# STOR-i Conference: 11th – 12th January 2024

## Titles and Abstracts

## Day 1

**Paul Harper, Cardiff University**

***Transforming Emergency and Urgent Care Services with OR***

In this talk I will discuss several related research projects, broadly within emergency and urgent care services. This includes working with ambulance providers and the Indonesian Government to help them make critical decisions on the optimal types, capacities and geographical locations of response vehicles. Such factors directly impact on the probability of patient survival, ability to respond to major disasters, and the overall quality of care provided. There are however many challenges faced in Indonesia, including vast geographical areas, traffic congestion, inadequate numbers of ambulances and a lack of a co-ordinated service.

**Christina Pagel, University College London**

***Simple maths for better decisions***

Much of my work involves careful analysis of large data sets to extract useful information for decision making. But this talk is not about that. Instead, I want to concentrate on two examples where very simple mathematics can provide a framework for supporting large scale health policy decisions. Specifically, I will cover two examples: 1) the benefits of vaccination against Covid in young children, and 2) whether to change the threshold for diagnosing a disease in the general population.

**Roberto Rossi, University of Edinburgh**

***On the Stochastic Inventory Problem Under Order Capacity Constraints***

We consider the single-item single-stocking location stochastic inventory system under a fixed ordering cost component. A long-standing problem is that of determining the structure of the optimal control policy when this system is subject to order quantity capacity constraints; to date, only partial characterisations of the optimal policy have been discussed. An open question is whether a policy with a single continuous interval over which ordering is prescribed is optimal for this problem. Under the so-called "continuous order property" conjecture, we show that the optimal policy takes the modified multi-(s, S) form. Moreover, we provide a numerical counterexample in which the continuous order property is violated, and hence show that a modified multi-(s, S) policy is not optimal in general. However, in an extensive computational study, we show that instances violating the continuous order property do not surface, and that the plans generated by a modified multi-(s, S) policy can therefore be considered, from a practical standpoint, near-optimal. Finally, we show that a modified (s, S) policy also performs well in this empirical setting (https://doi.org/10.1016/j.ejor.2023.06.045).

**Maddie Smith, STOR-i PhD student**

***How to combine forecasts? A puzzle***

Forecast combination is a ubiquitous problem for decision makers. The need to combine several expert forecasts into a single aggregate forecast can arise in many different settings, including macroeconomic forecasting, anticipating trends in retail and combining environmental forecasting models. Although this may appear a simple task on the surface, the combination of forecasts can prove a challenging problem due to features such as correlation, little historical data and subjective assessment. In this talk, we discuss popular point forecast combination methods in the literature and their associated challenges. We consider the forecast combination puzzle, and the bias-variance trade off when selecting appropriate linear combination weights. We propose an alternative forecast combination procedure based on Dynamic Linear Models, which enables weights to adapt dynamically based on past performance, and accommodates for correlations between forecasters. We show that this method can outperform a selection of benchmark forecast combination methods when applied to simulated and real data.

**Matthew Davision, STOR-i PhD student**

***Exploring a Many-objective University Timetabling Problem***

Many possible objectives to optimise arise as part of the timetabling process at universities. In this presentation, we give an overview of some of these objectives, and why they are important. Then we present a goal-programming inspired multi-objective approach to explore trade-offs between these objectives. To conclude we discuss the benefits and weaknesses of this approach and offer insights on how this work will progress in the future.

**Niall Adams, Imperial College London**

***Multiview learning and industry problems***

Multiview learning is concerned with extracting common signal from heterogeneous feature representations of instances. A common example, and that considered in this talk, is images and text. We will discuss three real examples, and outline the challenges of Multiview learning. Our recently developed method will be presented, and some strengths (and weaknesses!) discussed. Issue of practical relevance arising from real-world applications will be stressed to contrast between the challenges and thinking of industry and academia.

**Roland Langrock, Universität Bielefeld**

***Periodic variation in hidden Markov models***

The class of hidden Markov models (HMMs) is a popular tool for modelling time series driven by underlying states. HMMs can be used for example to relate animal movement processes to underlying behavioural modes, to infer the disease state of a patient from biomarkers, or to predict extreme share returns as a function of the underlying nervousness of the financial market. After a general introduction to this versatile class of time series models, I will focus on one particular extension of HMMs, namely model formulations that attempt to capture periodic variation in the state-switching dynamics (e.g. diel variation in animal behaviour, or seasonality in economic markets). Such periodic variation is commonly modelled using trigonometric functions, but can alternatively and more flexibly be incorporated using cyclic splines. I will showcase these modelling approaches in case studies, in which I will also illustrate a positive consequence of such periodic modelling, namely that for periodic HMMs the distributions of the dwell times in the different system states can deviate substantially from a geometric distribution as it would be implied by a homogeneous HMM.

## Day 2

**Chris Dent, University of Edinburgh**

***Digital Twin based decision support: the Climate Resilience Demonstrator project***

This talk will describe the Climate Resilience Demonstrator project, which ran under the National Digital Twin programme in 2021-2, and since then has been hosted by Connected Places Catapult and supported by the energy and water industries. As well as describing the specific project and use case, this will also be used as a basis for describing a wide range of issues in practical decision support modelling, including interoperability of data and models between organisations; uncertainty management in complex systems; modelling with limited conventional data; and deployment at enterprise scale within organisations.

**Thomas Newman, STOR-i PhD student**

***Probabilistic Inversion Modelling of Gas Emissions: A Gradient-Based MCMC Estimation of PDE Parameters***

In response to the escalating global concerns regarding air quality and the environmental impact of greenhouse gas emissions, detecting and quantifying sources of emissions has become critical. To understand and mitigate these environmental impacts, atmospheric gas emission studies must accurately estimate the location and rate of emission sources, background gas concentration, and sensor measurement error variance. In practice, these estimations are often biased when parameters in the gas dispersion model are assumed to be known. Here, we present an innovative approach to address the challenging task of estimating the parameters of the dispersion model, whilst simultaneously estimating sources, background, and measurement error. Estimation uncertainty is quantified through probabilistic inversion using gradient-based MCMC methods. Our focus lies on simultaneously estimating a large number of parameters which have complex relationships whilst enforcing sparsity, positivity, and computational efficiency.

**Katie Howgate, STOR-i PhD student**

***Optimizing Markdown Strategies: A Hierarchical Gaussian Process Approach for Demand Modelling in Retail***

When selling products, demand for each product does not remain consistent throughout its lifetime. As time progresses a product is deemed less desirable by customers due to factors such as declining quality or newer improved products being released. We often wish to maximise revenue and keeping prices consistent while demand is decreasing is not likely to achieve this. Therefore, we need to use historic sales data to determine a discount pricing strategy for products towards the end of their saleable lifetime, known as markdowns. We can use this data to model the relationship between demand and price which we can use to determine what set of discounts will be most profitable. Within available historic sales data, there are often only a limited number of discount percentages that have been applied to products. Additionally, some products have much more sparse data than others. Consequently, we can have a large amount of uncertainty in our demand. To tackle this, we can group products to enable data sharing among similar products. Here rather than grouping by physical characteristics, as is commonly done, we determine similar products are those that have similar demand behaviour. We cannot assume the demand for a luxury and a basic version of the same product follows the same demand function but grouping by physical characteristics can accidentally do this. We propose a hierarchical Gaussian process model and present some numerical results based on this. The model takes advantage of data sharing which also can capture the demand behaviour without a known underlying demand function. The model also gives a measure of uncertainty based on data availability. This work was undertaken in collaboration with Tesco.

**Carla Pinkney, STOR-i PhD student**

***Understanding Neuronal Synchronisation in High-Dimensions***

Understanding the complex dynamics of the brain is one of science's greatest challenges. In this work, we are interested in characterising electrical activity in the brain using neuron level data obtained from micro electrode array recordings. Statistically, these data can be thought of as observations from a marked multivariate point-process. We present novel methodology for the estimation of high-dimensional inverse spectral density matrices, which encode information regarding interactions between point-processes. This estimate can subsequently be used to infer networks of neuronal interactions. This project is in collaboration with world-leading statisticians and neuroscientists at the University of Washington in Seattle.

**Ben Lowery, STOR-i PhD student**

***Inventory control of end-of-sales-life products with integrated online and offline sales channels***

In a business with integrated online and offline storefronts, the management and control of stock for end-of-sales-life products facing uncertainty in demand poses a great deal of operational challenges. Historically there has been much research into inventory control under uncertainty. However, inventory control at the end-of-sales-life remains an under-researched area within the literature and practice. A key challenge is optimising stock levels to ensure remaining stock is sold, as well as preventing stockouts and unsatisfied customers. We develop a novel inventory model with integrated online and offline sales channels, to determine order quantities during the end-of-sales-life period. This work is done in partnership with Tesco Mobile.

**Ruth Misener, Imperial College London**

***OMLT: Optimization and Machine Learning Toolkit***

OMLT (https://github.com/cog-imperial/OMLT) is an open source software package incorporating surrogate models, which have been trained using machine learning, into larger optimization problems. We discuss the advances in optimization technology that made OMLT possible and show how OMLT integrates with Pyomo. We demonstrate how to use OMLT for solving larger decision-making problems in both computer science and process systems engineering. We close with discussing the open research questions related to this research stream. This work is joint with the Imperial Computational Optimisation Group (Francesco Ceccon, Alexander Thebelt, Calvin Tsay), Sandia National Laboratories (Jordan Jalving, Joshua Haddad), ad Carnegie Mellon (Carl Laird).