Nonlinear Time Series Analysis
Thresholding and Beyond

in honour of

Professor Howell Tong's 70th Birthday

19 -20 September 2014

Abstracts of Talks

(in the alphabetic order of the presenting authors' surnames)

Change-points in High Dimensional Settings

John Aston, Cambridge University (jada2@statslab.cam.ac.uk)

While there is considerable work on change-point analysis in univariate time series, more and more data being collected comes from high dimensional multivariate settings. One way to develop an asymptotic framework for such data is to use a Panel data setting where the number of dimensions increases with the sample size. Using contiguous alternatives in such a setup we can compare the asymptotic power of such projection procedures (including an oracle projection on the one hand and a random projection on the other hand) with a Panel statistic that uses the full multivariate information. If information is available to constrain the search region of the test, corresponding projections can lead to a considerable gain in power. All procedures depend on the unknown covariance structure between components whose estimation is very problematic in high dimensional situations, and the possible presence of change-points only further increases these difficulties. If the covariance assumptions made are violated, it not only leads to huge size problems for the Panel statistics but also to a massive power loss, where by looking at contiguous alternatives, the asymptotic power effectively becomes equivalent to choosing a random search direction to apply a univariate test. The size of projections procedures on the other hand is robust with respect to the unknown covariance structure between channels. At the same time while the power can also be affected by misspecification, the impact is often considerably smaller. The theoretic results will be accompagnied by small sample simulations and real data analysis.

(Joint work with Claudia Kirch)

Dynamic Factor Models, Cointegration, and Error Correction Mechanisms

Matteo Barigozzi, LSE (m.barigozzi@lse.ac.uk)

In this paper we study Dynamic Factor Models when the factors are I(1) and singular, By combining the classic Granger Representation Theorem with recent results by Anderson and Deistler on singular stochastic vectors, we prove that, for generic values of the parameters, the factor has an Error Correction representation with two unusual features: (i) the autoregressive matrix polynomial is finite, (ii) the number of error-terms is equal to the number of transitory shocks plus the difference between the dimension and the rank of the factor. This result is the basis for the correct specification of an autoregressive model for the factor. Estimation of impulse-response functions is also discussed. Results of an empirical analysis on a US quarterly database support the use of our model.

(Joint work with Marco Lippi, Matteo Luciani)

Prediction of Lévy-driven CARMA processes

Peter Brockwell, Colorado State University (peter.brockwell@gmail.com)

Is talk we consider the problem of determining the conditional expectations, $\mathbb{E}(Y(h)|Y(u), -\infty < u \leq 0)$ and $\mathbb{E}(Y(h)|Y(u), -M \leq u \leq 0)$ where $h > 0$, $0 < M < \infty$ and $(Y(t))_{t\in\mathbb{R}}$ is a continuous-time ARMA (CARMA) process driven by a Lévy process L with $\mathbb{E}|L(1)| < \infty$. If the driving Lévy process satisfies $\mathbb{E}(L(1)^2) < \infty$ then these are the minimum mean-squared error predictors of $Y(h)$ given $(Y(t))_{t\leq 0}$ and $(Y(t))_{-M\leq t\leq 0}$ respectively. In the course of the derivations we establish conditions under which the sample-path of L can be recovered from that of Y, both when Y is causal and strictly stationary and, without these assumptions, when L is a pure-jump Lévy process. When $\mathbb{E}(L(1)^2) < \infty$ and Y is causal and strictly stationary we also determine the best linear predictors $P(Y(h)|Y(u), u \leq 0)$ and $P(Y(h)|Y(-n\Delta), n \in \mathbb{N})$, comparing their performance with that of $\mathbb{E}(Y(h)|Y(u), -\infty < u \leq 0)$. Finally we use the expression for $P(Y(h)|Y(-n\Delta), n \in \mathbb{N})$ to establish a very simple algorithm for determining the parameters of the ARMA process obtained by sampling the CARMA process at regular intervals.

Threshold Modelling – Some Recent Advances and Challenges

Kung-Sik Chan, University of Iowa (kungsik.chan@gmail.com)

In this talk, I shall give an overview of some recent breakthroughs in threshold modelling, with illustrations drawn from financial time series and other areas. Some interesting future research problems and challenges in threshold modelling will be highlighted.

LASSO Estimation of Threshold Autoregressive Models

Ngai Hang Chan, Chinese University of Hong Kong (nhchan@sta.cuhk.edu.hk)

This paper develops a novel approach for estimating a threshold autoregressive (TAR) model with multiple-regime and establishes its large sample properties. By re-framing the problem in a regression variable selection context, a least angle regression (LARS) procedure is proposed to estimate a TAR model with an unknown number of thresholds, where the computation can be performed efficiently. It is further shown that the number of thresh- olds and the location of the thresholds can be consistently estimated. A near optimal convergence rate of the threshold parameters is also established and an improved version of the algorithm that incorporates LARS and stepwise regression variable selection technique is introduced. Simulation studies are conducted to assess the performance in finite samples. The results are illustrated with an application to the quarterly U.S. real GNP data over the period 1947-2009.

High-dimensional panel data segmentation

Haeran Cho, University of Bristol (haeran.cho@bristol.ac.uk)

In this paper, we discuss the problem of detecting multiple change-points in the mean of (possibly) high-dimensional panel data. CUSUM statistics have been broadly adopted in segmenting both univariate and multivariate data, which often estimate the likely location of a change-point in a single panel as where the series of CUSUM statistics is maximised. For n-variate data, several methods have been proposed to simultaneously segment the multiple panels by aggregating the n CUSUM series, thereby avoiding the necessity to prune the change-points estimated from individual panels.

One of the challenges in panel data segmentation is the fact that the detectability of a change-point is affected by a range of factors, such as its "sparsity" across the panels, the magnitude of jumps at the change-point and the "unbalancedness" of its location. However, the CUSUM statistic aggregation methods in the literature do not cater to all possible change-point configurations in the sense that e.g., taking the pointwise average of CUSUM statistics fails to detect a change-point which is very sparse across the panels.

Therefore we propose a novel approach to CUSUM statistic aggregation which can handle a wide spectrum of change-point scenarios in a data-driven way. To this end, the use of *double* CUSUM statistic is explored, where the key step is to partition the CUSUM statistics at each time point into "signals" (those corresponding to panels with a change-point in the neighbourhood) and "noise" (those corresponding to panels without a change-point). We show that the double CUSUM statistic, combined with a binary segmentation algorithm, achieves consistent change-point detection in terms of the total number of detected change-points and their locations, and its good practical performance is demonstrated in a comparative simulation study.

Volatility Decomposition and Online Volatility-Estimation with Nonlinear Market Microstructure Noise Models

Rainer Dahlhaus, University of Heidelberg (dahlhaus@statlab.uni-heidelberg.de)

A technique for online estimation of spot volatility for high-frequency data is developed. The method uses a price model with time shift in combination with a nonlinear market microstructure noise model. A benefit of the model is that it leads to an identifiable decomposition of spot volatility into spot volatility per transaction and the trading intensity - thus highlighting the influence of trading intensity on volatility. The online algorithm uses a computationally efficient particle filter. It works directly on the transaction data and updates the volatility estimate immediately after the occurrence of a new transaction. It also allows for the approximation of the unknown efficient prices. For volatility estimation a nonparametric recursive EM algorithm is used. We neither assume that the transaction times are equidistant nor do we use interpolated prices. For the theoretical investigations of the estimates we present a theoretical framework with infill asymptotics.

Comparing the Accuracy of Copula-Based Multivariate Density Forecasts in Selected Regions of Support

Cees Diks, University of Amsterdam (C.G.H.Diks@uva.nl)

This paper develops a testing framework for comparing the predictive accuracy of copula-based multivariate density forecasts, focusing on a specific part of the joint distribution. The test is framed in the context of the Kullback-Leibler Information Criterion, but using (out-of-sample) conditional likelihood and censored likelihood in order to focus the evaluation on the region of interest. Monte Carlo simulations document that the resulting test statistics have satisfactory size and power properties for realistic sample sizes. In an empirical application to daily exchange rate returns we find evidence that the dependence structure varies with the sign and magnitude of returns, such that different parametric copula models achieve superior forecasting performance in different regions of the support. Our analysis highlights the importance of allowing for lower and upper tail dependence for accurate forecasting of common extreme appreciation and depreciation of different currencies.

Switch time modeling for gene expression: An overview

Bärbel Finkenstädt, University of Warwick (B.F.Finkenstadt@warwick.ac.uk)

Time series relating to gene expression are now routinely measured in various important biological experiments such as microarrays, nanostring, etc as well as experiments based on bioluminescent imaging. One of the most important aim is to gain an understanding of the transcriptional regulation of genes, i.e. what determines their activation. A natural model is to assume that gene activation is a constant rate birth process but that the rate may change to different levels at unknown time points leading to the piecewise linear switch model. Statistical inference for such a model poses interesting and challenging problems, in particular since experiments can only measure events downstream whereas the processes of interest remain unobserved. We will give an overview of our experience with fitting such switch models to real data where, depending on the type of experiment, our assumptions range from stochastic differential equations to the use of ordinary differential equations, along with realistic stochastic formulations of the measurement processes. We will also present further results on extending this nonlinear approach to the multivariate case of identifying networks of interacting genes. Here we find that the concept of thresholding constitutes a simple yet realistic and very effective modelling device to help identifying network connections.

(Joint work with Dafyd Jenkins, Kirsty Hey, George Minas and David Rand)

The use of randomness in time series analysis

Piotr Fryzlewicz, LSE (p.fryzlewicz@lse.ac.uk)

This is an exploratory talk in which we describe different potential uses of randomness in time series analysis.

In the first part, we talk about Wild Binary Segmentation for change-point detection, where randomness is used as a device for sampling from the space of all possible contrasts (change-point detection statistics) in order to reduce the computational complexity from cubic to just over linear in the number of observations, without compromising on the accuracy of change-point estimates. We also discuss an interesting related measure of change-point certainty/importance.

In the second part, we use random contemporaneous linear combinations of time series panel data coming from high-dimensional factor models and argue that this gives the effect of "compressively sensing" the components of the multivariate time series, often with not much loss of information but with reduction in the dimensionality of the model.

In the final part, we speculate on the use of random filtering in time series analysis. As an illustration, we show how the appropriate use of this device can reduce the problem of estimating changes in the autocovariance structure of the process to the problem of estimating changes in variance, the latter typically being an easier task.

The mathematical structure of genetic information: a (nonlinear) time series perspective

Simone Giannerini, University of Bologna (simone.giannerini@unibo.it)

In this talk I will provide a brief account of the mathematical model of the genetic code introduced in Gonzalez (2008) and further studied in Giannerini et al.(2012). In particular I will focus on dichotomic classes, quantities derived from the mathematical properties of the model that can be seen as nonlinear functions of the information contained in a dinucleotide. Interestingly, such classes have a precise biochemical counterpart. We use the dichotomic classes as a coding scheme for protein coding DNA sequences and study their mutual dependence. We obtain meaningful tests by using an entropy based measure possessing many desirable properties together with suitable resampling techniques. We find universal correlation/optimization structures that point to the existence of irreversible processes and that might be related to possible mechanisms of error correction in the management of genetic information.

(Joint work with Diego Luis Gonzalez and Rodolfo Rosa)

IntegratedARCH and ARmodels: origins of long memory

Liudas Giraitis, Queen Mary University (l.giraitis@qmul.ac.uk)

Although propertiesof ARCH(∞) model arewellinvestigated, existenceof long memory FIGARCHand IARCH solutions was not established intheliterature. These two popular ARCH type models which are widely usedin appliedliterature, werecausing theoretical controversybecause ofsuspicionthat other solutionsbesides thetrivial zero one do not exist. SinceARCHmodelswith non-zerointercept have a unique stationary solution and exclude longmemory,existence of finite varianceFIGARCH and IARCH models and, thus, possibility of long memoryin ARCH settingwasdoubtful. The present papersolvesthiscontroversy by showing thatFIGARCH andIARCH equations havea non-trivialcovariance stationary solution,andthat such solutionexhibits long memory. Existence and uniqueness of stationary Integrated AR(∞) processes is also discussed, and long memoryasinheritedtheir feature isestablished.Summarizing, weshowthat covariancestationary IARCH, FIEGARCH andIAR processes exist, theirclassiswide, and they exhibit longmemory.

(Joint work withDSurgailisandASkarnulis)

High dimensional and Banded Vector Autoregression

Shaojun Guo, LSE and Chinese Academy of Sciences (guoshaoj@amss.ac.cn)

The vector autoregressive (VAR) model is a powerful tool in modeling multiple time series and has been widely used in various fields. Due to high dimensionality, traditional estimation procedures are no longer applicable directly since the number of the coefficient matrices grows dramatically. To overcome this difficulty, one often supposes that the coefficient matrices are sparse and fit the VAR model by Lasso-type penalized approach.

The sparsity assumption is essentially good but too 'universal'; it does not take into account some specific or prior information when facing a given dataset. In this paper, we consider the VAR model that the coefficient matrices are banded. The banded structure can be achieved by 'ordering' the data either based on some relevant subject knowledge or a linear transformation. We propose a modified marginal Bayesian information criterion to estimate the bandwidth and show that it is asymptotically consistent. We provide a simple ridge procedure to recover coefficient matrices and establish the convergence rate of this new estimator. Estimating auto-covariance matrices Σ_0 and Σ_1 is of also fundamental importance. We show that under certain conditions, Σ_0 and Σ_1 are 'approximately' banded in the sense that their spectral norms are very close to those of some banded matrices, respectively. Based on this, their banded estimators are shown to be consistent in the spectral norm under regularity conditions. Moreover, their convergence rates in the spectral norm are $\log(n/\log^{-1}(p))\sqrt{\log p/n}$. Simulation studies are conducted and support the relevant theories.

(Joint work with Yazhen Wang and Qiwei Yao)

Quantile Spectral Processes Asymptotic Analysis and Inference

Marc Hallin, Free University of Brussels and Princeton University (mhallin@ulb.ac.be)

Quantile- and copula-related spectral concepts recently have been considered by various authors. Those spectra, in their most general form, provide a full characterization of the copulas associated with the pairs (X_t, X_{tk}) in a process $\{X_t\}$, and account for important dynamic features, such as changes in the conditional shape (skewness, kurtosis), time-irreversibility, or dependence in the extremes, that their traditional counterpart cannot capture. In this paper, we provide a detailed asymptotic analysis of a class of smoothed rank-based cross-periodograms associated with the copula spectral density kernels introduced in Dette et al. (2011). We show that, for a very general class of (possibly non-linear) processes, properly scaled and centered smoothed versions of those cross-periodograms, indexed by couples of quantile levels, converge weakly, as stochastic processes, to Gaussian processes. A first application of those results is the construction of asymptotic confidence intervals for copula spectral density kernels. The same convergence results also provide asymptotic distributions (under serially dependent observations) for a new class of rank-based spectral methods involving the Fourier transforms of rank-based serial statistics such as the Spearman, Blomqvist or Gini autocovariance coefficients.

(Joint work with Holger Dette, Stanislav Volgushev, and Tobias Kley)

Specification in time series regression models

Javier Hidalgo, LSE (F.J.Hidalgo@lse.ac.uk)

We propose a specification test in regression models with data that follow a generalized linear model in the sense of Hannan (1970) or data exhibiting, possibly, long memory dependence. We show that, contrary to the case with weakly dependent data, test has some interesting features and properties. In particular, in some scenarios the test follows an asymptotically normal random variable.

Bayesian Inference for partially observed SDEs Driven by Fractional Brownian Motion

Kostas Kalogeropoulos, LSE (k.kalogeropoulos@lse.ac.uk)

We consider continuous-time diffusion models driven by fractional Brownian Motion (fBM), with observations obtained at discrete-time instances. As a prototypical scenario we will give emphasis on a a stochastic volatility (SV) model allowing for memory in the volatility increments through an fBM speci?cation. Due to the non-Markovianity of the model and the high-dimensionality of the latent volatility path, estimating posterior expectations is a computationally challenging task. We present novel simulation and re-parameterisation framework based on the Davies and Harte method and use it to construct a Markov chain Monte-Carlo (MCMC) algorithm that allows for computationally e?cient parametric Bayesian inference upon application on such models. The algorithm is based on an advanced version of the so-called Hybrid Monte-Carlo (HMC) that allows for increased e?ciency when applied on high-dimensional latent variables relevant to the models of interest in this paper. The inferential methodology is examined and illustrated in the SV models, on simulated data as well as real data from the S&P500/VIX time series. Contrary to a long range dependence attribute of the SV process (Hurst parameter H ¿ 1=2) many times assumed in the literature, the posterior distribution favours H ¡ 1=2 that points towards medium range dependence.

(Joint work with Joseph Dureau and Alexandros Beskos)

Baxter's inequality and sieve bootstrap for random fields

Jens-Peter Kreiss, Technical University of Braunschweig (j.kreiss@tu-bs.de)

The concept of autoregressive sieve bootstrap for time series is extended to random fields. Given a finite data sample of rectangular shape, the procedure fits a finite-order autoregressive model to the sample using Yule-Walker-type estimators. The residuals of this fit are resampled, which allows for construction of a bootstrap sample in order to approximate the distribution of the statistic of interest. The distinctive feature of the sieve bootstrap is that the order of the AR fit is chosen depending on the sample size; in particular, it increases to infinity as the sample size tends to infinity, but at a much slower rate. A general check criterion is presented which allows for a large class of statistics to determine whether the proposed bootstrap procedure works or not. This work depends largely on two general results for random fields which may be of interest of its own: The first one is a one-sided autoregressive representation of the underlying spatial process with summable autoregressive coefficients (which goes back to the early work of Whittle), while the other one is a kind of Baxters inequality for random fields.

(Joint work with Marco Meyer and Carsten Jentsch)

Neural network sieve bootstrap for nonlinear time series

Michele La Rocca, University of Salerno (larocca@unisa.it)

Sieve bootstrap schemes have proved to be effective resampling techniques. These methods are easy to implement, since they retain the simplicity of the classical residual bootstrap, and are robust to model misspecification, since they should be viewed as nonparametric techniques. In this talk, we discuss a bootstrap scheme suitable for nonlinear processes, developed in the same spirit as the AR-sieve bootstrap which uses the class of feedforward neural networks as sieve approximators. This seems to be justified for several reasons. First, feedforward neural networks are popular models in nonlinear time series analysis for their good forecasting accuracy. Secondly, this approach does not suffer for the so-called curse of dimensionality, since the approximation form does not appear to be so sensitive to the increasing dimension, at least within the confines of particular classes of functions. Therefore, extension of the neural network sieve bootstrap procedure to high- dimensional models is expected to be more straightforward than that of other nonparametric approaches. Moreover, neural networks are global nonparametric methods, and their use could stress different features and data structures when compared with the local nonparametric methods such as the kernel one. Finally and most importantly, under quite general conditions, this class of models provides an arbitrarily accurate approximation to an unknown target function of interest. If the network model is fitted to the data in such a way that the complexity of the network is allowed to increase at a proper rate with the sample size, the resulting function estimator can then be viewed as a nonparametric sieve estimator. The resampling scheme from the residuals of feedforward neural networks is shown to be asymptotically justified. Moreover, a Monte Carlo study shows that it has comparable performances to the AR-sieve bootstrap, when the process is linear, but it delivers better results when the process is nonlinear, both in terms of bias and variability.

(Joint work with Cira Perna and Francesco Giordano)

Nonparametric Eigenvalue-Regularized Precision or Covariance Matrix Estimator

Clifford Lam, LSE (c.lam2@lse.ac.uk)

Recently there are numerous works on the estimation of large covariance or precision matrix. The high dimensional nature of data means that the sample covariance matrix can be ill-conditioned. Without assuming a particular structure, much efforts have been devoted to regularizing the eigenvalues of the sample covariance matrix. We introduce nonparametric regularization of these eigenvalues through subsampling of the data. The subsampling idea for covariance matrix estimation is originally introduced in Abadir, Distaso and ike (2010). We improve on their covariance estimator, and for the first time provides vigorous proof that our version enjoys asymptotic optimal nonlinear shrinkage of eigenvalues with respect to the Frobenius error norm. Coincidentally, this nonlinear shrinkage is asymptotically the same as that introduced in Ledoit and Wolf (2012). One advantage of our estimator is its computational speed when the dimension p is not extremely large. Our estimator also allows p to be larger than the sample size n, and is always positive semi-definite. We prove that with respect to the Steins loss function, the inverse of our estimator is the optimal precision matrix estimator. We also showed that all the aforementioned optimality holds for data with a factor structure as well, which can be useful in portfolio allocation. Our method avoids the need to estimate the unknown factors and factor loadings matrix first, and directly gives the covariance or precision matrix estimator. We compare the performance of our estimators with other methods through extensive simulations and a real data analysis.

Analysis of a Laser-Chaos Communication Experiment

Tony Lawrance, University of Warwick (A.J.Lawrance@warwick.ac.uk)

This talk is concerned with analysing the results from an experimental communication system using distributed feedback chaotic lasers. The system generates 3 series of light intensity waves which are instrumentally converted to 10m discrete voltage values. Ten such runs were set up by Atsushi Uchida of Saitama University, Japan, as part of a collaboration, and the analysis arbitrarily concerns data from one run. The structure of the system is that there is a transmitter laser generating a wave of light intensity values which includes binary messages and which passes through a 60m fibre-optic cable before arriving at a receiver station. Using the idea that chaotic lasers can be synchronized by a trigger signal, a second identical laser at the receiver station is synchronized with the transmitter laser before inclusion of the binary message. These two chaotic streams of signals available at the receiver allow the binary messages to be decoded, but not without error there is optical noise and synchronization error. The decoding system used is known as chaos shift-keying, not previously used in laser communication, and is correlation based. A number of interesting data aspects only available in experimental circumstances, will be displayed; these include optical noise and synchronization error which are both found to be non-Gaussian and dependent on the transmitter laser signal. An important communication feature is bit error rate, often obtained by model simulation, but which can here be obtained empirically based on the 10m voltage sequences. The analysis points to low rates and the effectiveness of the synchronization. The security of the systems is not based on encryption, rather on stenography, the idea that the existence of the message is hidden in the transmitted laser wave. A mathematical model of the system can be adapted from the authors earlier work when the transmitter laser is represented by chaotic map output and the optical noise and synchronization error are IIGaussian, but awaits development under the present more challenging laser circumstances.

Inference on Structural Breaks in Panel Data Models with Interactive Fixed Effects

Degui Li, University of York (degui.li@york.ac.uk)

In this paper we consider estimation and inference of common structural breaks in panel data models with interactive fixed effects which are unobservable. We introduce a penalized principal component estimation procedure via adaptive group fused LASSO to detect the multiple structural breaks. Under mild conditions, we show that with probability tending to one our method can correctly determine the unknown number of breaks and consistently estimate the common break dates. Furthermore, to improve the convergence rates, we estimate the regression coefficients through the post-LASSO method and establish the asymptotic distribution theory of the resulting estimators. We also propose a data-driven method to determine the tuning parameter involved in the penalized principal component estimation procedure. The Monte Carlo simulation results demonstrate that the proposed method works well in finite sample case.

Some Results on the Buffered Time Series Models

Wai Keung Li, University of Hong Kong (hrntlwk@hku.hk)

We consider a new class of threshold time series models known as buffered processes. In this new class of models switching back and forth between two regimes depends on two different thresholds. We first investigate the self excited buffered autoregressive (BAR) process to some extent including a sufficient condition for geometric ergodicity and the asymptotic properties of the least squares estimators. We then extend the class of models to cover conditional heteroscedasticty resulting in the buffered GARCH and buffered AR-GARCH models. Simulation studies and applications to real data are considered to illustrate the potential of this new type of threshold models.

Self-weighted LAD Estimation for Infinite Variance Threshold Autoregressive Models

Shiqing Ling, Hong Kong University of Science and Technology (maling@ust.hk)

When its variance is infinite, the least squares estimator of the threshold autoregressive (TAR) model may not be consistent. Neither theory nor methodology can be applied to model fitting in this case. This paper is to develop a systematic procedure of statistical inference for the infinite variance TAR model. We first investigate the self-weighted least absolute deviation estimation for the model. It is shown that the estimated slope parameters are root-n consistent and asymptotically normal, and the estimated threshold is n-consistent and its limiting distribution is the minimizer of a compound Poisson process. A sign-based portmanteau test is proposed for model checking. Simulations are carried out to assess the performance of our procedure and a real example is given.

(Joint work with Yang Yaxing)

Semiparametric Nonlinear Regression Models for Irregularly Located Spatial Time-series Data

Zudi Lu, University of Southampton (Z.Lu@soton.ac.uk)

Nonparametric and semiparametric approaches have been popular in nonlinear modelling of univariate or small number multivariate time series data. However, they become increasingly challenging when applied to nonlinear analysis of big spatial time-series network data.

Bigger time-series data with more complex structures collected at irregularly spaced sampling locations are becoming more prevalent in a wide range of disciplines. With few exceptions, however, practical statistical methods for modeling and analysis of such data remain elusive. Here, we propose a class of spatio-temporal autoregressive partially linear regression models, which permits possibly nonlinear relationships between responses and covariates. In addition, the dependence structure is stationary over time but nonstationary over space, while the sampling grids can be irregular in space. We develop a computationally feasible method for estimation and thus enable our methodology to be applicable in practice. Asymptotic properties of our proposed estimates will be established and comparison will be made, in theory and via simulations, between estimates before and after spatial smoothing. For illustration, our methodology is applied to investigate housing prices in relation to mortgage rates in the United States from 1991 to 2012, with threshold structure found to be helpful for prediction.

(Joint work with Dawlah Al-Sulami and Jun Zhu)

Modelling, Analysis and Forecasting of Locally Stationary Time Series

Guy Nason, University of Bristol (G.P.Nason@bristol.ac.uk)

In this talk we will consider the modelling, analysis and forecasting of locally stationary time series. We will begin by considering various tests for stationarity and then examine tools that might be used to elicit the time-varying second-order structure of a locally stationary time series including local autocorrelations and spectral estimates. We will also consider some recent developments for forecasting locally stationary series focusing on successful coverage of prediction intervals. We will consider these locally stationary techniques in the context of macroeconomic series analysis and forecasting.

Estimating Multivariate Non-stationary Time Series Models in bf the Fourier Domain

Sofia Olhede, University College London (s.olhede@ucl.ac.uk)

Multivariate time series are complex structures, where it is difficult to both specify simple, valid models, and also good and computationally efficient inference mechanisms. Such is complicated when the observed series are non-stationary. Identical problems exist for inhomogeneous random fields.We shall discuss the simple device of spectral blurring coupled with the local Whittle likelihood to implement inference. Spectral blurring enables us to deal with finite sample effects, and avoid bias issues.Secondly, we shall discuss complicated bivariate models, and how appropriate choices of modelling parameterisations can help us specify simple structures, that in normal modelling paradigms are hard to formulate.

(Joint work with Adam Sykulski)

Threshold Models for Count Time Series

Jiazhu Pan, University of Strathclyde (jiazhu.pan@strath.ac.uk)

We consider the autoregressive models with threshold nonlinearity or switching regimes for non-negative integer valued time series. Assuming the conditional distribution given historical information is Poisson distribution, and the link between the conditional mean and its past values as well as the observed values of the Poisson process is different when a lag state or an unobservable (hidden) variable is in different regimes, we get Poisson Threshold or Markovian Switching AR models. We prove that our models can be approximated by a geometrically ergodic process under mild conditions on parameters. We also discuss statistical inference for the proposed models. Simulation studies and application to modelling financial count time series are presented to support our methodology.

Shrinkage Estimation of the Dependence Structure of High Dimensional Time Series

Rainer von Sachs, Catholic University of Louvain (rainer.vonsachs@uclouvain.be)

In high-dimensional time series analysis, when (effective) sample sizes and dimensionality are of the same order of magnitude, estimating the dependence structure across a panel of time series, via the covariance matrix (in the static case) or via the cross-spectral density matrix (in the dynamic case), can pose severe problems: generally the nonparametric matrix estimators are close to being ill-conditioned, and hence numerically very unstable, due to the well-known phenomenon of potentially high linear correlation among the columns of the matrix. One possibility to regularize these estimators is shrinkage towards a well-conditioned target: similarly to ridge regression, this approach reduces the dispersion among the empirical eigenvalues of the matrices, leads to better conditioning numbers and even better mean-squared error properties of the resulting estimators. In this talk we give a small overview on this type of shrinkage methods before we consider the particular case of estimating sudden changes in the structure of high-dimensional financial time series: based on a hidden Markov model, we treat regime switching of a vector of asset returns from some large portfolio, switching e.g. from a low volatile market to a state with higher risk. In this context, shrinkage of the state-by-state covariance matrices combined with an EM algorithm for estimation of the transition probabilities, allows to clearly stabilize both the estimators and the filters for reconstruction of the hidden state variables.

Extending the scope of the cube root asymptotics

Myung Seo, LSE (m.seo@lse.ac.uk)

This article extends the scope of cube root asymptotics for M-estimators in two directions: allow weakly dependent observations and criterion functions drifting with the sample size typically due to a bandwidth sequence. For dependent empirical processes that characterize criterions inducing cube root phenomena, maximal inequalities are established to derive the convergence rates and limit laws of the M-estimators. The limit theory is applied not only to extend existing examples, such as the maximum score estimator, nonparametric maximum likelihood density estimator under monotonicity, and least median of squares, toward weakly dependent observations, but also to address some open questions, such as asymptotic properties of the minimum volume predictive region, conditional maximum score estimator for a panel data discrete choice model, and Hough transform estimator with a drifting tuning parameter.

(Joint work with T Otsu)

Thresholding and Beyond in Ecology

Nils Christian Stenseth, University of Oslo (n.c.stenseth@ibv.uio.no)

The talk will focus on threshold modelling for the hare-lynx data, but will also provide a summary of how threshold modelling have been of great help in many other ecological systems.

Nonstationary processes with a threshold

Dag Tjøestheim, University of Bergen (Dag.Tjostheim@math.uib.no)

A review is given on Markov chain recurrence theory in the context of time series models that are both nonlinear and nonstationary. This is applied to threshold unit root processes and to threshold cointegration models.

Threshold Models for Functional Time Series with Applications

Ruey S. Tsay, University of Chicago (Ruey.Tsay@chicagobooth.edu)

We consider a threshold model for the time evolution of functional time series. The goal is to study the time-varying effects of explanatory variables on the functional time series. We propose methods for handling multivariate explanatory variables. For applications, we consider the time evolution of income distribution of Taiwan and daily distribution of stock returns of the U.S. market.

Advances in Shrinkage Methods for Spectral Matrices

Andrew Walden, Imperial College (a.walden@imperial.ac.uk)

A large amount of research has gone into the development of shrinkage methods for real-valued covariance matrices. In spectral analysis of p-vector valued time series there is often a need for good shrinkage methods too, most notably when the complex-valued spectral matrix is singular. The equivalent of the Ledoit-Wolf covariance matrix estimator for spectral matrices can be improved on using a Rao-Blackwell estimator, and we show its form and give an example of the improvement possible. Of course such estimators can be used to better estimate inverse spectral matrices too. Another estimation method is possible for the latter involving projection methods and invertible $L \times L$ matrices ($L < p$). We describe the method and suggest a way of selecting the dimension L using a predictive risk approach.

(Joint work with Deborah Schneider-Luftman)

Estimation of High-dimensional Vector Auto-regressive Processes

Wei Biao Wu, University of Chicago (wbwuchicago@gmail.com)

We will present a systematic theory for high-dimensional linear models with dependent errors and/or dependent covariates. To study properties of estimates of the regression parameters, we adopt the framework of functional dependence measures. For the covariates two schemes are addressed: the random design and the deterministic design. For the former we apply the constrained L1 minimization approach, while for the latter the Lasso estimation procedure is used. We provide a detailed characterization on how the error rates of the estimates depend on the moment conditions that characterize the tail behaviors, the dependencies of the underlying processes that generate the errors and covariates, the dimension and the sample size. Our theory substantially extends earlier ones by allowing dependent and/or heavy-tailed errors and the covariates.

Whittle Likelihood Estimation of Nonlinear Autoregressive Models with Moving Average Errors

Yingcun Xia, Singapore National University (staxyc@nus.edu.sg)

The Whittle likelihood estimation (WLE; Whittle, 1953) played a fundamental role in the development of both theory and computation of time series analysis (Hannan, 1973). However, WLE is only applicable to models whose theoretical spectral density function (SDF) is known up to the parameters in the models. In this paper, we propose a residual based WLE, called XWLE, that can estimate models with their SDFs only partially available, including many popular time series models with correlated residuals. Asymptotic properties of XWLE are established. In particular, XWLE is asymptotically equivalent to WLE in estimating linear ARMA models, and is also capable to estimate nonlinear AR models with MA errors and even with exogenous variables, and thus it has a much wider application than WLE. The finite-sample performances of XWLE are checked by simulated examples and real data analysis.

(Joint work with Tianhao Wang)

Identifying Cointegration by Eigenanalysis

Rongmao Zhang, Zhejiang University, China (rmzhang@zju.edu.cn)

We propose a new and easy-to-use method to identify cointegrated components for a vector time series. The method boils down to a simple eigenanalysis for a positive definite matrix. It is applicable under the setting when each component series is an $I(d)$ process with d being a non-negative integer. Furthermore we allow d being different for different component series. Asymptotic properties of the proposed methods are investigated. Illustration with both simulated and real data sets is also reported.

(Joint work with Peter Robinson and Qiwei Yao)

An Iterative Estimation Procedure for Generalised Varying Coefficient Models with Unspecified Link Functions

Wenyang Zhang, University of York (wenyang.zhang@york.ac.uk)

In this talk, the generalised varying-coefficient models with unspecified link functions will be addressed. A very weak identification condition will be presented for the generalised varying-coefficient models when their link function is unknown. Under the identification condition, I will introduce an iterative estimation procedure for the generalised varying-coefficient models with unspecified link function. An algorithm will also be introduced to implement the proposed estimation procedure. I will also show the asymptotic properties of the nonparametric estimators obtained by the proposed iterative estimation procedure, and some simulation study results. Finally, I will use the generalised varying-coefficient models with unspecified link function to analyse a real dataset.