STOR-i ANNUAL CONFERENCE: 7TH – 8TH JANUARY 2027

Titles and Abstracts

Day 1

Statistical data science for environmental systems

Claire Miller, University of Glasgow

Monitoring and managing environmental systems is becoming increasingly complex due to evolving pressures arising from climate change and human activity. For example, pressing challenges related to water quality, quantity (flooding and drought) and environmental sustainability require evidence to be integrated across environmental systems using diverse data sources, computational and statistical models and decision support tools. These challenges motivate further statistical innovation through questions related to monitoring design, combining information across heterogeneous data sources, quantifying uncertainty and supporting robust decision making. Through examples from current environmental collaborations, this presentation will illustrate the related statistical and data science approaches we are developing and highlight emerging research opportunities at the interface of statistics, data science and environmental science. This research is partly funded by NERC and EPSRC.

Automated Satellite-Based Detection and Quantification of Worldwide Methane Super Emissions

Thomas Newman, STOR-i Alumni, Netherlands Institute for Space Research

Since 2018, the TROPOspheric Monitoring Instrument (TROPOMI) aboard the Sentinel-5P satellite has provided daily global observations of atmospheric methane concentrations. Building on these observations, the Space Research Organisation Netherlands (SRON) has developed and operates an automated machine learning pipeline for the Copernicus Atmosphere Monitoring Service (CAMS), enabling the detection and quantification of methane emissions from TROPOMI data. The operational system processes new satellite observations on an hourly basis, combining data preprocessing, machine learning-based plume detection and emission quantification. Since becoming operational in January 2023, the pipeline has identified more than 8,000 emission events across 62 countries, covering sources including oil and gas operations, coal mines and urban areas. This presentation provides an overview of the end-to-end system, from satellite data processing and the challenges of uncertain training labels to the development and evaluation of detection algorithms and the delivery of emission alerts to international stakeholders.

Why Don't We Get It Right First Time? What Operational Research Can Teach Us About Making AI Deliver

Gilbert Owusu, Vanguide

Organisations are investing heavily in AI, yet few see a return. PwC's 2026 Global CEO Survey of more than 4,400 chief executives found that 56% have seen no significant financial benefit from AI so far, and only 12% report gains in both revenue and cost. Why does so much AI fall short of its promise?

This talk argues that the problem is rarely the algorithm alone. More often, AI fails because the problem was framed badly, the objective was wrong, real-world constraints were ignored, or no one planned how the model would be used, monitored and trusted once deployed. These are precisely the questions operational research has tackled for decades.

Drawing on experience leading AI and data transformation at scale in industry, I will show how OR approaches, including problem structuring, simulation, optimisation and decision analysis, can strengthen each stage of the AI lifecycle: from defining the right problem, through design and evaluation, to deployment and the management of risk and impact. I will close with open research questions where the OR community, and today's doctoral researchers in particular, can shape how AI is built and used.

Using relaxed problems to find exact solutions
Lauren Durrell, STOR-i PhD student
Many optimisation problems, such as locating warehouses or assigning jobs to machines, are computationally challenging to solve exactly, particularly as their size increases. Relaxation techniques offer a way to simplify such problems, but often at the expense of obtaining solutions that are feasible for the original problem.

This talk explores how semi-Lagrangian relaxation can be used to obtain provably optimal solutions while retaining some of the computational advantages of relaxation.

Through two examples, the Simple Plant Location Problem and the Generalised Assignment Problem, we illustrate how semi-Lagrangian relaxation works, the challenges involved in applying it, and how exploiting problem structure can improve its computational efficiency. We also highlight results from computational experiments, including faster solution times, stronger lower bounds, and optimal solutions for previously unsolved benchmark instances.

Dynamic Guardrails in Reinforcement Learning for Pricing Problems
Adam Page – STOR-i Impact Fellow
Reinforcement learning is used by vendors to dynamically price perishable products to maximise their revenue. These methods allow the pricing problem to be solved without knowledge of demand. When choosing exploration schemes, one usually focuses on learning fast and correctly, regardless of actions taken in the process. Unrestricted exploration in pricing has clear issues: setting extreme prices leads to financial or reputational damage. Safe exploration is addressed through the introduction of guardrails that dynamically adapt the action space by following a two-stage truncated-beta policy.

Robin Mitra, University College London

Rebecca Hamm – STOR-i Alumni, HEOR Cardiff

Who Gets Help in a Parallel Universe?
Matt Thomas, British Red Cross
Two analysts use the same dataset to answer the same question. Both make perfectly reasonable analytical choices, yet they arrive at different conclusions. What happens when those conclusions determine which communities receive humanitarian assistance - and which miss out?

Every analysis involves a multitude of decisions, from data preparation and model specification to interpretation. Together, these create a multiverse of plausible analyses, each potentially telling a different story.

Drawing on real-world examples from the British Red Cross, I'll explore how seemingly innocuous analytical choices can skew your results in surprising ways. We'll venture into the garden of forking paths, exploring specification curve and multiverse analysis as ways of exposing this hidden uncertainty. I'll share practical approaches to conducting these analyses, including how generative AI can help us build the tools and infrastructure to make them more accessible.

Finally, we'll tackle a difficult question: once we've revealed just how uncertain our conclusions really are, how do we communicate that uncertainty to decision-makers without undermining trust or paralysing action?

On the accuracy and accuracy-implications of inventory records
Aris Syntetos, Cardiff University
Inventory records are often inaccurate, and this is known to be the source of major cost and service inefficiencies in retailing. However, the practices of measuring and reporting errors in inventory records vary considerably across the sector. Further, and while the advantages of error-free stock records are apparent, retailers may tolerate minor discrepancies as inconsequential, blurring the line between 'accurate' and 'inaccurate' records. We set out, for the first time, to systematically research and understand inventory record accuracy measurement. We recognise that the implications of such accuracy for business practices are crucial, so we also consider how accuracy translates (or not) to bottom-line improvements. Our empirically grounded research involves detailed data from, and counterfactual analysis in, a number of major European retailers. We put forward a set of desirable attributes for the error measures and a protocol to assess the sales-lift and other implications of fixing the records. We offer an agenda for further research and operationalised suggestions that should be valuable for any professional involved in inventory decisions making.

Day 2

Laura Boyle, Queen's University Belfast

Online changepoint detection: moving beyond piecewise constant
Harry Ellingham – STOR-i PhD student
Detecting abrupt changes in the level of a time series is an important problem in applications ranging from industrial monitoring to environmental science. In many settings, the underlying signal may exhibit gradual variation even in the absence of changepoints. For example, measurements of temperature or pressure may evolve slowly over time, while sudden shifts in their levels can indicate events of practical significance. Distinguishing these abrupt changes from the natural evolution of the series, particularly in the presence of noise, presents an important statistical challenge.

Many existing methods for detecting changes in the level of a time series assume stationarity between changepoints, potentially limiting their applicability in practice. In this talk, we relax this

assumption by allowing the underlying signal to vary slowly between abrupt changes in level. We propose a methodology for detecting these changes in real time and demonstrate its effectiveness on both simulated and real-world data.

A Graph-Based Approach for Scaling Warehouse Order Batching
Mark Holcroft – STOR-i PhD student
The first process in many warehouses is the grouping of customer orders into batches such that pickers may collect items for several orders in a single trip. A good set of batches can ensure that each picker need only traverse a small portion of the warehouse, whereas a poor set can introduce redundancies and overlapping routes. As such, the quality of batches is a key driver of cost. Although exact models exist, their computational difficulty means that practitioners regularly rely instead on heuristics to generate good but often suboptimal solutions.

In this talk, we discuss how constraints on the direction aisles can be traversed by pickers can make exact methods more tractable by simplifying the routing aspect of the problem. In particular, by modelling the warehouse as a directed graph, we derive a closed-form representation of the minimum-distance route, and use this to construct a mixed-integer program to assemble batches. We examine how this scales as the number and size of orders increases, and compare its total picking distance with that of a batching heuristic.

Structured Multivariate Extreme Value Models for Flood Risk Estimation
Kristina Grolmusova – STOR-i PhD student
We exploit the new framework of multivariate geometric extreme value theory to develop statistical methodologies that will allow analysis of extremes of river flows at multiple locations on a river network simultaneously. Current methodologies within this framework are limited to a relatively low number of dimensions. This is insufficient for the purposes of flood risk estimation, since the number of locations on a river network where we get the river flow measurements from is often of the order 10-20. In order to create a parsimonious model in higher dimensions, we use recent work to define the gauge function which describes the limit shape of the scaled sample cloud in a structured way, using the river network. We apply this model to extreme flows at 10 locations on a river network around Preston, in north-west England. The model is shown to fit well and indicates strong extrapolation performance.