



2016 JOINT NZSA+ORSNZ CONFERENCE

**NEW ZEALAND STATISTICAL ASSOCIATION
OPERATIONS RESEARCH SOCIETY OF NEW ZEALAND**

ABSTRACT BOOKLET

**27th – 30th November 2016
Auckland University of Technology
New Zealand**



PLENARY SPEAKERS

Professor Di Cook

Monday 09:20 – 10:20

Di Cook is a Fellow of the American Statistical Association. Her research is in data visualization, exploratory data analysis, multivariate methods, data mining and statistical computing. She has developed methods for visualising high-dimensional data using tours, projection pursuit, manual controls for tours, pipelines for interactive graphics, a grammar of graphics for biological data, and visualizing boundaries in high-d classifiers. Di has experimented with visualizing data in virtual environments, participated in producing software including *xgobi*, *ggobi*, *cranvas* and several R packages. Her current work is focusing on bridging the gap between statistical inference and exploratory graphics. Di and her colleagues are doing experiments using Amazon's Mechanical Turk, and eye-tracking equipment. Some of the applications that Di has worked on include backhoes, drug studies, mud crab growth, climate change, gene expression analysis, butterfly populations in Yellowstone, stimulus funds spending, NRC rankings of graduate programs, technology boom and bust, election polls, soybean breeding, common crop population structures, insect gall to plant host interactions, soccer and tennis statistics.



Data Journalism Panel

Monday 16:05 – 17:05

The panellists will discuss the role and presentation of data in story-telling, and how this is influenced by modern, interactive visualisation techniques. The impact of the statistical literacy (or lack thereof) of the readership will also be discussed. The current and future training of data journalists will be debated. The panel will then be open to engage with conference delegates to address any questions and comments raised by the audience. The panel discussion will be moderated by [Professor Thomas Lumley](#). The panel members include: [Harkanwal Singh](#) (NZ Herald), [Keith Ng](#) (On Point), [Lillian Grace](#) (Figure.NZ), [Allan Lee](#) (School of Communications, AUT) and Professor Di Cook (Monash University).

Dr Mike O'Sullivan

Tuesday 09:00 – 10:00

Dr Michael O'Sullivan has a Ph.D. (Management Science and Engineering) and MS (Engineering-Economic Systems and Operations Research) from Stanford University and a MPhil (Distinction, Operations Research) and BSc (First Class Honours, Mathematics and Computer Science) from the University of Auckland. Dr O'Sullivan's research specialty is Operations Research (OR) and, in the recent years, combining OR with Analytics. He formed the research group ORUA and this group specialises in utilising OR and Analytics to develop intelligent systems. ORUA's research into OR and Analytics provides intelligence in many application areas including:



1. Healthcare – ORUA researchers are investigating models for providing OR/Analytics for healthcare systems including data-driven optimised rosters for General Medicine, optimal rostering and dispatch for Patient Transit, and simulation and optimisation for surgery scheduling;
2. Cloud Computing – ORUA's intelligent cloud programme is pioneering concept of OR/Analytics-based intelligence modules and these modules combine with cloud computing modules such as Compute, Store, and Connect to provide intelligent clouds;
3. Government – ORUA researchers are analysing data and forecasting future demand for government services. They will use this data and these forecasts to determine the best way to provide these services and identify the need for future investment.

ORUA's research programmes also realise new tools for OR and Analytics. This work provides innovative tools for use across all the application areas. Recently Dr O'Sullivan has been appointed a theme leader for Precise and Timely Healthcare as part of the Precision Driven Health Research Partnership.

PLENARY SPEAKERS

Dr Kevin Ross

Tuesday 14:00 – 15:00

Kevin Ross is Director of Research at Orion Health, a global healthcare platform managing health records for over 100 million patients worldwide, and General Manager of the Precision Driven Health partnership with District Health Boards and the University of Auckland. He founded and chairs the New Zealand Analytics Forum, a group of over 1700 professionals committed to learning and sharing best practice analytics. Prior to joining Orion, he was Chief Scientist of Optimisation Modelling at Fonterra, where he led the development of analytics capability for product mix, asset investment and forecasting for New Zealand's largest company. He was previously Associate Professor of Technology and Information Management at the University of California, Santa Cruz, and a Principal Consultant for PA Consulting Group. He has worked across multiple sectors, including for NASA, Bell Labs, Eli Lilly and London Councils. Kevin holds a PhD from Stanford University in Management Science and Engineering, and a BSc(Hons) from the University of Canterbury in Mathematics.



Professor Rhema Vaithianathan

Tuesday 16:05 – 17:05

A health economist with a passion for translational research, Rhema Vaithianathan is internationally recognised for her ambitious research using linked administrative data. Using data for social good is a common thread; the research projects that interest her most are those that exist in the junction between 'big data' and 'unsolved social problems'. A recent example is the Allegheny County project, where Rhema led an international research team to develop a child welfare predictive tool for Allegheny County, Pennsylvania, US. The tool, which is now live, makes use of linked data to support better decision making by front line staff as they decide whether or not to open child maltreatment investigations. In New Zealand, Rhema and her team recently completed several impact evaluations commissioned by the Ministry of Social Development, which was seeking empirical evidence on the effectiveness of interventions including the Family Start home visiting programme. In both her domestic and international research, she regularly collaborates with colleagues at AUT and other New Zealand universities, as well as researchers from University of Southern California, University of California, Berkeley and Stanford University. In 2016 she co-founded the Centre for Social Data Analytics at Auckland University of Technology. The Centre is a growing hub for research, education and collaboration on social data analytics, in New Zealand and internationally.



Professor David Morton

Wednesday 13:00 – 14:00

David Morton is a Professor of Industrial Engineering and Management Sciences at Northwestern University. His research interests include stochastic and large-scale optimization, where he seeks to formulate tractable models with appropriate fidelity, develop computational tools of practical value, and provide insights in applications and methodology. His applied interests include security, public health, and energy systems. He received a B.S. in Mathematics and Physics from Stetson University and an M.S. and Ph.D. in Operations Research from Stanford University. Prior to joining Northwestern, he was on the faculty at the University of Texas at Austin, worked as a Fulbright Research Scholar at Charles University in Prague, and was a National Research Council Postdoctoral Fellow in the Operations Research Department at the Naval Postgraduate School. He currently directs Northwestern's Center for Optimization and Statistical Learning <http://osl.northwestern.edu>.



MONDAY 09:20 – 10:20

MA Plenary 1

WG403

Statistics on street corners

Dianne Cook. Monash University, Australia. with Heike Hofmann, Andreas Buja, Deborah Swayne, Hadley Wickham, Eun-kyung Lee, Mahbulul Majumder, Niladri Roy Chowdhury, Lendie Follett, Susan Vanderplas, Adam Loy, Yifan Zhao, Nathaniel Tomasetti.

Perceptual research is often conducted on the street, with convenience sampling of pedestrians who happen to be passing by. It is through experiments conducted using passer-bys that we have learned about the effect of change-blindness (<https://www.youtube.com/watch?v=FWSxSQspiQ>) is in play outside the laboratory.

In data science, plots of data become important tools for observing patterns, making decisions, and communicating findings. But plots of data can be viewed differently by different observers, and often provoke skepticism about whether what you see “is really there”. With the availability of technology that harnesses statistical randomisation techniques and input from crowds we can provide objective evaluation of structure read from plots of data.

This talk describes an inferential framework for data visualisation, and the protocols that can be used to provide estimates of p-values, and power. I will discuss the experiments that we have conducted that

1. show that the crowd-sourcing provides results similar to statistical hypothesis testing,
2. how this can be used to improve plot design, and
3. p-values in situations where no classical tests exist.

Keywords: data visualisation; data mining; data science; statistical graphics; statistical inference; hypothesis testing; exploratory data analysis.

MONDAY 10:40 – 12:25

MB Ordination

WG403

On comparing quantiles of models for count data

John Maindonald. Statistics Research Associates, New Zealand.

Count data commonly follows a distribution that is mildly or strongly positively skewed. Fitted values from a generalized linear type model may be well above the median. An example will be shown where they lie close to the 75th percentile. This makes nonsense of a direct comparison, such as in a paper that is widely quoted in the ecological literature, between such fitted values and fitted values from a least squares fit to log

transformed data. The function `centiles.split()` in R’s `gamlss` package can be helpful in this connection. It allows a comparison, broken down in as much detail as seems useful, between the quantiles of the fitted distribution and the quantiles of the data.

In the choice between generalized linear type models (including the negative binomial) and least squares models fitted to log or power transformed counts, there can be good practical reasons, when counts are large, for preferring to work with log or power transformed counts. A strong case has been made for this approach to work with RNA-Seq counts, used as measures of gene expression.

Keywords: models for count data; negative binomial; quantiles; logarithmic transformation.

A goodness-of-fit test for the ordered stereotype model

Daniel Fernandez[†] & Ivy Liu[‡]. [†]SUNY, University at Albany, USA; [‡]Victoria University of Wellington, New Zealand.

This talk presents a new goodness-of-fit test for an ordered stereotype model used for an ordinal response variable. The proposed test is based on the well-known Hosmer-Lemeshow test and its version for the proportional odds regression model. The latter test statistic is calculated from a grouping scheme assuming that the levels of the ordinal response are equally spaced which might be not true. One of the main advantages of the ordered stereotype model is that it allows us to determine a new uneven spacing of the ordinal response categories, dictated by the data. The proposed test takes the use of this new adjusted spacing to partition data. A simulation study shows good performance of the proposed test under a variety of scenarios. Finally, the results of the application are presented.

Keywords: goodness-of-fit; Hosmer-Lemeshow test; ordinal data; ordered stereotype model; uneven spacing.

Extending factor analysis to binomially distributed data

John Holmes. University of Otago, New Zealand.

Factor analysis is a popular method of identifying a reduced set of latent variables (factors) from a larger set of correlated observed variables. Traditional factor analysis assumes observed variables are drawn from a multivariate normal distribution. Extensions of factor analysis to cases where observed variables are non-normal using a generalised linear model framework tends to be computationally costly due to a lack of closed form solutions. Here we will explore factor analysis when the observed variables are binomially distributed and show how the results in Polson, Scott and Windle on the plya gamma distribution can be used to obtain closed form solutions.

Keywords: factor analysis; Pólya gamma distribution; binomial data.

Maximum likelihood estimation in constrained contingency tables with binary response

Markus Stein. University of Auckland, New Zealand.

Likelihood-based approaches for combining aggregated data with an individual-level data in contingency tables will be considered in this presentation. Such data structure emerges in a wide body of literature, e.g. to reduce bias in “ecological inference” a reliable way consists in supplementing marginal distributions by unit-record information. Similarly this problem arises when researchers working on complex survey designs may take advantage of available marginal population information to improve efficiency. In these cases inferences via maximum likelihood can be computationally infeasible because calculating the likelihood function involves summing over the set of all possible tables. To avoid its calculations a binomial approximation to the true likelihood is presented. We also consider the stochastic EM algorithm and related methods that involve sampling the possible tables. Points of interest are accuracy and computational burden particularly in rare outcome problems. Moreover, we show that good approximations to the likelihood function can be carried out by sampling possible tables and maximising a sample version of this. Simulation results for estimation in a linear logistic model show large improvements by using the estimated likelihood.

Keywords: constrained contingency tables; case-control sampling; stochastic EM algorithm; binomial approximation likelihood; estimated likelihood.

Should we give a damn? The role of data transformation/distance used in analysing count data

Katharina Parry. Massey University, New Zealand.

As in all areas of life, good decision-making depends on being well-informed, and the information in a plot is worth a thousand words as they say. But what if that plot is misleading?

In this talk, you will be introduced to the world of ordinations, a method that allows us to “look” at high-dimensional data. An ongoing discussion regarding these methods is the magnitude of impact on the resulting plots from the data manipulation that usually takes place before the analysis.

This talk will be a visual study into the effects of data handling, in particular differences due to the way in which distances between points in space are defined.

Keywords: ordination; distance measure; procrustes.

multiple non-fatal faults are accumulated and then simultaneously reported and repaired. The reporting process is modelled as a stochastic process dependent on the underlying stochastic process generating the faults. The joint distribution of the reporting times and numbers of reported faults is derived.

Keywords: reliability; warranty analysis; stochastic process.

Strategic bidding in discrete accumulating priorities queue

Raneetha Abeywickrama[†], Moshe Haviv[‡] & Ilze Ziedins[†]. [†]University of Auckland, New Zealand; [‡]The Hebrew University of Jerusalem, Israel.

We consider a single server M/G/1 queue in which customers accumulate priority linearly while waiting. There are a number of priority classes, each of which accumulates priority at a different rate. Upon arrival, without knowing the current state of the system, each customer decides which priority class she should select out of the finite number of options. The greater the accumulation rate, the greater the cost of joining that priority class. When the server becomes idle, the customer with the greatest accumulated priority is chosen for service. If there is a tie, the first come first serve (FCFS) rule is applied. We will discuss the existence, uniqueness and stability of equilibrium strategies for this system.

Keywords: queueing; equilibrium strategies; accumulated priority; stability.

A queueing model for nurse staffing in the patient-at-risk team and intensive care unit

Ali Haji Vahabzadeh & Valery Pavlov. University of Auckland Business School, New Zealand.

The critically ill patients in a hospital’s wards may receive sub-optimal care because their deterioration is not identified or not acted upon timely by the ward nurses. On the contrary, a low-risk patient might be wrongly identified by a junior nurse in the ward as an acutely ill patient who needs to be admitted to an intensive care unit (ICU). To detect the early signs of critical health deterioration among patients in the ward and also avert preventable ICU admissions, hospitals establish a patient-at-risk (PAR) team. However, today’s global nursing shortage causes a serious and complex challenge for hospital managers on determining nurse staffing levels in hospital’s medical units. In this paper, we propose a queueing model to examine the impact of nurse staffing level between PAR and ICU on the ICU patient mortality rate. Using operational flow of patients of a large hospital in New Zealand, we build a discrete-event simulation model. Our preliminary results obtained from optimisation provide us with some insights into how the precious critical care nurses should be staffed with the PAR team and ICU.

Keywords: healthcare operations; patient flow; rapid response team; critical care outreach team.

Delayed reporting of faults in warranty claims

Richard Arnold[†], Stefanka Chukova[†] & Yu Hayakawa[‡]. [†]Victoria University of Wellington, New Zealand; [‡]Waseda University, Japan.

We present a model for the delayed reporting of faults:

Hidden Markov models with discrete bivariate emissions

Rolf Turner[†] & **Jonathan Williams**[‡]. [†]*Department of Statistics, University of Auckland, New Zealand;* [‡]*Northland District Health Board, New Zealand.*

This talk is on joint work with Dr. Jonathan Williams of the Northland DHB.

We seek to analyse the association between sequences of observations of levels of psychosis of patients and observations of their blood cell counts. Both sequences are discrete-valued, with the intrinsically discrete cell counts further discretised (on a 1 to 5 scale). The long-term goal is to investigate the possibility of predicting psychotic episodes from the values of blood counts or changes in these.

We fitted a hidden Markov model to the data (1258 subjects), by adapting my R package `hmm.disnp` to handle bivariate data. The major finding so far is that there is *highly* significant evidence of statistical dependence between the two observation sequences.

It was necessary to introduce a numerical optimisation procedure into the M-step of the EM algorithm employed. This necessity was in turn induced by the presence of missing values, intrinsic to the nature of the data. The maximum likelihood calculations used are valid only if these values are missing at random (MAR) – a dubious assumption in current circumstances. Simulation studies will be carried out to assess the sensitivity of the results to the MAR assumption.

Keywords: psychosis; blood cell counts; EM algorithm; missing at random.

User equilibria in systems of processor sharing queues

Niffe Hermansson, **Ilze Ziedins** & **Mark Holmes**. *University of Auckland, New Zealand.*

We consider properties of user equilibria in parallel processor sharing queues. Previous work has shown that these systems exhibit interesting, and sometimes perplexing, behaviours. In this presentation we will see some examples of surprising system behaviour, but also some encouraging properties of the system at equilibrium.

Keywords: processor sharing; user equilibria; queueing; networks.

ferences and conclusions. A basic problem is the unique identification of individuals. We consider two different MCMC algorithms for determining the true genotypes of animals based on observed genotypes. We can show that both algorithms produce irreducible chains which implies other properties that indicate the existence and uniqueness of a stationary distribution.

In this talk, we compare the convergence properties of both algorithms using an illustrative example which allows visualize the behaviour of the chains. Fascinatingly, both chains look extremely different but they converge to the same stationary distribution.

Keywords: irreducibility; convergence; missing data; Markov chains.

One step, two step, big step, small step: A comparison of MCMC methods for non-identifiable models

Sarah Pirikahu, **Geoff Jones** & **Martin Hazelton**. *Massey University, New Zealand.*

A model which is identifiable possesses the desirable property, that as you collect more data you are able to more accurately estimate your parameters of interest. For the majority of *frequentist* methods, identification is considered a minimum requirement. Over time though, non-identifiable models have become increasingly common in the field of epidemiology. The obvious solution to analysing these non-identifiable models has been to adopt a Bayesian approach; because as long as you have your model and appropriate prior information you can turn the Bayesian crank to come up with a solution. However, well-known MCMC samplers, such as Gibbs, Metropolis-Hastings and Metropolis-adjusted Langevin may not work well.

In this talk we discuss why these well-known MCMC samplers fail to converge to the required stationary distribution for large data sets, or require such a large amount of compute time that they become inefficient. Then through an example, that analyses the population attributable risk of leptospirosis in New Zealand (and its corresponding uncertainty), we compare the performance of these samplers; and show that an alternative importance sampling method performs much better for large data sets and can be used to explore the limiting posterior given infinite sample size.

Keywords: MCMC; population attributable risk; measurement error; limiting posterior.

Statistical principles on validating formative concepts: An annotated example

Ishani Soya, **Nihal Jayamaha** & **Nigel Grigg**. *Massey University, New Zealand.*

In science, a concept (e.g. pressure, temperature, viscosity) is conventionally viewed as a phenomenon that exists irrespective of the particular measurement scale or instrument (e.g. a pressure gauge) that is being used to observe the variation of the phenomenon. This principle holds true in social and behavioural sciences also. The key difference between a behavioural science measurement scale and a hard science measurement

MB Mixture Models / MCMC WA224

Markov chains for inferring true identity of individuals in a sample

Paula Bran, **Richard Barker** & **Matthew Schofield**. *University of Otago, New Zealand.*

In many situations there is uncertainty of the accuracy of the information in a dataset. Especially because the observations may be either duplicates, or wrongly reported, or missing. All of these result in error-prone data collection which seriously affect the posterior in-

scale is that the former usually requires multiple measures because the concept/phenomenon under consideration is not (normally) directly observable. The measures of conventional behavioural science concept are required to be strongly correlated with concept (convergent validity) and the measures themselves should covary because they mere reflections or manifestations of the concept. The statistical principles of validating a conventional behavioural science measurement scale are well established in the psychometric literature. In management science, however, on occasions, scientists need to form concepts, whose very meaning is dependent upon the measures being used to define the concept! Such concepts are known as formative concepts (Fornell and Bookstein, 1982; Gefen et al., 2000). Although formative concepts have a reasonably long history, the statistical principles of validating formative concepts are emerging (Hair et al., 2014).

In this presentation we explain how the formative construct 'Financial Health' in relation to a nonprofit organisation is validated using data that the first author collected for their research.

Keywords: formative constructs; validity; statistical principles.

Performance of Watanabe-Akaike Information Criterion (WAIC) and importance sampling methods in leave-one-out cross-validation approximation

M.A.C.S. Sampath Fernando, James M. Curran & Renate Meyer. University of Auckland, New Zealand.

Statistical models are widely used in explanation, prediction, or making inferences on real-world phenomena and various models can be derived using different modelling approaches.

In the Bayesian context, evaluation of the goodness-of-fit and comparison of the fitted models are essentially performed with posterior predictive checks. Cross-validation (CV) and information criteria are the mostly used tools for estimating unbiased measures of out-of-sample prediction error which assesses the predictive model accuracy. Even though many versions of cross-validations are available for evaluating prediction error, they all essentially take the same form. Leave-one-out cross-validation (LOO-CV) is the most computationally expensive variant among other CV methods, since it fits as many models as the number of observations. Utilising the existing MCMC results, importance sampling (IS-LOO), truncated importance sampling (TIS-LOO), and Pareto-smoothed importance sampling (PSIS-LOO) LOOs are defined as approximations to the exact LOO-CV. While IS-LOO, TIS-LOO, and PSIS-LOO are considered to be approximations of the exact LOO-CV, some drastic deviations were observed in the results. However, there are also some interesting relationships among the results of these approximation methods.

Keywords: cross-validation; importance sampling; information criteria; predictive accuracy.

A new cost recovering, revenue adequate in expectation single settlement scheme for electricity markets

Ryan Cory-Wright, Andy Philpott & Golbon Zakeri. Engineering Science, University of Auckland, New Zealand.

The penetration of intermittent renewable sources in electricity markets is increasing, because of the COP-21 Paris agreement, and because renewable sources have a marginal cost of 0. This intermittency increases the uncertainty in electricity markets, and the amount of expensive re-adjustment required from other generators when renewables deviate significantly from their forecast capacity.

The expected cost of deviating can be reduced by explicitly modelling the uncertainty in the capacity of renewable generators using a probability distribution, and dispatching generators according to the optimal solution of a stochastic program, which considers the expected cost of deviating. The first and second stages of this stochastic program correspond to policies for day-ahead and real-time dispatches under different realisations of a random variable.

An open question is how generators should be paid for their day-ahead and real-time dispatches. We contribute to this debate by deriving a single-settlement, discriminatory payment scheme which is revenue adequate in expectation and provides generator cost recovery in every scenario.

We discuss the attributes of this settlement scheme, and compare it with other schemes in the literature, in both theoretical terms and in the context of the New Zealand Electricity Market under April 2015 conditions.

Keywords: electricity markets; stochastic programming; pricing scheme.

Water behind capacity

Ellen Krohn Aasgard & Pål Henrik Roti. Norwegian University of Science and Technology, Norway.

This paper considers optimal scheduling for hydropower producers that deliver energy and certain types of system services. The focus is on the Nordic and European power system, but results may be generalized to other regions where hydropower is an important resource for electricity supply. A future power system based on intermittent renewables will depend on system services being supplied by conventional energy sources. As interconnections grow stronger, Nordic hydropower may deliver balancing services to the rest of the European power system. However, even if instant capacity is available from hydropower units, available energy is not always guaranteed as the consequences of real-time activation are often not analyzed for the entire multi-reservoir system. Water behind capacity is a term that describes the specific challenge of hydropower producers to comply with all technical, hydro-

logical and environmental constraints in the river chain in the case of activation.

In this work we analyze the effect of activation using a combined optimization-simulation approach. We illustrate by simple case study examples that the amount of reserves which can actually be delivered is limited, even though more capacity can easily be reserved from the hydropower units.

Keywords: hydropower; renewable energy; reserves; system services.

Electricity contracts and price premia

Hamish Mellor & Andy Philpott. University of Auckland, New Zealand.

In New Zealand, electricity spot prices fluctuate drastically due to a reliance on hydro-power. Such fluctuation requires participants in the market to trade hedge contracts to mitigate this risk. These contracts have prices agreed upon before the spot price is known. In a perfectly competitive market, the average value of a contract for a set period of time should equal the expected value of the average spot price over this time period. It has been observed that this may not be the case in NZ.

This project sought to quantify these differences, referred to as premia. It involved calibrating an existing mathematical model to recreate spot prices observed in 2008. This model was then simulated out of sample in 2012, a hydrologically comparable year, to produce a distribution of expected spot prices. This was then compared to contract prices for the same period to find evidence for premia existing.

The key conclusions: prices do not follow trajectories that would be expected from a perfectly competitive market; there is evidence of direct impacts on prices from the contracting positions of generators; factors other than hydrology are impacting the price of electricity; contracts are trading at significant premia over price expectations.

Keywords: electricity markets; hedge contracts; spot prices.

The impact of real world variation on the commercial optimisation application, pump scheduling

Cathy Daly. SUEZ (Derceto Ltd.), New Zealand.

For a water utility the use of a pump scheduling application can lead to cost savings in the millions of dollars per year. However, optimising energy costs without considering real world variation can lead to unsatisfactory schedules that conflict with required operation.

In this presentation we'll discuss the impact of real world variation in the form of required variation, variable network conditions, and operational interventions. The following examples will be used to convey issues that arise when real world variation is not accounted for, such as:

- ▷ water quality concerns caused by poor turnover of water in storage tanks;
- ▷ tank levels outside preferred operating levels due to discrepancies between observed and predicted water demand; and
- ▷ tank levels and pipe pressures outside preferred operating levels/pressures that occur when part of the network is manually operated and in mismatch to the optimal schedule.

Where required, variation, variable network conditions and operational interventions are accounted for in the optimisation application, the risk of unsatisfactory schedules can be mitigated and a near optimal schedule can be achieved.

Keywords: water; sustainability; commercial; real-world; pump; scheduling; variation.

Transmuted exponentiated modified Weibull distribution with an application to bladder cancer data

Muhammad Shuaib Khan, Robert King & Irene Lena Hudson. The University of Newcastle, Callaghan, Australia.

This research investigates the potential usefulness of the transmuted exponentiated modified Weibull (TEMW) distribution for modelling bladder cancer data. This distribution, formed using the quadratic rank transmutation map technique, contains twenty three lifetime distributions as the special cases. We obtain the analytic shapes of the density and hazard functions. Some structural properties of the transmuted exponentiated modified Weibull distribution are discussed, with special emphasis on its moments. We discuss estimation of the model parameters by the method of maximum likelihood and provide an application for bladder cancer data.

Keywords: exponentiated modified Weibull distribution; moment estimation; maximum likelihood estimation.

The Young Practitioner Prize is sponsored by



MONDAY 13:15 – 15:00

MC Spatio-temporal Modelling WG403**Practically relevant academic programmes for training data scientists***Leo Paas. Massey University, New Zealand.*

The failure to meet the growing need for data scientists in the new big data driven economy has been pointed out at many forums and conferences. In this presentation an avenue will be proposed to mitigate this growing need. The solution is based on academia and business breaking out of their silos and collaborating in multi-disciplinary programmes aimed at educating students to become analysts and data scientists that can work in a practically relevant manner for business.

Keywords: academic education; data science; big data.

Spatio-temporal modelling of air quality using R-INLA

Nokuthaba Sibanda[†], Munyaradzi Mujuru[‡] & Inos Dhau[‡]. [†]Victoria University of Wellington, New Zealand; [‡]Tshwane University of Technology, South Africa; [‡]University of Limpopo, South Africa.

The impact of poor air quality on the environment and subsequently to human health is now well recognised throughout the world. Air quality is determined by measuring concentration levels of several pollutants in the air, and these can be compared to WHO thresholds for acceptable pollutant levels. In Metropolitan Harare, Zimbabwe, air quality is a source of concern. To assess pollutant concentration levels over an entire region, environmental agencies need models to predict these levels at unmonitored sites. We use spatio-temporal modelling of four pollutant levels monitored at eight different sites to predict these levels throughout the city, focusing on the class of Bayesian hierarchical models. The aim is to develop a statistical model for determining pollutant concentration maps, with accompanying uncertainty measures. The limited number of monitoring stations means there is a greater level of uncertainty for some predictions. The models are fitted using the R software package R-INLA.

Keywords: air quality monitoring; spatio-temporal modelling; Bayesian hierarchical models.

Statistical methods and considerations to model growth and clearance of malaria parasitemia in human challenge studies

Louise Marquart, Leesa Wockner, Peter O'Rourke & James McCarthy. QIMR Berghofer Medical Research Institute, Australia.

Continuing emergence of antimalarial drug resistance

and renewed focus on worldwide eradication of malaria underlie research efforts to use induced blood stage malaria (IBSM) human challenge models to evaluate candidate vaccines and medications for malaria. Critical to the success of IBSM is the ability to monitor parasitemia with high levels of accuracy at low parasite densities, such as by real time quantitative polymerase chain reaction (qPCR). Ten IBSM studies with 127 participants, conducted since 2012, provide a rich source of pre- and post-treatment parasite count data to assess parasite growth and clearance.

Statistical models that can account for the malaria parasite life-cycle are fitted to evaluate growth rates and the pharmacodynamic effect of candidate antimalarials. Parasite growth and clearance models need to account for the cyclical nature of the sequestering parasite, and any lag and tail decay curve phases in parasitemia following antimalarial treatment.

This study describes the statistical considerations required when implementing non-linear mixed effect models to estimate parasite growth rates, and describes an analytic approach to estimate optimal parasite clearance. Furthermore, special considerations are required when using qPCR derived data, such as detecting outliers and handling values below the limit of detection and non-detects.

Keywords: statistical models for clinical studies; malaria parasite growth and clearance; non-linear mixed effect models.

A spatio-temporal model for outbreak surveillance

Jonathan Marshall. Massey University, New Zealand.

Most cases of food and water-borne illness in New Zealand are due to isolated events where only one or two people get sick, rather than being due to large outbreak events with common exposures.

The recent water-borne outbreak of campylobacteriosis in Havelock North, however, is a reminder that this is not always the case. Identifying cases likely to be related to a common exposure is important in order to inform public health guidelines and prevent illness in the future. When such outbreaks are large they are typically identified easily as they stand out clearly above trends in sporadic cases. When outbreaks are smaller, identifying them can be difficult.

One technique is to model the underlying spatial and temporal trends in risk, with a spatio-temporal field of outbreak indicators over and above these trends. A Bayesian model will be presented and applied retrospectively to cases of campylobacteriosis in the Manawatu region.

Keywords: public health; outbreak; disease surveillance; spatio-temporal model.

Testing for changing patterns of spatial relative risk*Martin Hazelton. Massey University, New Zealand.*

The spatial relative risk function is a useful tool for examining geographical patterns of disease. It is defined as the ratio of bivariate densities of the spatial coordinates of cases and controls for the disease of interest, and can be implemented in practice using kernel density estimates constructed from case-control data.

Estimation of relative risk functions from datasets collected over a sequence of time periods provides an opportunity to assess whether the spatial pattern of disease is changing. Informal inference can follow by simply inspecting the succession of risk functions using heat maps.

In this talk I will describe how formal testing can proceed based on various statistics measuring the temporal variation in the estimated relative risk functions. I will pay particular attention to the choice of smoothing bandwidth for the kernel density estimates, a perennial problem in this kind of kernel-based testing procedure.

Keywords: bandwidth; epidemiology; density estimation; kernel-based test; spatial risk.

Control charts scheme for high dimensional monitoring*Nurudeen Adegoke, Mat Pawley & Adam Smith. Massey University, New Zealand.*

Multivariate control charts typically assume the knowledge of in-control covariance structure. In practice, this assumption does not always hold, and a suitable estimate is obtained from the in-control reference sample of preliminary observations in Phase I, using the empirical covariance estimate or its related maximum likelihood estimate. The variability of the estimated parameter can result in chart performance that differs from the actual performance. However, using a representative subgroup size from Phase I would reduce the variability introduced into the system by parameter estimation, and there is a consensus on the importance of large subgroup size m for the estimation process from Phase I analysis. The restrictions on the choices of m and process characteristics p have limited the applications of the schemes to small p and large m data sets, which make the methods unsuitable for high-dimensional monitoring.

In this study, we propose an improvement to the monitoring procedures of the conventional memory-type control charts: multivariate exponentially weighted moving average (MEWMA) and multivariate cumulative sum control (MCUSUM) charts, for high dimensional monitoring. We propose the use of a shrinkage covariance estimate from Phase I, and we use simulation to determine the efficacy of this new procedure.

Keywords: covariance; estimation; empirical; shrinkage.

MC Statistical Theory 1 WG126**Contracted confidence intervals: a frequentist technique for greater efficiency when an existing estimate is sufficiently accurate***Robin Willink. Ministry of Business, Innovation and Employment, New Zealand.*

This talk presents a general construction of a confidence interval with reduced mean length when a prior estimate lies within a few standard errors of the true value. Examples show that when the prior estimate is accurate, a reduction in mean length of 10 to 40 percent can be observed depending on user-defined parameters of the procedure. The cost is increased mean length when the estimate is inaccurate. The limits of the interval are formed by moving one or both of the limits of the standard, unadjusted, interval towards the prior estimate.

A proof is given of the validity of the confidence interval, i.e. that the confidence coefficient is at least the stated value and so repeated use of the procedure achieves at least the implied proportion of correct inferences. No prior distribution or agreed prior information is required. The only type of prior opinion involved is a point estimate, and there does not need to be any agreement about this estimate.

Keywords: statistical inference; interval estimation; confidence interval; prior information; shrinkage estimation.

A functional central limit theorem on random fields for nested sampling domains*Leshun Xu, Alan Lee & Thomas Lumley. The University of Auckland, New Zealand.*

A functional central limit theorem is often called a Donsker's theorem or an invariance principle.

In this talk, we prove a functional central limit theorem on non-stationary random fields with strong mixing properties. We suppose the sampling domains are nested as they are growing. Then the standardized time-dependent process of the total convergences to a Brownian motion in distribution.

Keywords: functional central limit theorem; strong mixing; random field; nested sampling; non-stationary.

Anisotropic spatial modelling with Template Model Builder*Andrea Havron & Russell Millar. University of Auckland, New Zealand.*

Recent advancements in spatial modelling have been made using the new platform, Template Model Builder (TMB). This R package interfaces with a C++ automated differential modelling environment to simultaneously model fixed and spatial random effects within a

full likelihood-based framework. The likelihood of the multivariate normal spatial random effects are calculated directly by using Q , the inverse of the covariance matrix Σ . Model computation time is improved by the sparse property of this precision matrix. Q is estimated by a Gaussian Markov Random Field using a Stochastic Partial Differentiation Equation, which has a Gaussian Field with Matern covariance function as its solution. Lindgren et al. 2011 showed that this solution can be scaled by a d -dimensional H matrix for an anisotropic extension.

I will review this solution along with TMB C++ code developed by Kristensen et al. 2016. I will then apply these methods to a case study of surfclam biomass data from the Manawatu coast.

Keywords: spatial modelling; Template Model Builder; anisotropy.

An improved two-stage approach for joint modelling of longitudinal and time-to-event data

Huong Thi Thu Pham & Nur Darfiana. Flinders University, Australia.

Joint models for longitudinal and time-to-event data have been applied in many different fields. However, the main difficulty these models have to face is the computational problem. The requirement for numerical integration has become severe when the dimension of random effects increase. In this paper, a modified two-stage approach has been proposed to estimate the parameters in the joint models. This approach can reduce the computational time and allow for the application of extended joint models with a high dimension of random effects in the longitudinal submodel. In particular, in the first stage, the LMEs (linear mixed effects models) and BLUPs (best linear unbiased predictors) are applied to estimate parameters in the longitudinal submodel. Then, in the second stage, the approximation of the fully joint log-likelihood is proposed using the estimated values of these parameters from the longitudinal submodel. The survival parameters are estimated by maximizing the approximation of the fully joint log-likelihood. Simulation studies show that the modified two-stage approach performs well and is comparable with the existing shared random effects. Moreover, this approach also shows good results when the dimension of random effects increases.

Keywords: survival data; longitudinal data; two-stage approach; shared random effects approach; joint models.

populations. In stage-duration distribution models, parameter estimation for the models use the Laplace transform method. This method involves assumptions such as known constant shapes, known constant rates or the same overall hazard rate for all stages. These assumptions are strong and restrictive. The main aim of this paper is to weaken these assumptions by using a Bayesian approach. In particular, a Metropolis-Hastings algorithm based on deterministic transformations is used to estimate parameters. We will use two models, one at which has no hazard rates, and the other has stage-wise constant hazard rates. These methods are validated in simulation studies followed by a case study of cattle parasites. The results show that the proposed methods are able to estimate the parameters comparably well, as opposed to using the Laplace transform methods.

Keywords: multi-stage models; stage duration; stage frequency data; Bayesian analysis; destructive samples.

Bayesian non-parametric spectral density estimation using B-spline priors

Matthew Edwards[†], Renate Meyer[†] & Nelson Christensen[‡]. [†]University of Auckland, New Zealand; [‡]Carleton College, USA.

We present a new Bayesian non-parametric approach to estimating the spectral density of a stationary time series. A non-parametric prior based on a mixture of B-spline distributions is specified and can be regarded as a generalization of the Bernstein polynomial prior of Petrone (1999a,b) and Choudhuri et al. (2004). Whittle's likelihood approximation is used to obtain the pseudo-posterior distribution. This method allows for a data-driven choice of the smoothing parameter as well as the number and the location of the knots. Posterior samples are obtained using a parallel tempered Metropolis-within-Gibbs Markov chain Monte Carlo algorithm. We conduct a simulation study to demonstrate that under default noninformative priors, the B-spline prior provides more accurate Monte Carlo estimates in terms of L_1 -error and uniform coverage probabilities than the Bernstein polynomial prior. Finally, we demonstrate the algorithm's ability to estimate a spectral density with sharp features, using real gravitational wave detector data from LIGO's sixth science run, recoloured to match the Advanced LIGO target sensitivity.

Keywords: B-spline prior; Bernstein polynomial prior; Whittle likelihood; spectral density estimation; LIGO; gravitational waves.

MC Bayesian Statistics

WA224

A Bayesian approach for parameter estimation in multi-stage models

Hoa Thi Thu Pham & Alan Branford. Flinders University, Australia.

Multi-stage time evolving models are common statistical models for biological systems, especially insect

Bayesian option pricing framework with stochastic volatility for FX data

Rachel Wang[†], Boris Choy[‡] & Hoi Ying Wong[†]. [†]The Chinese University of Hong Kong, Hong Kong; [‡]The University of Sydney, Australia.

The application of stochastic volatility (SV) models in the option pricing literature usually assumes that the market has sufficient option data to calibrate the

model's risk-neutral parameters. When option data are insufficient or unavailable, market practitioners must estimate the model from the historical returns of the underlying asset and then transform the resulting model into its risk-neutral equivalent. However, the likelihood function of an SV model can only be expressed in a high-dimensional integration, which makes the estimation a highly challenging task. The Bayesian approach has been the classical way to estimate SV models under the data-generating (physical) probability measure. We propose an SV model that enables us to simultaneously and conveniently perform Bayesian inference and transformation into risk-neutral dynamics. Our model uses the Student-t and variance gamma distribution to replace the normality assumption on innovations of both return and volatility processes, and an empirical study shows that the estimated option prices generate a realistic implied volatility smile shapes. In addition, the volatility premium is almost flat across strike prices, so that adding a few option data to the historical time series of the underlying asset can greatly improve the estimation of option prices.

Keywords: option pricing; volatility smile; Student-t; variance gamma; MCMC.

DNest4: Diffusive nested sampling in C++ and Python

Brendon Brewer. Department of Statistics, University of Auckland, New Zealand.

About a decade ago, the novel Bayesian computation algorithm “Nested Sampling” (NS) was introduced by physicist John Skilling. The method has two main strengths that distinguish it from many of its competitors:

1. it can cope with awkward structures in the posterior distribution that defeat most other methods; and
2. it computes marginal likelihoods needed for model selection.

I will discuss a variant of called Diffusive Nested Sampling and its implementation in the DNest4 software package, through examples from astrostatistics.

Keywords: Bayesian computation; Markov chain monte carlo; nested sampling; phase transitions.

Non-parametric Bayesian spectral density estimation of stationary time series using a generalized Whittle likelihood

Renate Meyer[†], Claudia Kirch[‡], Matthew Edwards[†] & Alexander Meier[‡]. [†]University of Auckland, New Zealand; [‡]University of Magdeburg, Germany.

Non-parametric Bayesian inference has seen a rapid growth over the last decade but only very few non-parametric Bayesian approaches to time series analysis have been developed. Most notably, Carter and Kohn (1997), Gangopadhyay (1998), Choudhuri et al. (2004), and Rosen et al. (2012) used Whittle's likelihood for Bayesian modelling of the spectral density

as the main non-parametric characteristic of stationary time series. As shown in Contreras-Cristan et al. (2006), the loss of efficiency of the non-parametric approach using Whittle's likelihood can be substantial. On the other hand, parametric methods are more powerful than non-parametric methods if the observed time series is close to the considered model class but fail if the model is misspecified. Therefore, we suggest a non-parametric correction of a parametric likelihood that takes advantage of the efficiency of parametric models while mitigating sensitivities through a non-parametric amendment. We use a non-parametric Bernstein polynomial prior on the spectral density with weights induced by a Dirichlet process and prove posterior consistency for Gaussian stationary time series. Bayesian posterior computations are implemented via a MH-within-Gibbs sampler and the performance of the non-parametrically corrected likelihood is illustrated in a simulation study and LIGO gravitational wave data.

Keywords: Bayesian non-parametrics; stationary time series; Whittle likelihood; spectral density estimation; laser interferometric gravitational wave data.

YPP Session 2 / Transport WA220

MC

Risk-taking behaviour of public transport users

Chapa Jayalath[†], Anthony Downward[†] & Subeh Chowdhury[‡]. [†]Engineering Science, University of Auckland, New Zealand; [‡]Civil Engineering, University of Auckland, New Zealand.

Globally, public transport (PT) agencies are implementing integrated PT systems to increase ridership through modal shift from private cars. Transfers are an integral component of such integrated systems; they generate connectivity among modes thus allowing the system to function as one unit. However, it is well known that transfers can cause users to feel anxious due to the possibility of missed connections and being physically out of the vehicle. The PT network in Auckland is undergoing a transformation into an integrated system, meaning fewer direct routes and more connected routes. Utilising a previous survey done to understand the PT users' risk perceptions, this research project attempts to model users' risk aversion and forecast the use of the proposed transfer routes. Risk-averse commuters value reliability over savings in travel time. Transfer routes that provide better reliability and shorter travel time are seen as a more attractive option than a direct route, despite the inconvenience of having to change vehicles. In this study we find that females are more risk-averse than males and would prefer certain routes. Auckland Transport therefore has an opportunity to identify routes that have higher proportion of a particular gender and adopt targeted measures to enhance their commute.

Keywords: public transport; transfers; risk aversion; route choice; service quality.

Real-time prediction of bus arrival using joint models of vehicle and road states

Tom Elliott & **Thomas Lumley**. *University of Auckland, New Zealand.*

Studies have shown that an arrival-time countdown at bus stops can reduce passengers' perceived waiting times; of course, this assumes some level of accuracy in the predictions. For a variety of reasons, this is not often the case.

We present an approach to arrival time prediction that, rather than modelling only a single vehicle or route, takes into account all other transit vehicles in the area. Using a particle filter (PF) in real-time to model the state (location and speed) of all busses in Auckland, we generate a "map" of current traffic conditions for all bus routes, which is updated using a Kalman filter (KF). This in turn is used to refine the particle filter estimates for subsequent busses.

To make arrival time predictions, we use the KF state estimates to compute a vehicle's arrival time at all remaining stops, independently for each particle. This generates a distribution of plausible arrival times accounting for uncertainty, such as from traffic lights and time spent at stops dropping off and picking up passengers. To communicate uncertainty to passengers, we consider using prediction intervals (2 – 5 minutes) rather than point estimates (3 minutes), both of which are straightforward to obtain from the PF sample.

Keywords: real-time information; public transport; arrival time prediction; particle filter.

Mapping electric vehicle range

Danielle Gatland & **Andrea Raith**. *Engineering Science, University of Auckland, New Zealand.*

Electric vehicles have been growing in popularity in recent years, however in 2016 they still make up only a small proportion of New Zealand's vehicle fleet. One of the biggest barriers for drivers considering transitioning towards electric vehicles is range anxiety. An interactive map which accurately shows the full range of different models of electric vehicles is a powerful tool, which might increase the confidence and awareness of drivers in the capabilities of their electric vehicle. The reachable region of an electric vehicle can be found by constructing a shortest path tree, using an energy consumption formulation to determine path weights. The formulation used was based on theories of physics, and appears to reasonably estimate the reachable region of electric vehicles.

We developed an interactive tool to allow a user to investigate the range of a return trip of various models of electric vehicles. The user selects a starting point anywhere in New Zealand, which is used to determine the reachable region of a return trip of the vehicle.

Keywords: electric vehicle; range anxiety; energy consumption; visualisation.

Modelling traceability in the dairy industry

Melissa Welsh. *Victoria University of Wellington, New Zealand.*

Traceability is the capability to trace goods throughout the distribution chain. Traceability has become an increasingly important research area in recent years. It has always been an important aspect of production, but recent contamination events have highlighted its significance. The Fonterra botulism scare of 2013 in particular exposed a need for fast accurate product tracing in the New Zealand Dairy industry.

We present a Markov chain model for the flow of milk through the early stages of the dairy supply chain. The state of the Markov chain is the value of product at each location in the production chain, in this case the milk tanker, factory reception or processing.

The model incorporates parameters for product testing and tracing upon arrival in each state of the model. By varying these parameters we are able alter the precision of the traceability system, and gain an understanding of where and when traceability has the greatest impact. By analysing the results of simulations under various scenarios we are able to estimate the value traceability can contribute to the output of the production chain.

Keywords: traceability; dairy; supply chain; Markov chains.

A solution technique for rostering in hospital departments

Thomas Adams, **Michael O'Sullivan** & **Cameron Walker**. *Engineering Science, University of Auckland, New Zealand.*

Rosters in hospital departments directly affect the number of patients physicians are caring for each day, because the roster dictates who admits new patients. Therefore balance of workload across physicians can be controlled to a degree by how the physicians are rostered. To build a roster exhibiting a good workload balance across physicians a mixed integer programme is solved that utilises patient admission and discharge numbers.

Traditional rostering models do not include patient admissions and discharges. Including this data results in a model that is larger and more difficult to solve. A more difficult model to solve is not a concern for small problems. In particular, although the "vanilla" formulation can be solved for small departments (including admissions and discharge data over moderate time windows), it does not scale well, and so is less applicable for larger problem instances.

To reduce the time it takes to solve the model a new solution technique is developed, within the context of creating a roster for the General Medicine Department at Waitemata District Health Board.

Keywords: mixed-integer programming; rostering; healthcare modelling.

MONDAY 15:20 – 16:00

MD Software Packages

WG403

The web browser as an R graphics device

Paul Murrell. Department of Statistics, University of Auckland, New Zealand.

This talk will describe and demonstrate the **DOM** package for R. This package allows an R session to connect to a web browser via a two-way websocket. It is then possible to manipulate the content of a web page in the browser from the R session. It is also possible to call back to R in response to user events in the web browser. Possible applications include simple interactive graphics.

Keywords: R; statistical graphics; HTML; SVG; CSS.

DBSCAN clustering use in the electricity distribution network. VBA to R package

Garry Julian. Powerco, New Zealand.

The algorithm DBSCAN (Density Based Spatial Clustering for Applications with Noise) has now been used for a year for detecting load backfeeds in an electricity distribution network. This talk will explain the need for a clustering algorithm such as DBSCAN in the electricity industry. The presentation covers the use of a VBA implementation of DBSCAN in Excel then finishes by talking about the writing of an R package called **odcR** (Outlier Detection using Clustering in R) encapsulating an R package implementation of DBSCAN. A practitioner's journey.

Keywords: electricity distribution network; clustering; R package; DBSCAN.

MD Stochastic Modelling 2

WG126

A new structural approach in modelling hazard rate with covariates

Budhi Arta Surya. School of Mathematics and Statistics, Victoria University of Wellington, New Zealand.

There are two main approaches in modelling hazard rate: reduced form and structural approaches. In reduced form, occurrence of an event is caused by some exogenous factors represented by covariates. The hazard rate associated to the first time until an event observed is given by Cox proportional model (Cox, 1972). In contrast to reduced form, occurrence of an event in structural approach is determined by an underlying stochastic process. The observation of an event continues until the first time the underlying process enters a stopping region, which is a stopping time. The hazard rate is given by the ratio between the density and survival function of this stopping time. Aalen (1995) and Aalen and Gjessing (2001) introduced phase-type

distribution of homogeneous Markov chains and distribution of first-passage of a class of stochastic processes. However, because of the underlying process is Markovian, the hazard rate does not depend on the past information of underlying process (covariates). In this paper, we present following the recent work of Surya (2016) a new structural approach for modelling hazard rate which is given in closed form, and has the ability to capture heterogeneity covariates. The result closes the gap between reduced form and structural approach.

Keywords: reduced form approach; structural approach; hazard rate modeling; Markov mixture process; generalized phase-type distribution; survival analysis.

Markov processes: Dynamic visualizations to enhance student understanding

Maxine Pfannkuch & Stephanie Budgett. University of Auckland, New Zealand.

Finding ways to improve introductory students' understanding of probability ideas and theory is a goal of many first-year probability courses. In this presentation we explore the potential of a prototype tool for Markov processes using dynamic visualizations to develop in students a deeper understanding of the equilibrium and hitting times distributions. The tool and tasks were piloted on six introductory probability students using a two-person protocol. The talk will focus on the students' interactions with the Markov processes tool. The main findings of our exploratory study suggested that the tool and tasks have the potential to enhance students' probabilistic reasoning. The tool seemed to assist students to engage with and develop some intuition for Markov processes, to enhance their distributional ideas, to work flexibly between representations and to see mathematical structure.

Keywords: statistics education; introductory probability students; equilibrium distribution; hitting times distribution; animations.

Transportation Modelling

WA224

MD

Network aggregation with applications to evacuation planning

Bob Grün. University of Auckland, New Zealand.

For the safety of cities and events, the construction of evacuation plans is very important. These mathematical constructed plans often base on often too large graphs. We therefore considered two methods to reduce the graph size. The idea of both approaches is to reduce crossovers and roundabouts to a single node. Thus, in the first approach, we contract "small" edges such that the two connected nodes of an edge are merged. In the second approach, "small" edges are removed if there is a reasonable detour of this edge. It is clear that the contraction as well as the removal of an edge cause an error. In order to have information on the impact of the total error, we measure the quality of the graph aggregation through the average shortest path length change. From the theoretical point of

view, we consider worst case graphs and thus show error bounds on the average error. Even if both aggregation methods have worse theoretical error bounds, we have shown that for real world graphs, both methods have good results. Since the motivation of the aggregation rise from street networks, we consider the average aggregation error for 280 real world cities with different graph sizes.

Keywords: aggregation; evacuation planning; graph reduction.

Modelling fuel consumption in traffic assignment problems

Andrea Raith[†], *Clemens Thielen*[‡] & *James Tidswell*[†].
[†]Engineering Science, University of Auckland, New Zealand; [‡]Katholische Universität Eichstätt-Ingolstadt, Germany.

The Traffic Assignment (TA) Problem models route choices of transport network users assuming fixed demand between origins and destinations. An assumption of TA is that network users selfishly minimise their own travel time: TA solutions follow the so-called user equilibrium (UE). UE flow does not necessarily follow a system-optimal (SO) travel pattern (with respect to travel time or generalised cost). A SO travel pattern can be computed and congestion pricing theory shows that it can be enforced in a UE by charging network users appropriate congestion tolls.

We develop a model of fuel consumption within the TA framework. Our aim is to derive an SO distribution of traffic with respect to fuel consumption and to devise congestion pricing and speed limits to encourage traffic flow to follow this SO distribution. Unfortunately, fuel consumption is not an increasing function of speed or arc flow – an assumption required to ensure that TA models converge to UE or SO. We analyse the fuel-consumption TA model from a theoretical point of view and are able to show that, under appropriate assumptions, SO traffic patterns can be derived and enforced by congestion pricing and appropriate speed limits. We present results for standard TA instances.

Keywords: traffic assignment; fuel consumption; congestion pricing.

MD Dairy Modelling

WA220

POWDER: A strategic planning tool for New Zealand dairy farmers.

Oscar Dowson, *Andy Philpott*, *Andrew Mason* & *Anthony Downward*. University of Auckland, New Zealand.

Dairy farmers in New Zealand are required to make a vast number of sequential decisions in the face of long term environmental uncertainty and price volatility. This is largely due to the heavy reliance on outdoor pastoral farming and the exposure to international volatility in the price of milk. These factors run in contrast to the subsidised, indoor, controlled systems of Europe and the United States of America.

Moreover, many New Zealand farmers are deeply indebted, and therefore highly risk averse. Despite this, little research has been done in New Zealand to examine different management strategies, and how these strategies differ with risk, from a stochastic optimisation perspective. In this paper, we present POWDER: the Milk Production Optimiser incorporating Weather Dynamics and Economic Risk. POWDER is a multi-stage stochastic programme that breaks the dairy farming season into weeks and links these weeks by a system of linear dynamics. We use POWDER to show how management policies respond to differing levels of risk aversion, given different farming systems and geographic locations. POWDER shows promise as a strategic level planning tool to help participants in the New Zealand dairy industry understand, and plan for, farming in a stochastic world.

Keywords: dairy; stochastic; multi-stage; risk.

What happens in the container, stays in the container

Lisa Hall. Fonterra, New Zealand.

This year Fonterra took a problem to the second Maths in Industry New Zealand (MINZ) study group in Wellington. Milk powder shelf life is reduced by exposure to heat and moisture. We already know that refrigeration extends shelf life, while sitting a container of milk powder on a dock in the tropics decreases shelf life, but what about the conditions in between? The first component of this project explored complaints data, linking shelf life to various factors including product age, destination country and factory of origin. The second component involved designing a statistically robust experiment to test shelf life across five powder types and over fifty factories using temperature and humidity controlled containers. We want to ascertain what happens in the container, so that our customers receive the highest quality milk powder outside the container.

Keywords: shelf life; experimental design; Fonterra.

MONDAY 16:05 – 17:05

Data Journalism Panel

WG403

ME

Data Journalism Panel

Harkanwal Singh (NZ Herald), *Keith Ng* (On Point), *Lillian Grace* (Figure.NZ) & *Allan Lee* (School of Communications, AUT)

The panelists will discuss the role and presentation of data in story-telling, and how this is influenced by modern, interactive visualisation techniques. The impact of the statistical literacy (or lack thereof) of the readership will also be discussed. The current and future training of data journalists will be debated. The panel will then be open to engage with conference delegates to address any questions and comments raised by the audience.

The panel discussion will be moderated by Professor Thomas Lumley.

TUESDAY 09:00 – 10:00

TA Plenary 2

WG403

Analytics for understanding and improving health delivery services

Michael O'Sullivan. Engineering Science, University of Auckland, New Zealand.

I have been utilising Analytics to understand, evaluate, and improved health delivery services since 2008. My colleagues and I have employed a number of techniques / tools including Statistical Modelling, Data Visualisation, Simulation and Optimisation. In this presentation I will describe our experience working with New Zealand health providers and government to provide better outcomes delivered as efficiently as possible. I will present a number of relevant case studies, explain how we utilised Analytics approaches / tools in each case, and summarise the results of the research, particularly how it affected the health delivery service(s) involved.

Keywords: health, analytics, simulation, optimisation.

TUESDAY 10:20 – 12:05

TB Health Analytics 1

WG403

Can a brief telephone intervention for problem gambling impact on co-existing depression?

Nick Garrett, Maria Bellringer, Jussi Ranta & Max Abbott. Auckland University of Technology, New Zealand.

This study investigated the outcomes of a brief telephone intervention for problem gambling. A total of 150 callers were recruited and followed for 36 months. After giving consent, participants received a baseline assessment followed by a manualised version of the helpline's standard care. Eight-six percent of participants were re-assessed at three months, 79% at six months, 66% at 12 months, and 40% at 36 months. There was no evidence of major differential attrition consequently no adjustments were made to the data.

Depression is found to often be associated with problem gambling behaviour, and analysis was undertaken to examine the impact of a brief telephone intervention for problem gambling on rates of depression using logistic regression. At baseline depression was found to be associated with gender, problem gambling risk (PGSI), and deprivation (NZiDep). A multiple variable model found that PGSI and mental health medication best explained depression at baseline. A repeated measures logistic regression utilising all 36 months of data found that PGSI, NZiDep, and mental health medication were the best variables to explain the change over time.

The conclusion was that the intervention's impact on problem gambling behaviour also changed depression rates, however deprivation and mental health medication also contributed.

Keywords: problem gambling; depression; logistic regression.

Risk adjusted statistical process control methods and their applications in simulations and a real-life case study of ED intubations

Christin Coomarasamy[†], Irene Zeng[‡], Andrew Brainard[#] & Vanessa Thornton[‡]. [†]Auckland University of Technology, New Zealand; [‡]Counties Manukau Health, New Zealand; [#]Emergency department, University of Auckland; [‡]Emergency Medicine, Counties Manukau Health, New Zealand; [‡]Emergency Medicine, Counties Manukau Health, New Zealand.

Monitoring performances, adverse effects and mortality in hospital or community care are considered critical processes for quality assurance. The monitor includes the process of gathering cumulative information and statistical process control analyses. Results of analyses are compared to the expected values on a regular basis such that any abnormality will be detected and alarms can be given to the relevant stakeholders. The conventional approaches used in statistical process control are cumulative sum chart (CUSUM), exponential weighted moving averages chart (EWMA) and sequential probability ratio test. These approaches have been advanced in combination with a risk-adjusted method to account for the variability of patient baseline risk, but few evaluations have been conducted on these methods for different types of outcomes.

Our study aims to evaluate these methods in various monitoring settings. Simulations were used to create two types of scenarios: hospital indicator monitoring and community care surveillance. Outcomes with different statistical properties (continuous, binary or multiple categories) and baseline risk measurements of patients were simulated. Two different risk-adjusted methods: risk adjusted CUSUM and EWMA were applied to simulations and an ED intubation case study. Sensitivity, false alarm rates and performance in rare events are used as the criteria for the evaluations.

Keywords: CUSUM; EWMA.

ART: Adaptive Research Technologies

Ella Hodgson, Tasleem Hussein & Cliff Ashford. Information Communication Technologies Services, Auckland University of Technology, New Zealand.

In response to feedback from the AUT research community ICT embarked on a programme of work to deliver a suite of tools, services and resources under the collective name of Adaptive Research Technologies, or ART.

ART was designed and built through working directly with various research teams to provide a solution that would make things easier without restrict-

ing, or slowing them down. The solution provides data capture facilities, randomisation functionality, automated workflows, a Research Data Repository with features for cleaning, monitoring, and enhancing the data, a full range of analytical tools and reporting facilities, project management and document storage, high-performance computing access, data storage from the trivial to the titanic, online collaboration tools, as well as a bespoke development facility to integrate with external systems.

ART is now being actively used by seventeen research projects primarily in the area of clinical science, with another dozen across a wide variety of disciplines being initiated. The service is being made available to all post-grad students so they can secure their document and data while getting access to a broad range of tools.

Each feature was developed in direct partnership with research teams. The close working relationship allowed for fine-grained prioritisation and creative solution development based on real-world challenges. The ART service looks to provide a mobile-first form solution which can be used in difficult environments such as hospitals or out in the field.

Continuing professional development – what can NZSA do?

Ian Westbrooke[†], *Nokuthaba Sibanda*[‡] & *Richard Penny*[‡]. [†]NZ Department of Conservation, New Zealand; [‡]Victoria University of Wellington, New Zealand; [‡]Statistics New Zealand, New Zealand.

We presented a plan for continuing professional development (CPD) to the NZ Statistical Association executive in August (http://www.stats.org.nz/site/uploads/Continuing-Professional-Development_Oct-2016.pdf).

Our vision is that NZSA is recognised as the key facilitator / provider outside the formal education sector of professional development for statisticians and data analysts. This would be complementary to the formal education sector, activities of the NZSA Education committee and any NZSA mentoring programme.

The plan is to:

- ▷ Seek nominations and volunteers to provide workshops, courses or facilitate interactive sessions endorsed or organised by NZSA.
- ▷ Identify suitable presenters who may be visiting or prepared to visit NZ.
- ▷ Consult on how to assess which activities and presenters are appropriate for NZSA to organise or endorse.
- ▷ Explore options for providing CPD in association with NZSA or SSAI conferences.
- ▷ Explore possibilities of sponsoring or promoting activities that encourage good statistical practice through forums such as software user groups, NZ Analytics forums, or organisations that wish to increase skills of staff in data analysis and statistics.

While we see CPD activities as self-financing, we ask NZSA to agree in principle to underwriting some activities, potentially with funding from the Campbell fund. The audience will have time to respond to our proposals.

Keywords: continuing professional development.

Genetics

WG126 TB

Density approximation in population genetics

Wei Zhang[†], *Jing Liu*[‡], *Jesse Goodman*[†] & *Rachel Fewster*[†]. [†]Department of Statistics, University of Auckland, New Zealand; [‡]Shanghai Jiao Tong University, China.

The linkage disequilibrium coefficient r^2 is a measure of statistical dependence of the alleles possessed by an individual at two genetic loci. It has been widely used in association studies to search for the locations of disease-causing genes on chromosomes. However, so far most emphasis has been put on seeking the quantitative properties of r^2 or other measures of linkage disequilibrium. In this work, we propose a method to approximate the population-wide probability distribution of r^2 at stationarity by the maximum entropy principle, given a finite sequence of the distribution's moments. For the model considered in this work, we generalize an elegant analytic method proposed by Song and Song (2007) to compute the stationary moments of r^2 . Our method is based on the diffusion approximation, under which we demonstrate that for large numbers of models in population genetics, moments of the stationary distribution can be obtained without knowing the probability distribution itself. We illustrate our approach by recovering the stationary distribution of r^2 for a two-locus model incorporating mutation and recombination. Further, to show the robustness of the maximum entropy principle, we apply the approximate density on simulated data to estimate the mutation and recombination rates of a population.

Keywords: linkage disequilibrium; maximum entropy principle; stationary distribution.

Disease module detection (via graph clustering)

Matthieu Vignes. Massey University, New Zealand.

This work is an excuse to discuss the topic of graph clustering (aka graph partitioning, module detection or community detection). The data itself is a graph, i.e. a list of edges which represent some association to be considered between nodes. Nodes stand for statistical entities to be grouped. Most of the approaches we shall consider are based on a measures of 'compactness' or 'modularity' of the identified subgroups as compared to the rest of the graph. We will also present the motivating example of the detection of gene modules associated to a disease and our attempt at the subsequent DREAM11 challenge.

Keywords: module detection; graph clustering; modularity.

Bayes meets baby poo: Bayesian estimation of infant gut microbial diversity using metagenomic data

Roy Costilla, Ivy Liu & Richard Arnold. Victoria University of Wellington, New Zealand.

Metagenomics is the study of microbial communities using samples found in the physical environment. Through the sequencing of whole genomes, it aims to identify microbial species and their functions within their natural environment.

In this talk we identify variant strains, ie dominant genotypes, of *Bacteroides Faecis* (*B. faecis*) using shotgun metagenomic data from baby stool samples. Overall, the data comprises an ordinal response (equal to reference site, segregating site, and variant site) for 1992 sites measured over 25 occasions for one infant. We use a finite mixture model that includes random effects by strain and occasion and estimate the model parameters using a Bayesian approach. Model comparison is carried out using the Widely Applicable Information Criterion (WAIC).

The results show that there are at least three strains of *B. faecis* due to different levels of variation across occasions. Although using an ordinal response instead of counts, these results agree with a previous analysis of this dataset.

Keywords: ordinal data; finite mixtures; model-based clustering; Bayesian estimation; shotgun metagenomic data.

Estimation of quantitative genetics parameters in partially-genotyped populations

Alastair Lamont. University of Otago, New Zealand.

Quantitative genetics is the study of inheritance of quantitative traits. In livestock, two primary goals are the identification of both individuals with valuable DNA, and also variations in DNA that are valuable. For populations where the relationships among animals are known, there is an accepted modelling approach for this first goal, detailed in Henderson (1984).

Commercial genotyping now allows for direct measurement of the DNA of an animal at an efficient cost. This provides more information about the DNA that animals in a herd have in common. However, there is currently no consensus on how to estimate the parameters of interest in a partially genotyped population. We compare the accuracy of several modern approaches against that of Henderson, as well as an approach presently in development.

Keywords: quantitative genetics; model selection.

Parentage identification in New Zealand farmed deer using genotyping-by-sequencing

Ken Dodds, Suzanne Rowe, John McEwan, Tracey van Stijn, Rudiger Brauning & Shannon Clarke. AgResearch, New Zealand.

New Zealand deer farmers have used genetic markers to help them assign parentage for over 20 years. Know-

ing parentage is essential for breeders wanting to make rapid genetic gain in their livestock, and to avoid mating between close relatives. The current test is based on parentage likelihoods and marker exclusions using a set of microsatellite markers, but this technology is being phased out by the genotyping provider. An alternative technology is genotyping-by-sequencing (GBS). GBS, applied in a cost-competitive manner, does not provide definitive genotypes, but does provide information on several thousand-fold as many genomic positions. However, methods have been developed for relatedness estimation using GBS, and these can be used as the basis for a parentage test. Issues relating to this change in methodology will be discussed.

Keywords: parentage; genotyping; genetics; relatedness.

Demography

WA224

Genetics

TB

Dynamics of global remittances: a graph-based analysis

Felipe Lillo, Leidy García & Valentín Santander. UCM, Chile.

Human migration is an increasing worldwide phenomenon by which people leave their homeland to find better living conditions. This phenomenon entails several financial issues for a migrant. One of these issues are the remittances. The study of remittance dynamics at a global scale is still an emerging research topic.

This work contributes to the study of remittances by presenting a graph modelling approach that describes the dynamics of worldwide remittances. Some theoretical findings involving a probabilistic analysis of the model are also developed. Remittance data obtained from the World Bank is the main input of the model for which degree distribution, vertex degrees and two vertex cycles are analyzed for four different years. Results indicate a stable behavior throughout the years of both a preferential attachment process and the composition of the largest trading remittance communities.

Keywords: remittances; graph model; preferential attachment.

Redeveloping the confidentiality method for Statistics New Zealand business demography data

Alistair Ramsden & Caleb Moses. Statistics New Zealand, New Zealand.

Statistics NZ produces official business demography statistics, including counts of businesses and employees, disaggregated by categorical variables such as geography, industry, and employee count size group.

Statistics NZ has recently redeveloped how it confidentialises business demography statistics, to meet identified customer priorities, and also statutory responsibilities. This aligns with Statistics NZ's generic statistical business process model, other Statistics NZ strategic priorities, the NZ open government

Genetics

Demography

information and data programme, and the NZ data and information management principles.

We are analysing prospective confidentiality methods, determining which method is best, and implementing it. The proposed method uses a series of retained random numbers aligned with the series of raw interior cell data, modulo calculations on these retained random numbers to generate consistent exterior cell random numbers, random rounding, and bounded random noise, to automate producing confidentialised counts and count magnitudes. This method arguably produces limited data risk and data utility loss.

The proposed method relies on a corpus of previous methodological work. The proposed method is a test case for other work in progress developing a general Automated Confidentiality System (ACS), which seeks to automate key aspects of confidentiality within Integrated Data Infrastructure (IDI) and NZ Census.

Keywords: New Zealand; business; demography; confidentiality; method; automated confidentiality system; counts; count magnitudes.

Dynamic choropleths for poverty mapping

Geoff Jones. Massey University, New Zealand.

Spatial estimates of deprivation measures are often presented as a choropleth, ie a map of the country or region with sub-regions coloured differently to represent the estimated level of deprivation. Such a presentation tends to be taken as the exact truth by users, who may be basing important decisions such as aid allocation on the maps. However the estimates may have considerable uncertainty attached to them, so that ranking or pairwise comparisons of sub-regions are not as precise as the map may suggest. It is difficult to represent this uncertainty to the user. A separate choropleth showing the standard errors does not adequately convey the message. One possibility, if the map is to be viewed on-screen, is to make the choropleth dynamic, with the colours changing according to the uncertainty of the estimates.

An example is presented showing small-area estimates of sub-divisional poverty rates in Bangladesh. Some of the difficulties and issues in producing a suitable dynamic map are discussed.

Keywords: data visualization; Uncertainty; small area estimation.

Tree-based methods for poverty mapping

Penelope Bilton[†], Geoff Jones[‡], Ganesh Siva[§] & Steve Haslett[¶]. [†]Institute of Fundamental Sciences, New Zealand; [‡]Massey University, New Zealand; [§]AgResearch, New Zealand; [¶]Massey University / ANU, New Zealand.

Poverty mapping facilitates efficient allocation of \$Billions of aid resources, with the objective of eliminating poverty, the first of the United Nations Sustainable Development Goals.

Small area estimation techniques are used to generate

estimates of deprivation, such as poverty and malnutrition, at low geographical level, which are then displayed on a poverty map. The current technique used by organisations such as the World Food Programme consists of a unit-level regression model. An alternate methodology using tree-based methods is presented, which offers some practical advantages. Unit-level predictions from tree models are amalgamated to small area level, the trees adapted for complex survey design and a resampling scheme to generate standard errors of prediction is devised to overcome the inherent instability in the tree algorithm.

Real and simulated data based on a study of poverty mapping in Nepal are used to evaluate the methodology. Simulations indicate that approximately unbiased estimates result by amalgamating probability-based tree predictions, and standard errors obtained by employing a cluster bootstrap with cluster effects incorporated into tree predictions. The tree based methodology is used to generate small area estimates of poverty for a region in Nepal, which are comparable to published results from the standard method.

Keywords: small area estimation; complex survey design; sustainable development goals.

Spatially balanced sampling for socio-economic surveys

Naeimeh Abi. University of Canterbury, New Zealand.

In social science studies often socio-economic characteristics are influenced by the location of where people live. Therefore considering the geographical properties of sampling units in the sample selection process should improve survey design.

Spatially balanced sampling methods were designed initially for studies of natural resources and environmental processes. Their application for sampling human populations is relatively new. One of the reasons for the lag in uptake of these methods is that in socio-economic studies, target populations typically consist of separate discrete units such as mesh blocks or households. In these situations, simply imposing a geo-referenced grid over the study area, as is done in spatially balanced sampling methods, may lead to some areas being selected which have no sampling units.

In this presentation, we will extend the Balanced Acceptance Sampling (BAS) which is a new spatially balanced sampling design (Robertson et al., 2013), so it can be used for selecting discrete units. Our motivation is socio-economic studies. In this new version of the BAS method a spatial frame of discrete population units is created, and then BAS is implemented to select the sampling units. The results of implementing this new method for selecting mesh blocks are compared with other methods.

Keywords: socio-economic studies; spatially balanced sampling; balanced acceptance sampling; BAS-Frame.

TB Combinatorial Optimisation WA220

Finding sources of flexibility within tightly constrained operating room resources to reduce the impact of acute surgeries on elective surgery schedules

Daniel Cheng. University of Auckland, New Zealand.

When urgent acute surgeries arrive at a hospital, an elective surgery may have to be cancelled to create a vacant operating room. This cancelled elective surgery will need to be rebooked within the very near future. Since hospitals fully book their elective schedules to maximize utilization, vacant slots are rare. Instead, another elective surgery will be cancelled to create a timely vacancy for the previously cancelled elective surgery. This process of cancelling and rescheduling creates a domino effect that cascades throughout the booking list. The surgery schedule's inability to cope with disruptive events indicates rigidity, fragility and a lack of robustness. We explored two sources of hidden flexibility from within the system to aid in the schedule robustness. First, we incorporated the intangible costs of schedule disruptions as a component of our mixed-integer program's objective function. This allowed us to create an updated schedule that was as similar as possible to the original schedule. Secondly, we explored the use of Generalist surgeons who assume a flexible role. These surgeons are less averse to schedule disruptions and are able provide an absorbing buffer capacity, which will dampen the domino effect and allow the other surgeons' schedules to remain largely unaffected.

Keywords: robust scheduling; flexibility; operating room scheduling; operating room; healthcare; domino effect; linear programming; hospital scheduling.

Dynamic ambulance redeployment

Samuel Ridler, Andrew Mason & Andrea Raith. Engineering Science, University of Auckland, New Zealand.

Emergency Medical Services are commonly required to maintain a minimum level of service that specifies the speed with which different emergency medical calls should be responded to. Dynamic ambulance redeployment can help in ensuring that these service levels are met. Dynamic redeployment involves moving available ambulances to different ambulance stations or stand-by locations in real-time, in order to cover for ambulances that are currently busy responding to emergencies. Creating strategies that generate effective redeployments can be difficult to do well due to the stochastic nature of the system, e.g., call arrivals are random. An ambulance simulation model has been developed in the programming language Julia, in order to test the performance of different dynamic redeployment strategies, and for improving these strategies through simulation-optimisation methods. Details of the simulation will be presented, along with initial results of various redeployment strategies.

Keywords: ambulance; simulation; optimisation; redeployment.

Binary optimisation models for computing the frustration index in signed graphs

*Samin Aref[†], Andrew Mason[‡] & Mark C. Wilson[‡].
[†]University of Auckland, New Zealand; [‡]Engineering Science, University of Auckland, New Zealand; [‡]Computer Science, University of Auckland, New Zealand.*

Computing the frustration index of a signed graph is a key to solving problems in different fields of research including social networks, physics, material science, and biology. In social networks the frustration index, also referred to as correlation clustering, determines network distance from a state of structural balance. Although the definition of frustration index goes back to 1960, an exact algorithmic computation method has not yet been proposed. The main reason seems to be the complexity of computing the frustration index which is closely related to well-known NP-hard problems such as MAXCUT.

Quadratic and linear binary programming models are developed to compute the frustration index exactly. Gurobi Python interface is used to compute frustration index on real-world and synthetic datasets. The synthetic data involves Erdős-Rényi network, Barabási-Albert network, and specially structured random graphs. We also use well-known datasets from the sociology literature, such as Read's New Guinean tribes, and signed networks inferred from students' choice and rejection.

The results show that exact values of frustration index can be efficiently computed using our suggested optimisation models. We find that real-world networks exhibit a relatively low level of frustration which indicates that they are fairly close to balanced.

Keywords: binary programming; frustration index; correlation clustering; signed graphs; balance theory.

A liability allocation game

Robin Hankin. Auckland University of Technology, New Zealand.

Consider the following simple two-person zero-sum game. Two players are each given a fixed number of counters and have to allocate them among boxes labelled $1, 2, 3, \dots, k$. At times $t = 1, 2, 3, \dots$ a random box is chosen; the probability of choosing each box is known. If box i is chosen, and the player has at least one counter on box i , she removes a counter from that box; if there is no counter, she takes no action. The object of the game is to remove all the counters as quickly as possible. What allocation is optimal? This simple game is surprisingly difficult to analyze. In this short talk, I give some analytic and numerical results and suggest directions for further work.

Keywords: game theory; two-person zero-sum game; resource allocation; liability allocation; numerical method; optimisation.

Effectively solving 2D packing problems

Kumaran Ramakrishnan. Pacific Steel NZ.

This paper discusses the connections among published research work in two-dimensional packing problem (2DPP). Although there are many different variants in the 2DPP, the underlying principles are identical.

Based on the literature survey, similar approaches are applied to effectively solve any 2DPP. These are variations of sorting methods, multitude of deterministic heuristic placement routines, arrays of geometric techniques for overlap detection and the countless search methods (e.g. GA, SA, TS, EPSO) that are utilised to explore the solution space. On the other hand the solution space can either be the choice of next item chosen to be packed, the order of items in a packed layout, the target bin, etc. However any algorithm developed for one of the variants may provide a feasible solution to any other variant.

In addition, from the huge collections of research publications it can be gathered that researchers mainly focus on four groups of categories:

1. geometric techniques,
2. placement heuristics
3. meta-heuristics, and
4. lower bounds.

In this paper the author will examine the relationship between the 4 groups. It is hoped that the information provided will be a useful toolkit and enable the OR practitioners to effectively solve any 2DPP.

Keywords: cutting and packing; meta-heuristic.

TUESDAY 12:55 – 13:55

TC Health Analytics 2 WG403

An application of mediation analysis to randomised control trial data

Janet Pearson[†], Alain Vandal[‡] & Ema Tokolahi[†].

[†]Auckland University of Technology, New Zealand;

[‡]Department of Biostatistics and Epidemiology, Auckland University of Technology, Auckland and Health Intelligence and Informatics, Ko Awatea, Counties Manukau Health, New Zealand.

Pre-defined, and theoretically grounded mediation analysis of randomised control trial data can uncover interesting aspects of the workings of an intervention that conventional analysis usually cannot. For example, focus can turn to analysis of the mediating effect of adherence to the intervention (or knowledge gained), on the relationship between being in the intervention group and the outcome of interest. A brief introduction to mediation analysis and the paths involved will be given, along with an example of its use in random effects modelling as applied to data from a Cluster RCT with cross-over. Multivariate modelling, guided by Bauer & Preacher et al (2006); will be used to extract the covariance between coefficients for the different paths, to enable calculation of the standard er-

rors for the indirect and total effects. Complexities of the process in my example, looking at a mediator that varies only by cluster, and not by individuals within cluster, will be pointed out. Finally, a reflection on the utility of the mediation analysis in this setting will be given. Bauer D, Preacher J, Gil K. Conceptualizing and Testing Random Indirect Effects and Moderated Mediation in Multilevel Models: New Procedures and Recommendations. *Psychological Methods*, 2006, 11(2), 142-163

Keywords: mediation analysis; cluster randomised control trial; random effects; multivariate modelling; categorical IV in mediation analysis.

Feasibility study: Insulin therapy for diabetic foot ulcers

Diane Park. Auckland University of Technology, New Zealand.

Diabetic foot ulcers can have serious consequences for patients, and treatment incurs substantial costs to the health system. There is mounting observational evidence that improved glycaemic control will accelerate the healing of foot ulcers and reduce the incidence of extremity amputation, but no evidence exists from clinical trial. Feasibility of conducting a comprehensive clinical trial has not been undertaken successfully so far and is an important preliminary requirement. The primary objectives of this study are to investigate; the feasibility of recruitment of patients with diabetic foot ulcers for a future main study to address the effects of intensive insulin therapy; and the feasibility of intensive insulin therapy in a subset of those patients. The first five weeks of the overall 35-week accrual period was most active in terms of recruitment; the ratio of the recruitment rates in the first period compared to later period is 7.0 based on a Poisson exact test. Only 24 people among 78 satisfy all ten criteria (31% satisfaction), and the proportions of satisfied criteria differ significantly between periods. Kaplan-Meier analysis shows that approximately 17% of participants remained with a non-healed condition at Week 21. Hemoglobin-A1c (HbA1c) is shown as a sensitive covariate to model Ulcer-area-change.

Keywords: diabetic foot ulcer; insulin therapy; HbA1c.

Clinical decision rule development: a case study

Steve White, Alain Vandal, Peter McNair & Mark Laslett. Auckland University of Technology, New Zealand.

Clinical decision rules lead to further diagnostic or therapeutic action on the basis of results from a small battery of tests easily carried out in a clinical setting. Clinical decision rules must be stringently validated, but the initial development phase, while largely exploratory, can be informed by sound statistical practice. Typically, clinicians will find a logistic regression model for the gold standard decision based on dichotomous tests, using observed significance levels of the individual tests as criteria for model selection, then discard the fitted model and elaborate a rule based on

the number of successful tests. We bring in a few statistically based recommendations, nevertheless imbued with common sense, in regard to choosing an optimal number of tests and avoiding overfitting, selecting a good set of tests, and formulating a useful decision rule. We base these recommendations on information criteria, odds ratio magnitude and final model parameter estimates. These simple concepts are illustrated on data obtained from patients suffering from hip pain who underwent a guided anaesthetic injection into the hip joint as the reference standard for identifying intra-articular pathology.

Keywords: logistic regression; screening; clinical decision rules; information criteria.

TC Generalised Linear Models WG126

Focused model averaging in generalized linear models

Chuen Yen Hong, David Fletcher & Matthew Parry. University of Otago, New Zealand.

Under the frequentist framework, a model-averaged point estimate is a weighted mean of the estimates from each of the candidate models. The weights are often based on an information criterion, such as AIC or BIC. Focused model averaging is a method which tailors the weights to a focus parameter of interest. For a given focus parameter, model weights can be chosen in order to minimize the mean-squared error of the model-averaged estimator. Considering generalized linear models, we propose a focused model averaging procedure which uses standard results for maximum likelihood estimation when the model is misspecified. We use simulation to compare the proposed method with existing methods and find that the proposed method generally outperforms existing focused model averaging procedures.

Keywords: model averaging; focused model averaging; generalised linear models; mean squared error.

A set of novel link functions for modeling the mean of heavy-tailed distributions under the VGLM/VGAM approach

Victor Miranda Soberanis & Thomas Yee. Department of Statistics, University of Auckland, New Zealand.

Increasingly, distinct phenomena on rare events, e.g. in Geology, have revealed an important impact of heavy-tailed data at the point that uncertainty of predictions may grow sharply. Indeed, efforts to develop methodology and software to adequately address heavy-tailed data analyses have also grown correspondingly.

In particular, for discrete data, Vector Generalized Linear and Additive Models (Yee, 2015) (VGLM/VGAM) and the VGAM package for R allow to fit the Logarithmic (log-series), or the Zeta (Zifp) distributions; two prominent examples of statistical models to analyze rare events.

Ideally, in this fitting process each VGLM would correspondingly conform with an intrinsic link function. However, this is not generally true so that links from GLMs, as logit, are usually considered.

Developed under a common approach within the VGLM/VGAM framework, we introduce three novel link functions for the mean of some specific heavy-tailed distributions. Apart from naturally arising from the target distribution, these links are useful for event-rate data. That is, when the mean is modeled in terms of some rate of events and the time-frame observed. We will also show how the time and covariates, if any, can be incorporated in the analysis. Time permitting, some examples using the implemented software will be shown.

Keywords: VGLM/VGAM; link function; heavy-tailed distributions; event-rate data.

Parametric bootstrap and power analysis for GLMMs

Peter Green. The Agribusiness Group / Landcare Research, New Zealand.

Reliable inference from generalised linear mixed models often requires parametric bootstrapping. This means the simulation of a large number of replicates, which can be computationally expensive.

Power analysis for GLMMs is often done via simulation studies. This also requires a large number of replicates.

Combining the two can easily become computationally prohibitive, as nesting the bootstrap calculations within the simulation replications effectively squares the amount of computation needed.

But both procedures share some key features: they both involve simulating from a parametric model and refitting the model to the simulated data.

We attempt to solve the computational problem by using a single set of replicates to both estimate sampling distribution and calculate power.

Keywords: power; GLMM; simr.

Multicriteria Optimisation WA224 TC

Parallel implementation of a multi-objective linear programming solver

Mihai Bostan & Andrea Raith. Engineering Science, University of Auckland, New Zealand.

This project is focused on the design of a bi-objective optimisation solver for linear programs and its implementation in C++ using parallel programming techniques to take advantage of multi-core CPUs. The parallel solver can tackle a wide range of complex bi-objective problem instances, such as those arising in the fields of radiotherapy for cancer treatment and in network flow optimisation.

The parallel solver is based on the weighted sum scalarisation and ϵ -constraint method in multi-objective op-

timisation. Weighted sum scalarisation finds efficient solutions to linear biobjective problems, while the ϵ -constraint method finds efficient solutions to both linear and integer bi-objective problems.

Parallelisation of the solver code is achieved through use of the threading features present in the C++11 standard. The standard allows the creation of threads which are independent of one another and complete the most computationally expensive tasks in parallel. Both the weighted sum scalarisation approach and the ϵ -constraint method were parallelised.

Testing the weighted sum scalarisation approach against the ϵ -constraint method has produced significantly different results. Both approaches have seen considerable improvement in run time with parallelisation enabled, however the speedup is not proportional to the increase in the number of threads.

Keywords: multi-criteria; optimisation; parallel processing.

Knowledge-guided heuristics for radiotherapy treatment plan optimisation

Lisa Hillas & Andrea Raith. Engineering Science, University of Auckland, New Zealand.

Improvements to radiotherapy treatment techniques over the past two decades have resulted in complicated dose distributions being able to be planned and delivered accurately. However, the planning process itself remains to be a challenging and time consuming task, requiring constant input from a treatment planner. This research has investigated the possibility of automating the treatment planning process, and whether or not knowledge-based heuristics based on Data Envelopment Analysis can be used to identify high quality plans more efficiently. Multiobjective optimisation heuristics have been implemented and tested using the commercial radiotherapy treatment planning software RayStation. These heuristics aim to identify plans that are Pareto efficient in terms of particular plan quality metrics, such as percentage of tumour volumes receiving the prescribed dose, and the average dose received by healthy organs. The Pareto Set Pursuing algorithm (Shan and Wang, 2005) has been implemented. When tested on historical patient data, it has been found that the algorithm is able to identify plans of a similar quality to those actually used. It has been also been found that using Data Envelopment Analysis to guide the heuristic increases the efficiency of the heuristic.

Keywords: radiotherapy treatment planning; heuristic; multi-objective; optimisation; data envelopment analysis; knowledge-based optimisation.

Multi-objective navigation of external radiotherapy treatment plans using clinical criteria

Kuan-Min Lin[†] & Matthias Ehrgott[‡]. [†]University of Auckland, New Zealand; [‡]Lancaster University, United Kingdom.

This study considers a navigation method for finding the most preferable radiotherapy plans from a discrete

set using planner-defined clinical criteria. The method is based on interactively solving an optimisation model to identify a plan that best satisfies the aspiration values set by the planner. During navigation, the planner iteratively adjusts the aspiration values to match the preference learned from previous plans until the most preferable plan is identified. The aspiration values are treated as soft constraints in the optimisation model so that the model is always feasible. The use of soft constraints enables navigation among a discrete set and allows the planner to freely specify the aspiration values without producing an infeasible model. Application of the method on a prostate case shows that optimisation criteria, which are used in other navigation methods, may not correctly reflect plan quality and may mislead a planner to select a “sub-optimal” plan. Instead, the use of clinical criteria provides a more relevant measure of plan quality and aids the planner to quickly identify the most preferable plan.

Keywords: multi-criteria decision making; multi-objective optimization; navigation; radiotherapy treatment planning.

Census Collection

WA220

TC

Ex ante measures of census key performance indicators

John Crequer. Statistics New Zealand, New Zealand.

The traditional approach to measuring the quality of a census is to do a post-enumeration survey. Using a sample of the population we attempt to measure such things as coverage achieved in the census after the fact. However, we are now taking a new approach to running a census. Instead of delivery and collection by field staff, with resource allocation largely occurring before census day, the emphasis is on self-response and targeting collection where needed. This targeting requires some way of allocating resources to maximise the quality benefit of actions to achieve the objectives of the census during the census operation. This paper will discuss the development of measures to employ during census to provide some guidance on resource allocation.

Keywords: census; measurement; resource allocation.

Optimisation for the 2018 Census

Andrew Mason[†], John Crequer[‡], Geoff Leyland[#], Isaac Cleland[†] & Thomas Adams[†]. [†]University of Auckland, New Zealand; [‡]Statistics NZ, New Zealand; [#]Incremental, New Zealand.

For the 2018 census, Statistics NZ will encourage New Zealanders to submit census forms online. However, before the census, every street in the country has to be visited by a field officer to identify potential dwellings and to deliver paper forms to selected areas. After the census, non-responding households will need to be visited to encourage them to respond.

To plan the census, the country is divided into about 60,000 meshblocks, areas each containing very roughly

50 households. For each phase of the census, a field officer needs to be assigned to every meshblock of interest. Meeting the quality goals for the census within budget will require this fieldwork to be meticulously planned, so that the each field officer is assigned the best possible meshblocks.

We discuss how we are approaching the planning problems the census presents. We'll discuss how meshblocks interact with road networks, lakes and mountain ranges, and how we rapidly partition meshblocks into day-sized chunks of work for field officers, and how we make our solutions robust to the differences between estimated and actual workloads in a meshblock, the different work rates of different officers, and the comings and goings of a dynamic workforce.

Keywords: census; districting; optimisation.

Survivable electricity distribution network design using batteries

Ruth Dominguez Martin[†] & Andy Philpott[‡].
[†]University of Castilla-La Mancha, Spain; [‡]University of Auckland, New Zealand.

Improvements in the quality of batteries and a decline in their costs raise the possibility of installing batteries throughout distribution networks, to be used to supply power when lines fail. We propose a two-stage stochastic-programming problem to design networks that will survive line-failure events. This model determines optimal investments in new distribution lines and batteries. Particularly, we analyse the cost efficiency of building storage capacity under different cost scenarios, in comparison with the cost of building new distribution lines considering that the network is operated using a radial configuration. The proposed model is solved using the Julia programming language under JuMP. In order to solve large-scale problems, we apply a Dantzig-Wolfe decomposition technique using column generation. Numerical results in terms of investment costs and execution times are provided.

Keywords: column generation; distribution network; stochastic programming; storage capacity expansion.

TUESDAY 14:00 – 15:00

TD Plenary 3

WG403

Precision Health – New Zealand's golden analytics opportunity

Kevin Ross. Orion Health, New Zealand.

Health has lagged behind in allowing data to truly transform how we think and operate. We recognise more and more factors that influence our health, from family history and genomics to social and environmental factors, but our clinical experts have less and less time with a patient to diagnose and act. These factors can be tracked from medical records, devices and sensors, and the rapid increase in volume and complexity of this data presents a golden opportunity to trans-

form an industry that values evidence, consumes enormous resources, and touches every life. New Zealand is uniquely positioned to lead in transforming research into practice for data-driven health, due to the nature of our health system, population, and technology. I will introduce Precision Driven Health, a major public-private venture between industry, government and academia launched in 2016 to create the capability to optimize the health of each individual and their whanau by combining and learning from all available data. The research partnership addresses four key themes to enable this to become a reality: new data, predictive modelling, precise and timeline healthcare, and empowering patients. I will introduce the research partnership plans, and highlight this major opportunity for New Zealand.

Keywords: health; analytics.

TUESDAY 15:20 – 16:00

Health Economics

WG403

Plenary 3

TE

An econometric analysis of the patient-at-risk team and intensive care unit

Ali Haji Vahabzadeh & Valery Pavlov. University of Auckland Business School, New Zealand.

This paper explores the impacts of an early intervention of a patient-at-risk (PAR) team on an intensive care unit (ICU) patients. ICUs currently challenge with the rationing of nurse capacity, largely because they lack enough expertise to monitor the health conditions of life-threatening patients in the ICU. To reduce the workload of the ICU, hospitals construct the PAR team that is responsible for averting the preventable ICU admissions and detects the early signs of health deterioration among the critically-ill patients in the ward prior to admission to the ICU. Nevertheless, the efficacy of the PAR team on the ICU patient outcome such as mortality rate has not been well quantified. Using a patient-level dataset of a large hospital in New Zealand, our preliminary analyses show that the early intervention of PAR team on the ward may have no effect on the ICU mortality. We also discuss the impact of high occupancy level of the ICU nurses on the ICU LOS. Notably, we find that the ICU patients are likely to be bumped out of the ICU when the ICU nurses are highly occupied. However, the effect of PAR team on the prolonged ICU LOS may not be significant.

Keywords: healthcare delivery; capacity allocation; critical care outreach team; medical emergency team; patient flow.

Cost of traumatic brain injury in New Zealand: What can we learn from cost-of-illness studies?

Braden Te Ao. Auckland University of Technology, New Zealand.

We aimed to estimate from a societal perspective the 1-year direct and indirect costs of traumatic brain injury (TBI) for New Zealand (NZ) in 2010.

Health Economics

Census Collection

Plenary 3

An incidence-based cost of illness model was developed using data from the Brain Injury Outcomes New Zealand in the Community Study. Details of TBI-related resource use during the first 12 months after injury were obtained for 725 cases using resource utilization information from participant surveys and medical records. Total costs are presented in NZ dollars year 2010 value.

In 2010, 11,301 first-ever TBI cases were estimated to have occurred in NZ; total first year cost of all new TBI cases was estimated to be NZ \$71 million with total prevalence costs of NZ \$151 million. The average cost per new TBI case during the first 12 months and over a lifetime was NZ \$8,8724 (95% confidence interval [CI] \$7,117 – \$11,708), varying from NZ \$6,908 (95% CI \$5,596 – \$8,286) for mild cases to NZ \$54,606 (95% CI \$24,359 – \$97,372) for moderate/ severe cases.

Keywords: burden of disease; cost of illness; health services utilization; traumatic brain injury.

TE Poster Session

WG128

Spatially adaptive kernel density estimation: exploring global bandwidth selection techniques

Tilman Davies[†], *Martin Hazelton*[‡] & *Claire Flynn*[‡].
[†]University of Otago, New Zealand; [‡]Massey University, New Zealand; [‡]MetService NZ, New Zealand.

The highly heterogeneous nature of planar point patterns typically observed in many different fields of research renders non-parametric methods for estimation of the assumed bivariate density function particularly attractive. One often turns to kernel density estimation – whether it is simply for data exploration, or for the construction of key components of more complicated spatial models. Most often deployed with a constant amount of smoothing, recent advancements in KDE have highlighted clear advantages of allowing the amount of smoothing to vary spatially – where the adaptive bandwidths reduce to capture more detail in densely populated sub-regions, and increase in sparsely populated areas to ‘smooth over’ less important features. As with any smoothing problem, however, bandwidth selection is critical to the reliability of the results – and we cannot expect bandwidth selection methods developed for the fixed estimator to simply ‘work out of the box’ for choosing the requisite global smoothing parameter for the adaptive estimator. Sound data-driven solutions to this problem are nonexistent at present and, backed up by some novel theoretical and numerical explorations, we provide some initial recommendations for implementation of the flexible adaptive estimator.

Keywords: kernel smoothing; bandwidth selection; variable bandwidth; density estimation.

Shewhart, EWMA and CUSUM charts for monitoring fractional non-conformance in short-run production process

Xin Zhou, *Kondaswamy Govindaraju* & *Geoff Jones*.
 Massey University, New Zealand.

Considerable measurement errors are involved in the measurement of food quality characteristics such as fat percentage because routine testing in the shop-floor is based on high performance liquid chromatography (HPLC) technique. To adjust for the effect of measurement errors, Govindaraju and Jones (2015) proposed a fractional nonconformance probability measure for known measurement error distributions.

Exponentially weighted moving average (EWMA) and cumulative sum (CUSUM) chart accumulate information selectively from the past observations. Past research established that these charts are sensitive to detect of small shift levels (on the order of about 1.5 SD or less) compared to the ordinary Shewhart chart, but slower in detecting large shift in the process mean.

Fractional nonconformance was implemented on EWMA and CUSUM framework under various process models for short-run production monitoring. We found that EWMA and CUSUM chart are superior to Shewhart charting for detection of both small and big shifts (up to 3SD) only when the false alarm rate is small ($\alpha < 1\%$). When false alarm rate is of the order 5%, the performances of EWMA, CUSUM and Shewhart charts are indistinguishable. For short production runs, a very small false alarm rate is not feasible and hence the simpler Shewhart chart is sufficient.

Keywords: control chart; short run production; fractional nonconformance.

Bayesian non-parametric approaches to multivariate time series analysis

Alexander Meier[†], *Renate Meyer*[‡], *Claudia Kirch*[†] & *Matthew Edwards*[‡].
[†]OvGU Magdeburg, Germany; [‡]University of Auckland, New Zealand.

While there is an increasing amount of literature about Bayesian time series analysis, only few non-parametric approaches to multivariate time series exist. Most notably, Rosen and Stoffer (2007) and Shibin Zhang (2016) rely on Whittle’s likelihood, involving the second order structure of the time series by means of the spectral density matrix.

The latter is modeled with a smoothing splines prior for the components of the Cholesky decomposition. While these approaches are shown to perform well in many applications, their theoretical asymptotic behavior in terms of posterior consistency is not known. Consistency results are typically restricted to parameters outside a prior null set, which is unsatisfactory in infinite dimensions, or even fail entirely (Diaconis and Freedman, 1986).

Having a posterior consistency result in mind, we investigate multivariate extensions of the Bernstein-Dirichlet prior from Choudhuri et al. (2004), for which consistency under the Gaussianity assumption has been shown in the univariate case. We also consider a multivariate extension of the corrected parametric likelihood, a generalization of Whittle’s likelihood, which

has recently been developed by Kirch, Meyer et al. (2016).

Keywords: Bayesian non-parametrics; multivariate; time series analysis; spectral density; consistency.

Determining the probability distribution of earthquake forecasts

David Harte. GNS Science, New Zealand.

The Epidemic Type Aftershock Sequence (ETAS) model is probably the most successful model yet at forecasting earthquake aftershocks. It is a branching process using an analogy to the spread of an epidemic. All events in the history of the process increase the intensity at the current time at all points in space, the size of the effect being related to the space-time distance back to the historical event.

Earthquake forecasts based on a self-exciting model, like ETAS, are often produced by simulation. From these simulations, an empirical probability distribution can be derived for a forecast in a specified space-time-magnitude volume.

We will show that the forecast distribution can be characterised by probability generating functions. This shows how deeply complex the dependency structure is in such a model. We then consider whether the forecast distribution can be approximated, using less computation than that required for simulation, by a “standard” multi-parameter probability distribution. The multiple parameters gives us the ability to at least fit a distribution with comparable mean and variance to that of the forecast distribution. One of the main questions is how to determine the forecast mean and variance.

Keywords: ETAS model; stochastic process; branching process; earthquake forecasts.

Expectation-maximization of baby poo: Frequentist estimation of infant gut microbial diversity using metagenomic data

Roy Costilla, Ivy Liu & Richard Arnold. Victoria University of Wellington, New Zealand.

Metagenomics is the study of microbial communities using samples found in the physical environment. Through the sequencing of whole genomes, it aims to identify microbial species and their functions within their natural environment.

In this poster we identify variant strains, ie dominant genotypes, of *Bacteroides Faecis* (*B. faecis*) using shotgun metagenomic data from baby stool samples. Overall, the data comprises an ordinal response (“equal to reference site”, “segregating site”, and “variant site”) for 1992 sites measured over 25 occasions for one infant. We use a finite mixture model that includes random effects by strain and occasion and estimate the model parameters using the EM algorithm. Model comparison is carried out using several information criteria: the AIC, BIC and ICL. The results show that there are at least three strains of *B. faecis* due to different

levels of variation across occasions. Although using an ordinal response instead of counts, these results agree with a previous analysis of this dataset.

Keywords: ordinal data; finite mixtures; model-based clustering; EM algorithm; shotgun metagenomic data.

Modelling risk in multistage stochastic programming problems

Regan Baucke, Anthony Downward & Golbon Zakeri. Engineering Science, University of Auckland, New Zealand.

Objective function choice in multistage risk averse stochastic programming problems directly influences the policies generated from these optimisation models. We introduce and discuss necessary conditions for objective functions to be optimal when placed in a dynamic framework. We show how different models of risk are connected as members of a larger class of risk measures on these uncertainty spaces. We relate these problems to a general class of dynamic maximin programming problems. Finally, we introduce an algorithm for the efficient solution of general maximin stochastic dynamic programmes.

Keywords: multistage; risk; robust; dynamic; stochastic.

Electricity Markets 1

WA220

Poster Session

TE

Co-optimization of energy and reserve for a strategic major consumer in the New Zealand electricity market

Mahbubeh Habibian, Golbon Zakeri & Anthony Downward. University of Auckland, New Zealand.

We study demand-side participation in an electricity market for a major consumer. We construct models of this consumer’s optimal bidding strategies under two different assumptions. First, we treat all agents as price takers, leading to a competitive equilibrium in both energy and reserve. Second, the consumer is considered to be a price maker, anticipating how its bidding may affect the market prices. This model is a stochastic optimization problem for a strategic major consumer who has flexibility over its consumption and can offer reserve. Our model is a bi-level optimization model (reformulated as a mixed-integer program) that embeds the optimal power flow problem, in which electricity and reserve are co-optimized. We implement this model for a large consumer of electricity in the New Zealand Electricity Market. To reduce the solution time of the large mixed integer program, we explore the specific properties of the optimal power flow to reformulate the model. We show that by adding further constraints we manage to improve the performance of the model; however, a compromise between the number of scenarios and the accuracy of the solution computed is needed for practical purposes.

Keywords: stochastic optimization; electricity markets; demand response; integer programming.

Poster Session

Electricity Markets 1

Electricity offer strategy with constrained transmission

Keith Ruddell, Anthony Downward & Andy Philpott. Engineering Science, University of Auckland, New Zealand.

The NZ electricity market operates under locational marginal pricing for energy. In dispatching generation to meet load, a distinct price is assigned to each of over 200 nodes around the country. Price differences between nodes reflect loss of energy in transmission, and also constraints on flow of electricity to keep within the capacity of the transmission network. Price differences from normal losses are modest, however those resulting from constraints on the flow can be very large. Generators, in choosing their offer stacks, can cause these transmission constraints bind, thus reducing the amount of competition they face on the spot. As there are only a handful of large generators in the NZ market, there are often opportunities for this sort of market manipulation.

I will present a supply-function equilibrium model of an electricity spot market with a simple transmission network, in which generators cause lines to congest in order to increase their local market power. The model will also show how financial transmission rights and the location of retail commitments can attenuate the incentives for generators to cause congestion, thus restoring workable competition.

Keywords: supply-function equilibrium; electricity markets; industrial organization; optimal control.

TUESDAY 16:05 – 17:05

TF

Plenary 4

WG403

Translational research in the data sciences

Rhema Vaithianathan. Auckland University of Technology, New Zealand.

One problem with academic research can be that findings are placed on a shelf and gather dust once the journal article has been published. Translational research is gaining traction in healthcare and social sciences as a method of ensuring theoretical discoveries are converted into practical impact. Alongside traditional research methodologies, translational research combines input from a range of disciplines including ethicists, domain experts and data scientists to provide actionable recommendations for the application of findings. This has the result of (1) reducing the time lag between scientific discoveries and policy change, and (2) improving outcomes and well-being at a community level.

Data science and *Big Data* is one area ripe for translational research. Yet the area is challenging for a number of reasons – chiefly that the agents who can make the change do not want big data – they want big intelligence. Other challenges include the lack of social license for sharing data, and potential ethical traps in how data is actioned.

Drawing on recent examples, this talk will discuss opportunities and challenges in using predictive analytics and quasi-experimental methodologies for translational research. This will include a discussion on the development and implementation of the first Child Welfare predictive modelling tool, which is currently “live” in Allegheny County, US. A second predictive risk model will identify risk factors in California birth record data and discover communities where children have fared significantly better than expected. The talk will conclude with guidance on how to take data from the data lab and into actionable change.

Keywords: linked data; quasi-experimental research; translational research; predictive modelling.

WEDNESDAY 09:00 – 10:20

Environment / Biodiversity WG403 WA

Model selection and parameter inference in phylogenetics using Nested Sampling.

Patricio Maturana. University of Auckland, New Zealand.

Phylogenetics is the study of evolutionary relationships among organisms and as in any other field, statistics plays an important role. Under a Bayesian paradigm, we analyse molecular data to select among evolutionary models and to infer the tree which better fits the data. Bayes factors are a common method for statistical model selection. The computation of such factors is based on the marginal likelihood – an integral that can be hard to estimate depending on the model complexity. We carry out the analysis via Nested Sampling, a Bayesian algorithm that provides the means to estimate the marginal likelihood which also yields posterior tree samples in a single run. We present some examples in real data sets.

Keywords: model selection; nested sampling; phylogenetics.

Hierarchical modelling of occupancy and indices of abundance for New Zealand national biodiversity monitoring data

*Paul van Dam-Bates[†] & Andrew Gormley[‡].
[†]Department of Conservation, New Zealand; [‡]Landcare Research, New Zealand.*

The Department of Conservation measures the relative abundance of birds and mammals across New Zealand Public Conservation Land on an 8 km grid. Five repeated measurements are taken at each individual site to allow for estimation of occupancy. A state-space formulation of occupancy in a Bayesian framework is then used to get an estimate for each species across the PCL within a 5-year window. In the model, by conditioning on occupancy, abundance indices such as residual trap catch (RTC) and faecal pellet index (FPI) can be modeled as zero-inflated Negative Binomials. Birds

are modeled using a multi-species occupancy model and species richness is derived as the number of birds present at each site within an MCMC iteration. Single species occupancy models are run for possums, ungulates, hares, and rabbits. When the detection rate is high as it is for faecal pellet counts, the models provide very little new information over observed occupancy. In the next round of measurements site colonisation and extinction will be able to be incorporated to look at trend.

Keywords: occupancy; Bayesian; state-space; biodiversity.

Modeling environmental gradients in species composition using multivariate spatial factor analysis.

Russell Millar. University of Auckland, New Zealand.

An irregularly spaced AR1 process is formulated for a Gaussian Markov random field representation of the weights in a multivariate factor analysis. The AR1 process is computationally efficient thanks to specification via the sparse precision matrix. The model is applied to fish communities over a depth gradient and it is shown that environmental gradients and species turnover can be jointly modeled.

Keywords: environmental gradient; factor analysis; multivariate; spatial analysis; species composition.

Wasp population dynamics in their native and invaded ranges: plus ca change

John Haywood & Phil Lester. Victoria University of Wellington, New Zealand.

Introduced species often experience different population dynamics in their introduced and native ranges. We examined the long-term population dynamics of the invasive common wasp, *Vespula vulgaris*, in its native (English) range and its (New Zealand) invaded range. Wasp population time series were examined using partial rate correlation functions. Gompertz population models and multivariate autoregressive state-space (MARSS) models were fitted, both incorporating climatic variation.

Density dependence in wasp populations was similar in both countries, with previous-year wasp abundance the most important variable in predicting intrinsic rate of increase. No evidence of cyclic population dynamics was observed.

Both Gompertz and MARSS models highlighted the role of weather conditions in each country as significant predictors of annual wasp numbers. The temporal evolution of wasp populations at all sites was best modelled jointly using a single latent dynamic factor for local trends, with the inclusion of a latent spring weather covariate. That same parsimonious multivariate model structure was optimal in both the native and invaded ranges.

Spring weather in both countries has a major influence on wasp numbers, probably through impact on wasp colony initiation and early development. Invasive

species may not exhibit different population dynamics, despite considerable variation in abundance throughout their distribution.

Keywords: density dependence; invasive species; multivariate time series analysis; partial rate correlation functions; population dynamics.

Statistical Theory 2

WG126 WA

Optimal design measures via nonlinear optimization under asymmetric errors and second-order least squares estimation

Rahul Mukherjee[†] & S Huda[‡]. [†]Indian Institute of Management Calcutta, India; [‡]Kuwait University, Kuwait.

In recent years, it has been well recognized that the second-order least squares (SLS) estimator outperforms the ordinary least squares (OLS) estimator when the error distribution is asymmetric. The SLS estimator minimizes a quadratic form in the deviations of the observations as well as their squares from their respective expectations, the quadratic form being dictated by an appropriately chosen positive definite weighting matrix. It is applicable even when the error distribution is not exactly known. Because of these advantages of SLS estimation, the corresponding optimal design problem gains importance. From this perspective, we characterize optimal design measures via the use of directional derivatives, taking due cognizance of the nonlinearity of the underlying information matrix in the design measure. In contrast to OLS estimation, such nonlinearity poses a new challenge. Still, our findings have a broad spectrum of applicability allowing analytical as well as algorithmic results. This is illustrated with reference to binary design points which may also have to meet a string or consecutive ones property from practical considerations. The issue of robustness to the unknown skewness parameter of the error distribution is explored. Finally, we present procedures which entail N-run designs that are highly efficient, if not optimal.

Keywords: directional derivative; information matrix; robustness; second-order least squares; string property.

EM algorithm and stochastic control

Steven Kou[†], Xianhua Peng[‡] & Xingbo Xu[#]. [†]National University of Singapore, Singapore; [‡]Hong Kong University of Science and Technology, Hong Kong; [#]Goodman Sachs, USA.

We propose an algorithm called EM-Control (EM-C) algorithm to solve multi-period finite-time horizon stochastic control problems, where the optimal policy is not necessarily stationary. Generalizing the idea of the EM algorithm, the EM-C algorithm sequentially updates the control parameters in each time period in a time-backward manner. Similar to the EM algorithm, the EM-C algorithm has monotonicity of performance improvement in every iteration, and hence has good convergence properties. We apply the EM-C algorithm to solve stochastic control problems in real business cy-

cle and monopoly pricing of airline tickets, showing the effectiveness of the algorithm.

Keywords: EM Algorithm; stochastic control; economics; revenue management.

Computation of adaptive kernel intensity estimates of planar point process data

Tilman Davies[†] & *Adrian Baddeley*[‡]. [†]University of Otago, New Zealand; [‡]Curtin University, Australia.

The fixed-bandwidth kernel estimator of spatially continuous intensity functions on \mathbb{R}^2 can be formulated as a convolution of the kernel and discretised grid weights, allowing access to efficient fast-Fourier techniques to evaluate the offending expressions. However, employing constant smoothing can have adverse consequences when the observation locations are highly heterogeneous. An elegant solution is to allow the bandwidth attached to each observation to vary, leading to an adaptive estimator; certain forms of which have been shown to have important theoretical and practical advantages over fixed-bandwidth estimation. Its uptake in the applied literature remains limited because implementation is not straightforward – variable bandwidths prevent a directly analogous expression to the fixed-bandwidth convolution, meaning explicit evaluation on anything but relatively coarse grids with modest sample sizes becomes computationally prohibitive. We give an overview of more computationally feasible methods for evaluating the adaptive kernel intensity estimator, from simple bandwidth partitioning to designing trivariate convolutions that do possess the required periodic properties for subsequent 3D fast-Fourier options. Our proposed methods are substantially faster than direct implementation, include the non-trivial capability to perform adaptive edge-correction, and are generalisable to higher dimensions.

Keywords: kernel smoothing; variable bandwidth; fast evaluation; discrete Fourier transform; boundary bias.

Hybrid estimators for a general class of estimation problems

Mehdi Soleymani[†] & *Stephen Lee*[‡]. [†]University of Auckland, New Zealand; [‡]University of Hong Kong, Hong Kong.

We introduce a general framework for the combination of parametric and non-parametric estimators, which is termed a hybrid estimator. The hybrid estimator is a weighted average of the parametric and non-parametric estimates where the weights are computed automatically. When the convergence rates of parametric and non-parametric methods are different, we show that overall the new hybrid estimator has better performance compared to its parametric / non-parametric constituents. In the context of density estimation, we will illustrate a resampling method to compute the weights for weighted average of the parametric and non-parametric estimates. We show the asymptotic properties of this hybrid estimator in details. Later, we propose a new algorithm to calculate weights in a

simpler and more elegant way. We show that the new algorithm reduces the computational complexity of the hybrid estimator and brings more desirable properties to the hybrid estimator.

Keywords: bootstrap; hybrid estimator; kernel density estimator; convergence rate.

Electricity Markets 2

WA220

WA

Forward contracting in markets with private information

Andy Philpott[†] & *Eddie Anderson*[‡]. [†]University of Auckland, New Zealand; [‡]University of Sydney, Australia.

We consider the use of forward contracts to reduce risk for firms operating in a spot market. Firms have private information on the distribution of prices in the spot market. We discuss different ways in which firms may agree on a bilateral forward contract: either through direct negotiation or through a broker. We introduce a form of supply function equilibrium in which two firms each offer a supply function and the clearing price and quantity for the forward contracts are determined from the intersection. In this context a firm can use the offer of the other player to augment its own information about the future price.

Keywords: supply function equilibrium; electricity markets; contracts for differences.

Risk averse equilibrium in electricity market

Henri Gerard^{†‡}, *Andy Philpott*[‡] & *Vincent Leclère*[†]. [†]CERMICS, France; [‡]Electric Power Optimization Centre, University of Auckland, New Zealand.

New-Zealand electricity market is a complex system involving multiple agents either producers or consumers that have to deal with uncertainties. We choose to model uncertainty modeled by a scenario tree in a discrete time framework. We are looking for equilibria in multistage risk averse stochastic problems. This can be solved through extended mathematical programming formulation. However, the problem becomes intractable when the number of time steps increases and the number of scenario explodes.

One way to tackle such a difficulty is to decompose the original problem into subproblems which can be solved in parallel. In the risk neutral case with an additive criterion, it has been shown that a competitive equilibrium, even in incomplete market is equivalent to a decomposable optimization problem thanks to the linearity of the expectation. We investigate here the case of a risk averse social planner. We use price decomposition to handle the non linearity and decouple the problems. We are hence able to compute an approximation of the equilibria even for large scale problems using an Uzawa algorithm. After some recall of useful mathematical tools, we present the result of our numerical simulations and give some clues on how to select one equilibrium among several.

Keywords: equilibrium; risk; decomposition.

Identification of precursor patterns for prediction of frequency fluctuations in electrical power generation data

Md. Shahidul Islam, Russel Pears & Boris Bacic. School of Engineering, Computer and Mathematical Sciences, Auckland University of Technology, New Zealand.

Precursor pattern identification addresses the problem of detecting warning signals in data that herald an impending event of extraordinary interest. In the context of electrical power systems, identifying precursors to fluctuations in power generation in advance would enable engineers to put in place measures to mitigate against the effects of such fluctuations.

In this research we use the Haar wavelet to transform a time series defined on electrical power generation frequency which was sampled at intervals of 30 seconds, to identify potential precursor patterns. The power spectrum that results is then used to select high coefficient regions that capture a large fraction of the energy in the spectrum. We then subjected the high coefficient regions together with a contrasting low coefficient region to a non-parametric ANOVA test and our results indicate that one high coefficient region dominates by predicting an overwhelming percentage of the variation that occurs during the subsequent fluctuation event. These results suggest that the wavelet is an effective mechanism to identify precursor activity in electricity time series data.

Keywords: precursor pattern; frequency fluctuation; Haar wavelet transform; Kruskal Wallis test.

Furthermore, the model provides feasible consistent results using different solvers, while the same can not be said for the results of other approaches. The model was tested using a data set of 95 banks across 11 years and compared with the interval approaches.

Keywords: production economics; scale elasticity; returns to scale; most productive scale size; data envelopment analysis (DEA).

Advanced monitoring of project performance data

Amir Salehipour[†] & Leila Moslemi Naeni[‡]. [†] University of Newcastle, Australia; [‡] University of Technology Sydney, Australia.

The well-known earned value technique measures and evaluates project performance by uncovering schedule and cost deviations from the baseline plan. However, it is not established to determine acceptable levels of deviations from the baseline. We develop an advanced tool that incorporated several concepts of statistical modeling and data analysis to overcome this limitation. The developed tool is capable of monitoring trends in the project performance behavior, and allows them to be detected before the project deviates much from the baseline plan. The tool does not take any assumption regarding the data. Generally speaking, the developed tool has three steps:

1. it performs a pre-processing in order to prepare the data for the analysis,
2. supplies the processed data into the control charts, and
3. performs an advanced analysis by calculating the probability of failure.

We applied the tool to a construction project, and monitored four well-known schedule and cost earned value indexes. The investigated case showed that the new implemented tool enhanced capability of the earned value technique. We concluded that implementing the developed tool together with the traditional tools noticeably improves the project controlling scheme and provides more information on project progress.

Keywords: project performance monitoring; cost and schedule indexes; Shewhart individuals control charts; autocorrelated earned value.

Service level agreement between a supplier and multiple customers

Babak Abbasi[†], Zahra Hosseini[‡], Osama Alamri[†], Doug Thomas[#] & James Minas[‡]. [†] RMIT University, Australia; [‡] Deakin University, Australia; [#] Penn State University, USA; [‡] SUNY New Paltz University, USA.

In this study, we analyze item fill rate distributions for a supplier that exercises a static base stock policy and supply to multiple retailers. In particular, we analyze the impacts of the duration of performance review period and correlation between retailers' demands on the average fill rate and the probability of overreaching the target performance measure, when a supplier has mul-

WEDNESDAY 10:40 – 12:00

WB Performance Measures WG403

A new metric for scale elasticity in data envelopment analysis

Maryam Hasannasab, Israfil Roshdi, Dimitris Margaritis & Paul Rouse. Accounting and Finance, University of Auckland, New Zealand.

Scale elasticity (SE) is an important feature of technology and plays an important role in decision making, e.g. mergers and acquisitions, and in shaping competitive market structure, e.g. increasing returns to scale is an important source of market power. While there have been several attempts to develop useful measures of SE for data envelopment analysis (DEA) models, due to the linear piecewise nature of the DEA technology, these have resulted in interval measures with upper and lower bounds with ensuing issues of indeterminacy and computational limits. In this paper, we propose an LP model that provides a unique measure of SE based on a simple proposition of closeness to most productive scale size (MPSS). The model is non-oriented, and the result not only provides the SE measure but also allows returns to scale to be classified.

tiple customers and their demands are fulfilled from a pooled inventory. The results are provided based on random service policy that fulfill the demands of retailers based on first come first served policy.

Keywords: SLA; supply chain; multiple customers.

Supply chain bottlenecks and their applications on flexible planning

Samin Aref[†] & *Iman Kazemian*[‡]. [†]University of Auckland, New Zealand; [‡]University of Tehran, Iran.

We present a flexible supply chain design implementing an idea from production planning. The idea of capacity bottlenecks is used to improve flexibility in a multi-echelon multi-product supply chain. We optimize capacity allocations to bottleneck points in order to enhance overall flexibility. The proposed mixed-integer linear programming model minimizes the total cost of facility establishment as well as their utilization and transportation cost.

The performance of the model is evaluated by several test problems with uncertainty in demand, cost, capacity, and product specifications. The results indicate the superiority of the suggested model to the previous flexibility formulation method. According to the numerical results, the proposed model decreases the total cost of flexible supply chain by up to 16% on average. Another advantageous feature of the proposed model is its capability of solving previously insoluble test problems by optimizing flexibility levels.

Kazemian, I. & Aref, S. *Glob J Flex Syst Manag* (2016) <http://dx.doi.org/10.1007/s40171-016-0130-8>.

Keywords: supply chain; network design; flexibility; bottleneck; uncertainty; multi-echelon.

Modelling temporal dependence of bivariate survival data using copulas

Jose Romeo[†] & *Renate Meyer*[‡]. [†]Massey University, New Zealand; [‡]University of Auckland, New Zealand.

Copula models have become increasingly popular for modelling the dependence structure in multivariate survival data. In this work we explore a Bayesian approach to estimate the temporal association between bivariate survival times.

We consider models belonging to the Archimedean family of copulas as dependence structure and piecewise exponential distributions for modelling the marginal distributions. The temporal dependence is defined so that the copula association parameter varies in time.

For the estimation, we assume a two-step procedure. In a first step, from a Bayesian approach, we estimate the marginals survival function considering the bivariate lifetimes independents and following piecewise exponential distributions. Next, the estimation of the copula dependence parameter is carried out using a temporal factorization of the likelihood function and an autoregressive distribution of first order as prior distribution. The dynamic dependence is quantified through the Kendall's tau coefficient.

We illustrate the methodology using paired data from a study of laser photocoagulation therapy for retinopathy in diabetic patients.

Keywords: multivariate survival analysis; copula models; piecewise exponential distributions.

R-like pseudocode for fitting mixtures of distributions

Murray Jorgensen. *Auckland University of Technology, New Zealand*.

I will present an argument for writing your own code for fitting a mixture of distributions to data rather than hunting around for a package or special-purpose software to do the job.

I will outline using pseudocode, handwaving, and numerical examples how to go about doing this in R.

I will show how to make the mixture fitting more robust by adding a uniformly distributed component to the mixture.

If time permits I will outline some situations in which a statistician might want to fit a mixture.

Keywords: R; finite mixture; outliers; robustness.

Reference prior for Gaussian mixture models

Kate Lee. *Auckland University of Technology, New Zealand*.

While mixtures of Gaussian distributions have been studied for more than a century, the construction of a reference Bayesian analysis of those models still remains unsolved, with a general prohibition of the usage of improper priors due to the ill-posed nature of such statistical objects. By creating a new parameterisation

Two principles for checking Bayesian methodologies

Robin Willink. *Ministry of Business, Innovation and Employment, New Zealand*.

This talk presents two principles for checking the performance of an overall Bayesian methodology, such as the use of a standard method of eliciting prior belief together with the use of standard sampling models. For the methodology to be logically acceptable, (i) the mean value of the posterior distribution function evaluated at the prior p -quantile must be close to p and (ii) the mean value of the prior distribution function evaluated at the posterior median must be close to 0.5. Adherence to these principles can be tested using records of prior and posterior distributions. The second principle can be related to the definition and estimation of a 'bias' associated with the methodology.

Keywords: Bayesian statistics; performance criterion; statistical inference; validation; posterior distribution; median.

cantered on the mean and variance of the mixture distribution itself, we are able to develop here a genuine non-informative prior for Gaussian mixtures with an arbitrary number of components.

We demonstrate that the posterior distribution associated with this prior is almost surely proper and provide MCMC implementations that exhibit the expected exchangeability.

Keywords: reference prior; mixture model; improper prior.

WB Model Selection

WA224

Using mixture prior distributions to understand Bayes factors

Matthew Schofield & Richard Barker. University of Otago, New Zealand.

It is well known that Bayes factors and posterior model probabilities are sensitive to the prior distributions assumed on the parameters. This has led to comments such as “Diffuse prior distributions ... must be used with care” (Robert 2014) and “We do not see Bayesian methods as generally useful for giving the posterior probability that a model is true, or the probability for preferring model A over model B” (Gelman and Shalizi 2013).

We consider the calculation of Bayes factors for nested models. We show this is equivalent to a model with a mixture prior distribution, where the weights on the resulting posterior are related to the Bayes factor. These results allow us to directly compare Bayes factors to shrinkage priors, such as the Laplace prior used in the Bayesian lasso. We use these results as the basis for offering practical suggestions for estimation and selection in nested models.

Keywords: Bayes factors; Jeffreys-Lindley paradox; mixture prior distribution; regularization; lasso; model selection.

Using APC (all possible comparison) criterion for the analysis of non-orthogonal screening designs

Abu Zar Md. Shafiullah & Arden Miller. University of Auckland, New Zealand.

The APC-criterion for analysing orthogonal screening designs controls the error rate at the specific level of IER (individual error rate), EER (experiment wise error rate) or FDR (false discovery rate).

In this talk I present the application of APC-criterion to the non-orthogonal designs using the `leaps` package (Lumley, 2015) in R. Our simulation studies show that the APC-penalty constants for the orthogonal designs perform well with near-orthogonal (D-optimal) designs.

In general the effect size and the proportion of active

effect contrasts affect the performance of the criterion but the trade-off between error rate and power is maintained nicely over a wide range of possible true models.

Keywords: APC-criterion; non-orthogonal designs; D-optimality; IER; EER; FDR.

The accuracy of the APC method of model selection for different types of error control

Arden Miller & Abu Zar Md. Shafiullah. University of Auckland, New Zealand.

The “all possible comparisons” method (APC) of analysing unreplicated orthogonal factorial experiments is an extremely flexible AIC-like criterion. The APC method can be adapted to provide strong error control for three common measures of error: the individual error rate (IER), the experimentwise error rate (EER) and the false discovery rate (FDR).

This talk explores the accuracy of APC for these different types of error control. The accuracy is influenced by a number of factors including the size of the experiment, the prevalence of active effects, and the size of active effects in addition to the type and level of error control. Results from an extensive simulation study that explores accuracy under different scenarios are presented.

Keywords: model selection; screening experiment; orthogonal design.

Tourism Analytics

WA220

WB

On proportional odds modelling and marginal effects of Sardinian hotels data

*Thomas Yee[†], Giulia Contu[‡] & Claudio Conversano[‡].
[†]University of Auckland, New Zealand; [‡]University of Cagliari, Italy.*

We assembled a data set containing over 500 satisfaction scores appearing on the TripAdvisor website between the years 2008 and 2016 regarding hotels in Sardinia, Italy. The satisfaction data refer to the reputation of hotel located along Sardinian coasts, as expressed by clients with respect to different services (e.g., breakfast, restaurant, swimming pool) offered by the hotel. The data has been placed in the `VGAMdata` R package.

The statistical analysis involves fitting a proportional odds model with smoothing, and computing the marginal effects so that individual hotels could be identified that would benefit most by improving certain services.

Keywords: marginal effects; proportional odds model; Sardinia hotels; vector generalized additive model; `VGAM` R package; `VGAMdata` R package.

Tourism analytics: a recommendation engine for itinerary planning

Anthony Downward[†], **Subeh Chowdhury**[‡] & **Cengyue Chen**[†]. [†]*Department of Engineering Science, University of Auckland, New Zealand;* [‡]*Department of Civil Engineering, University of Auckland, New Zealand.*

Tourism is one of New Zealand's largest export industries, second only to dairy. In the year to June 2015, total tourism expenditure was \$29.8 billion. Despite this significant revenue, planning for tourists is not an easy task, since information about tourism activities is not centralised. Moreover, information about tourists themselves is not available for tourism providers. Due to these factors the potential benefits of tourism and not being maximized.

In this project we seek to address these issues by creating a planning tool for tourists, accounting for both individual preferences and demographic information. This tool will provide an optimised itinerary which maximises the enjoyment of the tourist, subject to their time and budgetary restrictions. This is modelled as a variant of a vehicle routing problem.

The additional benefit of this tool is that it can become a source of real-time information on tourists and provide the tourism sector with detailed information on the types of tourists visiting New Zealand.

Keywords: vehicle routing problem; recommendation engine; mixed-integer programming; tourism.

Developing forecasts for New Zealand of regional spend by international tourists

Raphael Aggio, **Michael Webster** & **Peter Ellis**. *Ministry of Business, Innovation and Employment, New Zealand.*

Estimating the future expenditure of international tourists visiting the country is important for the New Zealand tourism industry. It allows for planning and investing resources in a more strategic manner. MBIE already publishes annual forecasts of tourism numbers and spend at the national level, and stakeholders frequently ask for a regional dimension. We have now developed forecasts for 2016 – 2022 of international visitor spend by country of origin, at Territorial Authority (TA), Regional Tourism Organisation (RTO) area, and Regional Council (RC) levels. Using TA-level data from MBIE's Monthly Regional Tourism Estimates for the period of 2009 to 2015, we have evaluated the performance of a range of different modelling techniques based on combinations of Autoregressive Integrated Moving Average (ARIMA) and Exponential Smoothing (ETS). Variants include one model or an ensemble of models; with or without Box-Cox transformations; and with or without explicitly taking the hierarchical nature of the data into account. The relative performance of each method is assessed by its ability to predict 2015 results.

Keywords: forecast; hierarchical analysis; hybrid model.

WEDNESDAY 13:00 – 14:00

Plenary 5

WG403 WC

Optimizing prioritized and nested solutions

David Morton. *Northwestern University, USA.*

A typical optimization model in operations research allocates limited resources among competing activities to derive an optimal portfolio of activities. In contrast, practitioners often form a rank-ordered list of activities, and select those with the highest priority, at least when choosing an activity is a yes-no decision. Ranking schemes that score activities individually are well known to be inferior. So, we describe a class of two-stage stochastic integer programs that accounts for structural and stochastic dependencies across activities and constructs an optimized priority list. We further discuss a class of optimization models, subject to a single *budget* constraint, that naturally leads to a family of optimal nested solutions at certain budget increments. We use several applications to both motivate the work and illustrate results, ranging from a stochastic facility location model to a hierarchical graph clustering problem. We also describe possible extensions.

Keywords: stochastic program; submodular optimization; graph clustering.

AUTHOR INDEX

Aasgard, Ellen Krohn	MB	Haji Vahabzadeh, Ali	MB TE	Park, Diane	TC
Abbasi, Babak	WB	Hall, Lisa	MD	Parry, Katharina	MB
Abbott, Max	TB	Hankin, Robin	TB	Parry, Matthew	TC
Abeywickrama, Raneetha	MB	Harte, David	TE	Pavlov, Valery	MB TE
Abi, Naeimeh	TB	Hasannasab, Maryam	WB	Pawley, Mat	MC
Adams, Thomas	MC TC	Haslett, Steve	TB	Pears, Russel	WA
Adegoke, Nurudeen	MC	Haviv, Moshe	MB	Pearson, Janet	TC
Aggio, Raphael	WB	Havron, Andrea	MC	Peng, Xianhua	WA
Alamri, Osama	WB	Hayakawa, Yu	MB	Penny, Richard	TB
Anderson, Eddie	WA	Haywood, John	WA	Pfannkuch, Maxine	MD
Aref, Samin	TB WB	Hazelton, Martin	MB MC TE	Philpott, Andy	MB MD TC TE WA
Arnold, Richard	MB TB TE	Hermansson, Niffe	MB	Pirikahu, Sarah	MB
Ashford, Cliff	TB	Hillas, Lisa	TC	Raith, Andrea	MC MD TB TC
Bacic, Boris	WA	Hodgson, Ella	TB	Ramakrishnan, Kumaran	TB
Baddeley, Adrian	WA	Holmes, John	MB	Ramsden, Alistair	TB
Barker, Richard	MB WB	Holmes, Mark	MB	Ranta, Jussi	TB
Baucke, Regan	TE	Hong, Chuen Yen	TC	Ridler, Samuel	TB
Bellringer, Maria	TB	Hosseinfard, Zahra	WB	Romeo, Jose	WB
Bilton, Penelope	TB	Huda, S	WA	Roshdi, Israfil	WB
Bostan, Mihai	TC	Hudson, Irene Lena	MB	Ross, Kevin	TD
Brainard, Andrew	TB	Hussein, Tasleem	TB	Roti, Pål Henrik	MB
Bran, Paula	MB	Islam, Md.Shahidul	WA	Rouse, Paul	WB
Branford, Alan	MC	Jayalath, Chapa	MC	Rowe, Suzanne	TB
Brauning, Rudiger	TB	Jayamaha, Nihal	WB	Ruddell, Keith	TE
Brewer, Brendon	MC	Jones, Geoff	MB TB TE	Salehipour, Amir	WB
Budgett, Stephanie	MD	Jorgensen, Murray	WB	Santander, Valentín	TB
Chen, Cengyue	WB	Julian, Garry	MD	Schofield, Matthew	MB WB
Cheng, Daniel	TB	Kazemian, Iman	WB	Shafiullah, Abu Zar Md.	WB
Chowdhury, Subeh	MC WB	Khan, Muhammad Shuaib	MB	Sibanda, Nokuthaba	MC TB
Choy, Boris	MC	King, Robert	MB	Siva, Ganesh	TB
Christensen, Nelson	MC	Kirch, Claudia	MC TE	Smith, Adam	MC
Chukova, Stefanka	MB	Kou, Steven	WA	Soleymani, Mehdi	WA
Clarke, Shannon	TB	Lamont, Alastair	TB	Soysa, Ishani	MB
Cleland, Isaac	TC	Laslett, Mark	TC	Stein, Markus	MB
Contu, Giulia	WB	Leclère, Vincent	WA	Surya, Budhi Arta	MD
Conversano, Claudio	WB	Lee, Alan	MC	Te Ao, Braden	TE
Cook, Dianne	MA	Lee, Kate	WB	Thi Thu Pham, Hoa	MC
Coomarasamy, Christin	TB	Lee, Stephen	WA	Thi Thu Pham, Huong	MC
Cory-Wright, Ryan	MB	Lester, Phil	WA	Thielen, Clemens	MD
Costilla, Roy	TB TE	Leyland, Geoff	TC	Thomas, Doug	WB
Crequer, John	TC	Lillo, Felipe	TB	Thornton, Vanessa	TB
Curran, James M.	WB	Lin, Kuan-Min	TC	Tidswell, James	MD
Daly, Cathy	MB	Liu, Ivy	MB TB TE	Tokolahi, Ema	TC
Darfiana, Nur	MC	Liu, Jing	TB	Turner, Rolf	MB
Davies, Tilman	TE WA	Lumley, Thomas	MC MC	Vaithianathan, Rhema	TF
Dhau, Inos	MC	Maindonald, John	MB	van Dam-Bates, Paul	WA
Dodds, Ken	TB	Margaritis, Dimitris	WB	van Stijn, Tracey	TB
Dominguez Martin, Ruth	TC	Marquart, Louise	MC	Vandal, Alain	TC
Downward, Anthony	MC MD TE TE	Marshall, Jonathan	MC	Vignes, Matthieu	TB
Dowson, Oscar	MD	Mason, Andrew	MD TB TC	Walker, Cameron	MC
Edwards, Matthew	MC TE	Maturana, Patricio	WA	Wang, Rachel	MC
Ehrgott, Matthias	TC	McCarthy, James	MC	Webster, Michael	WB
Elliott, Tom	MC	McEwan, John	TB	Welsh, Melissa	MC
Ellis, Peter	WB	McNair, Peter	TC	Westbrooke, Ian	TB
Fernandez, Daniel	MB	Meier, Alexander	MC TE	White, Steve	TC
Fernando, M.A.C.S. Sampath	MB	Mellor, Hamish	MB	Williams, Jonathan	MB
Fewster, Rachel	TB	Meyer, Renate	MC TE WB WB	Willink, Robin	MC WB
Fletcher, David	TC	Millar, Russell	MC WA	Wilson, Mark	TB
Flynn, Claire	TE	Miller, Arden	WB	Wockner, Leesa	MC
García, Leidy	TB	Minas, James	WB	Wong, Hoi Ying	MC
Garrett, Nick	TB	Miranda Soberanis, Victor	TC	Xu, Leshun	MC
Gatland, Danielle	MC	Morton, David	WC	Xu, Xingbo	WA
Gerard, Henri	WA	Moses, Caleb	TB	Yee, Thomas	TC WB
Goodman, Jesse	TB	Moslemi Naeni, Leila	WB	Zakeri, Golbon	MB TE TE
Gormley, Andrew	WA	Mujuru, Munyaradzi	MC	Zeng, Irene	TB
Govindaraju, Kondaswamy	TE	Mukherjee, Rahul	WA	Zhang, Wei	TB
Green, Peter	TC	Murrell, Paul	MD	Zhou, Xin	TE
Grigg, Nigel	WB	O'Rourke, Peter	MC	Ziedins, Ilze	MB
Grün, Bob	MD	O'Sullivan, Michael	MC TA		
Habibian, Mahbubeh	TE	Paas, Leo	MC		