*, James Jordon*, and Mihaela van der Schaar. "Estimating the Effects of Continuous-valued Interventions using Generative Adversarial Networks." Advances in Neural Information Processing Systems (2020)

4. Daniel Jarrett*, Ioana Bica*, and Mihaela van der Schaar. "Strictly Batch Imitation Learning by Energy-based Distribution Matching." Advances in Neural Information Processing Systems (2020)

5. Jeroen Berrevoets, James Jordon, Ioana Bica, and Mihaela van der Schaar. "OrganITE: Optimal transplant donor organ offering using an individual treatment effect." Advances in Neural Information Processing Systems (2020)

Continued at turing.ac.uk/people/doctoral-students/ioana-bica

Thesis title

Causal inference methods for supporting, understanding, and improving decision-making

Supervisor

Mihaela van der Schaar

Main outcomes of my research

My research contributes to the area of causal inference by proposing machine learning methods that can be used for supporting, understanding, and improving decision-making, with a focus on the healthcare setting. To achieve this, I have developed causal inference methods capable of estimating the individualised effect of treatments from observational data for discrete treatments over time, but also for treatments with associated dosage in the static setting. Moreover, I have also worked on integrating counterfactual reasoning into batch inverse reinforcement learning to develop methods for better understanding the decision-making behaviour of experts, and on improving the robustness of decision-making by incorporating causal structure into learnt imitation policies. Overall, my PhD research resulted in 15 peer-reviewed publications and I was also involved in organising three workshops relevant to my work at top-tier machine learning conferences.

Impact of my work

My work introduces methodological advances in machine learning capable of reasoning about cause-and-effect relationships that can enable us to improve the delivery of personalized care for patients, support clinical decision-making and build a more transparent account of clinical practice. In particular, by being able to better estimate the causal effects of treatments, we can decide which intervention or sequence of interventions will result in the best outcome for each patient. In addition, methods for understanding expert preferences through counterfactual reasoning are important for finding the limitations of existing clinical practices, potential biases and variation in practice and overall, for building a more transparent account of clinical practice. Finally, by learning imitation policies that incorporate causal structure we can obtain more robust clinical decision guidelines.

Being a Turing Doctoral Student

Being a Turing student has allowed me to develop as a researcher, work on cutting-edge research problems and have access to amazing resources and support. To begin with, through the Turing I was able to obtain the guidance of leading experts in the field, expand my network, and collaborate with amazing researchers. Moreover, I am grateful that the Turing has provided me with crucial resources needed for my research, such as access to Microsoft Azure virtual machines which I used to perform experiments and thus validate and evaluate the methods developed throughout my doctoral studies. In addition, the travel support from the Institute enabled me to attend and present my work at top-tier conferences and learn more about state-of-the-art methods which has significantly contributed to the development of my own research.

Plans following my PhD completion

Following the completion of my PhD, I will join DeepMind as a Research Scientist.

Josh Cowls

Research areas

Ethics

Selected papers

J Cowls, P Darius, D Santistevan and M Schramm. 2022. Constitutional metaphors: Facebook's "supreme court" and the legitimation of platform governance. New Media & Society.

J Cowls, Andreas Tsamados, Mariarosaria Taddeo & Luciano Floridi. 2021. The AI gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations. AI & Society.

J Cowls, A Tsamados, M Taddeo, L Floridi. 2021. A definition, benchmark and database of AI for social good initiatives. Nature Machine Intelligence.

J Morley, J Cowls, M Taddeo, L Floridi. 2020. Ethical guidelines for COVID-19 tracing apps. Nature.

Ignacio Perez-Pozuelo, Dimitris Spathis, Jordan Gifford-Moore, Jessica Morley and Josh Cowls. 2020. Digital phenotyping and sensitive health data: Implications for data governance. Journal of the American Medical Informatics Association

J Cowls. 2019. The Silicon Cage: "Legitimate" governance 100 years after Weber. Presented

at Data Power 2019, University of Bremen, September 2019.

Thesis title

The political legitimacy of app store governance: The case of Parler

Supervisor

Luciano Floridi, Mariarosaria Taddeo

Main outcomes of my research

I use a case study of the removal of the fringe social media platform Parler from smartphone app stores in the aftermath of the riots at the US Capitol in January 2021 to identify the governance of app stores as an important aspect of internet governance. I find that Parler likely should have been removed sooner, owing to the extremist nature of the platform, and explore the reasons why it was retained for so long.

Impact of my work

My work contributes to a broader research agenda regarding how social media platforms should govern responsibly, identifies the human rights that should be upheld by online service providers, and provides recommendations for public policymakers to ensure that these responsibilities are met and rights safeguarded..

Being a Turing Doctoral Student

It gave me the opportunity to take a research project from start to finish, giving me the time and space to identify a topic of interest, engage with existing literature, and make theoretical and empirical contributions to the socially significant sphere of internet governance.

Karel Devriendt

Research areas

- Social networks
- Combinatorics
- Computing networks
- Applied mathematics
- Graph theory
- Dynamical systems and differential equations

Selected workshops

As an organiser of the Networks Seminar at Oxford, I invited Mateo Neira (another Turing DPhil student) to give a presentation on his work. This was in January 2020 and he presented his work on "Generative models and representational learning on street networks".

Otherwise, none of my papers or talks were facilitated directly by being at the Turing, but of course by being on a Turing studentship all my scientific output during my DPhil was indirectly supported by the Turing.

Thesis title

Graph geometry from effective resistances

Supervisors

Renaud Lambiote

Main outcomes of my research

My research focuses on developing new mathematical tools to study networks and complex systems. In practice, these can be social networks, transportation networks or even neuronal networks in the brain, which all rely on a shared set of models and concepts. During my DPhil, I looked specifically at how we can use distances between the individual parts of a network when analyzing it. We developed a method to measure the variance of a distribution defined on a network by making use of these distances, and as a more theoretical contribution, we used one particular distance - the effective resistance - to derive a new notion of "discrete curvature" on networks.

Impact of my work

Our measure of network (co)variance addresses a methodological gap in the analysis of networks and associated data. This has found applications in scientometrics, social networks and economics and we expect that this simple and intuitive measure will further develop to be an important practical tool. With our new approach to discrete curvature on the other hand, we contribute a new perspective to the research on discrete curvatures with some promising unique features, such as an associated Ricci flow (a way of rewiring a network based on its distances) and some new theoretical results for positively curved graphs (this is a special type of networks with interesting properties).

Being a Turing Doctoral Student

Through The Alan Turing Institute I got in contact with people I otherwise would never have met, such as students from universities all over the UK, and people outside of academia who found me through my Turing affiliation. These encounters really broadened my DPhil experience in a unique way. Thanks to the Turing DPhil studentship I was furthermore able to focus on my research and had a lot of extra freedom.

Plans following my PhD completion

I have just started a postdoc at the Max Planck Institute for Mathematics in the Sciences in Leipzig. I will continue working on some projects that started during my DPhil and broaden into some related topics.

Mateo Neira

Research areas

- Social data science
- Social networks
- Machine learning
- Applications (machine learning)

Selected papers & workshops

I co-organised the event: "Creativity and design in AI " along with Stephen Law.

I co-authored a paper titled: "An unsupervised approach to geographical knowledge discovery using street level and street network images" with Stephen Law, which was presented in GeoAI workshop of SIGSPATIAL.

I presented a paper titled "Urban Segregation on Multiplex Networks: a random walk approach" in the conference on complex systems 2019 in Singapore.

Thesis title

Modelling spatial segregation: A multiscalar approach to study segregation and its relation to transport network structure

Supervisors

Elsa Arcaute, Stephen Marshall

Main outcomes of my research

A theoretical framework and associated measures for analysing socio-spatial segregation and its relationship to transport networks across multiple scales using concepts from network theory and information theory. The framework and measures allow for policy-relevant analysis of the impact of transport infrastructure on socio-spatial segregation that avoids the issues of identifying appropriate spatial scales.
The framework was applied to study social-spatial segregation in Ecuador to demonstrate the insights we can gain with the proposed methodology.

Impact of my work

The methods and techniques developed as part of my work can be used to assess the dynamical aspects of segregation by taking into account transport network connectivity, and the results can be used to inform urban policy and planning.

In particular, as part of the research null-model that helps disentangle the role of spatial distribution of population and amenities and network connectivity, and a method to analyze how these interact at different scales were developed. The research provides new insights into the complex dynamics of segregation..

Being a Turing Doctoral Student

It helped me to develop as a researcher and learn from world leading academics in the area of network science and urban analytics. The training and support I received as a Turing Doctoral Student has been invaluable. During my time here I gained a deeper knowledge on the subjects I was interested in and developed a solid understanding of the theoretical part of my research thanks to the supportive and vibrant environment it provided.

Being part of a group of exceptional peers and working with outstanding and supportive supervisors and staff helped me grow as a researcher and as a person. During my time at the Turing I had the opportunity to collaborate with other researchers and develop meaningful professional and personal relationships.

Plans following my PhD completion

I am currently working as an associate data scientist with the urban design group of Foster and Partners. Within the team, I work on implementing novel frameworks and methods to understand cities and urbanisation to inform design strategies.

Naya Yerolemou

Research areas

- Geometry & topology
- Neuroscience
- Mathematical modelling
- Graph theory

Selected papers

(Pre-print) Morse theory for complexes of groups. Naya Yerolemou & Vidit Nanda

(The Australasian journal of combinatorics) Filtered matchings and simplicial complexes. Daniele Celoria and Naya Yerolemou

(The electronic journal of combinatorics) A discrete Morse perspective on knot projections and a generalised clock theorem. Daniele Celoria and Naya Yerolemou

(Symmetry) A topological selection of folding pathways from native states of knotted proteins. Agnese Barbensi, Naya Yerolemou, Oliver Vipond, Barbara Mahler, Pawel Dabrowski-Tumanski, Dimos Goundaroulis

(Proceedings: Research in Computational Topology 2) Graph pseudometrics from a topological point of view. Ana Lucia Garcia-Pulido, Kathryn Hess, Jane Tan, Katherine Turner, Bei Wang, and Naya Yerolemou

Thesis title

Topics in applied topology: Actions, knots and proteins

Supervisor

Vidit Nanda, Ulrike Tillmann

Main outcomes of my research

The development of a version of discrete Morse theory compatible with complexes of groups; the simultaneous generalisation of theorems from knot theory and combinatorics; and, the development of a quantitative way to systematically compare knotted proteins via a global shape statistic summarising their entanglement.

Impact of my work

The development of tools in pure and applied topology that have potential. For example, my work on complexes of groups could be used to perform computations on well-behaved compressed representations of symmetric spaces.

Being a Turing Doctoral Student

It allowed me to do a PhD, to work with some amazing researchers at Oxford and beyond, to travel to conferences and workshops, and to meet some great students from different research areas.

Plans following my PhD completion

Quant researcher at a hedge fund.

Réka Ágnes Kovács

Selected research areas

- Complexity (Algorithms)
- Combinatorics
- Convex programming
- Optimisation
- Operations research
- Applied mathematics

Selected papers & workshops

In May 2019 I attended the IPCO 2019 conference and presented a poster on Low Rank Binary matrix Factorisation.

In January 2020 I attended the Aussois Combinatorial Optimisation Workshop and I gave 5-min pitch talk on my plans for the PhD.

In June 2020 I attended online IPCO 2020.

In January 2021 I presented online a conference paper titled Binary Matrix Factorisation via Column Generation at AAAI2021, which is co-authored by my supervisors Raphael Hauser and Oktay Gunluk.

In June 2021 I attended online IPCO 2021.

In September 2021 I gave an online talk at INFORMS 2021 on Binary Matrix Factorisation via Column Generation.

In May 2022 I presented a conference paper titled On Minimally Non-Firm Binary Matrices at ISCO 2022 (no co-authors).

Thesis title

Exact and rank-k binary matrix factorisation under Boolean arithmetic

Supervisors

Raphael Hauser, Oktay Günlük

Main outcomes of my research

My research focuses on two problems in binary matrix factorisation, an applied and a theoretical one. Rank-k binary matrix factorisation (the applied problem) can be used to compress and complete datasets that contain responses to yes-or-no questions. I look at this problem through the methods of integer programming and developed an optimisation-based algorithm (AAAI 2021).

The second, theoretical part of my thesis is related to a problem in exact binary matrix factorisation. I try to understand the structure of firm binary matrices (a set of binary matrices that satisfy some nice properties) via constructing minimal counterexamples. I use techniques from graph theory and combinatorics for this and constructed four infinite classes of such counterexamples (ISCO 2022).

Impact of my work

My optimisation-based algorithm can be used to get accurate low-rank factorisations and completions of binary matrices and by this identifying discrete patterns in binary data. I hope that it may also help to more robustly benchmark heuristics for rank-k binary matrix factorisation.

My theoretical work starts the investigation into minimally non-firm binary matrices. I hope that it may encourage others to look at this problem and understand better the property of firmness.

Being a Turing Doctoral Student

Being a Turing Doctoral Student has been an amazing experience. It allowed me to meet many very talented, hard-working, and kind people from all over the world. It also offered me the unique opportunity to be able to work both from my home university, Oxford and the London Turing office, which I really appreciated.

I must also mention the generous financial support and the great range of training opportunities the Turing offered. Furthermore, the academic administration team has always been very kind and supportive even through the difficult times of COVID.

Plans following my PhD completion

I would like to go into industry and do an applied maths job.

Ryan Chan

Selected research areas

- Probabilistic programming
- Differential privacy
- Probability
- Monte Carlo methods

Selected papers & workshops

(All work was done with my supervisors Murray Pollock and Gareth Roberts unless otherwise stated)

Papers:

Divide-and-Conquer Monte Carlo Fusion (submitted and on arXiv: 2110.07265)

Divide-and-Conquer Generalised Bayesian Fusion (currently writing - hopefully submit by end of July 2022)

Conferences presented at:

O'Bayes 2019 (June 2019)

Greek Stochastics 2019 (August 2019)

The Turing Research Showcase 2020 (September 2020)

2021 World Meeting of the International Society for Bayesian Analysis (June 2021)

Bayes at Cirm 2021 (October 2021)

2022 World Meeting of the International Society for Bayesian Analysis (June 2022)

Thesis title

Monte Carlo methods for the unification of distributed statistical analyses

Supervisors

Murray Pollock, Gareth Roberts, Petros Dellaportas

Main outcomes of my research

Combining statistical analyses is a common challenge within statistics and data science. For instance, in distributed ‘big data’ problems, or when working under multi-party privacy constraints. During my PhD, I developed exact Monte Carlo methods for tackling this problem. The resulting methodology can be shown to be robust in a number of practical settings.

Impact of my work

Prior to my PhD, existing methodologies for tacking the problem of combining distributed statistical analyses were typically approximate methods which can perform poorly in many settings. The so-called 'Fusion' approaches were unique since they offered an exact Monte Carlo approach to this problem, but unfortunately these algorithms were not practical in several realistic scenarios. Over the course of my PhD, I have further developed these Fusion approaches which are now much more robust.

Being a Turing Doctoral Student

It allowed me to be a well-balanced data science researcher since we were able to engage with people with diverse backgrounds and research interests. Consquently, this made it very easy for me to be aware of what other researchers from different academic areas were studying.

Plans following my PhD completion

I will be a Research Software Engineering at the Turing within the Research Engineering (REG) Team.

Sam Miller

Selected research areas

Social data science	Applied (Machine learning)

Selected papers

Sam Miller, Helen Susannah Moat, and Tobias Preis. Using aircraft location data to estimate current economic activity. Scientific Reports, May 2020

Sam Miller, Abeer El-Bahrawy, Martin Dittus, Mark Graham, and Joss Wright. Predicting drug demand with Wikipedia views: Evidence from darknet markets. Proceedings of The Web Conference, April 2020

Sam Miller, Tobias Preis, Giovanni Mizzi, Leonardo Soares Bastos, Marcelo Ferreira

da Costa Gomes, Flavio Code co Coelho, Claudia Torres Code co, and Helen Su-sannah Moat. Faster indicators of chikungunya incidence using Google searches.

PLOS Neglected Tropical Diseases, June 2022

Thesis title

Faster socieconomic indicators using novel data sources

Supervisors

Suzy Moat, Tobias Preis

Main outcomes of my research

We show improved speed and accuracy for a range of socio-economic indicators such as GDP, illicit drug consumption, and the incidence of chikungunya.

Impact of my work

Chapters from the thesis were published in top venues like Nature Scientific Reports, PLOS Neglected Tropical Diseases and The Web Conference. Our results attracted significant interest from policymakers including the Office for National Statistics, the Department for International Trade, and the Brazilian Ministry of Health.

Being a Turing Doctoral Student

Being a Turing Doctoral student allowed me to broaden my horizons beyond my background in economics. I published work with a range of collaborators from computer science to social anthropology. I learned more about the tech sector, particularly how AI can be applied in a commercial setting. Turing's access to high-performance computing resources massively sped up my research.

Plans following my PhD completion

I've been involved with a start-up commercialising AI in the fitness sector - I'm keen to keep going with this.

Victor-Alexandru Darvariu

Selected research areas

- Algorithms
- Deep learning
- Artificial intelligence
- Multi-agent systems
- Reinforcement learning

Selected papers

Victor-Alexandru Darvariu, Stephen Hailes, and Mirco Musolesi (2021). Planning Spatial Networks with Monte Carlo Tree Search. Submitted for publication. Pre-print available on arXiv:2106.06768.

Victor-Alexandru Darvariu, Stephen Hailes, and Mirco Musolesi (2021). Solving Graph-based Public Good Games with Tree Search and Imitation Learning. In Proceedings of 35th Conference on Neural Information Processing Systems (NeurIPS 2021).

Victor-Alexandru Darvariu, Stephen Hailes, and Mirco Musolesi (2021). Goal-directed graph construction using reinforcement learning. Proceedings of the Royal Society A. 477(2553):20210168.

Thesis title

Learning to optimise networked systems (provisional)

Supervisors

Mirco Musolesi, Stephen Hailes

Main outcomes of my research

My thesis work is about the optimisation of networks using artificial intelligence techniques. An intuitive example is optimising the resilience of a transportation network such as the tube: how should we extend it so that it is as robust as possible to failures? Such problems are out of reach for classic mathematical optimisation methods but can be formulated and approached using the decision-making formalisms of AI and the generalisation capabilities of deep learning. The main outcomes of this research are algorithms for combinatorial optimisation problems as well as case studies that examine, empirically, particular classes of applications.

Impact of my work

The work has applications in transportation, infrastructure, and computer networks. It has the potential to translate to more efficient systems that consume less resources to achieve a given goal. Such methods are very early in the technological adoption cycle – research is only emerging – but, in my opinion, they have potential for transformative impact. I am hoping to pursue applications of my work to real-world problems in the future.

Being a Turing Doctoral Student

It has allowed me to interact with people that I may not have met otherwise in the typical insular academic department. This has definitely broadened my view of data science and AI research and has enabled me to kickstart collaborations and forge friendships. I also was able to attend interesting seminars (the Multi-Agent Systems series comes to mind) by being part of the Turing ecosystem. Finally, given the generous financial support I was able to focus fully on my research work – doing a PhD in London would have been impossible for me otherwise.

Plans following my PhD completion

My plan is to continue doing research in this area. I have not yet decided between academia and industry – actually, I am hoping I will never have to choose as I enjoy both conceptual and applied work.

turing.ac.uk
@turinginst