STOR-i
excellence with impact

Lancaster University

STOR-i ANNUAL CONFERENCE: 8TH – 9TH JANUARY 2026

Titles and Abstracts

Day 1

Experiences with dynamic pricing and reinforcement learning
Jeremy Bradley, Consulting Data Scientist
Over the last 5 years, my team has been exposed to a number of dynamic pricing problems across aviation, travel and tourism. We turned to RL (with some excitement initially) and have encountered some interesting experiences - some would say learning experiences - using tabular reinforcement learning in practice which helped and hindered us in equal measure. I will introduce the general dynamic pricing challenge and hopefully by the end, will have conveyed some of the pitfalls, misunderstandings and possibly even a few successes that we encountered in trying to tackle these.

A Bayesian Block Maxima Over Threshold Approach applied to Corrosion Assessment in Heat Exchanger Tubes
Jess Spearing – STOR-i Alumni, Shell
Remaining life estimation due to corrosion poses a hurdle for numerous industrial pro-cesses, and though corrosion can be measured directly, statistical approaches are often re-quired to either correct for measurement error or extrapolate estimates of corrosion severity where measurements are unavailable. This article considers corrosion in heat exchangers tubes, where corrosion is typically reported in terms of maximum pit depth per inspected tube, and only a small proportion of tubes are inspected, suggesting extreme value theory (EVT) as suitable methodology. However, in data analysis of heat exchanger data, shallow tube-maxima pits often cannot be considered as extreme; although previous EVT approaches assume all the data are extreme. We overcome this by introducing a threshold — suggesting a block maxima over threshold approach, which leads to more robust inference around model parameters and predicted maximum pit depth. The model parameters of the resulting left-censored generalized extreme value distribution are estimated using Bayesian inference, meaning parameter uncertainty is immediate from the estimation procedure. This framework also naturally extends to cases where there is significant measurement error or censoring of inspection results, which is a common challenge.

Risk-Aware Markov Decision Processes
Jefferson Huang – Naval Postgraduate School
Markov decision processes (MDPs) model sequential decision-making under uncertainty and serve as a foundation for the design of reinforcement learning algorithms. In the discrete-time setting, solving an MDP means finding a decision-making policy (e.g., mapping environment states to actions) that maximizes an objective that depends on the sum of the stochastic rewards earned over a finite or infinite planning horizon. For the “risk-neutral” setting, where the objective is the expected value of this sum, there exists a relatively complete theory on the existence and computation of optimal policies. In this talk, we present some recent results in the “risk-aware” setting where the expected value is replaced with the conditional value-at-risk of the total discounted rewards, which include connections to Markov games.

Optimization in transportation and logistics: trends and research ideas
Grazia Speranza – University of Brescia

Technological advancements over the past few decades have significantly transformed the way people travel and how goods are distributed. Today, a systemic approach to problem-solving and the use of advanced analytical methods are more essential than ever. In this talk, I will explore some key trends in the application of optimization models and algorithms within transportation and logistics with examples of projects from my own research experience. Optimization in transportation and logistics: trends and research ideas.

Statistical Modelling of Extreme Sea Surface Temperatures
Kajal Dodhia, STOR-i PhD student

Sea Surface temperature (SST) is gradually increasing around the world, causing significant negative environmental, ecological and socio-economic impacts. Existing methods that identify extreme SST are semi- or non-parametric and are primarily designed for ecological applications. These methods suffer from saturation and are heavily influenced by a specified climatological period. Although efforts have been made to mitigate saturation, it remains an open problem. We propose a new parametric method called the Varying Threshold Method for SST (VT4SST). This method uses a linear mixed model to identify the extremes above a varying threshold as values that are unusually high for a given time of year. Then, the excesses are modelled using the Generalised Pareto Distribution. Our method is driven by and applied to buoy data from the National Network of Regional Coastal Monitoring Programmes of England and Wales, which is rarely used in literature despite its high accuracy. We demonstrate its effectiveness in capturing extreme SSTs. We show that the VT4SST performs well, does not suffer from saturation, and can help provide risk estimates that can inform policy decisions aimed at managing coastal risks.

Towards the Integration of Strategic and Operational Decision Making
Luke Fairley – STOR-i Impact Fellow

Deterministic optimisation considers problems in which a decision-maker makes a single strategic decision, which are commonly formulated as mixed-integer programmes (MIPs). In contrast, stochastic dynamic optimisation considers problems involving sequential operational decisions in a random environment, typically modelled as Markov decision processes (MDPs). Frameworks that integrate strategic and operational decisions through MIPs and MDPs can be known as MDP Design frameworks (although this name is not set in stone), and the existing literature on this is limited to three prior studies each with their own approach, highlighting the complexity of merging these two disparate types of decision making. This talk covers these previous studies, introduces our own framework, and highlights the many research questions and the potential for application in this new and emerging field.

Data-driven modelling of cascades on networks
James Gleeson – University of Limerick

Network models may be applied to describe many complex systems, and in the era of online social networks the study of dynamics on networks is an important branch of computational social science. Cascade dynamics can occur when the state of a node is affected by the states of its neighbours in the network, for example when a Twitter/X user is inspired to retweet a message that she received from a user she follows, with one event (the retweet) potentially causing further events (retweets by followers of followers) in a chain reaction. In this talk I will review some mathematical models that can help us understand how social contagion (the spread of cultural fads and the viral diffusion of information) depends upon the structure of the social network and on the dynamics of human behaviour. Although the models are simple enough to allow for mathematical analysis, I will show examples where they can also provide good matches to empirical observations of cascades on social networks.

Deep extremal regression: Grey box models for univariate and multivariate extremes
Jordan Richards – STOR-i Alumni, University of Edinburgh
Despite some negative press, deep neural networks are a useful tool for statistical risk modelling. We present deep extremal regression models - neural networks designed to target descriptions of the conditional tails, such as extreme quantiles or tail indices. By combining deep neural networks with asymptotically-justified models from extreme value theory, we create grey box models which permit high-dimensional inference on extremes while retaining some of the interpretability of traditional statistical models. This talk covers a few examples of deep extremal regression models, which we use for extreme quantile regression, full (semi-parametric) conditional density estimation, and modelling of multivariate extremes via an angular-radial approach. Real data applications include US and Mediterranean wildfire extremes, and UK Metocean storms.

Uncovering Dependency Structures in Multivariate Markov Chain Models
Theodore Kypraios – University of Nottingham
Multivariate Markov Chain (MMC) models provide a powerful framework for capturing dependencies among multiple interrelated processes that evolve over time. They are widely used in finance, epidemiology, and genetics, where understanding the joint behaviour of several variables is crucial for prediction, risk assessment, and decision-making. When all chains share the same discrete state space, a natural parameterisation is through the joint transition probability matrix on the extended state space formed by the Cartesian product of the marginal spaces. However, the number of parameters grows rapidly with both the number of states and the number of chains, creating major challenges for estimation and inference. In this talk, I will introduce alternative, general parameterisations for MMC models that efficiently capture and characterise first-order dependence between chains—such as conditional independence, contemporaneous dependence, and Granger non-causality—while remaining compatible with likelihood-based inference. I will also present a new algorithm, the Generalized Individual Forward–Backward (GIFB) algorithm, which provides exact (in a Monte Carlo sense) inference for latent states in Hidden MMCs and reduces computational complexity from exponential to quadratic in the number of chains. The work is motivated by applications to multivariate Markov-modulated Poisson processes for detecting disruption on the National Rail network in Great Britain, using Twitter/X volume and content related to delays and disturbances. If time permits, I will illustrate how the proposed parameterisations can be applied to answer practically important questions in this context.

Day 2

Statistics and the law: improving probabilistic reasoning in criminal cases
Amy Wilson – University of Edinburgh
Criminal cases can feature multiple pieces of dependent evidence and multiple possible explanations for this evidence. It can be challenging to disentangle the correlations between pieces of evidence and to understand how to form a logically consistent argument that accounts for this evidence in a way that is probabilistically sound. There have been high profile miscarriages of justice that have resulted from failures in probabilistic reasoning and interpretation and there is currently controversy in the legal system regarding the use of statistics in the Lucy Letby case.

In this talk I will show how chain event graphs can be used to construct possible storylines for displaying the time evolution of events and evidence in criminal cases. These chain event graphs can be used to investigate possible arguments when drawing up a case, to make probabilistic assessments of the strength of evidence when prosecuting or defending or to reason about the likelihood of guilt at a whole case level. I will give two examples – a drugs on banknotes case and the case of the murder of Meredith Kercher. I will also discuss the relevance of graphical modelling approaches to the Letby case. To finish the talk I will explore the role that statistics and probabilistic reasoning can play in criminal cases and highlight where greater collaboration is needed between statisticians and those in the legal sector.

Multi-fidelity simulation optimisation using Gaussian Markov random fields
Graham Burgess – STOR-i PhD student
Simulation optimisation problems arise when one cannot directly evaluate an objective function but can only estimate it via stochastic simulation. Such problems are difficult to solve efficiently since a stochastic simulation can be computationally expensive to run and must be run multiple times. To tackle this, we propose an optimisation method which takes advantage of models of differing fidelity. We incorporate a low-fidelity alternative to the expensive "high-fidelity" stochastic simulation. A low-fidelity model is typically cheap but biased. We model the error of the low-fidelity model with respect to the high-fidelity model using a Gaussian Markov random field. Using the low-fidelity model and the model of its error, we reduce the reliance on the expensive high-fidelity model during optimisation.

Statistical Modelling of Human-Induced Seismicity: Magnitude Uncertainty and Spatio-Temporal Structure
Wanchen Yue – STOR-i PhD student
Human-induced seismicity, arising from industrial activities such as gas extraction and carbon storage, presents significant challenges for assessing seismic hazard. These events often display complex spatio-temporal patterns and are recorded under varying measurement conditions, making robust statistical modelling essential for understanding their behaviour and supporting risk management. Our research leverages statistical methodologies to address key challenges in seismic risk assessment. We develop novel techniques to estimate the upper tail of earthquake magnitude distributions, incorporating expert knowledge of the maximum possible magnitude and addressing measurement error in the magnitude distribution. This approach enhances the reliability of extreme magnitude predictions, which are vital for effective risk preparedness and mitigation. In addition, we investigate the spatio-temporal occurrence of induced earthquakes using extensions of the Epidemic-Type Aftershock Sequence (ETAS) model. By allowing the background intensity to vary with geomechanical or operational drivers, such as stress changes associated with subsurface activity, these developments aim to provide more realistic representations of induced earthquake clustering and its associated hazards.

Sequential Drifter Placement for Ocean Currents with BALLAST
Rui-Yang Zhang – STOR-i PhD student
Drifters are free-floating measuring devices of ocean currents that are being increasingly deployed by oceanographers. Such devices, once placed, will flow according to the currents and collect data accordingly. They are favoured by the practitioners due to their relatively low costs and abilities to capture both the spatial and temporal characteristics of the ocean currents. The deployment strategy of such sensors, however, is understudied. In this talk, we propose a formal Bayesian sequential decision-making framework for drifter placement that accounts for the structure of drifter observations and present BALLAST: Bayesian Active Learning with Look-ahead Amendment for Sea-drifter Trajectories. We observe noticeable benefits of BALLAST-aided sequential placement

strategies on both synthetic and high-fidelity ocean current models. This talk is based on https://arxiv.org/abs/2509.26005.

Interpretability and Explainability for Optimization

Marc Goerigk – University of Passau

Trust in computer systems is central to their practical use. This principle is well recognized in the artificial intelligence community, where interpretability and explainability have been long-standing research priorities and are now mandated by various regulations. Surprisingly, the operations research community has not accorded the same level of importance to trust, despite the widespread use of complex recommendations across diverse environments. Recently, however, this has begun to change, with new frameworks emerging to address these concerns. In this talk, I will present some of these ideas and discuss how they may transform the way we develop and apply optimization models and algorithms.