M.Sc. in Mathematical Modelling and Scientific Computing Dissertation Projects

February 2026

Contents

1 Industrial Projects

Descriptions of industrial projects have been omitted from the web version of the projects booklet. Please contact Kathryn Gillow, kathryn.gillow@maths.ox.ac.uk, if you believe you should have access to these.

2 Biological, Medical and Physical Application Projects

2.1 Active Fluid Models of Collective Cell Migration Guided by Chemical Cues

Supervisor: Giulia Celora
Contact: giulia.celora@maths.ox.ac.uk

Active fluid models have become a powerful framework for studying biological systems across scales, from individual cells to multicellular collectives and tissues. Despite their diverse physical and biological properties, these systems share key features: they consist of interacting units that operate out of equilibrium (are active) and exhibit emergent material properties such as surface tension and viscosity. Active fluid theories provide a unifying approach to investigate how activity and emergent material properties interact to shape the dynamics, organisation, and function of cellular collectives. In my research, I am interested in the applications and development of this framework to the study of collective migration of cell groups guided by chemical cues, a process known as self-generated chemotaxis.

Below are a few possible research directions that could develop into a project depending on the student's interests and background.

1. For students with a strong interest in numerical methods and interdisciplinary research, one possible direction is to extend our recent work [1] to characterise two-dimensional pattern formation arising from self-generated chemotaxis in a homogeneous cell population. This project would involve becoming familiar with the existing partial differential equation (PDE) model and developing its numerical implementation in two dimensions. By performing parameter sweeps over key model parameters, the aim would be to classify the resulting spatial patterns and compare them with available experimental observations from Chubb's lab (UCL); see Figure 1.

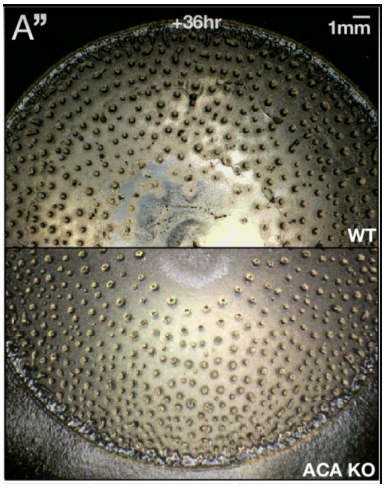

Figure 1: Example experimental observations.

2. For students with a strong interest in applied partial differential equations, one possible direction is to develop the theory of morphological transitions in chemotacting active droplets, as observed in [1]. This direction would involve singular perturbation techniques, exponential asymptotics, and boundary-layer analysis of modified thin-film equations.

3. For students with a strong interest in mathematical modelling, one possible direction is to work on the collective chemotaxis of heterogeneous cell groups. This project would involve developing a minimal theoretical model based on [1] to account for heterogeneity within migrating cell collectives (or tissues), composed of cells with distinct, specialised behaviours. Such heterogeneity is observed in several developmental systems, such as the zebrafish lateral line, where cells at the front and rear of the group specialise in either sensing or consuming the guiding chemical. By developing and analysing a minimal model, the project would investigate how such a type of cellular heterogeneity influences collective migration and under which conditions this leads (if possible) to more effective migration.

References

[1] H. Z. Ford, G. L. Celora, E. R. Westbrook, M. P. Dalwadi, B. J. Walker, H. Baumann, C. J. Weijer, P. Pearce, and J. R. Chubb. Pattern formation along signaling gradients driven by active droplet behavior of cell swarms. *Proc. Natl. Acad. Sci. U.S.A.*, 122(21):e2419152122, https://doi.org/10.1073/pnas.2419152122, 2025.

2.2 Molecular Dynamics Simulations of Bacterial Actin MreB Protein

Supervisor: Radek Erban
Contact: radek.erban@maths.ox.ac.uk
Collaborator: Rut Carballido-Lopez (Université Paris-Saclay, INRAE, France)

Actin is a well-studied protein in eukaryotic cells (animal and plant cells) that assembles into filaments to shape cells, enable movement, and perform other essential cellular functions. In bacteria, the actin structural homolog [1] MreB also forms filaments and is crucial for rod shape determination of many bacterial species [2]. While some properties and functions of the 'bacterial actin' MreB are similar to those of eukaryotic actin, many details remain unresolved (for example, how are filaments built and arranged in vivo?, how is their subcellular position controlled?, do MreB filaments nucleate?). In this project, we will use molecular dynamics (MD) simulations and experimentally determined structures of MreB from *Geobacillus stearothermophilus* [3] to study the role of ATP binding and hydrolysis in MreB polymerisation.

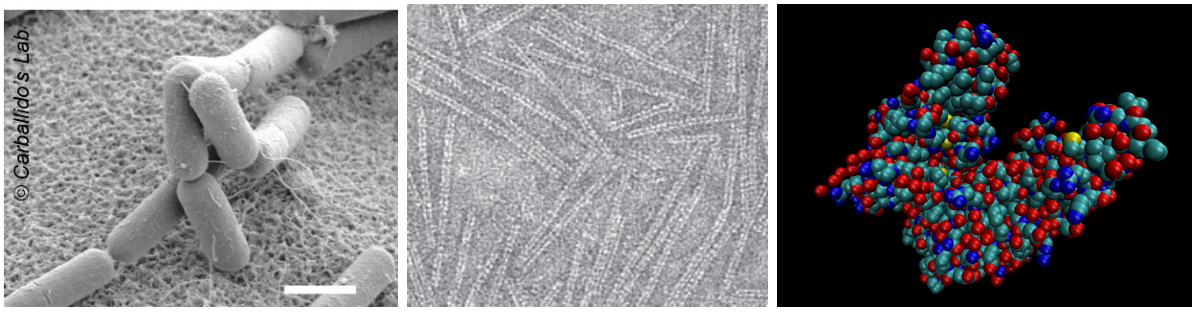

MD methods studied in this project will be based on classical mechanics and are sometimes called molecular mechanics approaches in the computational chemistry literature, [4, 5]. The growing importance of MD has been recognised by the award of the Nobel Prize in Chemistry to pioneers of this approach. In applications, MD approaches are used to study how complex biomolecules move, deform

and interact over time [6, 7, 8]. In mathematical terms, MD can be described by systems of ordinary differential equations (ODEs) or stochastic differential equations [9] (SDEs) for the positions and velocities of atoms or beads (representing collections of atoms), which can be also subject to algebraic constraints [10]. At the macroscopic level, biomolecules and their interactions can also be modelled by coarse-grained models, including approaches based on Brownian dynamics, Langevin dynamics and coarse-grained MD [9, 11, 12, 13].

References

[1] F. van den Ent, L. A. Amos, and J. Löwe. Prokaryotic origin of the actin cytoskeleton. *Nature*, 413:39–44, https://doi.org/10.1038/35092500, 2001.

[2] Laura J. F. Jones, Rut Carballido-López, and Jeffery Errington. Control of Cell Shape in Bacteria: Helical, Actin-like Filaments in *Bacillus subtilis*. *Cell*, 104(6):913–922, https://doi.org/10.1016/S0092-8674(01)00287-2, 2001.

[3] W. Mao, et al. On the role of nucleotides and lipids in the polymerization of the actin homolog MreB from a Gram-positive bacterium. *eLife*, 12:e84505, https://doi.org/10.7554/eLife.84505, 2023.

[4] E. Lewars. "Computational Chemistry", 3rd edition, Springer, 2016.

[5] D. Frenkel and B. Smith. "Understanding Molecular Simulation", Academic Press, 2nd edition, 2002.

[6] A. Colavin, J. Hsin, and K. C. Huang. Effects of Polymerization and Nucleotide Identity on the Conformational Dynamics of the Bacterial Actin Homolog MreB. *PNAS*, 111(9):3585–90, https://doi.org/10.1016/j.bpj.2013.11.961, 2014.

[7] H. Shi, et al. Sensing the shape of a surface by tightly surface-bound filaments. *PNAS*, 122(52):e2526131122, https://doi.org/10.1073/pnas.2526131122, 2025.

[8] R. Erban and Y. Togashi. Asymmetric Periodic Boundary Conditions for All-Atom Molecular Dynamics and Coarse-Grained Simulations of Nucleic Acids. *Journal of Physical Chemistry B*, 127(38):8257–8267, https://doi.org/10.1021/acs.jpcb.3c03887, 2023.

[9] R. Erban and S. J. Chapman. "Stochastic Modelling of Reaction-Diffusion Processes", Cambridge Texts in Applied Mathematics, Cambridge University Press, 2020.

[10] B. Leimkuhler and C. Matthews. "Molecular Dynamics. With deterministic and stochastic numerical methods", Springer, 2016.

[11] R. Erban. From molecular dynamics to Brownian dynamics. *Proceedings of the Royal Society A*, 470:20140036, 2014.

[12] R. Erban. Coupling all-atom molecular dynamics simulations of ions in water with Brownian dynamics. *Proceedings of the Royal Society A*, 472:20150556, 2016.

[13] R. Erban. Coarse-graining molecular dynamics: stochastic models with non-Gaussian force distributions. *Journal of Mathematical Biology*, 80:457–479, 2020.

2.3 Mathematical Modelling of Cell-Cycle and Spatial Dynamics from FUCCI Imaging Data

Supervisors: Dr Rebecca Crossley and Dr Stéphanie Abo
Contact: rebecca.crossley@maths.ox.ac.uk and stephanie.abo@maths.ox.ac.uk

Background and problem statement
Understanding cell cycle progression and proliferation dynamics is central to many areas of biology, including cancer research and developmental biology. Fluorescent Ubiquitination-based Cell Cycle Indicator (FUCCI) systems enable live-cell imaging of cell cycle phases by encoding phase-specific fluorescence signals, providing rich spatio-temporal datasets at the single-cell level [1]. However, extracting quantitative information from FUCCI data is challenging due to noise, cell crowding, variability in fluorescence intensity, and the need to link observations across time. Another challenge is determining how reliably mathematical model parameters can be inferred from such processed imaging data, or how sensitive model conclusions are to image processing choices. This project will focus on developing a reproducible workflow to extract cell-cycle parameters from FUCCI microscopy data and using these measurements to calibrate mathematical models of cell-cycle dynamics.

Figure 2: FUCCI fluorescence images [3]

Proposed approach and techniques
The student will analyse time-lapse FUCCI microscopy data to extract cell trajectories, spatial organisation, and fluorescence-based cell-cycle phase information (primarily using FIJI [2]). These data will be used to quantify cell-cycle statistics, density evolution, and migration behaviour, and to calibrate population models formulated as partial differential equations structured by cell-cycle phase and, where relevant, space. A central mathematical focus will be the assessment of robustness and identifiability, examining how inferred parameters and model predictions depend on image-processing choices and vary across datasets. Programming will primarily be in Python for data analysis and parameter estimation; prior experience with biological image analysis is not required.

What you'd hope to achieve
The project aims to (1) develop a reproducible workflow for processing FUCCI imaging data and a documented database of tracked cells with quantified cell-cycle characteristics; (2) obtain parameter estimates, with uncertainty quantification, for structured models of cell-cycle dynamics; and (3) analyse which cell-cycle parameters are reliably identifiable from FUCCI data and which remain poorly constrained. This work will provide training in biological image analysis, structured population modelling, and data-driven parameter estimation.

References

[1] Asako Sakaue-Sawano, et al. Visualizing spatiotemporal dynamics of multicellular cell-cycle progression. *Cell*, 132(3):487–498, 2008.

[2] Johannes Schindelin, et al. Fiji: an open-source platform for biological-image analysis. *Nature methods*, 9(7):676–682, 2012.

[3] Sebastian J. Streichan, Christian R. Hoerner, Tatjana Schneidt, Daniela Holzer, and Lars Hufnagel. Spatial constraints control cell proliferation in tissues. *Proc Natl Acad Sci*, 111(15):5586–5591, 2014.

2.4 Mathematical Modelling of Confinement in Multicellular Systems

Supervisors: Dr Rebecca Crossley and Dr Carles Falcó
Contact: rebecca.crossley@maths.ox.ac.uk and carles.falco@sjc.ox.ac.uk

Background and problem statement
Collective cell migration and spatial organisation are governed by various interactions, including: cell–cell adhesion, cell–extracellular matrix (ECM) coupling, and phenotypic plasticity [4]. In many biological contexts, including development and cancer invasion, cells undergo epithelial–mesenchymal transition (EMT), during which adhesion strengths, motility, and mechanical responses change substantially. Continuum models based on partial differential equations (PDEs) have successfully captured aspects of adhesion-driven aggregation and ECM-guided migration, yet it remains unclear under what conditions these mechanisms lead to confinement of a tumour, either spatially or across phenotypic subpopulations. This project will address this gap by developing and analysing PDE models that combine cell–cell adhesion and cell–ECM interactions in heterogeneous cell populations, aiming to identify when confinement emerges and how it is altered by EMT-driven changes in cells.

Proposed approach and techniques
The project will focus on formulating coupled nonlinear PDE models for cell density evolution incorporating adhesion-mediated fluxes [1], ECM-dependent migration [2], and phenotype-dependent parameters. Multiple cell types or phenotypic states will be represented through coupled systems or internal variables [3], with EMT modelled via reaction or switching terms that modify adhesion and motility. Analytical techniques will include nondimensionalisation, steady-state and stability analysis, an exploration of pattern formation or travelling waves, as well as aggregation regimes, all complemented by numerical simulations that investigate confinement and segregation behaviour beyond analytically tractable cases. Computation will be carried out in Python, with emphasis on understanding qualitative mechanisms rather than fitting specific experimental data.

What you'd hope to achieve
The project aims to identify mathematical mechanisms through which combined cell–cell adhesion and cell–ECM interactions generate spatial or phenotypic confinement, and to characterise how EMT-associated parameter changes influence these outcomes. Outcomes will include the development of a well-defined PDE model, analytical results describing regimes of confinement, and numerical demonstrations of emergent spatial organisation in heterogeneous populations. The project will provide training in continuum modelling, nonlinear PDE analysis, and scientific computing, while contributing theoretical insight into how adhesion and phenotype transitions shape multicellular structure.

References

[1] Carles Falcó, Ruth E. Baker, and José A. Carrillo. A local continuum model of cell-cell adhesion. *SIAM Journal on Applied Mathematics*, 84(3):S17–S42, 2023.

[2] Rebecca M. Crossley, Philip K. Maini, Tommaso Lorenzi, and Ruth E. Baker. Traveling waves in a coarse-grained model of volume-filling cell invasion: Simulations and comparisons. *Studies in Applied Mathematics*, 151(4):1471–1497, 2023.

[3] Carles Falcó, Rebecca M. Crossley, and Ruth E. Baker. Travelling waves in a minimal go-or-grow model of cell invasion. *Applied Mathematics Letters*, 158:109209, 2024.

[4] C. Venkata Sai Prasanna, Mohit Kumar Jolly, and Ramray Bhat. Multiscale modeling predicts dependence of mesenchymally transitioned tumor niche fitness on cell-cell and cell-matrix adhesions. *bioRxiv*, https://doi.org/10.1101/2025.06.25.661486, 2025.

2.5 Modelling Transient Behaviour in Elastic Instabilities

Supervisors: Prof. Dominic Vella and Dr Andrea Giudici
Contact: dominic.vella@maths.ox.ac.uk and andrea.giudici@maths.ox.ac.uk

Elastic instabilities are fascinating phenomena that occur when an elastic material under stress suddenly relaxes to a new configuration. In biological systems, elastic instabilities are believed to sculpt a variety of complex shapes, from the wrinkles on the edges of leaves to the villi in the human gut. Moreover, because some elastic instabilities are accompanied by a sudden release of elastic energy, they can be harnessed to achieve fast dynamics. For example, insects such as click beetles use a snap-through instability to jump many times their body length by slowly storing elastic energy and then releasing it all at once.

The dynamics of an elastic system undergoing a fast instability may seem qualitatively simple: the system becomes unstable and jumps to a new equilibrium. However, surprises often hide in the transient dynamics. For example, the rate at which the bifurcation point is crossed can affect how long it takes for the system to jump (e.g. [1]) and, even more surprisingly, can bias the symmetry of the motion during the jump [2], as shown in Figure 3. Furthermore, recent work has shown that an unstable system can exhibit oscillations that look "stable" for long periods, before a sudden snap-through [3].

Understanding this richness is not just fundamentally interesting; it also has consequences in engineering. Engineers have turned to elastic instabilities to mimic nature [4] and achieve fast motion and jumps (for example in robots). However, the design of these mechanisms is often driven by trial-and-error, leaving optimisation to chance rather than a rational, predictive approach.

This dissertation will focus on developing a mathematical model of snap-through in a buckled arch, extending the work in [2, 3], to understand how loading history and initial conditions affect the transient behaviour during snap-through. We will use a range of theoretical and mathematical tools that may include multiple-scales analysis, eigenfunction decompositions of the governing PDEs, Galerkin reduction, the PyElastica code and time-integration via the method of lines. Crucially, the project admits both a theory-focused and a numerically focused approach, allowing flexibility and giving students the chance to choose the particular focus of most interest for them.

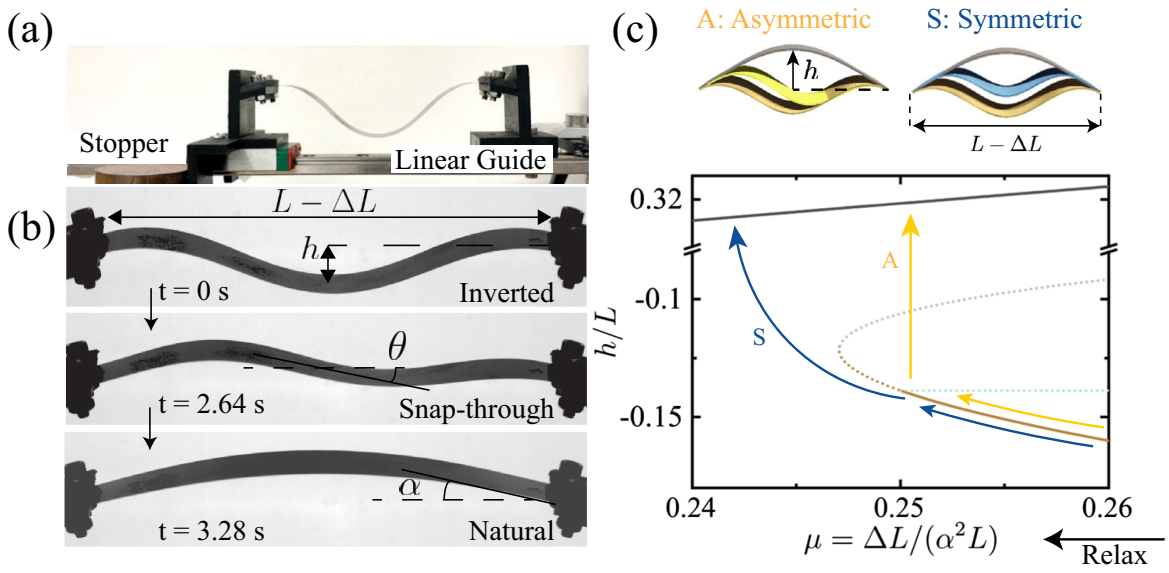

Figure 3:

Some knowledge of beam theory (for example from the Part C course "Elasticity and Plasticity") would be useful. However, any relevant background material can be taught as necessary.

References

[1] M. Liu, M. Gomez, and D. Vella. Delayed bifurcation in elastic snap-through instabilities. *Journal of the Mechanics and Physics of Solids*, 151:104386, 2021.

[2] Q. Wang, A. Giudici, W. Huang, Y. Wang, M. Liu, S. Tawfick, and D. Vella. Transient amplification of broken symmetry in elastic snap-through. *Physical review letters*, 132(26):267201, 2024.

[3] W. Simpkins, M. G. Hennessy, and M. Taffetani. Snap-through time of arches is controlled by slenderness and imperfections. *arXiv preprint:* arXiv:2508.10802, 2025.

[4] Y. Wang, Q. Wang, M. Liu, Y. Qin, L. Cheng, O. Bolmin, M. Alleyne, A. Wissa, R. H. Baughman, D. Vella, and S. Tawfick. Insect-scale jumping robots enabled by a dynamic buckling cascade. *Proceedings of the National Academy of Sciences*, 120(5):e2210651120, 2023.

2.6 Neural Field Pattern Persistence

Supervisor: Prof. Mohit Dalwadi
Contact: mohit.dalwadi@maths.ox.ac.uk

Neural field modelling provides a mathematical framework to understand how large-scale patterns of brain activity emerge from spatially structured neural interactions. By describing the evolution of averaged neural activity across space and time in a continuous manner, neural field models bridge the gap between detailed single-neuron dynamics and macroscopic measurements such as electroencephalography (EEG) and Functional Magnetic Resonance Imaging (fMRI). Unlike non-spatial models, neural

field models can naturally capture phenomena including travelling waves, localized activity bumps associated with working memory, and pattern formation arising from excitatory-inhibitory connectivity [1].

In this project, you will analyse a (nonlocal) integro-differential equation model for neural activity. In the absence of recurrent activity, the local dynamics decays to a rest state. Recurrent activity in the network, however, can drive pattern formation via a non-local Turing instability, as well as travelling and persistent localised activity (colloquially known as a 'bump attractor'). Here we will focus first on pattern formation, using numerical simulations as well as analytical tools such as stability and bifurcation analysis. In particular, we will focus on understanding how a more realistic time-varying firing function affects the onset of instability compared to constant a constant firing function. This will involve using novel techniques recently developed to understand rate-induced ('delayed') bifurcations and so-called R-tipping in PDEs [2], and extending them to analyse integro-differential equations. There are a number of different directions in which the project could go, depending on the student's interests.

This project aims to improve our understanding of neural activity by developing and applying novel techniques in multiscale applied mathematics. A successful project would provide a mathematical theory and biological interpretation of the importance of rate-induced effects in neural activity.

A key prerequisite for this project is a willingness to learn technical asymptotic analysis. This project involves working with Kyle Wedgwood from the University of Exeter.

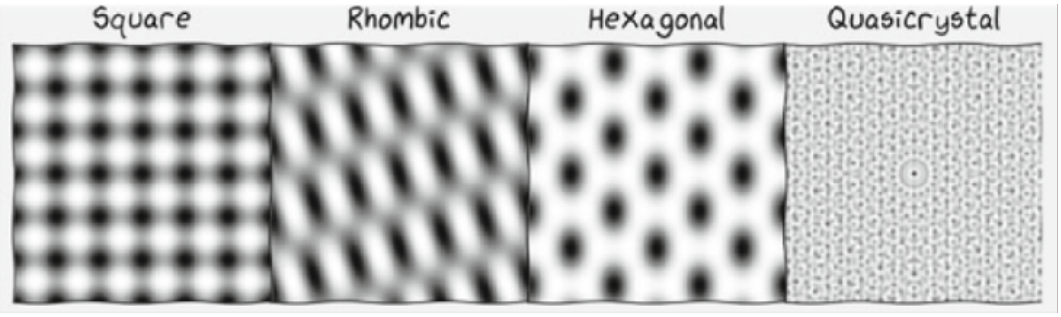

Figure 4: Examples of the pattern types observed in neural field models [1].

References

[1] S. Coombes and K. C. Wedgwood. *Neurodynamics*. Texts in Applied Mathematics. 2023.

[2] M. P. Dalwadi and P. Pearce. Universal dynamics of biological pattern formation in spatio-temporal morphogen variations. *Proceedings of the Royal Society A*, 479(2271):20220829, 2023.

2.7 A Mathematical Model to Understand Surging Glaciers

Supervisor: Prof. Ian Hewitt
Contact: ian.hewitt@maths.ox.ac.uk

Many of the world's glaciers and ice sheets undergo an intriguing behaviour in which they undergo quasi-periodic 'surges', during which they move much faster (∼100 times faster) than their usual

pace. Despite lots of attempts to understand what causes a surge, and why some glaciers surge while some do not, there are still no clear answers to those questions. An improved understanding of this behaviour would help us to better interpret the rapid glacier changes that are currently being observed in response to climate change.

A recent paper (Benn et al [1]) suggested a new theory for glacier surges, based on principles of mass and energy conservation. That paper described a 'lumped' ODE model for surging behaviour, which took the form of a dynamical system for the average ice thickness H and velocity u. The system can be analysed using phase-plane methods. The intention for this project is to extend the ideas of that study in one of a number of different ways, depending on the interest of the student.

One option is to extend the model to be spatially resolved, modelling the evolution along the length of the glacier. This would involve constructing a PDE model for the flow of a glacier and the evolution of its energy content. It would solve the model — probably numerically using either a finite element or a finite volume discretisation — and look to examine whether the behaviour seen in the lumped ODE model carries over to the spatially-resolved model.

Another option is to understand how surging behaviour interacts with the seasons. Glaciers melt predominantly during summer when a large amount of melt water is produced, which is known to affect their speed. This seasonal cycle is not included in the existing model, and the project would look to find a way to include it and investigate its effect.

This project would require interest in mathematical modelling, as well as (depending on the direction chosen) phase planes and bifurcation analysis, PDEs, numerical methods. It would benefit from some knowledge of fluid mechanics, though this is not essential.

References

[1] D. I. Benn, A. C. Fowler, I. Hewitt and H. Sevestre. A general theory of glacier surges. *J. Glaciology*, 2019. DOI: 10.1017/jog.2019.62

[2] D. I. Benn, I. J. Hewitt and A. Luckman. 2023 Enthalpy balance theory unifies diverse glacier surge behaviour. *Ann. Glaciol.*, 2023. DOI: 10.1017/aog.2023.23

[3] I. J. Hewitt. *C5.11 Mathematical Geoscience lecture notes (Chapter 3)*. Mathematical Institute, University of Oxford.

2.8 Pipe Formation in Wet Snow / Thermal Homogenisation for Porous Media Flow

Supervisor: Prof. Ian Hewitt
Contact: ian.hewitt@maths.ox.ac.uk

This project will develop a model to understand the phenomenon of 'piping' in a melting snowpack. Snow usually melts from the upper surface due to solar radiation. The water that is produced infiltrates the porous snow, where some of it refreezes and warms the subsurface. Some of the water may flow all the way through to the base. The infiltration is observed to be spatially heterogeneous, with much of the water flowing through narrow channels or 'pipes', while neighbouring snow is essentially devoid of water. It is of interest to understand firstly, how and why the pipes form, and secondly, what

determines the resulting macroscopic permeability. This is of relevance to climate scientists, because the depth to which meltwater infiltrates determines how likely it is to run-off and contribute to sea level.

The goal of the project is to use a continuum model to investigate the instability by which pipes form. One option would be to study the derivation of such a model using a formal homogenisation procedure. This would help to understand the conditions under which it is appropriate to treat the water and ice as having the same temperature, or when they should be treated as having two different temperatures. This would involve asymptotic methods applied to PDEs, and is expected to lead to a two-phase model of an unsaturated porous medium, with flow driven by capillary pressure and gravity, similar to the model described by Moure et al [3].

Another option is to start with the continuum model, and perform a linear stability analysis of one-dimensional solutions, to uncover the different modes of fingering that can occur. This could be compared with numerical solutions to study the non-linear development of pipes.

References

[1] C. R. Meyer and I. J. Hewitt. A continuum model for meltwater flow through compacting snow. *The Cryosphere*, 2017. DOI: 10.5194/tc-2017-128

[2] D. A. DiCarlo. Stability of gravity-driven multiphase flow in porous media: 40 Years of advancements. *Water Resources Research*, 49:4531–4544, 2013. DOI: 10.1002/wrcr20359

[3] A. Moure, N. Jones, J.,Pawlak, C. Meyer, and X. Fu. 2023 A Thermodynamic Nonequilibrium Model for Preferential Infiltration and Refreezing of Melt in Snow. *Water Resources Research*, 2023. DOI: 10.1029/2022WR034035

2.9 Visco-plastic Model for the Antarctic Ice Sheet

Supervisor: Prof. Ian Hewitt
Contact: ian.hewitt@maths.ox.ac.uk

Melting of the Earth's two great ice sheets, in Greenland and Antarctica, is likely to cause sea level rise of up to a metre over the next century; in total these ice sheets contain enough ice to raise sea levels globally by over 60m. The Antarctic ice sheet is thought to be particularly sensitive to increases in ocean temperature as well as air temperature, due to the fact that much of the ice sheet is grounded below sea level and therefore makes direct contact with the ocean. Modern numerical simulations of the ice sheet treat the ice as a non-Newtonian fluid, sliding over a rough bed. Different formulations and numerical discretisations of the problem appear to produce different results, and there is therefore significant uncertainty over the future behaviour of the ice sheet, even if the climatic forcing were assumed known.

This project will explore an alternative simpler model of the Antarctic ice sheet that treats the ice as a plastic material. The is somewhat different from the conventional models, but could provide a useful method to examine worst-case scenarios for how fast the ice sheet can collapse.

One option is to study the limit of a purely plastic model, in which case the problem reduces to solving a version of the eikonal equation for the ice thickness given the observed position of the boundary.

This problem is related to the problem of determining the shape of a sand-pile on a table of given shape, and involves solution of PDEs using a characteristic method. Starting with analytical solutions for simple geometries such as a circular or square ice sheet, an algorithm would be developed to solve for the ice thickness in more general geometries. If time allows, the method can be used to examine how the thickness of the ice sheet evolves over time in response to movement of the boundary.

A second option is to study a visco-plastic model in which the plastic flow is resisted by an effective viscosity. The limit of relatively small viscosity would be studied to understand how/when this model reduces to the purely plastic model, or to other approximate models. This option would make use of asymptotic expansions to derive simplified sets of equations that are appropriate for different climate forcings.

References

[1] J. F. Nye. The flow of glaciers and ice-sheets as a problem in plasticity. *Proc. Roy. Soc. A*, 207:554–572, 1951.

[2] N. Reeh. A plasticity theory approach to the steady-state shape of a three-dimensional ice sheet. *J. Glaciol.*, 28:431–455, 1982.

[3] L. Ultee and J. Bassis. The future is Nye: an extension of the perfect plastic approximation to tidewater glaciers. *J. Glaciol.*, 62:1143–1152, 2016. DOI: 10.1017/jog.2016.108

2.10 Rate-Induced Tipping of the Atlantic Meridional Overturning Circulation Under Linear and Periodic Freshwater Forcing

Supervisor: Prof. Irene Moroz
Contact: irene.moroz@maths.ox.ac.uk

A very recent paper, [1], investigates critical transitions of the Atlantic Meridional Overturning Circulation (AMOC) under non-stationary external forcing, based on a three-box salinity-advection model. The model of Neff et al [2] is analysed for the bifurcation structure of equilibria and their basins of attraction, clarifying how circulation regimes are connected through saddle-node bifurcations. Two types of time-dependent forcing are considered: a linear drift, in which the freshwater flux decreases linearly from an initial to a final value over finite time before remaining constant, and a cosine-type forcing representing seasonal variability. The evolution of trajectories, non-autonomous basins of attraction, basin sizes, and tipping probability matrices across different forcing rates are explored. The results show that AMOC multistability is highly sensitive to the forcing rate which reshapes basin sizes and governs cross-basin transfers, ultimately determining whether the AMOC undergoes irreversible collapse. The conclusion is that both the amplitude and the rate of external forcing must be considered when assessing AMOC stability and climate risks, rather that relying solely on static bifurcation structures.

The idea of the project is to investigate a greater range of forcing amplitudes and frequencies in the Neff et al model of AMOC, part of global thermohaline circulation. Then the analysis would be extended to the 3 cell Rooth model with both Temperature and Salinity variations, in which 2 of the cells are in the Arctic and Antarctic, the 3rd being at the Equator. The fate of the AMOC is a very pressing and topical issue.

References

[1] Jicheng Duan, Chris Budd, Irene Moroz, Zhouchao Wei, and Yiwen Chen. From Multistability to Collapse: Rate-Induced Tipping of the Atlantic Meridional Overturning Circulation Under Linear and Periodic Freshwater Forcing. *Journal of Nonlinear Science* (submitted), 2026.

[2] Alannah Neff, Andrew Keane, Henk A. Dijkstra, and Bernd Krauskopf. Bifurcation analysis of a North Atlantic Ocean box model with two deep-water formation sites. *Physica D*, 456:133907, 2023.

3 Numerical Analysis Projects

3.1 Multilevel Monte Carlo

Supervisor: Prof. Mike Giles
Contact: mike.giles@maths.ox.ac.uk

Multilevel Monte Carlo (MLMC) is a big research area for which I have many references available from my research webpage and my MLMC community webpage. I am open to suggestions for projects involving MLMC, but have the following specific ideas:

- use of MLMC for stochastic models of proton transport (for this medical treatment) based on this arXiv paper
- use of path branching in MLMC to address challenges with discontinuous functionals, based on this paper
- application of MLMC to stochastic PDEs, possibly co-supervised by Prof. Ana Djurdjevac, based in part on this paper
- use of adjoints and MLMC to compute sensitivities in mathematical finance — needs someone with prior knowledge of mathematical finance

3.2 High Performance Computing

Supervisor: Prof. Mike Giles
Contact: mike.giles@maths.ox.ac.uk

I also very willing to supervise projects involving High Performance Computing (HPC), using OpenMP on CPUs, or CUDA on GPUs, perhaps as part of a project with another supervisor.

One possible topic in HPC is parallel sorting which is very important in machine learning for randomising the training data. This is a heavily researched area, and yet I think there may be scope for improvement in the state-of-the-art under certain assumptions concerning the hardware being used, recognising that data movement is often the most time-consuming operation in a sorting algorithm.

3.3 Numerical Linear Algebra and Applications

Supervisor: Prof. Yuji Nakatsukasa
Contact: yuji.nakatsukasa@maths.ox.ac.uk

Description of proposal:
I would be happy to supervise MMSC dissertations on topics in (randomised) numerical linear algebra (NLA), and/or its applications in data science or scientific computing.

Depending on your preference, the focus can be theoretical, computational/algorithmic, or application oriented. Potential projects include

- Core randomised NLA, for example CUR and its applications, analysis and exploration of randomised SVD or least-squares problems, or Krylov subspace methods, possibly incorporating

mixed precision arithmetic [5].

- Parameter-dependent linear systems [7] and its use/analysis in e.g. pseudospectra computation or nonlinear eigenvalue problems. Relatedly, study the decay of singular values for matrices consisting of parameter-dependent vectors (e.g. [3, 8]). [with Dr Hussam Al Daas (STFC)]
- Efficient solution of the Helmholtz equation at many wave numbers [with Dr Carolina Urzúa Torres (TU Delft)]
- NLA techniques in data science and scientific computing, for example (i) kernel methods [6, 7], (ii) model order reduction [1] and rational approximation [9], (iii) NLA aspects of attention in transformers [11], (iv) studying double descent and overfitting [2, 4] from the perspective of Lebesgue constants [10].

References

[1] A. C. Antoulas, C. A. Beattie, and S. Güğercin. *Interpolatory Methods for Model Reduction.* SIAM, 2020.

[2] M. Belkin, D. Hsu, and J. Xu. Two models of double descent for weak features. *SIAM J. Math. Data Sci.*, 2(4):1167–1180, 2020.

[3] H. A. Daas and N. I. M. Gould. Extended-krylov-subspace methods for trust-region and norm-regularization subproblems. *arXiv preprint*, arXiv:2511.11135, 2025.

[4] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Ann. Stat.*, 50(2):949–986, 2022.

[5] N. J. Higham and T. Mary. Mixed precision algorithms in numerical linear algebra. *Acta Numerica*, 31:347–414, 2022.

[6] T. Hofmann, B. Schölkopf, and A. J. Smola. Kernel methods in machine learning. *Ann. Stat.*, 36(3):1171–1220, 2008.

[7] E. Jones and Y. Nakatsukasa. Subapsnap: Solving parameter-dependent linear systems with a snapshot and subsampling. *arXiv preprint*, arXiv:2510.04825, 2025.

[8] D. Kressner and C. Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. *SIAM J. Matrix Anal. Appl.*, 32(4):1288–1316, 2011.

[9] Y. Nakatsukasa and L. N. Trefethen. Applications of AAA rational approximation. *arXiv preprint*, arXiv:2510.16237, 2025. [To appear in Acta Numerica.]

[10] L. N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition.* SIAM, 2019.

[11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Adv. Neural Inf. Process. Syst.*, 30, 2017.

3.4 Automated Parallel Goal-Based Adaptivity

Supervisor: Prof. Patrick Farrell
Contact: patrick.farrell@maths.ox.ac.uk

Background and problem statement: When solving a mathematical problem numerically, it is crucial that the accuracy of the approximation be quantified. Scientists, engineers, and policymakers use numerical solutions for design and decisions; can they rely on those approximations?

Since the 1990s, a mathematical theory of *goal-oriented error estimation and adaptivity* for partial differential equations (PDEs) has been developed [1]. If the variational statement of our PDE problem is

$$u \in V : \quad F(u; v) = 0 \quad \text{for all } v \in V,$$

and our discrete approximation is

$$u_h \in V_h : \quad F(u_h; v_h) = 0 \quad \text{for all } v_h \in V_h,$$

for a discrete approximation space V_h, the aim of goal-based error estimation is to compute η such that

$$\eta \approx |J(u) - J(u_h)|$$

for a goal functional $J : V \to \mathbb{R}$. The aim of goal-based adaptivity is to efficiently design an adapted mesh with associated discrete function space V_h such that

$$\eta < \varepsilon$$

for a target approximation error $\varepsilon > 0$. As a concrete example, F might represent the Navier–Stokes equations, u the velocity, pressure, and temperature of the fluid, and J represent the drag coefficient of a candidate wing design.

In the literature goal-based adaptivity is essentially always applied on a case-by-case basis. In a previous (and very successful) MMSC project, we automated the application of goal-based adaptivity for stationary (i.e. not time-dependent) boundary value problems and eigenvalue problems in the finite element computing environment Firedrake [2]. In this project we now aim to extend this automation to time-dependent problems.

Planned approach: The dual-weighted residual method has been formulated for linear time-dependent problems in the classic review article of Becker & Rannacher [1]. This formulation relies on the use of Galerkin methods in both space and time; expressing Galerkin methods in time in Firedrake is now very convenient due to recent developments in its time discretisation module Irksome [3]. These methods require the solution of the corresponding adjoint problem, which is also automatically supported in Firedrake [4]. The setup and solution of the adjoint for time-dependent problems is highly nontrivial, as it propagates causality backwards in time and requires recording the entire forward trajectory for linearisation.

We will first verify that time-dependent problems expressed in Irksome can be automatically adjoined. We will then manually implement the estimators for linear problems discussed in Becker & Rannacher [1]. Once this is working, we will attempt to automate the logic for the derivation of the dual-weighted residual error estimator in time, and compare this to the manual approach. If time permits, we will then investigate applications to nonlinear problems, e.g. [5, 6].

Desired outcomes: One output of the MSc research project will be a software contribution to the Firedrake system making goal-based adaptivity for time-dependent problems available to non-expert users. An aspirational outcome would be to apply goal-based adaptivity to energy- and entropy-stable discretisations of the compressible Navier–Stokes equations.

Dependencies: This project depends on a good working knowledge of *Finite Element Methods for PDEs*.

References

[1] R. Becker, and R. Rannacher. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica*, 10:1–102,2001. DOI: 10.1017/S0962492901000010

[2] D. A. Ham, P. H. J. Kelly, L. Mitchell, C. J. Cotter, R. C. Kirby, K. Sagiyama, N. Bouziani, S. Vorderwuelbecke, T. J. Gregory, J. Betteridge, D. R. Shapero, R. W. Nixon-Hill, C. J. Ward, P. E. Farrell, P. D. Brubeck, I. Marsden, T. H. Gibson, M. Homolya, T. Sun, A. T. T. McRae, F. Luporini, A. Gregory, M. Lange, S. W. Funke, F. Rathgeber, G.-T. Bercea, and G. R. Markall. *Firedrake User Manual*, Imperial College London, University of Oxford, Baylor University, and University of Washington, 2023. DOI: 10.25561/104839

[3] P. E. Farrell, R. C. Kirby, and J. Marchena-Menendez. Irksome: automating Runge–Kutta time-stepping for finite element methods. *ACM Transactions on Mathematical Software*, 47(4), 2021. DOI: 10.1145/3466168

[4] P. E. Farrell, D. A. Ham, S. W. Funke, and M. E. Rognes. Automated derivation of the adjoint of high-level transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369–C393, 2013. DOI: 10.1137/120873558

[5] J. Roth, J. P. Thiele, U. Köcher, and T. Wick. Tensor-product space-time goal-oriented error control and adaptivity with partition-of-unity dual-weighted residuals for nonstationary flow problems. *Computational Methods in Applied Mathematics*, 24(1):185–214, 2023. DOI: 10.1515/cmam-2022-0200

[6] B. Endtmayer, U. Langer, and A. Schafelner. Goal-oriented adaptive space-time finite element methods for regularized parabolic p-Laplace problems. *Computers and Mathematics with Applications*, 167:286–297, 2024. DOI: 10.1016/j.camwa.2024.05.017

3.5 Time Discretisation of the Navier–Stokes–Onsager–Stefan–Maxwell Equations for Multicomponent Flow

Supervisor: Prof. Patrick Farrell
Contact: patrick.farrell@maths.ox.ac.uk

Background and problem statement:

Many fluids consist of mixtures; for example, air is a mixture of nitrogen, oxygen, carbon dioxide, and other species. In many situations, it is not necessary to resolve the motions of the individual species, such as when modelling the flow of air over an aircraft. However, in other contexts, detailed knowledge of the transport of individual species is required. Examples include biological applications, where one may be interested in the transport of oxygen and carbon dioxide in blood, in chemical

engineering, where one may be interested in separating or combining the constituents of petroleum, or in electrochemistry, where the performance of a lithium-ion battery is often limited by the transport of lithium ions within an electrolyte. We describe this situation as a *multicomponent flow*, where a fluid is composed of $2 \leq n \in \mathbb{N}_+$ distinct chemical species in a common thermodynamic phase.

If the mixture is dilute, the system of equations to be solved can be reduced to the Navier–Stokes equations coupled to n independent advection-diffusion equations. However, when the mixture is concentrated, then the model to be solved becomes the Navier–Stokes–Onsager–Stefan–Maxwell equations. These are substantially harder to solve (and also more interesting).

In a series of works my DPhil students and I have studied the formulation and discretisation of these equations [1, 2, 3, 4, 5, 6]. Each of these works have considered the stationary case, where the model is not time-dependent. With a reasonable understanding of the spatial discretisation now in hand, it is time to turn our attention to the time discretisation of the non-stationary model.

A key aspect of these equations is that they enjoy a great deal of *structure*. The mass-average constraint relates the barycentric velocity and the species fluxes; the Gibbs–Duhem relation constrains the diffusion driving forces; the partial densities should be nonnegative and sum to the total density. When we move to the time-dependent problem, two crucial thermodynamic principles arise: in the isothermal case, the energy of an unforced system should be dissipated; in the anisothermal case, the energy should be conserved, and entropy should be dissipated. Different formulations at the continuous level expose different opportunities and challenges for structure-preservation on discretisation.

In this project we will consider the time discretisation of the Navier–Stokes–Onsager–Stefan–Maxwell equations, with a view to devising time discretisations that preserve thermodynamic structure.

Planned approach: In recent work Boris Andrews and I have devised a framework for devising arbitrary-order time discretisations that preserve conservation and dissipation laws [7, 8]. This will be the main tool we use to devise structure-preserving discretisations. A lowest-order energy-stable discretisation for a closely related system was recently considered by Brunk et al. [9]; a possible question is to investigate whether our framework applied to this system yields the same discretisation at lowest order.

We should begin with the simplest possible setting, the Onsager–Stefan–Maxwell equations with a specified bulk velocity (decoupling from the Navier–Stokes equations) for an ideal mixture of ideal gases. Depending on progress here we could go to ideal mixtures of non-ideal gases, non-ideal mixtures, or to the coupled isothermal Navier–Stokes–Onsager–Stefan–Maxwell equations. If extremely successful we could proceed to consider anisothermal problems where both energy- and entropy-stability should be enforced.

A starting point is the formulation and discretisation in the stationary case. My DPhil students Aaron Baier-Reinio and Kars Knook have recently written a demo solving for a mixture of two hydrocarbons, benzene and cyclohexane, with the Firedrake finite element system [10]:

https://www.firedrakeproject.org/demos/multicomponent.py.html

Desired outcomes: A basic goal would be to successfully implement some standard time discretisations for a model. An extremely ambitious goal would be to devise and implement energy- and entropy-stable discretisations for the anisothermal Navier–Stokes–Onsager–Stefan–Maxwell equations.

Dependencies: This project depends on a good working knowledge of *Finite Element Methods for*

PDEs.

References

[1] A. J. Van-Brunt, P. E. Farrell, and C. W. Monroe. Augmented saddle point formulation of the steady-state Stefan–Maxwell diffusion equations. *IMA Journal of Numerical Analysis*, 42:3272–3305, 2022. DOI: 10.1093/imanum/drab067

[2] A. Van-Brunt, P. E. Farrell, and C. W. Monroe. Consolidated theory of fluid thermodiffusion. *AIChE Journal*, 68(5):e17599, 2021. DOI: 10.1002/aic.17599

[3] F. R. A. Aznaran, P. E. Farrell, C. W. Monroe, and A. J. Van-Brunt. Finite element methods for multicomponent convection-diffusion. *IMA Journal of Numerical Analysis*, 45(1):188–222, 2024. DOI: 10.1093/imanum/drae001

[4] A. Van-Brunt, P. E. Farrell, and C. W. Monroe. Structural electroneutrality in Onsager–Stefan–Maxwell models with charged species. *Electrochimica Acta*, 441: 141769, 2022. DOI: 10.1016/j.electacta.2022.141769

[5] A. Baier-Reinio and P. E. Farrell. High-order finite element methods for three-dimensional multicomponent convection-diffusion. *SIAM Journal on Scientific Computing*, 2025. DOI: 10.48550/arxiv.2408.17390

[6] A. Baier-Reinio, P. E. Farrell, and C. W. Monroe. Finite element methods for electroneutral multicomponent electrolyte flows. *arXiv preprint*, arXiv:2510.14923, 2025.

[7] B. D. Andrews and P. E. Farrell. Enforcing conservation laws and dissipation inequalities numerically via auxiliary variables. *SIAM Journal on Scientific Computing*, 47(6), 2025. DOI: 10.1137/25M1756673

[8] B. D. Andrews and P. E. Farrell. Conservative and dissipative discretisations of multi-conservative ODEs and GENERIC systems. *arXiv preprint*, arXiv:2511.23266, 2025. [Submitted to Computers and Mathematics with Applications.]

[9] A. Brunk, A. Jüngel, and M. Lukáčová-Medvidová. A structure-preserving numerical method for quasi-incompressible Navier–Stokes–Maxwell–Stefan systems. *arXiv preprint*, arXiv:2504.11892, 2025.

[10] D. A. Ham, P. H. J. Kelly, L. Mitchell, C. J. Cotter, R. C. Kirby, K. Sagiyama, N. Bouziani, S. Vorderwuelbecke, T. J. Gregory, J. Betteridge, D. R. Shapero, R. W. Nixon-Hill, C. J. Ward, P. E. Farrell, P. D. Brubeck, I. Marsden, T. H. Gibson, M. Homolya, T. Sun, A. T. T. McRae, F. Luporini, A. Gregory, M. Lange, S. W. Funke, F. Rathgeber, G.-T. Bercea, and G. R. Markall. *Firedrake User Manual*, Imperial College London, University of Oxford, Baylor University, and University of Washington, 2023. DOI: 10.25561/104839

3.6 Efficient Algorithms for Infinite-Dimensional Spectral Computations

Supervisor: Prof. Patrick Farrell
Contact: patrick.farrell@maths.ox.ac.uk

Background and problem statement: The task of computing the spectrum

$$\mathrm{Sp}(A) = \left\{ z \in \mathbb{C} \mid (A - zI)^{-1} \text{ does not exist as a bounded operator } \right\} \tag{1}$$

of a linear operator A on a Hilbert space H is a fundamental problem in mathematics. At least twenty Nobel Prizes in Physics have come from phenomena associated with eigenvalues, which are a subset of the spectrum [1].

Computing spectra in infinite dimensions is much, *much* harder than in finite dimensions. This is primarily because fundamental new phenomena arise in infinite dimensions that cannot occur in finite dimensions. Whereas in finite dimensions the spectrum of an $n \times n$ matrix consists only of n isolated points (eigenvalues), in infinite dimensions the spectrum can be continuous, and form e.g. entire intervals, the unit disk, or Cantor-type sets.

The standard approach to computing spectra of infinite-dimensional problems is to discretise them, and then to compute the eigenvalues of the discretised operator, hoping that they approach those of the underlying one [2]. However, this is fraught with difficulties. If the problem is not discretised very carefully, *spectral pollution* can occur, where the discretised operator has eigenvalues that converge to something *not* in the spectrum of the original one [3, 4]. There are other failure modes, like *spectral invisibility*, where parts of the spectrum are always neglected, no matter how fine the resolution is taken. Furthermore, with these approaches, it is extremely difficult or impossible to compute those continuous parts of the spectrum that are structurally excluded in finite dimensions.

In recent work, Matt Colbrook in Cambridge has proposed a new algorithm for computing the spectra of infinite-dimensional operators that (in certain cases) can be proven to converge to the correct spectrum, so that neither spectral pollution nor invisibility occurs, and that is able to correctly compute continuous spectra [4, 5, 6, 7, 8]. The key idea is to take a grid of points in $\mathbb{C}$, and for each point compute the lowest eigenvalue of the (square, non-negative, self-adjoint) operator

$$(A - zI)^*(A - zI). \tag{2}$$

This gives an upper bound on the distance from z to $\mathrm{Sp}(A)$. This approach is, I believe, the right outer loop for computing spectra. However, the algorithm is expensive, because there are many expensive inner subproblems, and the bound acquired can be poor for coarse discretisations.

Many open questions arise: what information from one z can be recycled for the eigensolves of other nearby points? Can multilevel solvers for eigenproblems [9, 10, 11, 12] be fruitfully applied in this context? Can we use adaptive eigensolver ideas to compute the distance to the spectrum with a guarantee of accuracy [13, 14, 15]? Algorithms for computing pseudospectra have a similar flavour, and many ideas here can be extended to this context [16].

Planned approach: We will implement Colbrook's algorithm and make a thorough literature review of ideas that might be useful. We will then consider different solver strategies for the subproblem at every z. The first problem to attack will be the famous Maxwell eigenproblem, which if discretised naïvely exhibits spectral pollution [3].

Desired outcomes:

One possible outcome would be an adaptive multilevel solver for the inner subproblem that automatically and efficiently computed the bound to the spectrum to within a guaranteed error tolerance.

Dependencies:

This project depends on a good working knowledge of *Finite Element Methods for PDEs* and *Numerical Linear Algebra*.

References

[1] L. N. Trefethen. Favorite eigenvalue problems. *SIAM News*, 44(10), 2011-12.

[2] D. Boffi. Finite element approximation of eigenvalue problems. *Acta Numerica*, 19:1–120, 2010. DOI: 10.1017/S0962492910000012

[3] D. N. Arnold. *Finite element exterior calculus*. Society for Industrial and Applied Mathematics, 2018. DOI: 10.1137/1.9781611975543

[4] M. J. Colbrook. *Infinite-Dimensional Spectral Computations: Foundations, Algorithms, and Modern Applications*. Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, 2026. [Forthcoming in August 2026.] https://www.cambridge.org/9781009382526

[5] M. J. Colbrook, B. Roman, and A. C. Hansen. How to compute spectra with error control. *Physical Review Letters*, 122(25), 2019. DOI: 10.1103/physrevlett.122.250201

[6] M. J. Colbrook, A. Horning, and A. Townsend. Computing spectral measures of self-adjoint operators. *SIAM Review*, 63(3):489–524, 2021. DOI: 10.1137/20m1330944

[7] M. J. Colbrook. On the computation of geometric features of spectra of linear operators on Hilbert spaces. *Foundations of Computational Mathematics*, 24(3):723–804, 2022. DOI: 10.1007/s10208-022-09598-0

[8] C. Drysdale, M. J. Colbrook, and M. T. M. Woodley. Computation and verification of spectra for non-Hermitian systems. *Physical Review Letters*, 135(17), 2025. DOI: 10.1103/hf8s-v93b

[9] Z. Cai, J. Mandel, and S. McCormick. Multigrid methods for nearly singular linear equations and eigenvalue problems. *SIAM Journal on Numerical Analysis*, 34(1):178–200, 1997. DOI: 10.1137/s1064827594261139

[10] A. V. Knyazev and K. Neymeyr. Efficient solution of symmetric eigenvalue problems using multigrid preconditioners in LOBPCG. *Electronic Transactions on Numerical Analysis*, 15:38–55, 2003.

[11] A. V. Knyazev, M. E. Argentati, I. Lashuk, and E. E. Ovtchinnikov. Block locally optimal preconditioned eigenvalue xolvers (BLOPEX) in Hypre and PETSc. *SIAM Journal on Scientific Computing*, 29(5):2224–2239, 2007. DOI: 10.1137/060661624

[12] H. Chen, H. Xie, and F. Xu. A full multigrid method for eigenvalue problems. *Journal of Computational Physics*, 322:747–759, 2016. DOI: 10.1016/j.jcp.2016.07.009

[13] R. Becker, and R. Rannacher. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica*, 10:1–102,2001. DOI: 10.1017/S0962492901000010

[14] T. Vejchodský. Three methods for two-sided bounds of eigenvalues—a comparison. *Numerical Methods for Partial Differential Equations*, 34(4):1188–1208, 2018. DOI: 10.1002/num.22251

[15] S. Giani, L. Grubišić, H. Hakula, and J. S. Ovall. A posteriori error estimates for elliptic eigenvalue problems using auxiliary subspace techniques. *Journal of Scientific Computing*, 88(3), 2021. DOI: 10.1007/s10915-021-01572-2

[16] L. N. Trefethen and M. Embree. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, 2005.

3.7 Thermodynamically-Consistent Closures of the Boltzmann Equation

Supervisor: Prof. Patrick Farrell
Contact: patrick.farrell@maths.ox.ac.uk

Background and problem statement: The Boltzmann equation

$$\partial_t f + v \cdot \nabla_x f = Q(f, f), \tag{3}$$

is the fundamental equation of statistical mechanics. Here $f(\vec{x}, \vec{v}, t) : \mathbb{R}^3 \times \mathbb{R}^3 \times [0, \infty) \to \mathbb{R}$ denotes the distribution function describing the probability density of finding a particle at position $\vec{x}$ with velocity $\vec{v}$ at time t, and Q is a quadratic operator modelling particle collisions. The Boltzmann equation models the evolution of this probability density as the molecules in a gas collide and approach thermodynamic equilibrium.

The Boltzmann equation is the foundation for deriving continuum models for fluids, like the famous Euler and Navier–Stokes equations. The main tool for doing this is the Chapman–Enskog expansion [1], which is a formal asymptotic expansion of the distribution function f in terms of a small parameter ε, known as the Knudsen number, i.e.

$$f = f_0 + \varepsilon f_1 + \varepsilon^2 f_2 + \cdots, \tag{4}$$

where f_0 is the so-called local Maxwellian distribution, which represents the thermodynamic equilibrium state of the system and is defined by $Q(f_0, f_0) = 0$, and f_1, f_2, ... are the non-equilibrium corrections to the distribution function.

When truncated at zeroth order $f \approx f_0$, the Chapman–Enskog expansion leads to the compressible Euler equations. When truncated at first order $f \approx f_0 + \varepsilon f_1$, it leads to the compressible Navier–Stokes equations. Both of these are very successful and widely used models. When truncated at second order $f \approx f_0 + \varepsilon f_1 + \varepsilon^2 f_2$, it leads to the so-called Burnett equations [2]. However, the Burnett equations suffer from a fundamental flaw: the rate of entropy production can be negative, and they are not thermodynamically consistent. This manifests as various kinds of instabilities and ill-posedness of the equations.

It is deeply unsettling that the asymptotic procedure that furnishes the Euler and Navier–Stokes equations next gives something that violates the laws of physics. Can we find an alternative construction of hydrodynamic equations that furnishes the Euler and Navier–Stokes equations, and then furnishes something physical?

Planned approach: We plan to explore different approaches for deriving thermodynamically-consistent higher-order continuum models. Our first approach will be to combine the Levermore moment-closure approach with the framework of Rajagopal and Srinivasa for deriving constitutive relations. The Levermore moment-closure approach [3] guarantees non-negativity of entropy production, but the resulting expression for the distribution function is very difficult to work with. The Rajagopal–Srinivasa approach [4] takes as input two scalar expressions for the entropy production and internal energy and uses these in a constrained optimisation problem to derive all constitutive relations. Can these approaches be hybridised in a Levermore–Rajagopal–Srinivasa method, where the Levermore moment closure is used only to compute the entropy production and internal energy, and the Rajagopal–Srinivasa optimisation problem is then used to derive the resulting constitutive relations?

Desired outcomes: A desired outcome would be a thermodynamically-consistent alternative to the Burnett equations. If successful, in future work it would be of great interest to study the numerical discretisation of these equations.

Dependencies: This course depends on a good working knowledge of *Finite Element Methods for PDEs*, and background in statistical mechanics and fluid mechanics. The student undertaking this project would ideally *love* computing tricky integrals.

References

[1] S. Chapman and T. G. Cowling. *The mathematical theory of non-uniform gases: an account of the kinetic theory of viscosity, thermal conduction and diffusion in gases.* Cambridge University Press, 3rd edition, 1970.

[2] D. Burnett. The distribution of molecular velocities and the mean motion in a non-uniform gas. *Proceedings of the London Mathematical Society*, s2-40(1):382–435, 1936. DOI: 10.1112/plms/s2-40.1.382

[3] C. D. Levermore. Moment closure hierarchies for kinetic theories. *Journal of Statistical Physics*, 83(5):1021–1065, 1996. DOI: 10.1007/BF02179552

[4] K. R. Rajagopal and A. R. Srinivasa. On thermomechanical restrictions of continua. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 460(2042):631–651, 2004. DOI: 10.1098/rspa.2002.1111

3.8 Finite Element Discretisations for Magnetic Advection–Diffusion Eigenvalue Problems

Supervisors: Prof. Kaibo Hu and Dr Jindong Wang
Contact: kaibo.hu@maths.ox.ac.uk and jindong.wang@maths.ox.ac.uk

Background and problem statement: For Maxwell's equations, a wide range of finite element discretisations have been developed and analysed [1, 2]. These methods play a central role in computational electromagnetics. In the context of the associated eigenvalue problem,

$$\nabla \times \nabla \times \boldsymbol{B} = \lambda \boldsymbol{B}, \tag{5}$$

substantial progress has been made in understanding its spectral approximation properties [3, 4]. It is well known that certain discretisations, such as those based on continuous Lagrangian finite elements, may lead to spurious eigenvalues and nonphysical modes.

In magnetohydrodynamics (MHD) and dynamo theory, a closely related but more challenging problem arises from the magnetic advection–diffusion eigenvalue problem. The issue becomes even more critical in the dynamo context, where the accurate computation of magnetic eigenmodes is directly linked to physically observable phenomena. The central question in dynamo theory is whether a given conducting flow can amplify and sustain a magnetic field against diffusion [5, 6]. This mechanism underlies the generation of magnetic fields in astrophysical and geophysical systems, such as stars, planets, and galaxies.

Mathematically, the problem is often formulated as

$$-\nabla \times (\boldsymbol{\beta} \times \boldsymbol{B}) + R_m^{-1} \nabla \times \nabla \times \boldsymbol{B} = \lambda \boldsymbol{B}, \tag{6}$$

where $\boldsymbol{\beta}$ denotes a prescribed velocity field. Owing to the presence of the advection term, the resulting operator is nonsymmetric, and its spectral properties are significantly more complicated than in the classical Maxwell case. There is also an open issue highlighted in [7],

The physically meaningful magnetic Reynolds numbers R_m are of order 10^8. The corresponding matrices are (and will remain) beyond the reach of any computer.

This highlights the computational difficulty of the problem at physically relevant parameter regimes. Recent studies also indicate that, if not discretised appropriately, this problem may suffer from similar difficulties, including spurious spectral components even for relatively small Reynolds numbers. However, the precise mechanisms underlying these numerical difficulties, as well as their dependence on discretisation choices, remain insufficiently understood.

Planned approach: The project will explore and compare various finite element discretisations for the magnetic advection–diffusion eigenvalue problem. Numerical experiments will also be conducted to investigate the spectral behaviour for different choices of the velocity field $\boldsymbol{\beta}$. By systematically comparing different discretisation strategies, the project aims to gain insight into how numerical choices influence the accuracy, stability, and reliability of the computed eigenvalues and eigenfunctions.

Expected outcomes: The expected outcome of this project is a systematic comparison of various finite element discretisations for the magnetic advection–diffusion eigenvalue problem under different velocity fields $\boldsymbol{\beta}$. The project will benchmark the numerical solutions and characterise the performance of different methods.

References

[1] Peter Monk. Finite element methods for Maxwell's equations. *Oxford University Press*, 2003.

[2] Ralf Hiptmair. Finite elements in computational electromagnetism. *Acta Numerica*, 11:237–339, 2002.

[3] Snorre H. Christiansen. Foundations of finite element methods for wave equations of Maxwell type. *Applied Wave Mathematics: Selected Topics in Solids, Fluids, and Mathematical Methods, Springer, Berlin, Heidelberg*, pp. 335–393, 2009.

[4] Daniele Boffi. Finite element approximation of eigenvalue problems. *Acta Numerica*, 19:1–120, 2010.

[5] S. Childress and A. D. Gilbert, Stretch, Twist, Fold: The Fast Dynamo, *Springer, Berlin, Heidelberg*, 1995.

[6] V. I. Arnold, Some remarks on the antidynamo theorem, Vladimir I. Arnold – Collected Works: Hydrodynamics, Bifurcation Theory, and Algebraic Geometry 1965–1972, *Springer, Berlin, Heidelberg*, pp. 387–396, 2013.

[7] V. I. Arnold and B. A. Khesin, Topological Methods in Hydrodynamics, Vol. 19, *Springer, New York*, 2009.

3.9 Dynamical Evolution of Topological Magnetic Configurations

Supervisors: Prof. Kaibo Hu and Dr Boris Andrews
Contact: kaibo.hu@maths.ox.ac.uk and boris.andrews@maths.ox.ac.uk

Magnetic relaxation, i.e. time evolution of magnetic fields in conducting fluids, is a fundamental problem in magnetohydrodynamics (MHD), with important applications in astrophysics and fusion science. When magnetic field lines possess nontrivial topological structures, such as knots, links, or braids, their evolution becomes highly constrained by topology. A central quantity is magnetic helicity

$$H = \int_\Omega A \cdot B \mathrm{d}x,$$

which measures the average linking of field lines and is conserved in ideal MHD. However, helicity captures only part of the topology: configurations such as the Borromean rings may have vanishing helicity despite nontrivial linking.

Recent structure-preserving finite element methods, particularly those based on finite element exterior calculus (FEEC), allow discretizations that preserve divergence constraints and helicity. Surprisingly, numerical schemes that preserve different levels of topology can produce qualitatively different relaxation behavior. It is therefore not clear how general topological configurations evolve under magnetic relaxation, nor how strongly the discrete enforcement of topology influences the computed dynamics.

The goal of this project is twofold. First, we investigate, through numerical experiments, the evolution of various knotted and linked magnetic configurations under the magneto-friction model, a gradient-flow approximation of MHD. Beyond classical examples such as Hopf fields and braids, we consider configurations including Borromean rings and other higher-order links. Second, we study how numerical algorithms that preserve different levels of topological structure behave on these configurations, and how this affects long-time dynamics and equilibrium states.

By systematically comparing divergence-preserving, globally helicity-preserving, and locally helicity-preserving schemes, we aim to clarify the interplay between topology, relaxation, and structure-preserving discretization in MHD.

Pre-requisites include knowledge of finite element methods and applied PDEs. Understanding some topology (particularly knot theory) is a plus, but not necessary.

References

[1] V. I. Arnold and B. A. Khesin, Topological Methods in Hydrodynamics, Vol. 19, *Springer, New York*, 2009.

[2] K. Hu, Y. J. Lee, and J. Xu. Helicity-conservative finite element discretization for incompressible MHD systems. *Journal of Computational Physics*, 436:110284, 2021.

[3] M. He, P. E. Farrell, K. Hu, and B. D. Andrews. Helicity-preserving finite element discretization for magnetic relaxation. *arXiv preprint*, arXiv:2501.11654, 2025.

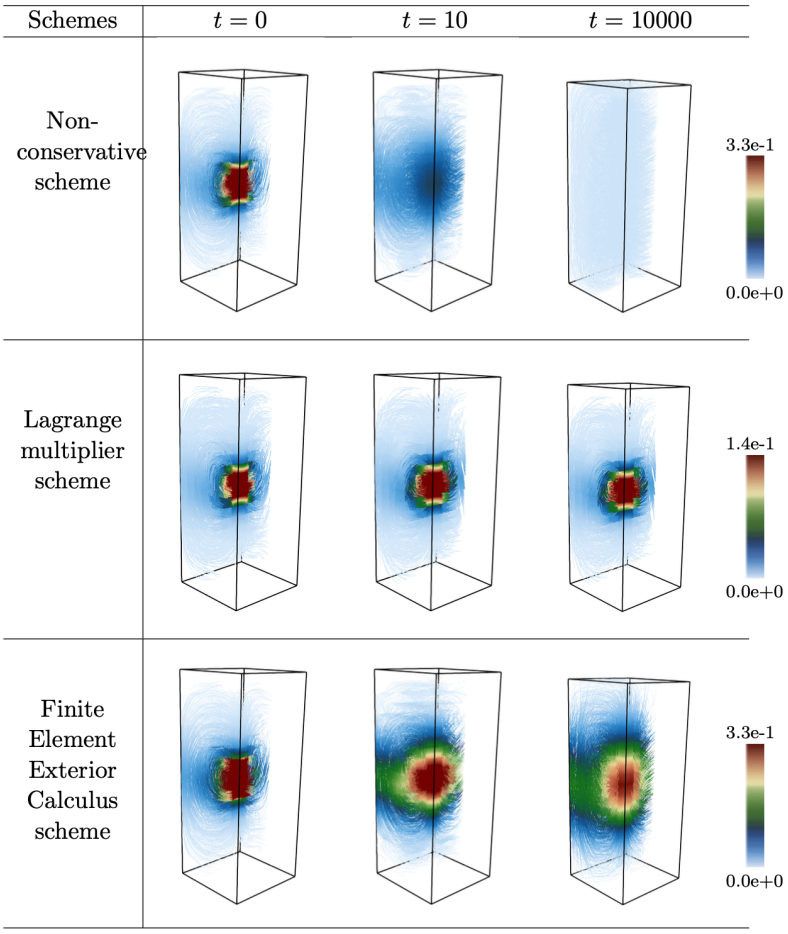

Figure 5: *Magnetic knots*: morphology comparison of magnetic field evolution under different topological constraints.

3.10 Numerical Simulation of Elastic Bouncy Objects

Supervisor: Dr Ioannis Papadopoulos
Contact: ioannis.papadopoulos@maths.ox.ac.uk

[Availability: I will be away from Oxford until early June. However, I will be happy to meet online before my return.]

Background and problem statement: Imagine throwing a tennis ball as hard as you can into the ground. If you could slow down time at the moment of impact, you would see the tennis ball deform, become less round, before returning to its original shape and bouncing back up. There are three physical behaviours one must model:

1. The deformability of the tennis ball, normally described by a theory of elasticity;
2. The non-penetration condition, i.e. the fact that the tennis ball cannot pass through the ground;

3. Newton's third law describing the forces exerted on the ball.

Once modelled, we want to numerically approximate the solution. This requires solvers for the nonlinear physics (which we can recast as constrained optimisation problems) coupled with a discretisation so that we can solve these problems on our computer.

Planned approach: As a first step we will model the ball with linear elasticity which will be discretised with finite elements in Firedrake [1] (a Python-based finite element library) using their Netgen [2] support and new SubMesh features. The non-penetration conditions will be modelled as a Signorini problem (the correct way to model them!). Signorini problems are actually constrained infinite-dimensional problems in disguise. Namely, one seeks a displacement $u : \Omega \to \mathbb{R}^d$ minimising

$$\min_{u \in H^1(\Omega;\mathbb{R}^d)} \int_\Omega \mu|\varepsilon(u)|^2 + \frac{\lambda}{2}|\nabla \cdot u|^2 \,\mathrm{d}x \;\text{ subject to }\; u \cdot n \geq 0 \;\text{ on }\; \Gamma_C + \text{bcs},$$

where $\epsilon = (\nabla + \nabla^\top)/2$ is the symmetrised gradient and Γ_C is the contact boundary. These optimisation problems will be then be solved via a powerful new technique called proximal Galerkin [3, 4]. The first goal is to numerical simulate a 2D linear elastic ball compressed into the ground. Next we can add the dynamics of bouncing. Once that is achieved, we would generalise this to 3D, consider multimaterial balls (like the empty air in a tennis ball), and even a slightly deformable ground. This project can also take a numerical linear algebra route, investigating fast linear system solvers for those that arise.

Desired outcomes: A desired outcome is a working simulation of a 3D ball bouncing on a rigid floor. If successful, we can consider how to speed up the solver, investigate other elasticity models, and perhaps publish in a scientific computing journal.

Prerequisites: Knowledge of the finite element method and Python would be very beneficial.

References

[1] D. A. Ham, P. H. J. Kelly, L. Mitchell, C. J. Cotter, R. C. Kirby, K. Sagiyama, N. Bouziani, S. Vorderwuelbecke, T. J. Gregory, J. Betteridge, D. R. Shapero, R. W. Nixon-Hill, C. J. Ward, P. E. Farrell, P. D. Brubeck, I. Marsden, T. H. Gibson, M. Homolya, T. Sun, A. T. T. McRae, F. Luporini, A. Gregory, M. Lange, S. W. Funke, F. Rathgeber, G.-T. Bercea, and G. R. Markall. *Firedrake User Manual*, Imperial College London, University of Oxford, Baylor University, and University of Washington, 2023. DOI: 10.25561/104839

[2] Jack Betteridge, Patrick Farrell, Matthias Hochsteger, Christopher Lackner, Joachim Schöberl, Stefano Zampini, and Umberto Zerbinati. ngsPETSc: A coupling between NETGEN/NGSolve and PETSc. *Journal of Open Source Software*, 9(104), 2024.

[3] Brendan Keith and Thomas M. Surowiec. Proximal Galerkin: A structure-preserving finite element method for pointwise bound constraints. *Foundations of Computational Mathematics*, Springer, pp.1–97, 2024. DOI: s10208-024-09681-8

[4] Jørgen S. Dokken, Patrick E. Farrell, Brendan Keith, Ioannis P. A. Papadopoulos, and Thomas M. Surowiec. The latent variable proximal point algorithm for variational problems with inequality constraints. *Computer Methods in Applied Mechanics and Engineering*, 445: 118181, 2025.

3.11 Large-Scale Topology Optimisation via Preconditioning

Supervisor: Dr Ioannis Papadopoulos
Contact: ioannis.papadopoulos@maths.ox.ac.uk

[Availability: I will be away from Oxford until early June. However, I will be happy to meet online before my return.]

Background and problem statement: Consider a box domain with inlets on one side and outlets on the other. You have been tasked with designing pipes to carry the fluid from the inlets to the outlets that minimises the fluid's energy loss. The catch is that you only have enough material to build pipes that can occupy a fraction of the space in the box. The question is, where do you place these pipes?

Engineers use *topology optimisation* to solve problems like the one above, as well as for finding optimal designs for car parts, airplane wings, engineering tools and many more besides. There are several modelling techniques and solvers for topology optimization, but at a minimum, one must repeatedly solve PDEs associated with the physics such as the Navier–Stokes or elasticity equations. When scaling to large 3D problems, these solves create computational bottlenecks that are difficult to overcome.

Planned approach: This project will model fluid topology optimisation problem via the density approach. One introduces a variable $\rho : \Omega \to [0, 1]$ such that where $\rho = 1$ indicates the presence of fluid and $\rho = 0$ is void. Intermediate value of ρ are penalised. We will then adapt a recent solver called SiMPL for density-based topology optimisation [1, 2]. This is a robust first-order solver that decouples the optimisation problem into a "forward" solve (i.e. solve the fluid equation) and an adaptive descent step. It is also discretisation agnostic allowing us to plug-and-play with various finite element discretisations inside Firedrake (a Python-based finite element library) [3]. This is advantageous because one can use *preconditioners* coupled with iterative solvers to significantly reduce the cost of the forward solve. As the Reynolds number of the flow increases, these preconditioners become harder to design. However, we will make use of augmented Lagrangian preconditioners designed by Farrell et al. [4, 5]. With all the the tools in hand, we will be able to solve three-dimensional large Reynolds number Navier–Stokes topology optimisation problems.

Desired outcomes: A desired outcome is a working pipeline for preconditioning large-scale 3D fluid topology optimisation problems. If successful, one may also consider time-dependence or preconditioners for other topology optimisation problems involving linear elasticity. Strong results may be published in a good scientific computing journal.

Prerequisites: Knowledge of the finite element method, numerical linear algebra, and Python would be very beneficial.

References

[1] Dohyun Kim, Boyan S. Lazarov, Thomas M. Surowiec, and Brendan Keith. A simple introduction to the SiMPL method for density-based topology optimization. *Structural and Multidisciplinary Optimization*, 68(4):74, 2025.

[2] Brendan Keith, Dohyun Kim, Boyan S. Lazarov, and Thomas M. Surowiec. Analysis of the SiMPL method for density-based topology optimization. *SIAM Journal on Optimization*, 35(2):1134–1164, 2025.

[3] D. A. Ham, P. H. J. Kelly, L. Mitchell, C. J. Cotter, R. C. Kirby, K. Sagiyama, N. Bouziani, S. Vorderwuelbecke, T. J. Gregory, J. Betteridge, D. R. Shapero, R. W. Nixon-Hill, C. J. Ward, P. E. Farrell, P. D. Brubeck, I. Marsden, T. H. Gibson, M. Homolya, T. Sun, A. T. T. McRae, F. Luporini, A. Gregory, M. Lange, S. W. Funke, F. Rathgeber, G.-T. Bercea, and G. R. Markall. *Firedrake User Manual*, Imperial College London, University of Oxford, Baylor University, and University of Washington, 2023. DOI: 10.25561/104839

[4] Patrick E. Farrell, Lawrence Mitchell, and Florian Wechsung. An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number. *SIAM Journal on Scientific Computing*, 41(5):A3073–A3096, 2019.

[5] Patrick E. Farrell, Lawrence Mitchell, L. Ridgway Scott, and Florian Wechsung. A Reynolds-robust preconditioner for the Scott-Vogelius discretization of the stationary incompressible Navier-Stokes equations. *The SMAI Journal of Computational Mathematics*, 7:75–96, 2021.

3.12 Polynomial Optimization Problems: Algorithms and AI-Integrated Approaches

Supervisors: Prof. Coralia Cartis and Kate Wenqi Zhu
Contact: coralia.cartis@maths.ox.ac.uk and wenqi.zhu@maths.ox.ac.uk

High-order tensor methods for solving both convex and nonconvex optimization problems have recently attracted significant research interest. These methods lead to algorithms that achieve *optimal global convergence rates* [7] while enjoying *local convergence rates faster than Newton's method*. A central ingredient of a p-th order method, with $p \geq 2$, is the (global) minimization of a *regularized p-th order Taylor subproblem* to calculate the step change from one iteration to the next.

The minimization of the regularized Taylor subproblem is closely connected to *global polynomial optimization*, as well as to classical questions such as *Hilbert's 17th problem* concerning nonnegativity and sum-of-squares (SoS) representations [2, 5]. Recent exciting developments reveal that *AI-integrated approaches*—including reasoning-based large language models (LLMs) [1] and neural sum-of-squares with transformers [4]—have shown competitiveness in certifying nonnegativity of polynomials and tackling global polynomial optimization problems.

This regularised Taylor subproblem lies at the heart of modern high-order optimization algorithms and has become an *active frontier of research*. We have recently developed several practical methods (e.g., CQR, DTM [6, 3]) to approximately solve this subproblem, and we have observed encouraging numerical performance in nonconvex settings. However, the *theoretical convergence properties* of these methods warrant further exploration.

In this project, students will have the chance to explore one or two of the following potential directions:

- Explore the relationship between global optimality conditions for the regularized Taylor polynomials and their SoS representations via various relaxations techniques.
- Investigate how modern AI tools can be integrated with classical optimization theory to design new solvers and certificates.
- Analyze the convergence behavior of high-order subproblem solvers using tools from *continuous optimization*, such as secular equation techniques (see trust-region methods in B2 Lectures).
- Leverage ideas from *numerical linear algebra (NLA)* to improve the scalability of subproblem solvers.

All these directions build on our previous work and so there exist well set foundations for the project on which to build extensions, in terms of algorithms, implementations and theory.

This project is ideal for students interested in optimization, numerical linear algebra and/or the intersection of mathematics and AI.

References

[1] Kechen Li, Wenqi Zhu, Coralia Cartis, Tianbo Ji, and Shiwei Liu. Sos1: O1 and r1-like reasoning llms are sum-of-square solvers. *arXiv preprint*, arXiv:2502.20545, 2025.

[2] Wenqi Zhu and Coralia Cartis. Sufficiently Regularized Nonnegative Quartic Polynomials are Sum-of-Squares. *arXiv preprint*, arXiv:2601.20418, 2026.

[3] Wenqi Zhu and Coralia Cartis. Global Optimality Characterizations and Algorithms for Minimizing Quartically-Regularized Third-Order Taylor Polynomials. *arXiv preprint*, arXiv:2504.20259, 2025.

[4] Nico Pelleriti, Christoph Spiegel, Shiwei Liu, David Martínez-Rubio, Max Zimmer, and Sebastian Pokutta. Neural Sum-of-Squares: Certifying the Nonnegativity of Polynomials with Transformers. *arXiv preprint*, arXiv:2510.13444, 2025.

[5] Jinling Zhou, Xin Liu, Jiawang Nie, and Xindong Tang. A Tight SDP Relaxation for the Cubic-Quartic Regularization Problem. *arXiv preprint*, arXiv:2511.00168, 2025.

[6] Coralia Cartis and Wenqi Zhu. Cubic-quartic regularization models for solving polynomial subproblems in third-order tensor methods. *Mathematical Programming*, Springer, pp.1–53, 2025.

[7] Coralia Cartis, Nicholas Ian Mark Gould, and Philippe L. Toint. Evaluation complexity of algorithms for nonconvex optimization. *MOS-SIAM Series on Optimization*, 2022.

3.13 Differential Privacy in Optimization: SGD and Beyond

Supervisor: Prof. Coralia Cartis
Contact: coralia.cartis@maths.ox.ac.uk

Problem statement: When supervised models are trained on sensitive data, such as medical or financial data, the resulting model may unintentionally reveal information about individual datapoints. Differential Privacy (DP) addresses this by ensuring that the training algorithm's output differs only slightly if any datapoint is modified. This can, for instance, be achieved by adding random noise to the data to mask their individual contribution, while still providing an overall contribution that is correct on average.

Planned Approach: The standard algorithm used in learning is Stochastic Gradient Descent (SGD). This algorithm randomly selects a subset of datapoints to use in each iteration, and performs a gradient step using these points. By additionally adding random noise to the selected data, we obtain Differentially Private SGD (DP-SGD) [1], for which one can theoretically prove the DP property and convergence to a solution of the problem.

This project aims to go beyond the study of SGD, by considering a broader class of algorithms known as *fixed-point methods* [2]. These generalize various optimization algorithms, one of which is SGD. The purpose is to add DP to such methods, and to show, numerically and if possible, theoretically,

that these methods still converge. By specializing the fixed-point operator, various algorithms may be recovered, now with provable DP [3].

Desired Outcomes: A theoretical goal of this project will be to have strong guarantees of the DP fixed-point methods, both from the DP and the convergence perspective. A numerical goal will be to specialize the methods to various cases, thus introducing new DP algorithms, and illustrating these numerically by running them on synthetic and/or real-life examples.

References

[1] Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private Stochastic Convex Optimization with Optimal Rates. In *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 32:11282–11291, 2019.

[2] Felix Lieder. On the Convergence Rate of the Halpern-iteration. *Optimization Letters*, 15(2):405–418, 2021. DOI: 10.1007/s11590-020-01617-9

[3] Edwige Cyffers, Aurélien Bellet, and Debabrota Basu. From Noisy Fixed-Point Iterations to Private ADMM for Centralized and Federated Learning. In *Proceedings of the 40th International Conference on Machine Learning*, PMLR, pp.6683–6711, July 2023.

4 Data Science

4.1 Non-Markovian Models for Interacting Systems

Supervisor: Prof. Renaud Lambiotte
Contact: renaud.lambiotte@maths.ox.ac.uk

Background and problem statement: Networks offer a powerful framework for representing the structure and dynamics of complex systems across different scales. In these models, direct interactions between components are depicted as edges, while indirect interactions are captured through paths and walks. A common assumption is that system flows follow sequences of independent transitions — as a Markov process. However, real-world path data often exhibit higher-order dependencies, necessitating more advanced modelling approaches.

Description of the approach planned and the techniques needed: To address this issue, researchers have introduced network models with memory, often under the form of a higher-order Markov model. Integrating Markov models of varying orders within a multi-order framework significantly enhances next-element prediction in real-world systems. However, enforcing a single fixed Markov order across an entire system imposes a rigid constraint, as real-world dependencies often extend across multiple orders. The uneven distribution of observations in real data can lead such models to overfit certain regions while underfitting others simultaneously. A fixed-order modelling is also often impractical in practice, as the state space explodes with higher-order models.

What you'd hope to achieve: The purpose of this project will be to explore the possibility to constructing compact network representations from path data, balancing accuracy and model complexity.

References

[1] Rohit Sahasrabuddhe, Renaud Lambiotte, and Martin Rosvall. Concise network models of memory dynamics reveal explainable patterns in path data. *arXiv preprint:* arXiv:2501.08302, 2025.

[2] Renaud Lambiotte, Martin Rosvall, and Ingo Scholtes. From networks to optimal higher-order models of complex systems. *Nature Physics*, 15(4):313–320, 2019.

4.2 Symmetry-Consistent Guidance for SE(3)-Equivariant Molecular Diffusion

Supervisor: Dr Sathya Subramanian
Contact: sathya.subramanian@cs.ox.ac.uk

Background and problem statement: Diffusion models enable high-dimensional conditional generation by steering the generation of new samples toward desired signals. This control often comes at the cost of sampling instability or reduced diversity. In structure-based drug design, the conditioning signal is often chosen to be a target of the drug discovery campaign, e.g. a protein pocket. The objective is to generate molecules (represented as tensors) that are chemically valid and geometrically compatible with the target. This setting imposes rigid-motion symmetry: the generative dynamics (and any guidance term) should respect SE(3) transformations of 3D space. Recent target-aware methods therefore combine SE(3)-equivariant backbones with energy- or force-inspired guidance to bias generation toward stronger pocket complementarity. Such guidance can, e.g. take the form of

vector-field guidance derived from binding-energy gradients. A central open problem is designing guidance that preserves SE(3)-symmetry while maintaining stable sampling and a good fidelity–diversity (coverage) tradeoff.

Suggested approach: The dissertation should present a mathematically precise account of guidance mechanisms used in practice, contrasting (i) conditional-score and classifier-free guidance, and (ii) explicit energy- or force-based guidance, using an SDE/ODE viewpoint to give a precise discussion of stability and bias effects. Interested students could then implement or reproduce a clear classical baseline for 3D generation under SE(3)-equivariance (e.g. an EDM/GeoDiff-style model) and a target-aware setup (e.g. docking/pose diffusion or a VFDiff-style guided generator). Once the student has developed theoretical insights into the problem, possible experiments include running one-factor-at-a-time comparisons to isolate the effect of: (a) backbone choice (e.g. equivariant message passing vs. transformer-style attention over geometric graphs), (b) guidance type and strength, and (c) sampler settings (number and size of steps), benchmarking the experiments against standard metrics in graph machine learning for drug discovery (e.g. validity, diversity, target-compatibility proxies such as docking scores or RMSD where appropriate), and explicit robustness diagnostics (failure rates, sensitivity to hyperparameters).

Desirable outcomes: Students can aim to achieve a reproducible, practically useful set of conclusions about which symmetry-consistent guidance strategies yield the best stability–compute–quality tradeoffs in target-aware 3D molecular diffusion. The dissertation should prioritise theoretical insights for potential improvements, and interpretable diagnostics, aiming to provide concrete recommendations for implementing guidance in SE(3)-equivariant molecular generative pipelines. A very good dissertation could give a simple set of practical rules for choosing guidance strength, sampler settings, and backbone architecture that generalise across small benchmarks. These outcomes would lay the groundwork for reliable SE(3)-equivariant guided diffusion pipelines in structure-based molecular design.

Pre-requisites: Familiarity with diffusion models and basics of rigid motions in 3D geometry. Some exposure to group equivariance, geometric deep learning, and PyTorch is helpful.

References

[1] J. Ho and T. Salimans. Classifier-Free Diffusion Guidance. *arXiv preprint*, arXiv:2207.12598, 2022.

[2] E. Hoogeboom, V. Garcia Satorras, C. Vignac, and M. Welling. Equivariant Diffusion for Molecule Generation in 3D. *arXiv preprint*, arXiv:2203.17003, 2022.

[3] M. Xu, L. Yu, Y. Song, C. Shi, S. Ermon, and J. Tang. GeoDiff: a Geometric Diffusion Model for Molecular Conformation Generation. *arXiv preprint*, arXiv:2203.02923, 2022.

[4] G. Corso, H. Stärk, B. Jing, R. Barzilay, and T. Jaakkola. DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking. *arXiv preprint*, arXiv:2210.01776, 2022

[5] L. Tan, J. Liu, G. Yue, Q. Zou, D. Cao, X. Zeng, and X. Fu. VFDiff: SE(3)-Equivariant Vector Field Guided Diffusion Model for Target-Aware Molecule Generation in 3D. 2024/25.

[6] S. Tang. A Complete Guide to Spherical Equivariant Graph Transformers. *arXiv preprint*, arXiv:2512.13927, 2025.

4.3 Minimal Hybrid Subroutines for Quantum Diffusion-Inspired Generation and Guidance

Supervisor: Dr Sathya Subramanian
Contact: sathya.subramanian@cs.ox.ac.uk

Background and problem statement: Recent work on *measurement-based quantum diffusion models* reframes forward and backwards diffusion as physically meaningful quantum operations (channels, stochastic measurement trajectories, and recovery maps). This creates a bridge between conceptual primitives used in generative modelling within classical computing (noise injection, reverse dynamics, score/flow, and guidance), and quantum computing (measurement, unitaries, and channel inversion). The central question for this dissertation is broad: *can one find and isolate a small, well-defined subroutine in a diffusion model (which could be enhanced by "guidance") that admits a meaningful quantum implementation, and analyse whether it yields any advantage over classical models, perhaps under explicit computational assumptions?* This avoids the unrealistic goal of end-to-end "quantum advantage" for a full application, and instead targets a modular, testable intermediate (e.g. sampling problems or estimation problems) whose complexity can be compared theoretically.

Suggested approach: The dissertation should first develop a concise mapping between classical diffusion mechanisms and their quantum counterparts, clarifying what objects correspond to scores, guidance signals, and reverse-time updates in the quantum formalism. Building on this, one possible direction is to select a minimal sub-task within a classical diffusion model as a candidate for speedup by offloading to a quantum co-processor, and study it rigorously within an idealised setting. This could include: (i) a quantum-assisted *estimation* primitive for expectations needed by a guidance signal (e.g. using classical shadows to estimate observables, with explicit sample complexity bounds), or (ii) a *recovery* primitive as a quantum analogue of reverse-time correction in a diffusion process (e.g. Petz recovery / approximate channel inversion in lieu of classical denoising). The work will combine (a) complexity theory (sample/query complexity, circuit depth, error models) with (b) small-scale numerical simulation to validate whether the hybrid ingredient behaves as intended.

Desirable outcomes: A research outcome could be either a concrete minimal hybrid scheme whose advantage is demonstrated in a clear metric (e.g. estimator variance or sample complexity under stated assumptions), or a principled negative result showing why common candidate quantum subroutines fail to help under realistic constraints. In both cases, the dissertation can aspire to be a reusable conceptual and technical bridge between guided diffusion and quantum diffusion formalisms, with a precise elaboration of the scope, underlying assumptions, and empirical checks.

Pre-requisites: Linear algebra, probability, and basic quantum computing (states, channels, and measurement). Familiarity with PyTorch and diffusion models is helpful.

References

[1] X. Liu, J. Zhuang, W. Hou, and Y.-Z. You. Measurement-Based Quantum Diffusion Models. *arXiv preprint*, arXiv:2508.08799, 2025.

[2] A. Cacioppo, L. Colantonio, S. Bordoni, and S. Giagu. Quantum Diffusion Models. *arXiv preprint*, arXiv:2311.15444, 2023.

[3] H.-Y. Huang, R. Kueng, and J. Preskill. Predicting Many Properties of a Quantum System from Very Few Measurements. *Nature Physics*, 16:1050–1057, arXiv:2002.08953, 2020.

[4] P. Holderrieth and E. Erives. An Introduction to Flow Matching and Diffusion Models. *arXiv preprint*, arXiv:2506.02070, 2025.

[5] O. Santoso. Brief Overview of the Petz recovery map and its applications. https://mcgreevy.physics.ucsd.edu/w23/final-papers/2023W-213-Santoso-Owen.pdf, 2023.

[6] J. Ho and T. Salimans. Classifier-Free Diffusion Guidance. *arXiv preprint*, arXiv:2207.12598, 2022.

4.4 Fourier Features for Learning High-Frequency Functions

Supervisor: Dr Paz Fink Shustin
Contact: paz.finkshustin@maths.ox.ac.uk

Deep neural networks are known to exhibit a "spectral bias", where the model prioritizes low-frequency components, leading to slow convergence for high-frequency functions. This limitation is tied to the rapid eigenvalue decay of the Neural Tangent Kernel (NTK). Random Fourier Features (RFF) have shown empirical success in mitigating this bias, as a positional encoding or input mapping. This project proposes a comprehensive investigation into the role of RFF as a mathematical and numerical preconditioner, exploring how frequency mappings reshape the NTK and influence the conditioning of the underlying optimization problem.

The approach involves both numerical analysis and practical implementation, allowing for an exploration of several interconnected research paths. These include a comparative analysis between RFF-based networks and classical kernel-based methods, such as Kernel Ridge Regression (KRR) and Nyström approximations. Furthermore, the project will evaluate the impact of various frequency-tuning strategies and bandwidth selection on the convergence rates and stability of Gradient Descent. Implementation will be conducted using Python, involving tasks such as high-fidelity signal reconstruction or Physics-Informed Neural Networks (PINNs) to validate the theoretical findings.

By bridging the gap between classical approximation theory and modern deep learning, the project aims to identify the most effective strategies for embedding prior spectral knowledge into neural architectures. Ultimately, this work seeks to provide efficient optimization strategies for multi-scale and high-frequency domains, improving the performance of neural models on complex mathematical tasks.

References

[1] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547. 2020.

[2] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31, 2018.

[3] S. Wang, X., Yu, and P. Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.

[4] A. Rahimi and B. Recht. Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems*, 20, 2007.

4.5 Mathematical Foundations of Operator Learning

Supervisor: Prof. Jared Tanner
Contact: jared.tanner@maths.ox.ac.uk

Machine learning techniques are actively being extended to scientific computing tasks such as solving differential equations and learning PDEs or operators from samples. This project would consider learning corrections to the idealised operator for ptychography. The nearby Diamond Light Source facility is the UK's national synchrotron science facility where, amongst other things, they determine structure of materials though ptychography. Ptychography uses far-field diffraction pattern measurements to determine the structure of materials. When the facilities are ideally calibrated the forward problem is well understood and the task becomes purely a large scale inverse problem. This project seeks to reduce the calibration time by allowing ML techniques to determine corrections from misscalibration and reduce the overall measurement process time in favour of added computational reconstruction complexity. This project will join an ongoing partnership and will engage with DPhil student Alan Murithi as well as two members of staff at Diamond Light Source, Jaroslov Fowkes and Benedikt Daurer. The project has capacity to include highly computational aspects as well ad delving into the foundational mathematical questions of operator learning.