



UNIVERSITY OF
BIRMINGHAM

IT Services



BEAR Case Studies - Vol 5

We advance
We activate

birmingham.ac.uk/bear

125
years

birmingham.ac.uk

Copyright & Acknowledgements

Copyright © The University of Birmingham 2025
Published in the United Kingdom by The University of Birmingham,
Edgbaston, Birmingham, B15 2TT, UK.

All rights reserved. Except as otherwise permitted under the Copyright Designs and Patents Act 1988, no part of the work may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher. Inquiries concerning reproduction outside these terms should be sent to the compiler at the following address:

Aslam Ghumra,
The University of Birmingham,
IT Services, Building G5
Edgbaston,
Birmingham, B15 2TT, UK.
email: bearinfo@contacts.bham.ac.uk

Acknowledgments

We would like to thank everybody who contributed to this publication.
To see these and newer case studies please visit
<https://blog.bham.ac.uk/bear/category/case-study/>

Cover Photo

The left hand side depicts the BlueBEAR cluster from 2025 and the right hand side is the IBM 3090-200S system of 1983

URL : <http://epapers.bham.ac.uk/id/eprint/4428/>
DOI : <https://doi.org/10.25500/epapers.bham.00004428>

Foreword

Professor Andrew Morris
Chair of Research Computing
Management Committee



This is the fifth edition of the BlueBEAR case study comprising 24 new projects. Once again the diversity of the projects is amazing: from quantum computing to tooth enamel; modelling FinTech markets to RNA sequencing.

Computing facilities, like BlueBEAR, Baskerville and Sulis do not just appear at Universities: to keep results like these studies, there's a lot happening behind the scenes. There is a dedicated team of Research Infrastructure and Software Engineers, outreach teams and academics making it happen. Their job includes keeping the hardware running, but also installing and upgrading software; training and coaching users how to get the best out of the systems; and attracting new users who would benefit from it. In the background there's also discussion with suppliers, like NVIDIA and Lenovo; future planning for expansion of the BlueBEAR and the data-centres themselves; along with analysis of what up-and-coming users of BlueBEAR will need in 1, 2, 5 years time.

The Research Computing Management Committee that I chair, (RCMC for short) is one of the ways that academics can help with the running of the facilities, including providing those future insights and making the argument for continued BlueBEAR funding.

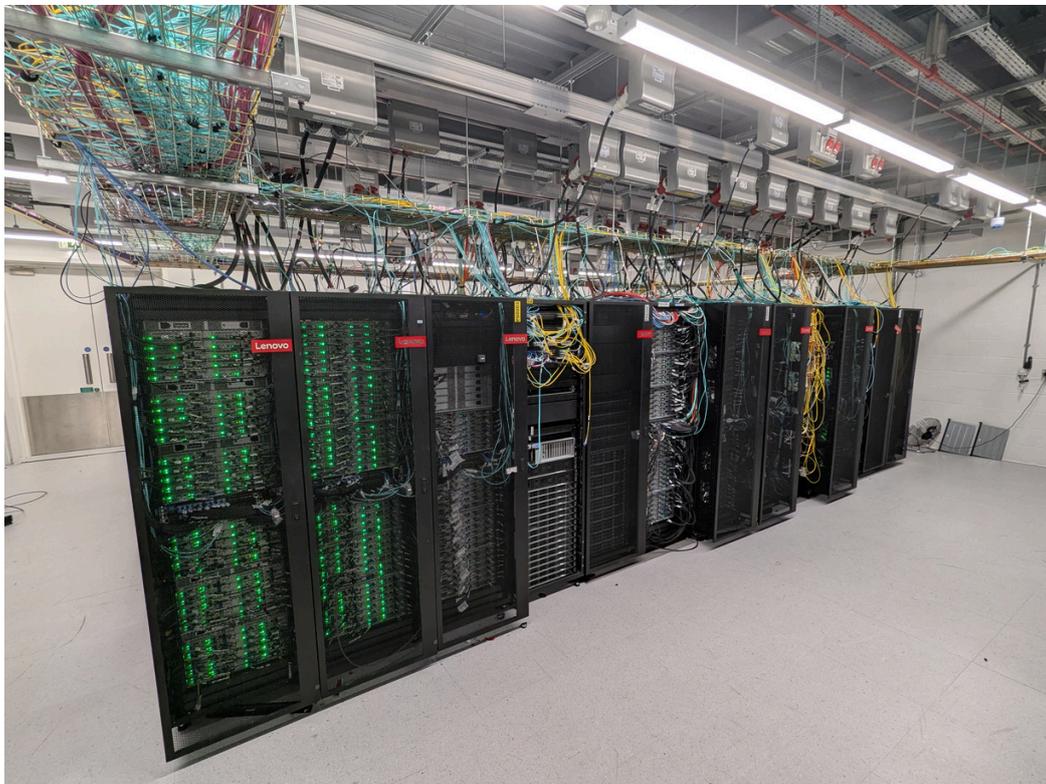
At many institutions free access to these sorts of resources is no longer a given. It is increasingly important that we highlight the excellent value-for-money that BlueBEAR provides in terms of the wealth of high-impact outputs we produce. Along with acknowledgements in publications and conference talks, an important way to do this is by publishing case studies demonstrating the breadth and quality of the work we carry out. These studies provide the details so that we can all boast, to anyone who might listen, about of your and BlueBEAR's success

Abstract

This collection celebrates the extraordinary breadth and ambition of research powered by the University of Birmingham's Environment for Academic Research (BEAR), which is led by the Advanced Research Computing team within IT Services. Across every College, researchers are accelerating discoveries, tackling complex challenges, and expanding the boundaries of knowledge thanks to BEAR's cutting-edge digital infrastructure.

From high-performance computing on BlueBEAR to secure research data storage and specialist software, BEAR empowers staff and students to work at a scale and speed that transforms what is possible. These case studies showcase not only the diversity of research undertaken at the University, but also the unique value that BEAR brings in enabling innovative, impactful, and world-leading work.

Together, they highlight how BEAR continues to be a driving force behind research excellence and an essential catalyst for future breakthroughs.



BlueBEAR Computer Cluster - 2025

Contents

Life Sciences & Health	Page
• Radiation transport simulations	• 10
• Human tooth enamel mineralisation	• 12
• RNA sequencing analysis	• 15
• Infant gut and malnutrition	• 21
• SMLM Data Exploration with Nano-org	• 26
• Telomere lengthening (ALT)	• 28
Data, AI & Methods	Page
• 10,12,14 to infinity	• 7
• Quantum chemical modelling	• 8
• Multimodal Analysis in Cognitive Translation Research	• 13
• Extraction and Analysis of Hydrological Spatial Data	• 14
• Extracting tender information with NLP	• 16
• Understand investor behaviour in FinTech markets	• 20
• High-Resolution Soil Hydrology Analysis	• 22
• Bayesian modelling of decision thresholds	• 25
• Energy Data Landscapes	• 27
Materials	Page
• Modelling multiphase flows for sustainable manufacturing	• 9
• Fluid dynamics in Scroll Expanders	• 11
• Shaking up things: Optimising mixing	• 17
• What makes things go BOOM!	• 18
• Modelling defects in inorganic phosphor crystals	• 19
• Bio-Inspired FSI and HPC-driven Aerodynamic Optimisation	• 23
• Atomic Simulation to Build Better Batteries	• 24
• Modelling accelerator-based neutron sources	• 29

Contributors

Name	Email	Page
Anil Taskin	axt126@student.bham.ac.uk	11
Anubrata Das	a.das.1@bham.ac.uk	28
Asim Abbas	axa2233@student.bham.ac.uk	16
Chenlu Yang	cxy348@student.bham.ac.uk	14
Dylan Owen	d.owen@bham.ac.uk	26
Edward Barbour	e.barbour@bham.ac.uk	27
Emily Allman	ela090@alumni.bham.ac.uk	17
Fahad Al-Gburi	fsa956@student.bham.ac.uk	9
Faysal Khaleel	fak917@student.bham.ac.uk	9
Grant Wilson	i.a.g.wilson@bham.ac.uk	27
Guilin Luo	gxl355@student.bham.ac.uk	22
Hibah Saddal	hxs619@student.bham.ac.uk	23
Jakob Hein-Paar	jmh421@student.bham.ac.uk	18
Jonathon Riddell	jonathon.riddell@kotharcomputing.com	7
Jonny Lewis	j.lewis.2@bham.ac.uk	15
Kiran Nutter	kjn718@student.bham.ac.uk	10

Contributors

Name	Email	Page
Leon Mijic	lxm697@student.bham.ac.uk	21
Louis Butt	ljb841@student.bham.ac.uk	29
Mahdi Tavakol	mahdi.tavakol@eng.ox.ac.uk	12
Maryia Shymanovich	mxs1756@student.bham.ac.uk	19
Mingjie Zhang	mxz140@student.bham.ac.uk	13
Mohammad B M A J Aldoub	mxs1542@student.bham.ac.uk	20
Oskar Soulas	ogs353@student.bham.ac.uk	24
Peter Russell	ppr466@student.bham.ac.uk	8
Ronan Broderick	r.a.broderick@bham.ac.uk	28
Sandeep Shirgill	s.shirgill@bham.ac.uk	26
Sophie Wetz	d.g.heinke@bham.ac.uk	25

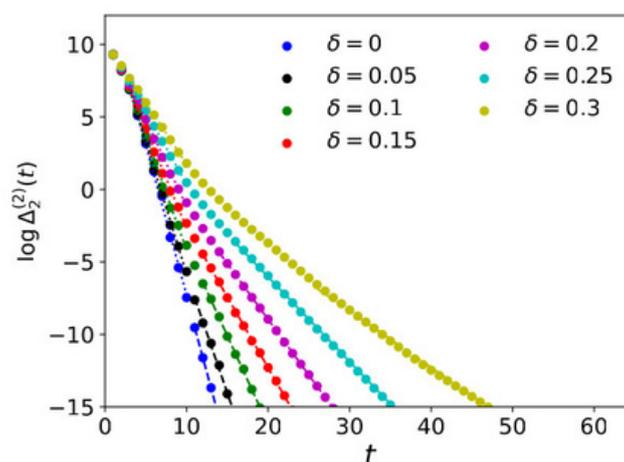
10,12,14 to infinity: quantum research is classically hard by Jonathon Riddell



Contexts for this research range from models of magnetism, metals or even simulating quantum circuits (those things that the quantum computers should end up running, but for now, we use classical computers). The key aspect that makes these problems computationally very challenging is really in the memory demands of the problems. The typical degree of freedom we study is the “qubit”. For every qubit we insert into the numerics, we have to either double or quadruple the memory of the program, depending on what task we are doing.

With memory scaling so poorly, having high memory (512 GB) nodes has been key in simulating the physics we care about. Once you are manipulating data as large as 100s of GB, you also need fast compute to properly update the memory as you proceed in an algorithm. To this end we are typically reserving 200-400 GB of memory along with 64+ cores, which are readily available on BlueBEAR.

A recent project we are quite excited about is creating efficient algorithms to prepare “random states” on quantum computers. Think of this as being equivalent to preparing random numbers on a classical computer, except, the community still isn’t quite sure how to efficiently do this. Here we’ve had to study the spectrum of a non-hermitian matrix for 32 qubits, as well as do dynamics. You can see this in the figure below. The parameter δ controls how quickly we prepare the random state on the quantum computer (data generated on BlueBEAR). Δ here quantifies how far we are from the statistics we are interested in.



Quantum chemical modelling of solids for optoelectronic applications by Peter Russell



Peter's research is all about exploring the weird and wonderful world of solid-state materials — the kind that could one day power your phone, your car, or even your solar panels. His group is on a mission to model these materials using computational and theoretical methods, with a strong focus on green energy applications. Think of it as chemistry meets cutting-edge tech, with a splash of sustainability.

Peter's current playground? Optoelectronics — materials that are both transparent and conductive. These are the unsung heroes behind solar panels, smartphones, aircraft windows, and more. Pretty cool, right? But here's the twist: modelling solid materials isn't exactly a walk in the park. Crystals, for example, are theoretically infinite (yes, infinite!). To make this manageable, scientists simplify the problem by zooming in on a single repeating unit — kind of like looking at one tile in a never-ending mosaic (Figure 1).

Even this tiny unit can be packed with atoms, and atoms come with electrons — lots of them. Peter's simulations often juggle between 150 and 300 electrons, each behaving like a tiny diva with its own quirks. Since electrons basically run the show in chemistry, modelling their behaviour is key.

To tackle this many-particle puzzle, Peter uses density functional theory — a well-established method that makes the problem solvable. But even with clever maths, the computational demands are sky-high. That's where BlueBEAR, the university's high-performance computing system, swoops in like a superhero. Thanks to BlueBEAR, Peter can crunch the numbers, simulate the materials, and push the boundaries of what's possible in optoelectronic research — all without frying his laptop.

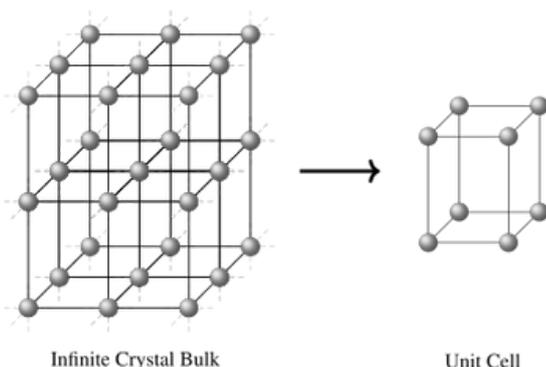


Figure 1: Simplifying an infinite crystal to its repeating unit

Modelling multiphase flows for sustainable manufacturing and thermal energy applications

by Fahad Al-Gburi and Faysal Khaleel

We are PhD researchers in the Schools of Engineering and Chemical Engineering at the University of Birmingham, working to make advanced materials manufacturing and electronics cooling more sustainable and efficient.

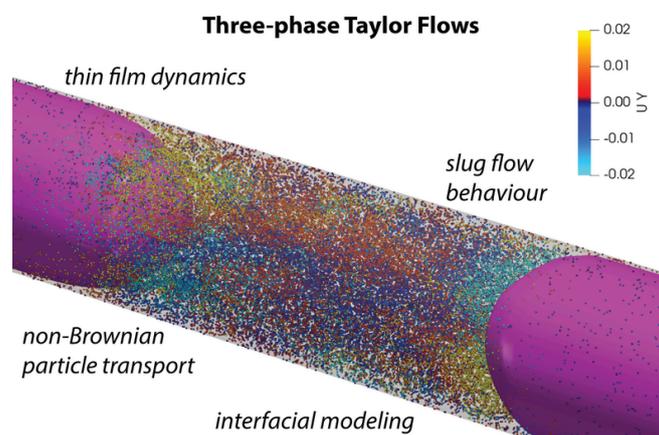


Our work focuses on complex fluid systems involving solids, liquids, and gases, which are key to creating modern materials. By understanding how heat, mass, and flow behave in these systems, we can reduce waste, save energy, and design materials with tailored properties.

To achieve this, we develop advanced computer models that simulate these processes across scales—from nanometres to metres. We've recently published research on three-phase flows and are now building an open-source solver within the OpenFOAM framework to predict multiphase flow and heat transfer at the nanoscale. These simulations require enormous computing power, which is made possible by the University's BlueBEAR high-performance computing (HPC) facilities.

Another part of our research tackles the growing challenge of cooling high-density electronics, such as those in data centres. By introducing innovative cooling methods like nanofluids and phase-change materials, we aim to improve energy and water efficiency—helping HPC systems become more sustainable.

Access to BlueBEAR has also enabled us to create detailed visualisations for public engagement on issues like clean air and reducing exposure to harmful particles, work that has been featured in major media outlets.

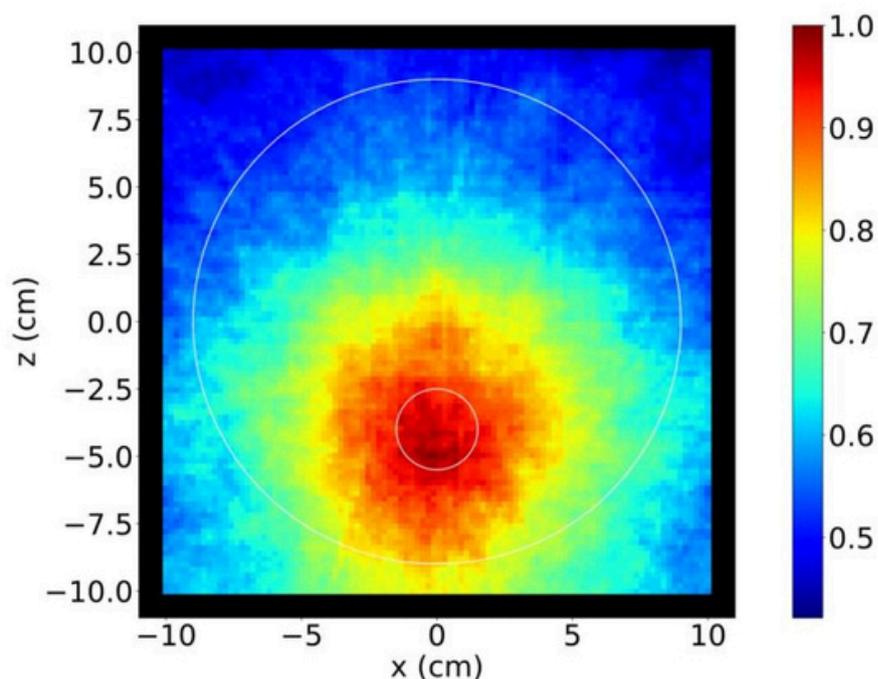


Utilising BlueBEAR for radiation transport simulations and computational fluid dynamics by Kiran Nutter



Boron Neutron Capture Therapy (BNCT) is a form of cancer therapy that utilises both biological and physical targeting of the tumour cells and can achieve high precision dose delivery. However, obtaining an accurate measure of the dose that is delivered to the patient during treatment is particularly challenging. One possible solution to this problem is to image the gamma rays that are produced during the neutron capture reaction on boron and use this to infer the dose delivered.

I proposed a design for an imaging camera that could be used for this purpose, and BlueBEAR provided me with the capacity to run detailed particle tracking simulations, with Geant4, to evaluate the feasibility of this camera design. The results from this study were very promising and were published in 2024; the image shown is presented as the main result in this paper, which was generated after a simulation running on BlueBEAR for approximately 30 days <https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2024.1347929/full>



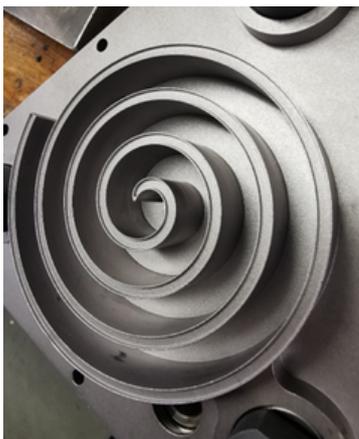
Fluid dynamics in Scroll Expanders and Static-Shaft Wankel Expanders by Anil Taskin



We're researching ways to turn low-temperature heat—such as industrial waste heat—into electricity using Organic Rankine Cycle (ORC) systems. These systems use special machines called expanders, which are key to making renewable energy and waste heat recovery more efficient and cost-effective.

Under the guidance of Dr. Raya Al-Dadah in Mechanical Engineering, I design and optimise expanders like Static-Shaft Wankel and Scroll Expanders. These devices can convert heat from sources as low as 60–90°C into usable power.

To achieve this, I've developed advanced numerical models in MATLAB and run detailed simulations using Ansys Fluent on the University's BlueBEAR high-performance computing (HPC) system. BlueBEAR allows us to test multiple designs quickly, saving time and cost compared to traditional experiments. This work has led to several high-impact publications and is helping us create more efficient, sustainable energy solutions.



Scroll Expanders



Components of wankel expander

Elucidating the Pathways for Human Tooth Enamel Mineralisation by Mahdi Tavakol

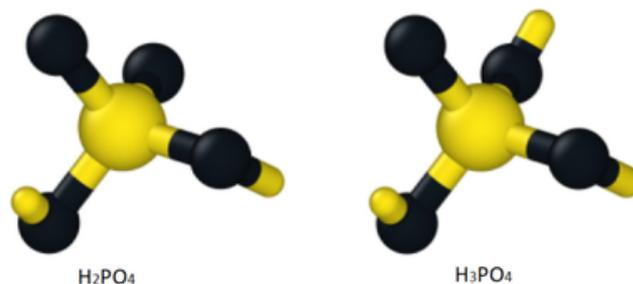
I'm Mahdi Tavakol, a computational scientist at the University of Oxford and visiting researcher at the University of Birmingham. My work is part of an EPSRC-funded project exploring how human tooth enamel forms and is broken down, using advanced computational techniques.



Our goal is to understand the molecular processes involved in tooth mineralisation and demineralisation, which could lead to new preventive and therapeutic solutions in dentistry. I use high-performance simulations—such as molecular dynamics and constant-pH modelling—to study the main tooth mineral, hydroxyapatite, and how adding ions like magnesium or fluoride can improve its stability and resistance to decay.

A key finding is that effective ion incorporation requires enamel to undergo cycles of demineralisation and remineralisation, highlighting the importance of natural repair mechanisms. This knowledge could help design treatments that strengthen teeth rather than relying on artificial fillings.

The University of Birmingham's BlueBEAR supercomputing facility has been essential for running large-scale simulations quickly and accurately, enabling breakthroughs that would take months on standard systems. These insights bring us closer to next-generation dental care.



Combined Molecules

Utilising BEAR workshops for Multimodal Analysis in Cognitive Translation Research by Mingjie Zhang

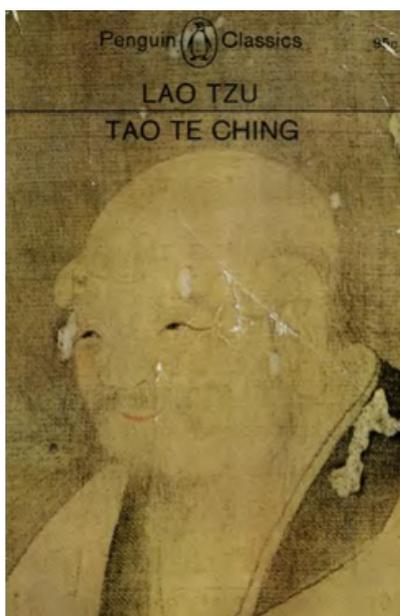
Mingjie Zhang, a postgraduate researcher in the Department of Languages, Cultures, Art History and Music at the University of Birmingham, is using digital humanities methods in her project “Translating between Minds: Cognitive Insights of the Dao De Jing.” Her work involves large-scale text analysis, including word embeddings, cluster analysis, and multimodal translation studies.



To build these skills, Mingjie attended BEAR’s Software Carpentry training series, covering tools like R, Python, MATLAB, and GitHub. These sessions helped her clean, analyse, and visualise data efficiently, model algorithms, and manage collaborative projects. She also learned image processing techniques that support her research on illustrated translations.

BEAR’s drop-in sessions provided hands-on support for troubleshooting issues such as software installation and debugging, enabling Mingjie to overcome technical challenges quickly. Thanks to these resources, she can now apply advanced computational techniques seamlessly, improving both the efficiency and quality of her research.

Mingjie strongly recommends the BEAR training series to researchers at all levels who want to enhance their digital skills and empower their projects.



The Book of Dao De Jing

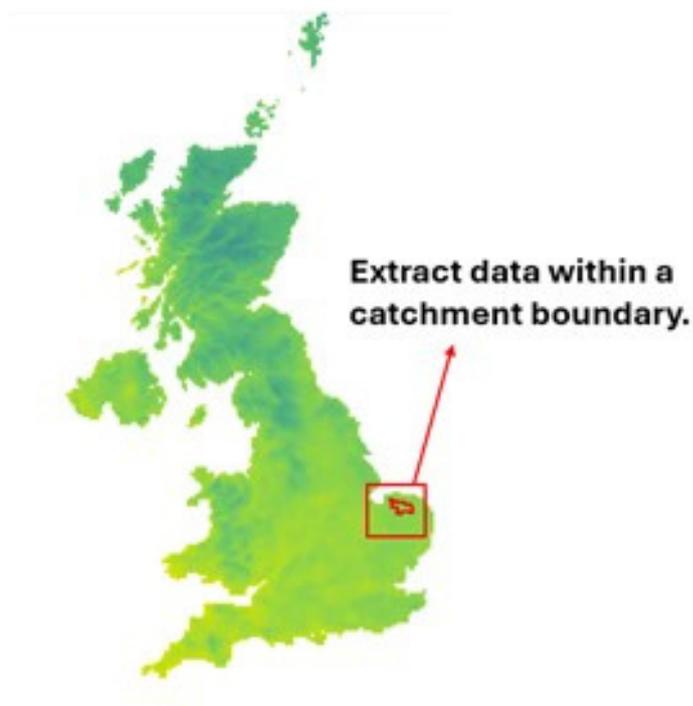
Extraction and Analysis of Hydrological Spatial Data by Chenlu Yang

Chenlu Yang, a PhD student in the School of Geography at the University of Birmingham, is researching flood behaviour and hydrological extremes across the UK. Her work uses high-resolution climate datasets and decades of catchment-scale observations to track changes in flood intensity and distribution.



This research involves processing large files—such as NetCDF, TIFF rasters, and shapefiles—which can be extremely time-consuming on a standard computer. By using the University’s BlueBEAR high-performance computing (HPC) system, Chenlu can run complex R and Python scripts in parallel, speeding up data extraction and analysis dramatically. Tasks that would take days on a laptop can now be completed in just hours using up to 72 cores.

BlueBEAR also provides secure, scalable storage for her datasets, ensuring efficient management of climate data. These resources allow Chenlu to focus on understanding how flood patterns are changing—work that is vital for improving resilience to extreme weather events.



RNA sequencing analysis made easy with BEAR

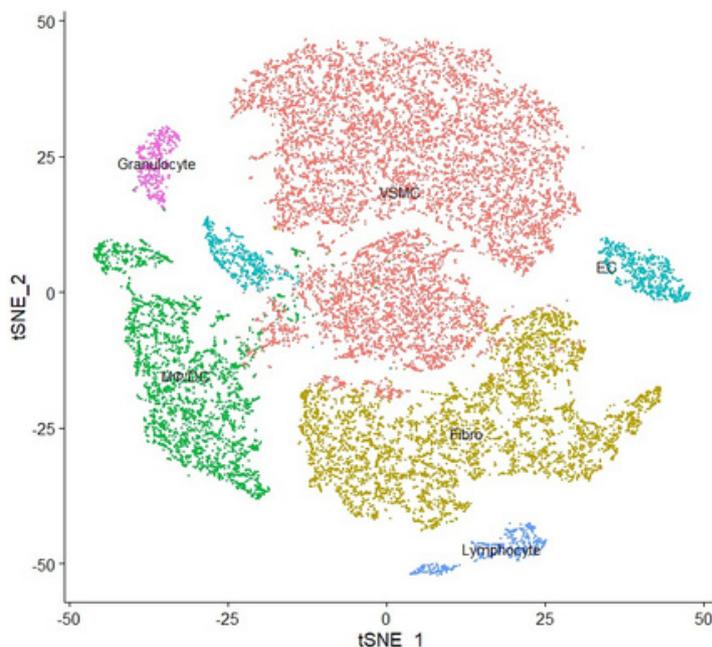
by Jonny Lewis



Jonny Lewis, a researcher in Inflammation and Ageing at the University of Birmingham, uses BEAR's high-performance computing (HPC) resources to study how blood flow patterns affect the cells lining blood vessels. When blood doesn't flow smoothly—a condition known as shear flow—it can change how these cells behave and what molecules they release, influencing processes like vessel growth and permeability.

To investigate this, Jonny and his team analysed bulk and single-cell RNA sequencing data, which involves processing huge datasets to understand gene expression changes. Downloading and analysing these files required over 170GB of storage and advanced computing power—tasks that would be impossible on a standard computer.

Using BlueBEAR, Jonny could run complex scripts, download large datasets, and process them with specialised tools like Cellranger and RStudio Server, all while continuing lab work. BEAR's secure storage and scalable computing allowed the team to uncover how molecules such as galectins respond to different blood flow conditions, providing insights into vascular health and disease.



TSNE plot of the different cell clusters found in the blood vessels

Extracting tender information with BlueBEAR: tackling complex documents with NLP and image processing by Asim Abbas

Asim Abbas, a PhD researcher in Artificial Intelligence at the University of Birmingham, is tackling the challenge of extracting structured information from complex tender documents. Unlike standard PDFs, these documents often include unstructured text, tables, diagrams, and non-standard layouts, making them difficult to process with traditional tools.



To solve this, Asim uses Natural Language Processing (NLP) and image processing techniques to convert document pages into images and apply advanced models for information extraction. This work requires significant computing power and high-performance GPUs, which the University's BlueBEAR HPC platform provides.

BlueBEAR's scalable resources and easy-to-use environment allow Asim to process large datasets efficiently, optimise workflows, and focus on research rather than technical limitations. His work aims to automate tender document analysis—an innovation with real-world applications for industry.

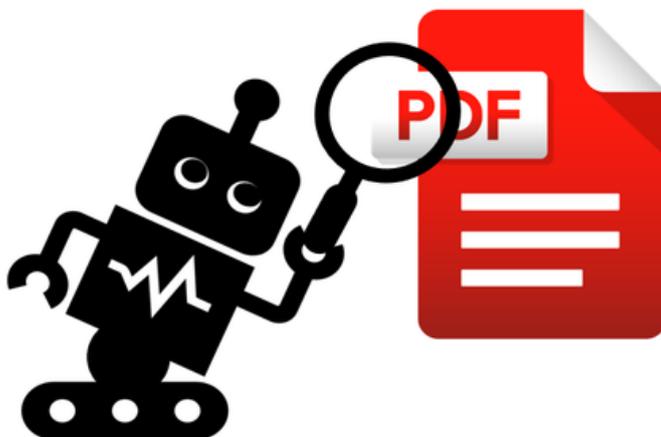


Image taken from
<https://www.linkedin.com/pulse/your-journey-towards-automation-can-robot-read-pdf-documents-hoh/>

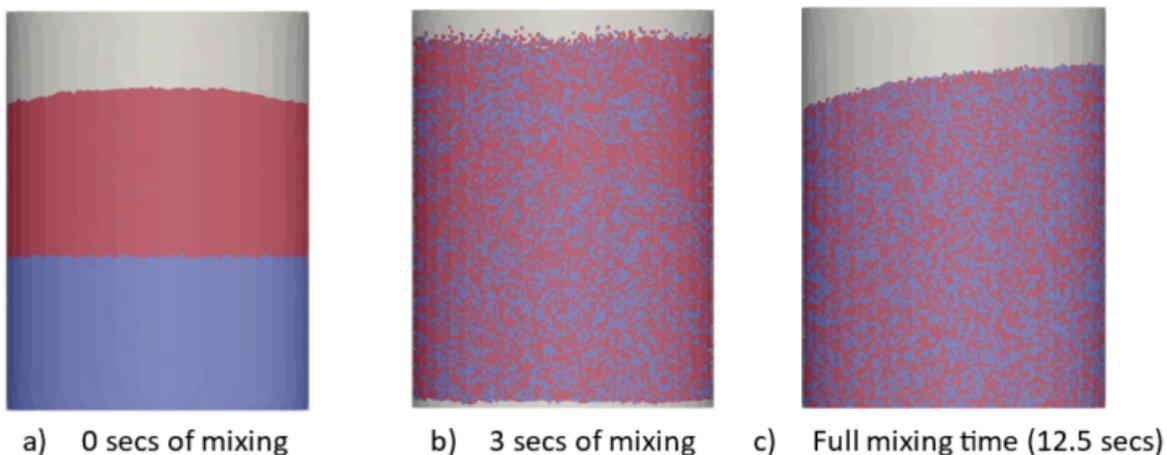
Shaking up things: Optimising mixing using BlueBEAR by Emily Allman



I'm a recent MEng Chemical Engineering graduate from the University of Birmingham. My final-year project investigated how particle properties affect mixing efficiency in Resonant Acoustic Mixers (RAM) – a vibration-based, blade-free mixing technology a new technique that is rapidly gaining attention in pharmaceuticals and other industrial sectors because it's efficient and cost-effective.

Using the BlueBEAR high-performance computing facility, I ran physics-based simulations (Discrete Element Method) to study the effects of particle friction, elasticity and fill height. Fill height was found to be the dominant factor, with optimal mixing around 80%, aligning with industry practice. Particle friction and restitution also influenced mixing, particularly at early stages.

The benefit of using BlueBEAR is not just being able to run large simulations, but also being able to run large numbers of distinct simulations in parallel, thus rapidly exploring a huge parameter space in a manner that would be (in terms of both time and money!) practically impossible in real life



Simulation results generated through ParaView for 80% fill height after a) 0 secs, b) 3 secs and c) 12.5 secs of mixing, with a sliding friction of 0.8.

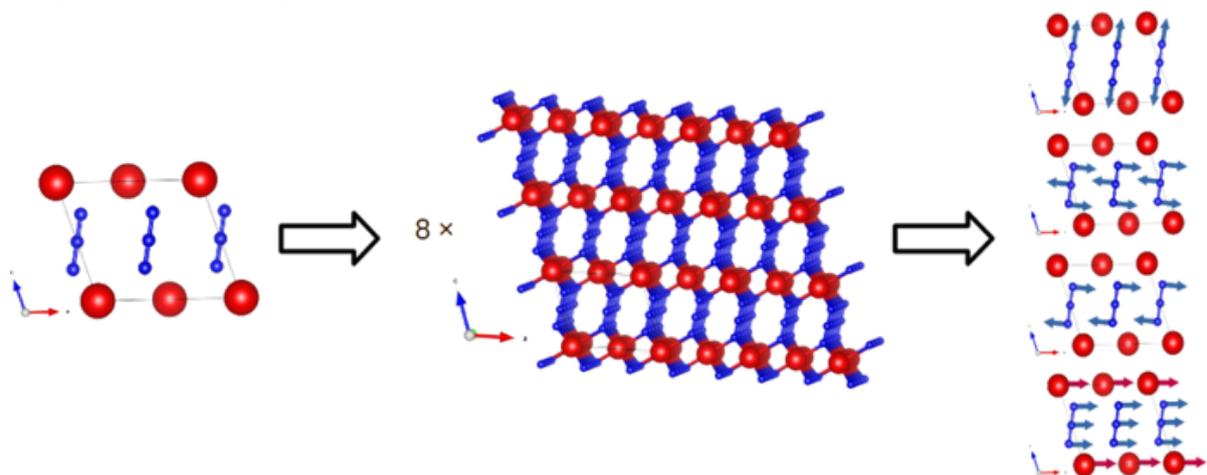
What makes things go BOOM!– from vibrations to detonation by Jakob Hein-Paar

My research focuses on vibrational metallization in metal azides, a class of energetic materials (explosives). This involves studying how energy moves through their crystal structures. Inside these crystals, atoms vibrate—these vibrations, called phonons, can interact with electrons and sometimes close the electronic bandgap, making the material behave like a metal and potentially triggering chemical reactions, including explosions.



To understand these processes, I run thousands of complex calculations to model atomic movements and energy transfers. This requires enormous computing power, which is why I use the University of Birmingham's BlueBEAR high-performance computing (HPC) system. BlueBEAR provides the memory and parallel processing needed to handle large-scale simulations quickly and efficiently.

By leveraging this technology, we aim to design safer, more environmentally friendly energetic materials through computer modelling—reducing the need for risky physical testing.



Workflow for the simple sodium azide (NaN₃, Na: red, N: blue). Supercells with different, small atomic displacements are created from the primitive cell. From the calculated interatomic forces, the dynamical matrix can be constructed and the eigenvectors, governing the atomic displacement of the phonon modes, obtained. Figure created with VESTA.

Modelling defects in inorganic phosphor crystals by Maryia Shymanovich

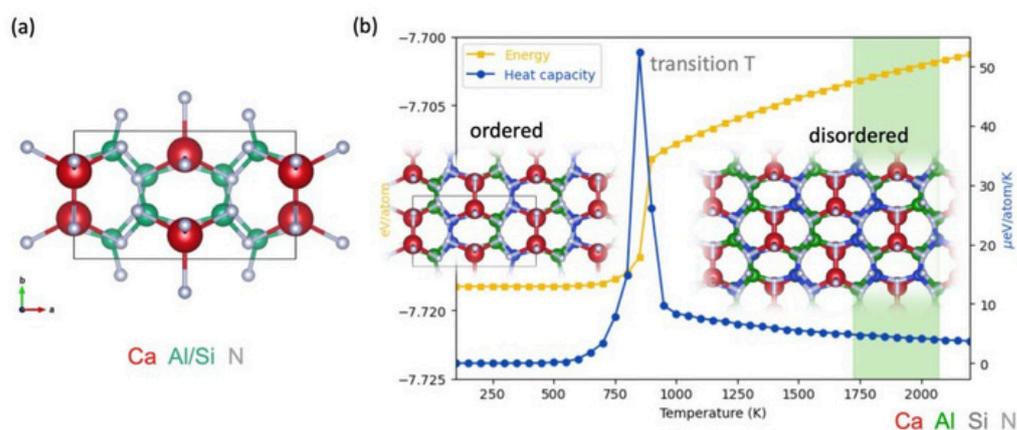
I'm a second-year PhD student in Chemistry at the University of Birmingham, studying defects in inorganic phosphor materials used in energy-efficient lighting. Phosphors are key components in LED bulbs—they absorb high-energy light and convert it into different colours, creating bright, tuneable white light for homes and industry.



Defects in these materials, such as irregularities in their crystal structure, can dramatically affect performance. My research uses advanced computational modelling to understand these defects and predict how they influence optical properties. This involves thousands of density functional theory (DFT) calculations and large-scale simulations, which would take weeks on a normal computer.

Using the University's BlueBEAR high-performance computing (HPC) system, I can run these calculations in parallel, reducing processing time from weeks to days. This allows me to explore complex phenomena like cation disorder in CaAlSiN_3 (a red-emitting phosphor for warm-white LEDs) and its impact on electronic properties.

Without BlueBEAR, this research would not be possible. These insights will help design better phosphor materials for next-generation lighting.



(a) 24-atom unit cell of CaAlSiN_3 showing tetrahedral sites occupied by Al and Si. (b) Energy and heat capacity from Monte Carlo simulations; Sharp peak represents the transition from ordered structure with space group Cc (9) to a disordered phase. The green area shows the range of synthetic temperatures indicating that experimentally observed structure of CaAlSiN_3 exists in the disordered state.

How I used BEAR to understand investor behaviour in FinTech markets

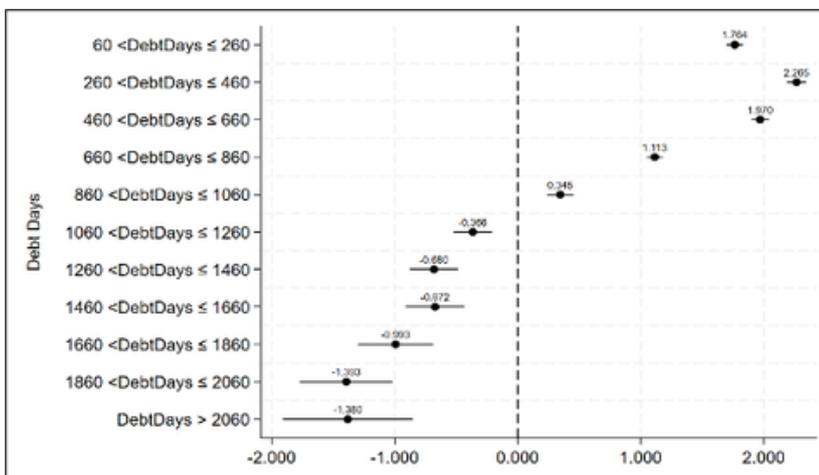
by Mohammad Aldoub



My research studies how investors respond to negative shocks on peer-to-peer lending platforms, focusing on changes in belief dispersion during periods of uncertainty. P2P lending refers to a digital platform where borrowers and lenders interact directly without traditional financial intermediaries. Using loan-level data from Bondora, a leading European P2P platform, the analysis covers more than 45 million transactions from 2016 to 2024 to examine investor behaviour in the secondary market.

The results show that investors initially disagree strongly about a loan's value, with some seeing opportunities and others perceiving high risk. As risk increases and the loan remains unpaid, this disagreement first widens and then narrows as investors converge on the view that the loan is too risky, leading to herding behaviour

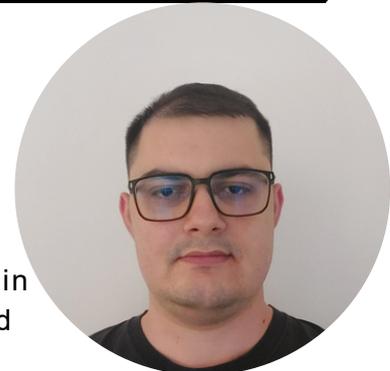
Processing such large datasets and running complex models would be impossible on a standard computer. The University's BlueBEAR high-performance computing system makes this work feasible by enabling fast data cleaning, large-scale analysis, and secure data storage. This computational capacity directly supports our study by enabling us to conduct the required large-scale analysis reliably and systematically, ensuring that research objectives can be addressed effectively.



Herding Behaviour in High-Risk Loans: Evidence from Debt Days Analysis

BEAR helps delve into the infant gut and malnutrition

by Leon Mijic

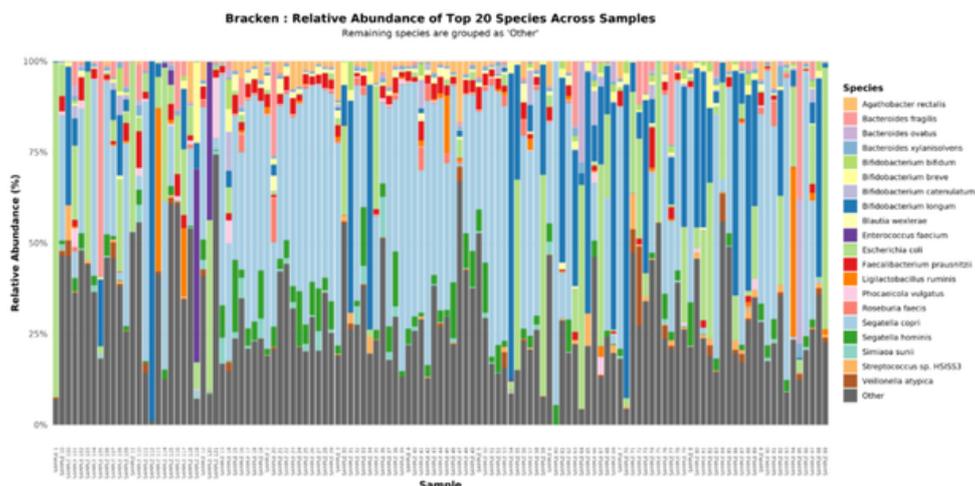


The human gut hosts trillions of microbes that play vital roles in health. Leon's research focuses on how the early-life gut microbiome interacts with child malnutrition in low- and middle-income countries, influencing growth and long-term health.

Understanding these interactions could lead to probiotic or dietary solutions to reduce the harmful effects of malnutrition during childhood. To study this, terabytes of metagenomic data alongside clinical information is being analysed using a custom bioinformatics pipeline for genome assembly, taxonomy, and functional profiling.

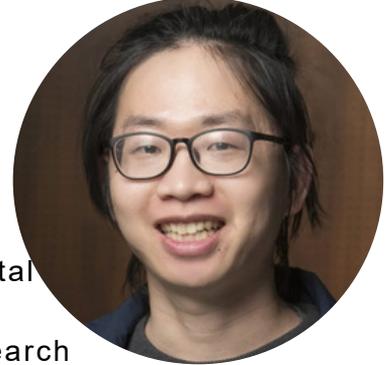
Such large-scale analysis requires enormous computing power. On a normal computer, one dataset would take decades to process. Using the University's BlueBEAR high-performance computing (HPC) system, we completed the equivalent of 24 years of CPU time in just 14 days. BlueBEAR also provides secure storage and easy-to-use tools, making advanced microbiome research possible.

These insights could help design targeted interventions to improve child health globally.



Stacked bar plot of relative abundances of bacterial species. The x axis represents the samples that have been analysed while the y axis represents relative abundances of bacterial species. Each colour corresponds to a different species of bacteria. The taxonomic abundances have been assigned using Kraken2 and Bracken with the standard database and visualized using ggplot2.

High-Resolution Soil Hydrology Analysis Enabled by BlueBEAR by Guilin Luo



I'm Guilin, a PhD student in Geography at the University of Birmingham. My research uses high-resolution environmental data to understand how temperate forest soils respond to climate change. At the Birmingham Institute of Forest Research (BIFoR), we continuously monitor atmospheric and soil conditions, creating detailed datasets that reveal how rainfall and other factors affect soil moisture at different depths.

To process this data, I use R and Python to clean sensor outputs, remove noise, and prepare them for analysis. This work supports better land-use planning, vegetation modelling, and resilience assessments for extreme weather.

Handling such large datasets and running advanced machine-learning models requires significant computing power. The University's BlueBEAR high-performance computing (HPC) system makes this possible, reducing processing times from days to hours and enabling interactive visualisations for publications and stakeholder engagement. Expert support and training from the BEAR team have been invaluable in optimising workflows and ensuring robust results.

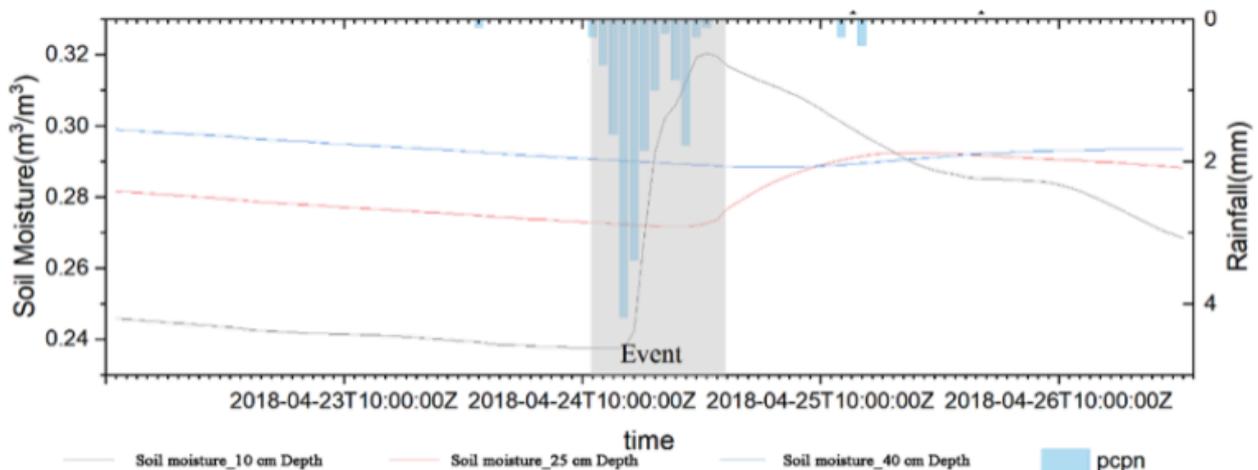


Fig. Soil moisture response at multiple depths during a concentrated rainfall event

Bio-Inspired FSI and HPC-driven Aerodynamic Optimisation

by Hibah Saddal



Learning from natural flyers, our research is aimed at improving the aerodynamic performance of Unmanned Aerial Vehicles (UAVs), including bio-inspired strategies to exploit coupled fluid-structure interactions (FSI). One research focus is the investigation of flexible flaps attached to the upper surface of airfoils, inspired by the covert feathers of birds, to increase lift and delay stall for high angles of attack flight. Another research direction explores morphing wings as an alternative to conventional rigid wing designs for improved aerodynamic efficiency and adaptability across varying flight conditions.

To model these ideas, we conducted two-way coupled FSI simulations using a partitioned strong-coupling approach. Simulations are performed using OpenFOAM for the fluid solver, while the structural response is computed using CalculiX. To facilitate data exchange between both our solvers for the FSI problem, we use the preCICE coupling library.

As these coupled simulations are computationally demanding, we use the University of Birmingham's BlueBEAR high-performance computing (HPC) system. The use of BlueBEAR HPC resources enabled us to run large batches of simulations in parallel, significantly reducing computation time. This work highlights how HPC and open-source software are driving innovation in aerospace technologies.



An example visual from our research, titled "Whispers of Vortices from a Covert-Inspired Leading-Edge Flap" by Hibah Saddal and Dr. Chandan Bose, shows a covert-inspired flexible flap interacting with the vortex shedding over a wing section under stalled conditions.

Atomic Simulation to Build Better Batteries

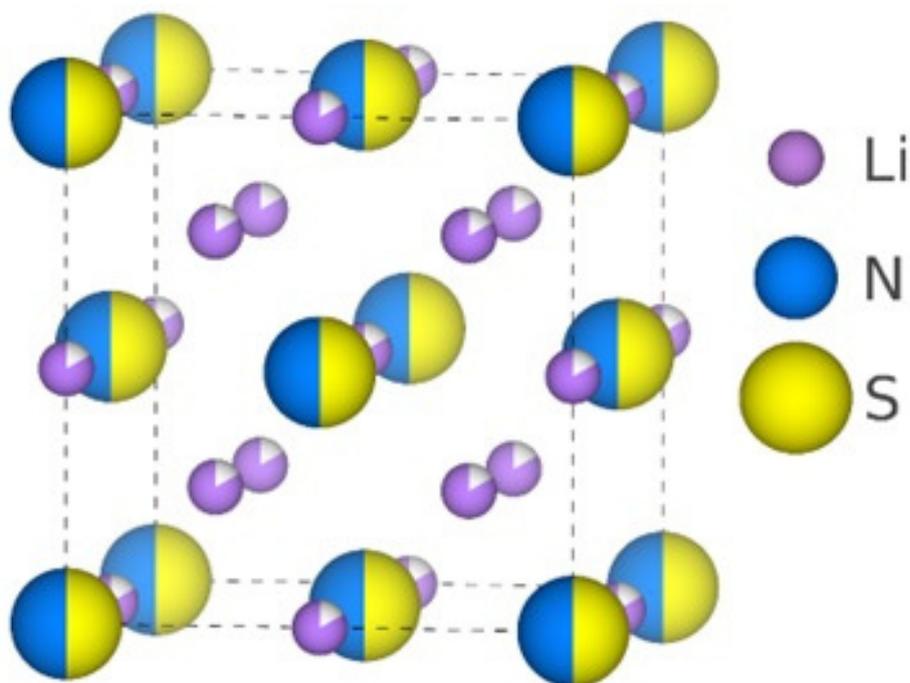
by Oskar Soulas

I'm in the Scanlon Materials Theory Group at the University of Birmingham, where we use advanced computational tools to design materials from the atomic level up. My research focuses on solid electrolytes—key components in next-generation batteries that could replace flammable liquid electrolytes in lithium-ion cells.



I study a family of compounds called Li-S-N antifluorites, which have the potential for high lithium-ion conductivity if their atomic structure is optimised. Using the University's BlueBEAR high-performance computing (HPC) system, I run thousands of simulations to understand how atomic ordering affects lithium mobility. These include quantum-mechanical calculations, cluster expansion models, and molecular dynamics to link structure with stability and conductivity.

Our findings show that disordered atomic arrangements enable faster lithium movement, while ordered structures restrict it. This insight helps guide experimental work toward the most promising candidates, accelerating the development of safer, high-performance energy storage.



Structure of Li-S-N solid electrolytes with disordered positions

Bayesian Modelling of Rising Decision Thresholds in DDMs using BlueBEAR

by Sophie Wetz

I'm a Master's student in Computational Neuroscience. My project is supervised by Dr Dietmar Heinke and tests whether dynamic decision boundaries improve the accuracy of the Drift Diffusion Model (DDM) compared to fixed boundaries in a perceptual decision-making task. Dynamic boundaries may occur when early evidence is unreliable, and delaying a response allows better decisions.



To analyse this, I used Bayesian inference with advanced sampling methods (DE-MCMC) to handle complex, high-dimensional parameter spaces. These computations are extremely demanding, so I ran them on the University's BlueBEAR high-performance computing (HPC) system. BlueBEAR reduced processing time from weeks on a single machine to just days, enabling rapid testing of model variants and thorough checks.

The findings provide partial support for dynamic boundaries in decision-making and highlight the need for more flexible models. This research shows how HPC resources accelerate progress in understanding the brain's decision processes.

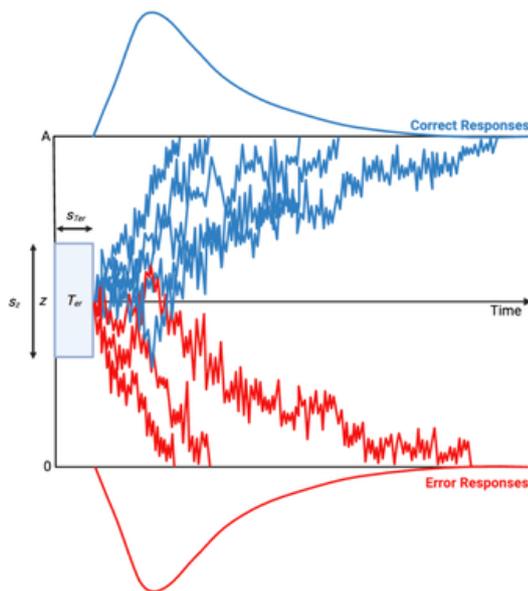


Illustration of the Drift Diffusion Model taken from Deakin et al. (2024). The blue lines show evidence accumulation towards a correct response while the red lines depict accumulation preceding an incorrect response. Probability distributions of response times are shown at each boundary.

Deakin, J., Schofield, A., & Heinke, D. (2024). Support for the Time-Varying Drift Rate Model of Perceptual Discrimination in Dynamic and Static Noise Using Bayesian Model-Fitting Methodology. In *Entropy* (Vol. 26, Issue 8, p. 642). MDPI AG. <https://doi.org/10.3390/e26080642>

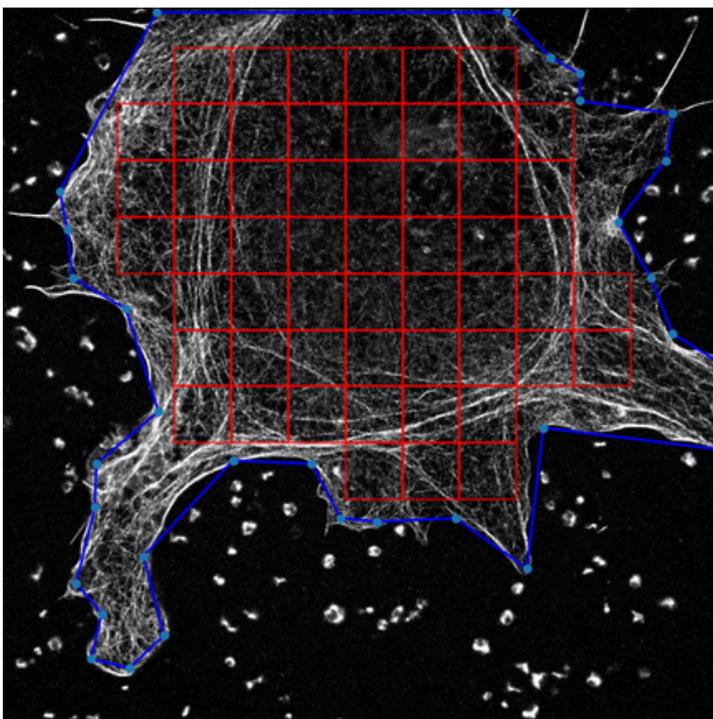
SMLM Data Exploration with Nano-org and BlueBEAR by Dylan Owen and Sandeep Shirgill

Nano-org is a cutting-edge platform that uses ultra-precise coordinate data to map the nanoscale distribution of proteins within cells. This enables researchers to explore how protein architectures vary across cell types and conditions, opening up a new field called spatial nano-omics.



The platform offers a user-friendly interface for uploading and analysing single-molecule localisation microscopy datasets. Behind the scenes, it uses advanced algorithms for similarity analysis and clustering to uncover biologically meaningful patterns. These large-scale computations are powered by the University of Birmingham's BlueBEAR high-performance computing (HPC) system, which provides parallel processing, virtual machines, and efficient data management.

Nano-org is built on a modern web stack and is continually evolving. Current developments include integrating machine learning models to detect nanoscale patterns linked to protein function and cellular states. This collaboration between the Research Software Group and Dylan Owen's lab demonstrates how HPC resources transform complex research into practical, high-performance tools for life sciences.



Rendered SMLM image showing actin filaments, key components of the cytoskeleton, within a cell. Nano-org's automated segmentation (blue polygon and red squares) enables region-specific analysis of nanoscale protein distributions

Energy Data Landscape – Challenges and Opportunities by Grant Wilson and Edward Barbour

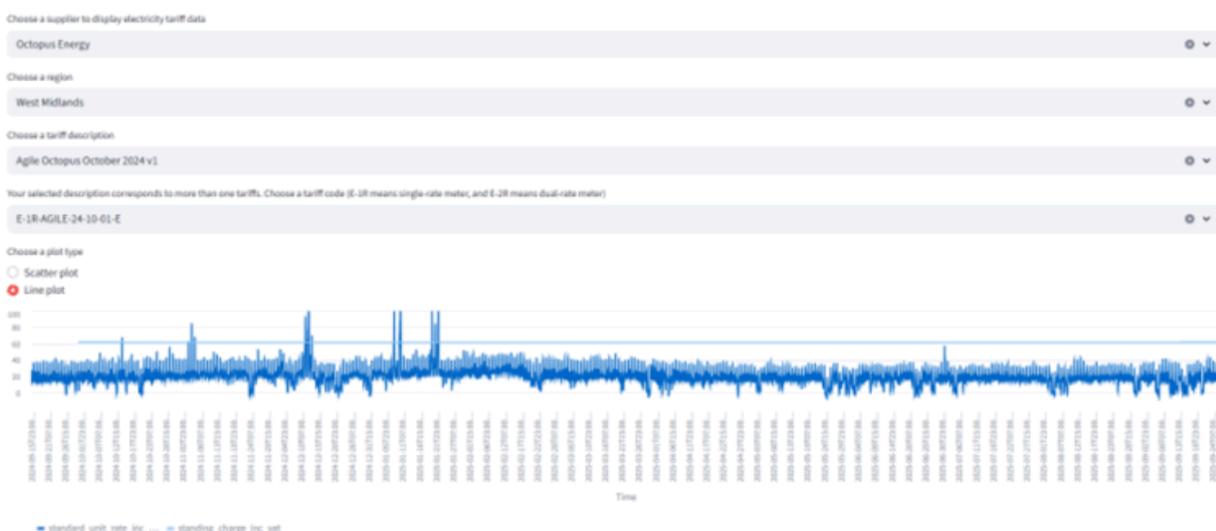


Researchers at the University of Birmingham are investigating the costs of domestic heat pump operation under different energy tariffs. While exploring tariff data, they encountered major challenges: inconsistent formats, limited access, reliance on PDFs, and a lack of standardised APIs. Only one supplier, Octopus Energy, offers usable API access. Most suppliers provide only snapshots of prices, making long-term analysis difficult.

To address this, the team developed a proof-of-concept visualisation tool to track tariff trends and highlight the fragmented data landscape. They recommend that Ofgem introduce standards to ensure tariff data meets FAIR principles—Findable, Accessible, Interoperable, and Reusable—so researchers and policymakers can analyse energy costs more effectively.

The next step is to present findings to Ofgem and advocate for improved data transparency, supporting better policy and consumer decision-making.

Electricity tariffs visualisation



Investigating the Alternative Lengthening of Telomeres (ALT)

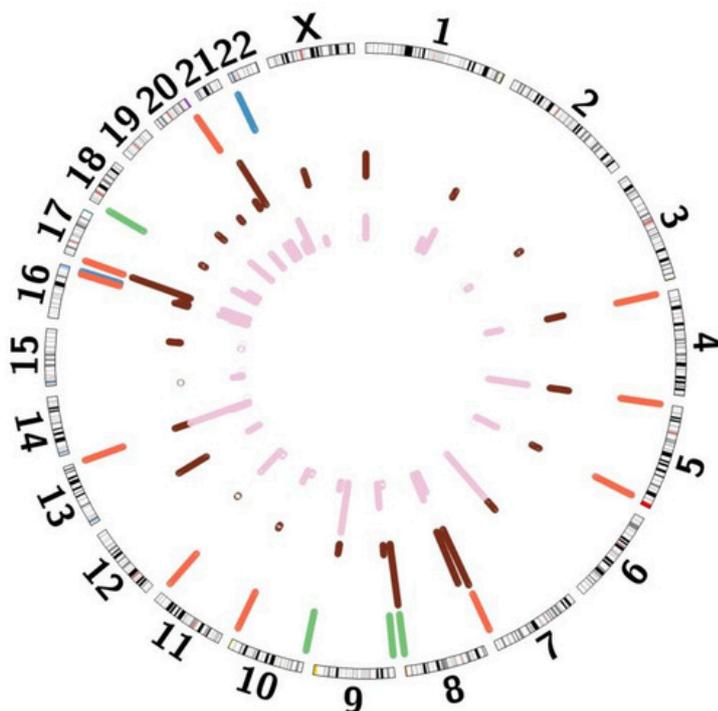
by Ronan Broderick and Anubrata Das

Our lab studies Alternative Lengthening of Telomeres (ALT), a mechanism used by 10–15% of cancers to survive by extending the DNA at chromosome ends. There are currently no effective treatments for ALT-dependent cancers, such as high-grade glioma, so understanding this process could lead to new therapies and diagnostic tools.



We use Oxford Nanopore long-read sequencing to analyse telomeric DNA and rely on the University of Birmingham's BlueBEAR high-performance computing (HPC) system for GPU-powered data processing. BlueBEAR also provides secure storage for large genomic datasets and tools for image analysis, including deep learning libraries like Cellpose and Ilastik. Collaboration is streamlined through BEAR GitLab, which hosts our DNA and image-processing scripts.

These resources allow us to process complex genomic and imaging data efficiently, accelerating research into cancer mechanisms and potential treatments.



Synchronous U2OS cells

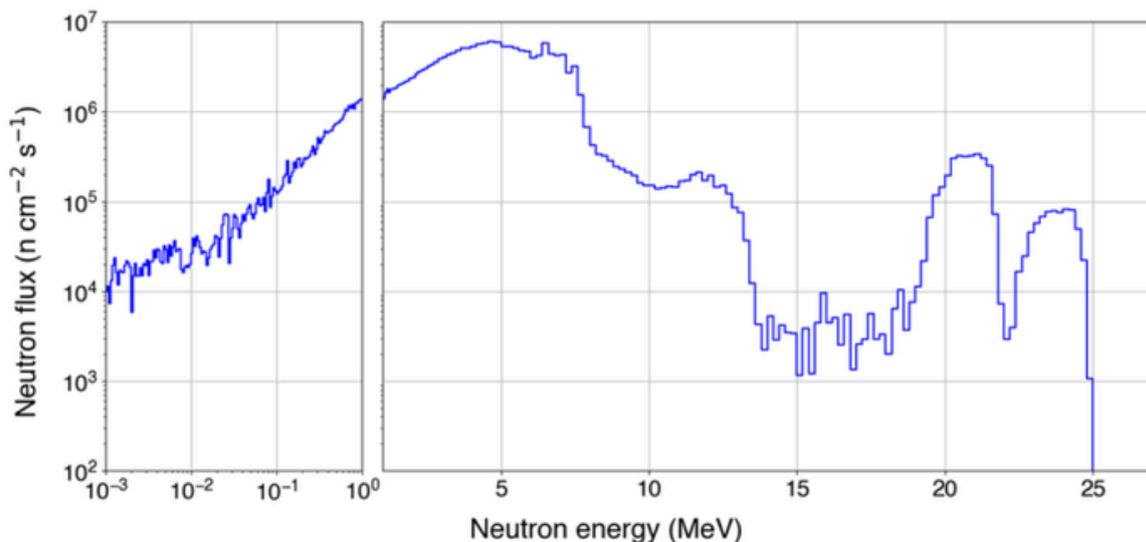
Modelling accelerator-based neutron sources using the BlueBEAR cluster by Louis Butt

I'm a PhD student in the School of Metallurgy and Materials at the University of Birmingham, researching accelerator based neutron sources. My work involves measuring neutron output from our source at the University's MC40 Cyclotron and comparing it with predictions from complex radiation transport models. This helps verify the accuracy of nuclear data libraries used in simulations for future technologies like fusion reactors and spallation sources.



Running these models requires Monte Carlo simulations of billions of particles, which would take weeks on a single computer. Using the University's BlueBEAR high-performance computing (HPC) system, I can run these simulations in parallel across multiple cores, reducing runtime to hours or days. This allows me to test many model variations while focusing on challenging experimental measurements.

BlueBEAR makes it possible to combine theory and experiment efficiently, accelerating research that supports the development of next-generation neutron facilities.



Neutron energy spectrum from an example BlueBEAR simulation of our accelerator-based neutron source. To generate this spectrum, the behaviour of 10,000,000,000 particles was simulated, and an average result across every simulation was taken.