Mathematical and Computational Foundations of Artificial Intelligence

Link to funding opportunity

Summary of Successful Outline Proposals (Ordered by grant number)

Grant No.	PI	Title
EP/Y007166/1	Professor Ruth Misener	FoundOpt: Foundations of Optimisation for AI
EP/Y007174/1	Professor Paul Fearnhead	Probabilistic AI Hub
EP/Y007484/1	Professor Michael Bronstein	Mathematical Foundations of Intelligence: An "Erlangen Programme" for AI
EP/Y007514/1	Dr Sidharth Jaggi	Information theory for distributed AI
EP/Y007603/1	Professor Richard Samworth	MSAI: Mathematical and Statistical underpinnings of AI
EP/Y007646/1	Professor Mihaela van der Schaar	Mathematics of Stable and Trustworthy AI Computation (MOSAIC)
EP/Y007727/1	Professor Michael Wooldridge	Decentralized AI: UKRI Hub in Computational & Mathematical Foundations of AI
EP/Y007735/1	Professor Aldo Faisal	AI Hub in Foundational Reinforcement Learning (RLHub)

Grant Reference Number: EP/Y007166/1

Title: FoundOpt: Foundations of Optimisation for AI

Team:

- Professor Ruth Misener, Imperial College London (PI)
- Dr Alain Zemkoho, University of Southampton
- Dr Anastasia Borovykh, Imperial College London
- Dr Calvin Tsay, Imperial College London
- Dr Dario Paccagnan, Imperial College London
- Dr Gah-Yi Ban, Imperial College London
- Dr Martin Haugh, Imperial College London
- Dr Michael Osborne, University of Oxford
- Dr Panos Parpas, Imperial College London
- Dr Sarah Filippi, Imperial College London
- Dr Selin Ahipasaoglu, University of Southampton
- Professor Paul Goulart, University of Oxford
- Professor Wolfram Wiesemann, Imperial College London

Contact: Ruth Misener – r.misener@imperial.ac.uk

Brief description of proposed work:

Computational optimisation is a foundational AI tool that practitioners need every time they train a machine learning model. Optimisation is also commonly used to test vulnerabilities in AI models through verification and to integrate the outputs of uncertainty quantification into optimal decision-making strategies.

But the current state-of-the-art in optimisation is not fit for purpose. There are many issues that motivate us to rethink the very foundations of optimisation theory, computation, and software. For example, in optimisation, one of the first things we teach students is that an optimisation algorithm has converged, or the minimal solution is found, when the algorithm has found a point on a surface that is at the bottom of a valley. In other words, if we're trying to minimise a function, we seek a point where small movements in any direction would yield a worse solution. But Zhang et al. (ICML 2022) find something unexpected: neural network weights do not necessarily converge to stationary points of the loss function. In other words, the very first thing we teach in optimisation modules is inconsistent with a typical outcome of training neural networks!

There are many other examples of issues going to the very core of computational optimisation:

* In black-box optimisation, or optimisation where only expensive function evaluations are available, the settings typically studied by researchers do not align with the needs of AI practitioners.

* Constraints designed to reduce algorithmic bias sometimes cut the optimisation space in a way that reduces our ability to explain the outcome.

* Typical practice in machine learning overparameterises neural networks: researchers including Belkin et al. (PNAS 2020) show how this practice improves the observed behaviour of trained neural networks. But adding the rules imposed by fairness principles may mean that this overparameterisation makes the subsequent decision-making process much more difficult.

We also see possibilities for rethinking optimisation as it applies to causality, scalable uncertainty quantification, and robustness/verifiability analysis in AI.

Optimisation, particularly as a sub-field of operational research, has always focussed on finding the best value(s) for objective function(s). But, with next-generation AI technologies, we need to move away from the simply-defined goals of faster computational methods and algorithms. Instead, we seek more tailored methods that actually address the needs of next-generation AI technology.

FoundOpt aims to develop a new generation of optimisation theories, algorithms, and open-source software that can address these pressing needs of AI:

* First, we have assembled a top team of computer scientists and mathematicians with expertise in operations research and statistical machine learning. The research of this transformative team, which spans from control theory to causality, is uniquely suited to deliver an exciting new suit of approaches getting to the core of the needs in modern AI.

* Second, we have identified several underpinning technologies that need to be reimagined by this team: these include hierarchical optimisation, Bayesian optimisation, mixed-integer optimisation, uncertainty quantification for optimisation, mixed-integer optimisation, and optimisation under uncertainty.
* Third, we have found many instances where the current optimisation technology is insufficient, for example applications to AI causality, interpretability, algorithm bias, and verifiability.

Our research plan centres on three strands: (i) developing the new algorithms needed for modern AI use cases, (ii) proving their performance through extensive computational tests, and (iii) developing, releasing, and maintaining open-source software giving access to our developments. We will also engage extensively with companies, charities, policy makers, and other key stakeholders across the UK.

Grant Reference Number: EP/Y007174/1

Title: Probabilistic AI Hub

Team:

- Professor Paul Fearnhead, Lancaster University (PI)
- Dr Aretha Teckentrup, University of Edinburgh
- Dr Richard Turner, University of Cambridge
- Professor Anthony Lee, University of Bristol
- Professor Catherine Powell, The University of Manchester
- Professor Gareth Roberts, University of Warwick

Contact: Paul Fearnhead – p.fearnhead@lancaster.ac.uk

Brief description of proposed work:

Probabilistic AI holds the potential to create the next generation of AI tools which can overcome many of the current shortcomings within AI, such as poor uncertainty quantification, overly confident predictions, and uninterpretable decision-making. At its core, probabilistic AI provides a framework for embedding probability models, probabilistic reasoning, and measures of uncertainty within AI methodology. The importance of utilising probabilistic reasoning to produce optimal decisions is well understood and has been shown theoretically. However, whilst much of the current AI toolbox is motivated by probability models, shortcuts in modelling and inference are often taken to deal with the computational challenges of data size and complexity.
The ability to fit probabilistic AI models fully, so as to appropriately quantify uncertainty, would be ground-breaking. For example, creating new AI tools with robust inferences and decision-making; a new ability to include physics and system knowledge, with data, to improve accuracy; being able to propagate uncertainty as we fuse AI models; and the ability to infer causal links and better predict the impact of policy changes. These are all important features of AI which will be impactful across diverse areas such as healthcare, climate science and weather prediction, drug discovery and material design.
Developing practical, general-purpose probabilistic AI methods requires overcoming substantial challenges, and at their heart many of these challenges are mathematical. This hub will bring together UK researchers across the breadth of Applied Mathematics, Computer Science, Probability and Statistics, with a range of non-academic and international partners, to solve these underlying challenges. This will lead to a new generation of mathematically-rigorous, scalable and uncertainty-aware AI algorithms, which can be applied and relied upon to benefit UK society within future technologies, such as AI-assisted drug discovery, autonomous vehicles and smart energy grids.
The hub will become a national resource for research and innovation in probabilistic AI and cement the UK's world-leading position in this area. The hub will support a range of activities open to the wider UK research community which are designed to facilitate engagement between mathematical sciences researchers and AI applications in science and industry. The hub will produce a culture change within the mathematical sciences in the UK, where cross-disciplinary mathematics research forms a core component of AI, rather than a peripheral extension to AI developments.

Grant Reference Number: EP/Y007484/1

Title: Mathematical Foundations of Intelligence: An "Erlangen Programme" for AI

Team:

- Professor Michael Bronstein, University of Oxford (PI)
- Dr Anthea Monod, Imperial College London
- Dr Nina Otter, Queen Mary University of London
- Dr Primoz Skraba, Queen Mary University of London
- Dr Yue Ren, Durham University
- Professor Alessandro Abate, University of Oxford
- Professor Heather Harrington, University of Oxford
- Professor Jacek Brodzki, University of Southampton
- Professor Jared Tanner, University of Oxford
- Professor Jeffrey Giansiracusa, Durham University
- Professor Ran Levi, University of Aberdeen

Brief description of proposed work:

In 1872, Felix Klein published his now famous Erlangen Programme, in which he treated geometry as the study of invariants, formalised using group theory. This radically new approach allowed tying together different types of non-Euclidean geometries that had emerged in the nineteenth century and has had a profound methodological and cultural impact on geometry in particular and mathematics in general. New fields of mathematics such as exterior calculus, algebraic topology, the theory of fibre bundles and sheaves, and category theory emerged as a continuation of Klein's blueprint. The Erlangen Programme was also fundamental for the development of physics in the first half of the twentieth century, with Noether's theorem and the notion of gauge invariance successfully providing a unification framework for electromagnetic, weak, and strong interactions, culminating in the Standard Model in the 1970s.

Now is the time for an "Erlangen Programme" for AI, based on rigorous mathematical principles that would bring better understanding of existing AI methods as well as a new generation of methods that have guaranteed expressive and generalisation power, better interpretability, scalability, and data- and computational-efficiency. Just as the ideas of Klein's Erlangen Programme spilled into other disciplines and produced new theories in mathematics, physics, and beyond, we will draw inspiration from these analogies in our AI research programme. By resorting to powerful tools from the mathematical and algorithmic fields sometimes considered "exotic" in applied domains, new theoretical insights and computational models can be derived.

Our "Erlangen Programme of AI" will study four fundamental questions that underlie modern AI/ML systems, striving to provide rigorous mathematical answers. How can hidden structures in data be discovered and expressed in the language of geometry and topology in order to be exploited by ML models? Can we use geometric and topological tools to characterise ML models in order to understand when and how they work and fail? How can we guarantee learning to benefit from these structures, and use these insights to develop better, more efficient, and safer new models? Finally, how can we use such models in future AI systems that make decisions potentially affecting billions of people?

With a centre at Oxford, and broad geographic coverage of the UK, the Hub will bring together leading experts in mathematical, algorithmic, and computational fields underpinning AI/ML systems as well as their applications in scientific and industrial settings. Some of the Hub participants have a track record of previous successful work together, while other collaborations are new.

The research programme in the proposed Hub is intended to break barriers between different fields and bring a diverse and geographically-distributed cohort of leading UK experts rarely seen together with the purpose of strong cross-fertilisation. In the fields of AI/ML, our work will contribute to the exploitation of tools from currently underexplored mathematical fields. Conversely, our programme will help attract the attention of theoreticians to new problems and applications.

Grant Reference Number: EP/Y007514/1

Title: Information theory for distributed AI

Team:

- Dr Sidharth Jaggi, University of Bristol (PI)
- Dr Amanda Prorok, University of Cambridge
- Dr Ayalvadi Ganesh, University of Bristol
- Professor Deniz Gunduz, Imperial College London
- Professor Ioannis Kontoyiannis, University of Cambridge
- Professor Jonathan Lawry, University of Bristol

Contact: Jonathan Lawry – J.Lawry@bristol.ac.uk

Brief description of proposed work:

Artificial intelligence (AI) is on the verge of widespread deployment in ways that will impact our everyday lives. It might do so in the form of self-driving cars or only in the form of navigation systems optimising routes on the basis of real-time traffic information. It might do so through smart homes, in which usage of high-power devices is timed intelligently based on real-time forecasts of renewable generation on the grid. It might do so by automatically coordinating emergency vehicles in the event of a major incident, natural or man-made. It might do so by coordinating swarms of small robots collectively engaged in some task, such as search-and-rescue. Much of the research on AI to date has focused on optimising the performance of a single agent carrying out a single well-specified task. There has been little work so far on emergent properties of systems in which large numbers of such agents are deployed, and the resulting interactions. Such interactions could end up disturbing the environments for which the agents have been optimised. For instance, if a large number of self-driving cars simultaneously choose the same route based on real-time information, it could overload roads on that route. If a large number of smart homes simultaneously switch devices on in response to an increase in wind energy generation, it could destabilise the power grid. If a large number of stock-trading algorithmic agents respond similarly to new information, it could destabilise financial markets. Thus, the emergent effects of interactions between autonomous artificial agents inevitably modify their operating environment raise significant concerns about the predictability and robustness of critical infrastructure networks. At the same time, they raise the prospect of optimising distributed AI systems to take advantage of cooperation, information sharing, collective learning and decision making.

The key future challenge is therefore to design distributed systems of interacting AIs that can exploit synergies in collective behaviour, while being resilient to unwanted emergent effects. Biological evolution has addressed many such challenges, with social insects such as ants and bees being an example of highly complex and well-adapted responses emerging at the colony level from the actions of very simple individual agents!

The goal of this project is to develop the mathematical foundations for understanding and exploiting the emergent features of complex systems composed of relatively simple agents. While there has already been considerable research on such problems, the novelty of this project is in the use of information theory to study fundamental mathematical limits on learning and optimisation in such systems. Information theory is a branch of mathematics that is ideally suited to address such questions. Insights from this study will be used to inform the development of new algorithms for artificial agents operating in environments composed of large numbers of interacting agents.

The project will bring together mathematicians working in information theory, network science and complex systems with engineers and computer scientists working on machine learning, AI and robotics. The aim is to translate theoretical insights into algorithms that are deployed on real systems; lessons learned from deploying and testing the algorithms in interacting systems will be used to refine models and algorithms in a virtuous circle.

Grant Reference Number: EP/Y007603/1

Title: MSAI: Mathematical and Statistical underpinnings of AI

Team:

- Professor Richard Samworth, University of Cambridge (PI)
- Dr Henry Reeve, University of Bristol
- Dr Timothy Cannings, University of Edinburgh
- Professor Ata Kaban, University of Birmingham
- Professor Carola-Bibiane Schönlieb, University of Cambridge
- Professor Ramji Venkataramanan, University of Cambridge

Brief description of proposed work:

We are at the early stage of the AI revolution. Scientific fields such as precision medicine, genetics and robotics stand on the verge of making discoveries that would have seemed unimaginable a decade ago. AI-driven technologies such as ChatGPT and driverless cars promise to have greater and greater impacts on our lives. At the core of these advances are algorithms, and our increasing dependence on their automated responses makes it vital that these methodologies are robust, fair and efficient. Our vision is to create an ambitious, large-scale, interactive hub, combining broad skills from across the mathematical, statistical, machine learning and information engineering spectrum to tackle foundational questions that underpin AI.

Our research will be grouped around three interconnected themes:

Robust statistical methodology for AI: Modern data sets arise in many different forms, and their underlying generating mechanisms can be extremely complex. There may be missing data, or corruption, or we may observe data from a different distribution from our target population. Recent breakthroughs in handling heterogeneous missingness and transfer learning suggest adaptive approaches for these new structures, while being robust to a wide range of departures from standard modelling assumptions; new techniques such as conformal inference allow uncertainty quantification in these general settings. Often, we want an interpretable model to explain a phenomenon; we may also seek to understand causal relationships, even from observational data.

Explainable machine learning (ML) with guarantees: ML has witnessed spectacular practical successes demonstrating what is possible; moreover, its interplay with statistics and computational learning theory yields insight into what is guaranteed, what is impossible, and which data and problem characteristics make one problem instance easier or harder than another. Our next goal is to open the black box and provide explainable machine decisions. Recent results in multi-output learning, model compression and heterogeneous ensembles pave the way towards general-purpose mathematical representations that can break down complex, nonlinear structures into intuitively appealing (including visual) thought processes with confidence guarantees.

Information theory and AI: Algorithms such as message passing play a key role in understanding the trade-offs between statistical accuracy and computational efficiency. Moreover, information-theoretic concepts and lower bounds provide crucial insight into the fundamental limits of AI techniques, and motivate new procedures, e.g. information directed sampling. The long-term goal is to answer questions such as why large language models trained with deep learning algorithms sometimes work so well and sometimes fail, and to predict their generalisation properties when applied outside their intended conditions. We will characterise the performance of modern AI algorithms on different test classes, including adversarial high-dimensional data.

Grant Reference Number: EP/Y007646/1

Title: Mathematics of Stable and Trustworthy AI Computation (MOSAIC)

Team:

- Professor Mihaela van der Schaar, University of Cambridge (PI)
- Dr Ajitha Rajan, University of Edinburgh
- Dr Hana Chockler, King's College London
- Dr Sheehan Olver, Imperial College London
- Professor Anders Hansen, University of Cambridge
- Professor Desmond Higham, University of Edinburgh
- Professor Eleni Vasilaki, University of Sheffield
- Professor Ivan Tyukin, King's College London
- Professor Nicholas Higham, The University of Manchester

Brief description of proposed work:

Artificial Intelligence (AI) is revolutionising society and impacting many areas of science. AI is capable of exceeding human capabilities in a range of well-defined tasks. However, like humans, AI systems have limitations. In particular, they can be extremely sensitive to the extent that, for example, a seemingly imperceptible change to an image of a chair can cause the image to be misclassified as a toaster. They can be unreliable, for example, a reconstruction of an MRI scan can include spurious features that appeared in previous, unrelated scans. They can be uninterpretable, for example an "explanation" of why the image is classified as a toaster may be opaque or misleading.

For these reasons, in order to make the most of the opportunities offered by modern AI technologies, we need to improve our understanding of their capabilities and limitations, and develop approaches that, where possible, improve on stability and trustworthiness.

The UK Government white paper "Establishing a pro-innovation approach to regulating AI", from July 2022, emphasizes the national need for better understanding by specifying cross-sectoral principles for AI regulation that include

Ensure that AI is used safely,
Ensure that AI is technically secure and functions as designed, and
Make sure that AI is appropriately transparent and explainable.

The basic tenet of the MOSAIC Hub is that significant progress in these directions can only be made by bringing together an appropriate and committed team of mathematical scientists and computer scientists, working alongside a range of non-academic collaborative partners who can implement, validate and provide feedback on the results that arise. The underlying challenges that we will address range from high-level, big picture issues such as "under what conditions is it inevitable that a class of AI tools will be vulnerable to targeted attack" to very practical issues of "can we design and test an algorithm that detects and guards against specific type of attack"?

So far, advancements in this area have been somewhat piecemeal and ad hoc, resulting in a conflict escalation between attack and defence strategies with no end in sight. The MOSAIC team has the breadth and depth to provide rigorous, concrete answers, and hence to determine, and where possible extend, the methodological boundaries inside which modern AI techniques provide the adequate robustness and trustworthiness. In addition to making Ai systems more reliable, results in this area will shape political and legal decisions for regulations and hence have a huge impact on the market for AI technologies.

MOSAIC will be led by a core team of nine researchers from leading universities across the UK. By employing and working with PhD-educated research assistants, MOSIAC will help to train the next generation of game-changers in this extremely important area. The hub will work in collaboration with a range of external partners from across a range of sectors. These partners use AI to make decisions and drive policies, and they are committed to providing research challenges, giving critical feedback and making immediate use of MOSAIC's outputs. The hub will also provide a leadership and awareness-raising role across the UK's mathematical and computer science communities. We will offer secondments to Early Career Researchers, train 10 new PhD students, and deliver outward-facing workshops that will be open to all scientists.

In particular, in order to maximize the value of the grant, the team will build on their current track record of leadership and leverage their network of connections with institutes and laboratories (including The Alan Turing Institute, The Isaac Newton Institute, The International Centre for Mathematical Sciences, The Edinburgh

Parallel Computing Centre, All Party Parliamentary Group on AI Data Governance) to ensure the optimal reach and impact.

Grant Reference Number: EP/Y007727/1

Title: Decentralized AI: UKRI Hub in Computational & Mathematical Foundations of AI

Team:

- Professor Michael Wooldridge, University of Oxford (PI)
- Dr Galit Ashkenazi-Golan, London School of Economics and Political Science
- Professor Graham Cormode, University of Warwick
- Professor Lukasz Szpruch, University of Edinburgh
- Professor Rahul Savani, University of Liverpool

Brief description of proposed work:

The history of computing has been marked by a steady transition from monolithic, isolated systems to decentralized, networked systems. At the time of writing, however, most contemporary AI systems (including IBM's Watson, DeepMind's AlphaGo, and OpenAI's GPT-3) are inherently isolated, typically interacting only with human users. Over the decades ahead, this will change: future AI systems will interact not just with humans but with other AI systems, and with many other conventional computing systems besides. History tells us this decentralized AI future is inevitable. Of course, this presents enormous opportunities - but realizing it efficiently, safely, and responsibly raises a raft of scientific questions at the intersection of mathematics, statistics, and theoretical computer science. Such questions include, for example:

* How should decentralized systems evaluate the trustworthiness of the information that they are presented with and make decisions in uncertain environments?

* How can systems of autonomous and semi-autonomous entities interact with humans safely and robustly?

* How can vast quantities of heterogenous information be combined to make informed choices?

* How can we provide rigorous, provable guarantees on the overall output of decentralized AI and, for example, prevent these systems from tacitly colluding?

To be able to answer these questions, we need to develop theoretical foundations for the design of safe and scalable algorithms & protocols with provable guarantees for controlling complex distributed systems comprising (semi-)autonomous entities with their own goals in data-rich environments. Current AI systems are brittle with respect to a variety of objectives (e.g., accuracy, safety). These issues will inevitably be amplified in decentralized AI systems. To provide mathematical guarantees on the overall outputs of decentralized AI, and hence enable accountability, we need to understand how uncertainty and robustness guarantees propagate through AI pipelines (potentially involving multiple AI systems using multiple AI technologies) to determine what overall certifiable robustness guarantees can be obtained.

To address these questions, our hub is structured around four key research themes:

Theme 1: Decentralized Data: This theme will develop novel algorithms for decentralized training with statistical privacy guarantees, generalizing initial efforts in federated learning, and drawing on new developments in multi-party computation, differential privacy, and zero-knowledge proofs to ensure validity without compromising privacy.

Theme 2: Decentralized Learning: This theme will develop the theory for decentralised learning, aiming to put this area on a rigorous mathematical basis with provable computational properties.

Theme 3: Decentralized Decision-Making: In this theme we will study decentralized AI through the lens of game theory and mechanism design, for example studying distributed, adversarial, and strategic robustness for distributed systems.

Theme 4: Decentralized Verification & Reasoning: Contemporary AI systems use a range of (symbolic) AI techniques in addition to ML, including search, constraint satisfaction/optimisation, planning, and symbolic reasoning (verification/model checking, synthesis/planning, theorem proving).

Tackling these and many other questions will require meaningful co-creation between experts in AI, theoretical computer science, mathematics, optimization, game theory, and related topics: the team is ideally placed for this, bringing together researchers from Oxford, Edinburgh, Liverpool, LSE, and Warwick who have expertise in core AI/ML, mathematics, optimisation, and game theory.

The hub will be open to new participants, with open calls for new research projects in Y3-4 and Y4-5.

Grant Reference Number: EP/Y007735/1

Title: AI Hub in Foundational Reinforcement Learning (RLHub)

Team:

- Professor Aldo Faisal, Imperial College London (PI)
- Dr Baharak Rastegari, University of Southampton
- Dr Ciara Pike Burke, Imperial College London
- Professor Giovanni Montana, University of Warwick
- Professor Robert Piechocki, University of Bristol
- Professor Samuel Kaski, The University of Manchester
- Professor Sebastian Stein, University of Southampton

Contact: Aldo Faisal – RLhub@imperial.ac.uk

Brief description of proposed work:

We propose establishing an Artificial Intelligence (AI) Hub in the mathematical, statistical and computational foundations underpinning Reinforcement Learning (Foundational RL Hub). Reinforcement Learning (RL) is one of the pillars of the AI revolution and an essential component for realising the full potential of AI to transform industry and society both in the UK and further afield.. Within AI, RL is the approach by which computers can learn how to act or interact through experience, e.g. to control robots, drive self-driving vehicles, or find better ways to treat patients. The great advantage of RL methods is that they do not require to be told at every step what to do; instead, they just need to receive feedback in good time if they are doing well or not, and the algorithms will work out how to improve themselves. Multiple estimates suggest that the value of implementing RL throughout the global economy is £5 trillion, or five-fold the combined value of supervised learning. To extract value from RL in real-world settings and leverage its potential for transforming UK services and industry, these algorithms must be built on solid foundations and extended to a level that makes them deployable at scale. To this end, it is essential to connect the underpinning mathematics, statistical and operations research-based theories to deepen our understanding and capability to develop RL algorithms. These foundational connections will allow us to create new AI RL methodologies that fulfil the requirements of safety, trust, explainability, fairness and privacy-respecting AI technology that regulators and users of AI will demand. Our vision is to co-create with the UK academic community a hub of world-leading standing that will drive the UK's RL innovation capability & capacity as a source of ideas, solutions, researchers, resources, and education for RL. Our goal is to recruit present and future foundational researchers to engage with the enormous potential and value that RL can realise and to transform RL from a field currently driven by clever hacks into a foundational field where principled understanding allows us to exploit deep connections to deliver better, faster, safer and society-wide more desirable algorithms.

The Foundational RL Hub will be driven by co-creation between our AI & RL communities and our four foundational communities (mathematics, statistics, operations research & computation). The hub's leadership team comprises both AI and foundational research leaders and a diversity of geographic sites, protected characteristics, and career stages. Our engagement and partnership strategy aims to mobilise the foundational communities to work on problems in RL and to get the AI communities to embrace the many advantages of deeper theoretical understanding and its more powerful methods bring about. To this end, the Foundational RL Hub will offer more than £5.5Million in open competition funding for postdoctoral researchers and resources for creating bespoke research and impact events. To make the Foundational RL Hub accessible to researchers across the UK, we will set up three physical sites in the North, West and South of the country that offer a seminar room, co-working and hotdesking spaces that will enable a new level of scientific cooperation on top of best-practice hybrid and online approaches. Hub projects connecting foundational and AI research will jump-start by making all the hub's PostDocs follow a paired supervision model requiring two supervisors from different domains (e.g. AI and Mathematics). The 35 PostDoc-years supported by the hub will be complemented by funding attracted through proactive partnership building. To further foster the hub's aim to become the national centre for foundations of RL, we planned over £1Mio in flexible funds to co-create symposia, and hackathons and support growth of foundational RL throughout the UK academic, industry and stakeholder communities.