STOR-i Conference: 11-12 January 2018

Abstracts

Day 1

Solving large fleet management problems in the face of uncertainty
Professor Enrique Munoz de Cote, Head of Multi-Agent Systems, prowler.io

A brief overview of some interesting unsolved problem in the field of field management from a computer science angle. The point of view discussed will be about how to use new collected data as evidence for decision making. The decision making problem to be addressed is to coordinate a large fleet to solve long-term objectives when the environment is uncertain. I will present some of our recent (centralised and decentralised) solutions, our simulation platform and future challenges in the field.

Influence maximization in stochastic and adversarial settings
Professor Po-Ling Loh, Department of Electrical and Computer Engineering, University of Wisconsin-Madisson

We consider the problem of influence maximization in fixed networks, for both stochastic and adversarial contagion models. Such models may be used to model infection spreads in epidemiology, as well as the diffusion of information in viral marketing. In the stochastic setting, nodes are infected in waves according to linear threshold or independent cascade models. We establish upper and lower bounds for the influence of a subset of nodes in the network, where the influence is defined as the expected number of infected nodes at the conclusion of the epidemic. We quantify the gap between our upper and lower bounds in the case of the linear threshold model and illustrate the gains of our upper bounds for independent cascade models in relation to existing results. In the adversarial setting, an adversary is allowed to specify the edges through which contagion may spread, and the player chooses sets of nodes to infect in successive rounds. Our main result is to establish upper and lower bounds on the regret for possibly stochastic strategies of the adversary and player. This is joint work with Justin Khim (UPenn) and Varun Jog (UW-Madison).

Sequential Monte Carlo and Dynamic Latent Space Networks
Kathryn Turnbull, STOR-i PhD student

The ubiquity of high-dimensional networks motivates the study of efficient inference schemes for this type of data. When considering networks, we note that high-dimensional may refer to the number of nodes, the number of observations in time, or both. We consider temporally evolving networks whereby we observe connections between a fixed set of nodes over time. Whilst a range of modelling approaches exist for this type of data, we focus on the latent space approach. In this framework we assume that nodes can be represented in a p-dimensional space such that the probability of a connection between each pair of nodes is determined by their relative positions. This framework can be adapted to the time varying scenario by using a state space model. In this setting inference is typically carried out via MCMC or Variational methods. We instead note that this presents a natural framework for the application of Sequential Monte Carlo techniques. However, here we have a high-dimensional and highly-dependent state space. This talk will discuss how SMC may be applied in this context.

Dynamic stochastic block models: parameter estimation and detection of changes in community structure

Matt Ludkin, STOR-i Alumni, Lancaster University

The stochastic block model (SBM) is widely used for modelling network data by assigning individuals (nodes) to communities (blocks) with the probability of an edge existing between individuals depending upon community membership. In this talk, we introduce an autoregressive extension of the SBM, based on continuous-time Markovian edge dynamics.

This allows for the movement of individuals between communities. An effective reversible-jump Markov chain Monte Carlo algorithm is introduced for sampling jointly from the posterior distribution of the community parameters and the number and location of changes in community membership. The algorithm is successfully applied to a network of mice.

Further details are available in the corresponding paper at:
https://link.springer.com/article/10.1007/s11222-017-9788-9

OR in Maritime Logistics

Dr Kerem Akartunali, Department of Management Science, University of Strathclyde

In this talk, I will first present a brief overview of maritime OR literature and the most classical problems studied in this area. Rather than focusing on traditional problems, I will present two different problems to provide a different perspective.

In the first part of the talk, I will discuss offshore wind turbine installation logistics optimization. Due to significant government incentives for renewables, the accessibility of the surrounding seas and the immense potential for energy generation, installing wind farms tens of miles offshore has become very attractive to many energy companies, although there are many unknowns in these big scale projects. I will discuss wind turbine installation logistics, where multiple vessels operate on the installation of tens, if not hundreds, of wind turbines with various constraints from precedence of tasks to varying degrees of limiting weather. I will also discuss some interesting extensions stemmed from these problems.

In the second part of the talk, I will discuss crew scheduling and re-scheduling problems, which originated from a company operating a fleet of Offshore Supply Vessels (OSVs) on a global scale. With the aim of assigning crew to different tasks or shift patterns during a predetermined planning horizon, crew scheduling often aims to minimize costs while attempting to satisfy various legal and contractual requirements, as well as physical constraints imposed by the geographical and temporal aspects of the tasks. Crew scheduling problems have been widely studied in various transportation sectors, such as airlines, railways and urban buses. However to date it appears that application of these problems in the maritime context has been very limited, where many unique challenges exist due to the variety of tasks to be carried out, the varying experience levels of crew, and long duty times. I will briefly discuss differences between different transport settings, and then present mixed-integer programming formulations. Some discussion on effective solution methods and future perspectives will conclude this part of the talk.

STOR-i excellence with impact

Lancaster University

Taking Action with Incomplete Information and Unknown Deadlines
Professor Mike Atkinson, Operations Research Department, Naval Postgraduate School

Optimal stopping problems balance the trade-off between a known immediate reward and unknown future rewards. We consider three different optimal stopping applications where the decision maker faces an unknown deadline. If the decision maker does not act (i.e., "stop") by the deadline, additional costs are incurred.

1. An investigator receives a stream of intelligence about the location of a target. At any time, the investigator can engage the location he thinks contains the target or wait for more intelligence. The investigator incurs costs when he engages the wrong location, based on insufficient intelligence, or waits too long in the hopes of gaining better situational awareness, which allows the target to disappear.

2. A time-critical operation is contingent on completing a preliminary set of actions in a checklist. The operation may be executed before the full checklist is completed but then it may fail. The question is when to abort the checklist and initiate the operation.

3. During an artillery battle, firing multiple rounds from the same location has two main benefits: a high rate of fire at the enemy and improved accuracy as the shooter's aim adjusts to previous rounds. However, firing repeatedly from the same location carries significant risk that the enemy will detect the artillery's location. Therefore, the shooter may periodically move locations to avoid counter-fire.

Scientific decision making in Royal Mail
Dr Jeremy Bradley, Royal Mail Group, Industry Speaker

In a large business, many people make decisions - operational and strategic - every day. Sometimes those decisions are based on experience and expertise and sometimes, on data and recent-trends. Often, though, individuals in a large organisation are implicitly given OR-style problems in resource allocation or planning with limited access to tools to help them solve them. In this talk, I will lay out the approach we, in the data science group, have been taking to help design and deploy such tools to help colleagues in Royal Mail.

Day 2

Strategic and Tactical Methods for Reducing Delays at Airports
Dr Rob Shone, Research Associate, Lancaster University

Air traffic growth is placing the resources of many of the world's busiest airports under seemingly unsustainable pressure. Flight delays are not only inconvenient for passengers, but also have serious financial consequences for airlines and other airport stakeholders. The expansion of airport infrastructure is usually not possible in the short-term, so airports must strive to improve the efficiency of their operations in order to cope with increasing volumes of traffic.

The OR-MASTER project is a joint collaboration between Lancaster University and Queen Mary University which aims to develop innovative and mathematically sophisticated methods for improving the management of scarce resources at busy airports. This talk will begin with an overview of the project's objectives and will go on to describe the slot allocation process, in which we aim to determine a schedule of operations at a single airport which broadly satisfies airlines' requirements without violating capacity and fairness constraints. The remainder of the talk will focus on a dynamic optimisation problem in which expected queueing delays must be minimised by controlling runway configurations and the balance of throughput rates between arrivals and departures in response to real-time observed queue lengths and weather conditions.

Stochastic and statistical modelling for engineering applications
Professor Darren Wilkinson, School of Mathematics & Statistics, Newcastle University

The ability to make credible simulations of open engineered biological systems is a necessary step towards the application of scientific knowledge to solve real-world problems in this complex engineering domain. One important application of this approach is in the design and management of wastewater treatment systems. However, the simulation of open biological systems is challenging because they often involve a large number of bacteria (up to 10^18 individuals in a wastewater treatment plant), and are physically complex. Since the models are computationally expensive, and due to computing constraints, the consideration of only a limited set of scenarios is often possible. A simplified approach to this problem is to use a statistical approximation of the simulation ensembles derived from the complex individual based (IB) models at a fine scale which will help in reducing the computational burden. A statistical emulator is also an effective tool for studying and incorporating micro-scale processes in a computationally efficient way into macro-scale models, leading to a genuinely multi-scale modelling approach.

Newcastle University Frontiers in Engineering Biology (NUFEB, research.ncl.ac.uk/nufeb) is a large collaborative project between Newcastle University and Northumbrian Water, funded by the EPSRC, concerned ultimately with improving the design of open engineered biological systems including wastewater treatment plants.

Systematic trading, noisy data and star-gazing
Dr Anthony Ledford, MAN Investments, Industry Speaker

In this talk I will explain how modelling noisy data underpins research and development within Man AHL's systematic trading business. We'll look at some examples including a Bayesian combination of classifiers model used elsewhere for supernovae detection, the application of a multi-armed bandit, and how to understand what your predictive model is doing before it has done it.

Online learning for online learning
Ciara Pike-Burke, STOR-i PhD Student

In recent years, technology has changed education with students today being able to supplement their studies with online resources and activities. This allows for much more personalisation and, in many cases, the online platform itself is responsible for selecting which questions to ask the individual students. A natural question that arises from this is how best to select questions for a particular student in order to maximize their learning? Online learning is a framework that allows us to use sequential data in order to help inform decisions and predictions. Within this, we are particularly interested in the family of sequential decision problems commonly conceptualised by the multi-armed bandit problem. In this talk, we discuss several developments in online learning and multi-armed bandits specifically motivated by the problem of selecting questions for an online education platform.