Recent Advances in Time Series Analysis (RATS 2026)

June 17 - 20, 2026

View from Protaras

Sponsors
Cyprus Academy of Sciences, Letters, and Arts
Technische Universität Braunschweig

Organizers
Jens-Peter Kreiss, Efstathios Paparoditis, and Dimitris N. Politis

Venue: Radisson Beach Resort, Larnaca

Time	Wednesday	Thursday	Friday	Saturday
08:50–09:00	**Welcome**			
09:00–10:00	Gregory Rice A functional ARCH model	George Michailidis High-dimensional matrix autoregressions: A unified framework for structured regularization	Carsten Jentsch Global estimation in tv envir: asympt and bootstrap results Dan Nordman A resampling method for transfor. long-memory times series	Suhasini Subba Rao Graphical models for multivariate non-Gaussian time series – via conditional specifications
10:00–10:30	**Coffee Break**			
10:30–12:30	Paul Doukhan A brief history of weak dependencies Moritz Jirak Weak dependence and optimal quantitative self-normalized ... Zhou Zhou Simult. inference for non-linear time series, a sieve M-regr apprch Georg Köstenberger Sharp oracle inequalities for covariate selection via AIC	Piotr Kokoszka Deep learning estimation of spectral density of fnctnl data... Anne van Delft Characterizing metric space valued process's: separating classes Siegfried Hörmann Testing for Isotropy of function-valued randm fields Hanlin Shang Interpretble models for forecasting high-dim. fnctnl time series	Patrice Bertail Splitting Markov chain in high dims with applications... Lionel Truquet On transfer of mixing rates for Markov chains with appl.... Miles Lopez Testing elliptical & independent component models in high dims Kejin Wu Distributional conformal prediction for Markov processes	Richard Davis Estimating quantile treatments without strict overlap Qiwei Yao Nonlinear independent component analysis for time series Yunyi Zhang Analysis of quadratic forms of high-dimensional non-station... Liudas Giraitis Estimation of partially observed AR(p) time series
12:30–14:00	**Lunch Break**			**Lunch Break**
14:00–15:30	**Time for Personal Discussions**			Daniel Rademacher Mercer expansions in Sobolev spaces and appl's Pavlos Zoubouloglou Breuer-Major thms for Hilbert space valued randm var
15:30–16:00	**Time for Personal Discussions**			**Farewell Coffee**
16:00–17:30	Andreas Anastasiou Spectral mean estimators of time ser: finite smple appr Sumanta Basu Estimation and inference of large sparse spec prec matrices Anne Leucht Gaussian approx'm for lag-window estimators of spectral dens	Haeran Cho Detection & identification of multiple change points in factor models Martin Wendler Change-point tests based on U-statistics Edward Eriksson The unseen species problem revisited	**Free Afternoom**	
17:30–18:00	**Coffee Break**			
18:00–19:30	Michael Neumann Nonparametric symmetry tests for int-valued time series Anna Bykhovskaya Generalized AR mult models: From binary to poisson Alexander Lindner Divergnce of sums of i.i.d. rand var with exponential weights	Christian Francq Dynamic CAPM with long memory factors Han Xiao Dynamic matrix factor model for counts data Likai Chen Valid inference for factor models with partially missing returns	**Free Afternoon**	
19:30			**Departure Conference Dinner**	

Workshop on Recent Advances in Time Series Analysis (RATS 2026) — Program Overview

Abstracts - Keynote Talks

A functional ARCH model

Gregory Rice
University of Waterloo, Canada

AutoRegressive Conditional Heteroscedasticity (ARCH) models are standard for modeling time series exhibiting volatility, with a rich literature in univariate and multivariate settings. In recent years, these models have been extended to function spaces. However, functional ARCH and generalized ARCH (GARCH) processes established in the literature have thus far been restricted to model "pointwise" variances. In this paper, we propose a new ARCH framework for data residing in general separable Hilbert spaces that accounts for the full evolution of the conditional covariance operator. We define a general operator-level ARCH model. For a simplified Constant Conditional Correlation version of the model, we establish conditions under which such models admit strictly and weakly stationary solutions, finite moments, and weak serial dependence. Additionally, we derive consistent Yule-Walker-type estimators of the infinite-dimensional model parameters. The practical relevance of the model is illustrated through simulations and a data application to high-frequency cumulative intraday returns.

High-Dimensional Matrix Autoregressions: A Unified Framework for Structured Regularization

George Michailidis
University of California, Los Angeles

While vector autoregressions (VARs) effectively capture lead-lag relationships among time series, they suffer from severe parameter proliferation in high-dimensional settings. This issue is particularly acute for matrix-variate data tracking multiple variables across numerous entities. Naively vectorizing these inherently two-way data obscures their underlying structural relationships, rendering estimation both statistically and computationally inefficient. The bilinear Matrix Autoregressive (MAR) model addresses this by preserving the intrinsic matrix structure, thereby substantially reducing the effective number of parameters. Imposing additional structural constraints, such as sparsity or low-rankness, further enhances parsimony but leads to a nonconvex optimization problem with significant computational and statistical challenges. To address these issues, we develop a novel two-stage estimation procedure for regularized MAR models. The method first constructs high-quality initial estimators and subsequently refines them via an alternating least squares algorithm. We establish rigorous theoretical guarantees for the resulting estimators under high-dimensional scaling, demonstrating that the final error decomposes into computational and statistical components. Notably, the statistical error matches the

performance of a global-minimizing oracle. Finally, the empirical performance of the proposed method is illustrated through extensive simulation studies and an application to macroeconomic indicators from Eurozone economies.

Graphical models for multivariate non-Gaussian time series – via conditional specifications

Suhasini SubbaRao
Texas A&M University, USA

Graphical models are ubiquitous for summarizing conditional relations in multivariate data. In many applications involving multivariate time series, it is of interest to learn an interaction graph that treats each individual time series as nodes of the graph, with the presence of an edge between two nodes signifying conditional dependence given the others. Typically, the partial covariance is used as a measure of conditional dependence. However, in many applications, the outcomes may not be Gaussian and/or could be a mixture of different outcomes.

In this talk, we propose a broad class of time series models for multivariate mixed-type time series, that includes the classical VAR model as a special case. For each node in the time series we model its conditional distribution with a distribution from the exponential family. We call this construction a conditionally exponential stationary graphical model CEStGM. This modelling approach has several potential advantages. The first is that due to the versatility of the exponential class, it allows one to stitch together several different variable "types". The second is that the univariate conditional specification allows for simple estimation procedures using standard GLM tools. Finally, the conditional specification is specifically designed to succinctly encode process-wide conditional independence in its parameters.

We derive conditions that ensure the model leads to a well defined strictly stationary time series and show that the model is geometrically beta-mixing. We propose an approximate Gibbs sampler for simulating sample paths from CEStGM. Finally, we conclude the talk with some numerical experiments and real data examples.

Abstracts - Talks

Spectral Mean Estimators of Time Series: Finite Sample Distributional Approximations

Andreas Anastasiou
University of Cyprus, Nicosia

In this talk, we consider general spectral means, based on a finite signed spectral Borel measure, and derive finite sample bounds for the Wasserstein distance of the distributions of the corresponding sample estimators (integrated periodograms) to their Gaussian limits. Our results also apply to absolutely continuous and discretized versions of spectral means, based on a function φ. The bounds obtained are given explicitly, they are computable and are applicable under non-restrictive conditions regarding Fourier coefficients, and for a broad class of short-range dependent processes. For n the sample size, the rates of convergence obtained are the closer to the optimal $\sqrt{n}$ rate the better the time series can be approximated by an m-dependent process and the smoother the function φ is. We illustrate our findings through some examples.
Joint work with Efstathios Paparoditis and Tobias Kley.

Estimation and inference of large sparse spectral precision matrices

Sumanta Basu
Cornell University

Graphical models offer a powerful framework to capture intertemporal and contemporaneous relationships among the components of a multivariate time series. For stationary time series, these relationships are encoded in the multivariate spectral precision matrix. In the first part of this talk, we will present adaptive penalization methods for estimation of these objects under suitable sparsity assumptions. We will discuss new optimization algorithms for Complex Lasso and Graphical Lasso (CLASSO and CGLASSO), and investigate estimation consistency under a double-asymptotic regime where the dimension of the time series increases with sample size. In the second part of the talk, we will introduce a debiased version of CGLASSO to conduct statistical inference on individual off-diagonal entries of the spectral precision matrix, and discuss its asymptotic properties.

Splitting Markov chain in high dimensions with applications to estimation and bootstrap

Patrice Bertail
University of Nanterre, Paris

Large-dimensional Markov chains appear in many models and many applications, including EEG or multicompartment models in food risk assessment which motivated this work. In this talk, we describe the Multidimensional Approximative Regenerative (MAR) algorithm designed for high-dimensional Markov chains and the associated Block Bootstrap method (MARBB). We focus on a vector autoregressive (VAR(1)) process with a low-rank structure, but the same ideas have been developed for multi(single) index non-linear models. We first use a reduction algorithm that transforms the original high-dimensional time series into a lower-dimensional Markov chain. Once the chain is reduced, we use the regenerative properties of Harris recurrent Markov chains to extend existing results from the low-dimensional settings to the high-multidimensional case. This approach enables the identification of approximate regeneration times of the lower-dimensional Markov chain. In turn, they lead to the splitting of the original high-dimensional Markov chain into (approximate) regenerative blocks. These blocks are then used to construct relevant statistics of interest and to a bootstrap method using (approximate) regenerative blocks. We give MARBB consistency results and apply our algorithm to simulated data.
Joint work with A. Dudek and K. Marek.

Generalized Autoregressive Multivariate Models: From Binary to Poisson

Anna Bykhovskaya
Duke University

We present a framework for binary autoregressive time series in which each observation is a Bernoulli variable whose success probability evolves with past outcomes and probabilities, in the spirit of GARCH-type dynamics, accommodating nonlinearities, network interactions, and cross-sectional dependence in the multivariate case. Existence and uniqueness of a stationary solution is established via a coupling argument tailored to the discontinuities inherent in binary data. A key theoretical result, further supported by our empirical illustration on S&P 100 data, shows that, under a rare-events scaling, aggregates of such binary processes converge to a Poisson autoregression, providing a micro-foundation for this widely used count model. Maximum likelihood estimation is proposed and illustrated empirically.
Link to the paper: https://arxiv.org/pdf/2604.14394 .

Valid Inference for Factor Models with Partially Missing Returns

Likai Chen
Washington University, St. Louis

We propose a new methodology for valid coefficient inference in financial factor models when the response panel of returns is partially missing. The leading application is alpha inference in a Fama-French factor regression, where a fund's or stock's risk-adjusted return α_i is the parameter of interest but returns are missing due to fund closure, delisting, sparse trading, or asynchronous markets. The framework combines a static factor backbone (estimated via PCA on observed returns) with an optional diffusion-based residual correction to construct a synthetic prediction $S_{i,t+1}$ for every unit, observed or missing. The prediction is never substituted for the response; instead, it is residualized against the factor-model design W_{it} on the full panel and enters a labeled-sample regression as an orthogonalized auxiliary control. We establish unbiasedness and asymptotic normality for $\hat{\textbf{beta}}$ (and therefore $\hat{\alpha}_i$), derive the efficiency gain $\text{ARE} = 1/(1-\pi_T\rho^2)$ where π_T is the missing fraction and ρ is the partial correlation between target and surrogate, and show through Monte Carlo simulation that the gain is realized only when the surrogate uses information beyond W.

Detection and identification of multiple change points in tensor factor models

Haeran Cho
University of Bristol

We study the problems arising from modeling high-dimensional tensor-valued time series under a Tucker decomposition-based factor model with multiple structural change points. First, we propose TFMseg, an algorithm for detecting the multiple change points, which utilizes the low-rank structure of the data for statistical and computational efficiency. Noting the unique challenge posed by the multi-dimensional array setting under which the changes in different modes interact, we also investigate the problem of identifying each change with different modes post-segmentation. To this end, we formalize the mode-identifiability of each change and propose an algorithm for detecting the modes at which the data are undergoing a mode-identifiable shift. We establish the consistency of both change point detection and mode-identification methods under a weak moment condition, and demonstrate their good performance on simulated datasets where, in particular, it is shown that the mode-identification step can improve the post-segmentation estimation of the mode-wise loading space. Additionally we analyze the datasets on New York City taxi usage and Fama-French portfolio returns using the proposed suite of methods.

Estimating quantile treatments without strict overlap

Richard A. Davis
Columbia University, New York

We consider the problem of estimating quantile treatment effects without assuming strict overlap , i.e., we do not assume that the propensity score is bounded away from zero. More specifically, we consider an inverse probability weighting (IPW) approach for estimating quantiles in the potential outcomes framework and pay special attention to scenarios where the propensity scores can tend to zero as a regularly varying function. Our approach effectively considers a heavy-tailed objective function for estimating the quantile process. We introduce a truncated IPW estimator that is shown to outperform the standard quantile IPW estimator when strict overlap does not hold. We show that the limiting distribution of the estimated quantile process follows a stable distribution and converges at the rate $n^{1-1/\gamma}$, where $\gamma > 1$ is the tail index of the propensity scores when they tend to zero. We also propose a new mean stabilized IPW estimator that uses truncated propensity scores as an alternative to the stabilized IPW. This allows for a trade-off between bias and variance that leads to an improved mean squared error, especially in the very heavy-tailed regimes. The limit distribution in this case is infinitely divisible. We also discuss a novel data-driven method for estimating an optimal truncation level. The performance of our estimators in numerical experiments and in a dataset that exhibits the presence of extreme propensity scores.
(This is joint work with Marco Avella Medina and Gennady Samorodnitsky.)

Characterizing metric-space-valued processes: separating classes and weak invariance principles for measure-theoretic inference

Anne van Delft
Columbia University, New York

This article investigates stochastic processes taking values in metric spaces that lack a topological vector space structure. Analyzing such processes is often complex, as it necessitates accounting for the interplay between the topological, geometric, and temporal dependence structures. Tractable separating classes are established based on the structural geometry of the underlying space. While traditional kernel methods suffer from geometric distortion when the underlying space cannot be isometrically embedded into a Hilbert space, we bypass these rigid constraints by exploiting a deeper topological property. This foundation provides a basis for the subsequently introduced measure-theoretic inference methodology.

A brief history of weak dependencies

Paul Doukhan
University Cergy-Pontoise, Paris

More than a purely formal exposition, the aim is to demonstrate the advantages and disadvantages of most of the most commonly used concepts, depending on the intended purpose and the types of models involved. Recently, Wu proposed asymptotic results that improve upon virtually everything imaginable, with speeds corresponding to those known to be optimal in the independent framework. The price to pay is the representation of such models, necessarily discrete-time, as functions of a sequence of independent variables, for which coefficients and couplings, technically somewhat unnatural, can be determined. I will attempt to highlight all the advantages of these techniques, as well as developments for models used in statistical practice.

The Unseen Species Problem Revisited

Edward Eriksson
Max Planck Institute for Mathematics in the Sciences, Leipzig

Imagine an ecologist who wants to predict how many new species a return expedition to a location would discover. Formally, given n i.i.d. samples from an unknown distribution over an unknown set, how many new outcomes will be observed if we collect m more samples? We consider also a version with set-values samples. This generalized version is equivalent to forecasting the growth of the vertex set in an edge exchangeable (hyper-) graph. The main results are (1) a new estimator which in the classical problem is, up to an explicit multiplicative constant, optimal in the class of estimators respecting a natural problem symmetry and (2) prediction guarantees for a regular-variation-adapted estimator when the sets are of bounded size.

Dynamic CAPM with long memory factors

Christian Francq
University of Lille

Factor models are widely used in finance, and recent advances allow slope coefficients (known as 'betas') to vary over time. However, estimation theory for these dynamic conditional betas typically relies on short-memory volatility models, which can be restrictive in empirical applications. Moreover, exogenous variables such as realized moments have proven useful in recent volatility modeling studies. In this paper, we introduce

a multivariate framework that allows for time-varying betas where covolatilities can exhibit higher persistence than standard exponential decay. We incorporate covariates in the dynamics of both conditional variances and betas. We establish stationarity conditions for the proposed model, prove the consistency and asymptotic normality of the quasi-maximum likelihood (QML) estimator, and propose goodness-of-fit tests. Monte Carlo experiments assess the finite-sample performance of the estimation procedure. Finally, we discuss the choice of relevant exogenous variables and illustrate the model's effectiveness through real data applications.
Joint work with Julien Royer and Jean-Michel Zakoian.

Estimation of Partially Observed AR(p) Time Series

Liudas Giraitis
London School of Economics

This paper builds on methodology that extends the standard estimation of stationary autoregressive (AR) time series to settings with unevenly spaced observations and missing data, providing appropriately corrected estimates of autoregressive model parameters. Using these methods, we develop time-series tools for the estimation and forecasting of autoregressive models with cyclically varying parameters, where the periodicity is assumed to be known.
The proposed approach helps uncover cyclical patterns in the data and is illustrated using the Central England Temperature (CET) time series (1772 - present) and daily stock market data.

Testing for Isotropy of Function-Valued Random Fields

Siegfried Hörmann
University of Technology, Graz

This project is motivated by the analysis of the morphology of paper-based materials. Paper consists of fibers packed closely together. We consider the paper's surface as a two-dimensional domain and the z-axis as the position within the thickness of the paper.
For each spatial location (x, y) on the paper, we have functional measurements

$$D(z|x, y),$$

describing the paper's density at depth $z \in [0, 1]$, where 0 and 1 correspond to positions on the two opposite surfaces of the paper. These measurements are collected at very high spatial resolution over the (x, y)-domain.

An important question for materials scientists is whether the internal structure of the paper is homogenous and invariant under rotations. For example, one aims to determine whether the rolling process induces a preferred orientation of the fibers.
Statistically speaking, we aim to figure out if the corresponding functional random field is isotropic.
We propose a testing approach, where the functional data are transformed into spherical data, and we reduce the problem to that of testing uniformity on the sphere. Because the observations exhibit spatial dependence and are basically sampled over a continuum, existing methods from directional statistics must be appropriately adapted to this setting.
This is joint work with Matthias Neumann and Julius Baumhakel.

Global estimation in time-varying environments: some asymptotic and bootstrap (in)consistency results

Carsten Jentsch
TU Dortmund

While most statistical methods for time series analysis typically rely on stationarity assumptions, these are often not fulfilled exactly in applications. That is, the mean, the autocovariance structure or even the whole distribution may change slightly as time evolves. However, the fact that stationarity does typically not hold exactly, but only approximately, is commonly ignored in practice. In this paper, we consider the (global) statistics of the sample mean and the sample autocovariances in time-varying (i.e. locally stationary) environments and derive asymptotic theory. While, as expected, the sample mean is converging to the average (integrated) mean of the process, the situation becomes more delicate for the sample autocovariance. Moreover, we discuss (in)consistency of a seemingly appropriate component wild block bootstrap in several settings.
Joint work with Matei Demetrescu and Christoph Hanck.

Weak dependence and optimal quantitative self-normalized central limit theorems

Moritz Jirak
University of Vienna

Consider a stationary, weakly dependent sequence of random variables. Given only mild conditions, allowing for polynomial decay of the autocovariance function, we show (nonuniform) Berry-Esseen bounds of optimal order $n^{-1/2}$ for studentized (self-normalized) partial sums, both for the Kolmogorov and Wasserstein (and L^p) distance. The results show that, in general, (minimax) optimal estimators of the long-run variance lead to suboptimal bounds in the central limit theorem, that is, the rate $n^{-1/2}$ cannot be reached. This can be salvaged by simple methods: In order to maintain the optimal speed of convergence $n^{-1/2}$, simple over-smoothing within a certain range is necessary and sufficient. The setup contains many

prominent dynamical systems and time series models, including random walks on the general linear group, products of positive random matrices, functionals of Garch models of any order, functionals of dynamical systems arising from SDEs, Glauber dynamics, iterated random functions and many more.

Sharp oracle inequalities for covariate selection via the AIC

Georg Köstenberger
University of Vienna

Prediction is a core task in time series analysis, and the choice of the forecasting model is of fundamental importance. To address this problem (among other things), Akaike introduced the AIC and FPE, and demonstrated their significant usefulness for prediction in two landmark papers. In subsequent seminal works, Shibata developed a notion of asymptotic efficiency and showed that both AIC and FPE are optimal, setting the stage for decades-long developments and research in this area and beyond.
Most of the literature on the usage of AIC for prediction focuses on the case of nested models. However, there is no fundamental information theoretic reason for this restriction, as the AIC is essentially an optimal estimator of the prediction error, which is based on minimizing the empirical Kullback-Leibler divergence between any two models. This point of view suggests that the AIC (and its variants) should be able to select an (optimal) forecasting model from any set of candidate models. In this work, we establish sharp, finite-sample oracle inequalities for the AIC in the non-nested case, subject only to a very general notion of weak dependence. This establishes a unversality property of the AIC, in the sense that it can not just be used to compare between nested models, but rather between arbitrary models. Our framework contains many prominent dynamical systems such as random walks on the regular group, functionals of iterated random systems, functionals of (augmented) Garch models of any order, functionals of (Banach space valued) linear processes, possibly infinite memory Markov chains, dynamical systems arising from SDEs, and many more. The proofs require a new notion of uniform integrability in the Wiener algebra, a new uniform version of the Wiener-Levy theorem and a uniform Baxter-type inequality, which may be of independent interest.

Deep learning estimation of the spectral density of functional time series on large domains

Piotr Kokoszka
Colorado State University

The talk is concerned with the estimation of the spectral density of a time series of surfaces or images defined on large grids of hundreds of thousands of points. Such sequences occur in biomedical and climate studies. The estimator is the output of a multilayer perceptron neural network. Existing estimators use sample autocovariance kernels represented as high-dimensional matrices whose manipulation is computationally demanding or even infeasible. We use the theory of spectral functional principal components to derive our deep learning estimator and prove that it is a universal approximator to the spectral density under general assumptions. It can be trained without computing the autocovariance kernels and can be parallelized to provide the estimates much faster than existing approaches. We validate its performance by simulations and an application to fMRI images.

Gaussian approximation for lag-window estimators of spectral densities

Anne Leucht
University of Bamberg

We develop a Gaussian approximation for the maximum of a suitably normalized lag-window estimator of the spectral density evaluated at all positive Fourier frequencies. This max-statistic plays an essential role to construct goodness-of-fit tests for the second-order structure of stationary time series and to derive simultaneous confidence bands for the spectral density. The Gaussian approximation provides the opportunity to verify asymptotic validity of a multiplier bootstrap procedure and, even further, to derive the corresponding rate of convergence. A small simulation study illustrates the properties of this bootstrap proposal for finite samples.

Divergence of sums of i.i.d. random variables with exponential weights

Alexander Lindner
Ulm University

It is well known that a random walk having finite expectation diverges almost surely to $+\infty$ if and only if the expecation of the increments is strictly positive, while it oscillates when the expectation is 0. Less well known is the divergence behaviour of random walks when the increments have infinite expectation. In 1973, Erickson obtained an integral criterion characterising when the corresponding random walk diverges to $+\infty$,

$-\infty$, or when it oscillates. In this talk we are interested in the divergence behaviour of $W_n = \sum_{k=1}^{n} c^k X_k$, where (X_k) is i.i.d. and $c \in (0,1)$ is a given constant. Such a sequence arises when expressing an AR(1) process in terms of its infinite past. It is well known that this sum converges almost surely if and only if X_1 has finite log-moment, but we are interested in the divergence behaviour when it has infinite log-moment. We give sufficient analytic conditions for W_n to exhibit an oscillating behaviour and sufficient conditions for it to diverge almost surely to $+\infty$. The talk is based on joint work in progress with Ross Maller and Maximilian Strobel.

Testing elliptical and independent component models in high dimensions

Miles Lopez
University of California, Davis

Due to the broad applications of elliptical models, there is a long line of research on goodness-of-fit tests for empirically validating them. However, the existing literature on this topic is generally confined to low-dimensional settings, and to the best of our knowledge, there are no established goodness-of-fit tests for elliptical models that are supported by theoretical guarantees in high dimensions. In this work, we propose a new goodness-of-fit test for this problem, and our main result shows that the test is asymptotically valid when the dimension and sample size diverge proportionally. Remarkably, it also turns out that the asymptotic validity of the test requires no assumptions on the population covariance matrix. In addition, we will discuss some more recent work extending these ideas to a goodness-of-fit test for independent component models in high dimensions. (A subset of this work will appear in JASA: https://doi.org/10.1080/01621459.2025.2518617. Joint work with Siyao Wang and Mingshuo Liu.)

Nonparametric symmetry tests for integer-valued time series

Michael H. Neumann
Friedrich-Schiller-Universität Jena

During the last years, there have been many proposals for modeling integer-valued time series. We propose tests of hypotheses related to certain symmetry and antisymmetry properties. For example, we consider the hypotheses that the conditional mean is an odd function or that the conditional variance is an even function. The proposed test statistics are nonparametric and have non-standard limit distributions. We show that the wild bootstrap offers a simple method of generating asymptotically correct critical values. The talk is based on joint work with Paul Doukhan and Christian Weiß.

A Resampling Method for Transformed Long-Memory Times Series

Dan Nordman
Iowa State University

Long-range dependent time series are often conceptualized by an unknown transformation of an underlying long-memory Gaussian process. The so-called Hermite rank of this transformation is a process parameter that critically impacts statistical inference, as sampling distributions change with the unknown rank. A compounding issue is that ranks can further vary between statistics computed from the same time series. Over the past 50 years, no approach has existed to generally approximate Hermite ranks from data. This talk describes a method for approximating both the Hermite rank as well as dependence parameter of the underlying Gaussian process, without knowledge of the underlying transformation that defines the observed long-memory time series. As a result, the estimation approach can then be coupled with a bootstrap method for approximating the sampling distribution of statistics in practice. The inference method is illustrated through numerical studies and examples.

Mercer Expansions in Sobolev Spaces and Applications to Stochastic Processes

Daniel Rademacher
University of Technology, Graz

We establish a fundamental extension of Mercer's celebrated theorem by introducing a novel class of higher-order kernel operators acting on Sobolev spaces $H^k(\Theta)$, where $\Theta \subset \mathbb{R}^d$ is a bounded domain and $k \in \mathbb{N}_0$ corresponds to the order of weak differentiability. These operators transcend the classical L^2-framework by explicitly incorporating the information encoded in the (weak) derivatives of the kernel. We prove that this class corresponds precisely to the Hilbert-Schmidt mappings between Sobolev spaces, thereby providing a complete operator-theoretic characterization. The spectral decomposition of these operators then yields Mercer-type expansions that are optimal in $H^k(\Theta \times \Theta)$. Notably, we show that for $k > d$, these expansions also converge uniformly without requiring the kernel to be positive definite - a result previously unknown in the literature for arbitrary symmetric kernels. For positive definite kernels, we confirm the nuclearity of these higher-order operators and establish a significant refinement of Mercer's Theorem, which ensures uniform convergence of all term-wise derivatives and provides explicit convergence rates (including derivatives) tied to the spectral decay. These results lead to novel spectral representations of Reproducing Kernel Hilbert Spaces and have subtle implications for stochastic analysis. Applied to the covariance kernels of weakly differentiable random fields, our theory provides refined Karhunen-Loève expansions that facilitate the simultaneous mean-square optimal approximation of both the process and its derivatives.

Interpretable models for forecasting high-dimensional functional time series

Hanlin Shang
Macquarie University, Sydney

We study the modeling and forecasting of high-dimensional functional time series, which can be temporally dependent and cross-sectionally correlated. We implement a functional analysis of variance (FANOVA) to decompose high-dimensional functional time series, such as subnational age- and sex-specific mortality observed over years, into two distinct components: a deterministic mean structure and a time-varying residual process. Unlike purely statistical dimensionality-reduction techniques, the FANOVA decomposition provides a direct and interpretable framework by partitioning the series into effects attributable to data-specific factors, such as regional and sex-level variations, and a grand functional mean. From the residual process, we implement a functional factor model to capture the remaining stochastic trends. By combining the forecasts of the residual component with the estimated deterministic structure, we obtain the forecasted curves for high-dimensional functional time series. Illustrated by the age-specific Japanese subnational mortality rates from 1975 to 2023, we evaluate and compare the accuracy of the point and interval forecasts across various forecast horizons. The results demonstrate that leveraging these interpretable components not only clarifies the underlying drivers of the data but also improves forecast accuracy, providing more transparent insights for evidence-based policy decisions.

On the Transfer of Mixing Rates for Markov Chains in Random Environments, with Applications to Time Series Analysis and Machine Learning

Lionel Truquet
University Rennes, ENSAI-CREST

For the analysis of time series or the study of the convergence of certain stochastic algorithms, such as those based on Langevin dynamics, discrete-time Markov processes are naturally perturbed by the introduction of stochastic disorder, which may be referred to as noise, an exogenous factor, or a random environment depending on the terminology used across different fields. The aim of this presentation is to show that one can study the stability of these Markov chains in random environments on unbounded spaces using adapted versions of more classical stability conditions (drift conditions, small set conditions, or contraction of random mappings). Sufficient conditions for the existence of stationary and ergodic measures, as well as control of the alpha-mixing coefficients for these processes, can then be obtained. From this, one can derive a limit theory for autoregressive processes with exogenous covariates, quantify mixing rates for the study of deep learning algorithms, and obtain qualitative estimates of the convergence rates of certain stochastic algorithms.

Change-point tests based on U-statistics: First versus full approach and first versus last approach compared

Martin Wendler
Otto-von-Guericke-Universitẗ Magdeburg

Many authors use U-statistics to generalize the CUSUM change-point statistic to changes beyond the mean. We compare two approaches of constructing CUSUM-type change-point tests, which we call the first-vs-full and first-vs-last approach. The question naturally arises if the two tests substantially differ and, if so, which of them is better in which data situation. In large samples, both tests are similar: under some general assumption on the time series, they are asymptotically equivalent under the null hypothesis and under sequences of local alternatives. In small samples, there may be quite noticeable differences, which is in line with a different asymptotic behavior under fixed alternatives. We examine the examples Gini's mean difference, the sample variance, and Kendall's tau.

Distributional Conformal Prediction for Markov Processes

Kejin Wu
Loyola University Chicago

We introduce the Markov Distributional Conformal Prediction (MDCP) method that extends the distributional conformal prediction (previously developed for regression) to the setting of a strictly stationary Markov process. Instead of relying on a specific model structure to do prediction, the idea of distributional conformal prediction interval aligns with the Model-Free (MF) Prediction Principle. In analogy to MF prediction of Markov processes, our method exploits the probability integral transform based on estimated transition distribution functions to transform the Markov data to an i.i.d. dataset. We show a non-asymptotic error bound of MDCPs unconditional coverage rate under a beta-mixing condition and other standard assumptions on the kernel estimators. The asymptotic validity of the conditional prediction interval is also verified. In addition, we show that our conditional prediction interval is still asymptotically valid with Markov processes being $L^p - m$-approximable instead of satisfying the mixing property. Numerical simulations and real data experiments are deployed to empirically illustrate the finite-sample performance of MDCP, and compare it with the MF bootstrap.

Dynamic matrix factor model for counts data

Han Xiao
Rutgers University

A dynamic factor model is proposed for matrix time series, where the observations are counts. The model is formulated as Poisson observations conditional on the rate matrices, which have log-normal distributions. The logarithm of the rate matrices has the form of a dynamic Gaussian factor model of matrix time series, where the dynamics are captured by the factor process. We use a log moment method to estimate the loading matrices, and the variational inference to estimate the factor process. An autoregressive model is imposed on the factor process to enable predictions. It is also considered to include a trend component in the log rate matrices to capture the possible nonstationarity. Theoretical and numerical analyses are conducted for the proposed model.

Nonlinear Independent Component Analysis for Time Series

Han Yan and Qiwei Yao
London School of Economics

Independent component analysis (ICA) is a fundamental tool for dimension reduction and recovering latent independent signals from multivariate observations with broad applications across various fields. Most existing ICA methods, however, are designed for linear mixtures and independently observed data, and are therefore inadequate for nonlinearly mixed data or the data exhibiting temporal dependence. To fill in the gap, we propose a nonlinear ICA for multivariate time series such that the transformed time series are independent across all time lags. We use an energy-distance to measure independence and employ invertible neural networks (INN) with coupling flows to represent nonlinear and invertible transformations. We establish theoretical guarantees by characterizing the approximation capacity and function class complexity of the INN class and deriving a convergence rate for the nonlinear ICA estimator. Numerical experiments demonstrate that the proposed method performs well in decomposing multivariate time series into mutually independent component series.

Analysis of quadratic forms of high-dimensional non-stationary time series, with application to ANOVA and independent testing

Yunyi Zhang
Chinese University of Hong Kong

Quadratic forms of time series are common in statistical applications, making it worthwhile to study these structures. This talk introduces theoretical results, including concentration inequality, Gaussian approximation theorem, and consistent variance estimator, for a quadratic form of high-dimensional non-stationary time series. Building on these results, we establish an ANOVA procedure and a hypothesis testing method for independence of two time series. We further develop distributional results of the test statistics, and propose dependent wild bootstrap algorithms to facilitate hypothesis testing through Monte-Carlo simulations. Numerical studies and real-life data applications demonstrate the good performance of the proposed test statistics.

Simultaneous Inference for Nonlinear Time Series, a Sieve M-regression Approach

Zhou Zhou
University of Toronto

This talk studies simultaneous inference of conditional distributions in nonlinear time series from a sieve M-regression perspective. Existing literature on sieve M-regression has primarily focused on pointwise asymptotics, leaving the development of uncertainty quantification over the entire predictor space unexplored. We address this gap by establishing a uniform Bahadur representation for the sieve M-estimator, accommodating dependent data and a growing number of sieve basis functions. A novel high-dimensional empirical process theory is developed for temporally dependent data, and a specifically designed M-decomposition method is utilized to control high-dimensional complexities. Building on this representation, we develop a convex Gaussian approximation to characterize the asymptotic behavior of the estimator and construct valid simultaneous confidence regions (SCRs). To facilitate practical implementation, we introduce a self-convolved bootstrap algorithm that accurately approximates the distribution of the maximal deviation. Our inferential framework is supported by rigorous error bounds and validated through numerical simulations and real data applications.

Breuer-Major Theorems for Hilbert space valued random variables

Pavlos Zoubouloglou
University of Münster

Breuer–Major theorems are central limit theorems for nonlinear transformations of stationary Gaussian processes. They describe when sums of dependent random variables behave, after normalization, like a Gaussian random variable. In this talk, I will discuss an infinite-dimensional version of this theory. Let $\{X_k\}_{k\in\mathbb{Z}}$ be a stationary Gaussian process taking values in a separable Hilbert space $\mathcal{H}_1$, and let $G : \mathcal{H}_1 \to \mathcal{H}_2$ be a possibly nonlinear transformation. We study the asymptotic behavior of the normalized partial sums

$$\sum_{k=1}^{n} G(X_k).$$

Under two different sets of sufficient assumptions, involving both the transformation G and the dependence structure of the process $\{X_k\}$, we prove a central limit theorem for these sums. The proof uses tools from the infinite-dimensional Malliavin–Stein method, which provides a way to prove Gaussian limits together with quantitative error bounds. I will also discuss several examples, including applications to statistics arising in functional data analysis.

Abstracts - Posters

Learning Graphical Models for Nonstationary Multivariate Time Series

Noah Bolanos
Texas A&M University

NonStGM is a general nonparametric graphical modeling framework for studying associations among the components of a nonstationary multivariate time series. The proposed framework captures conditional correlations in the form of an undirected graph. In addition, to describe the more nuanced nonstationary relationships among the components of the time series, we incorporate within the graph architecture the notion of conditional nonstationarity. This allows one to distinguish between direct and indirect nonstationary relationships among system components, and can be used to search for small subnetworks that serve as the "source" of nonstationarity in the system. In this work, we describe the process to estimate the graph from data. We study the sampling properties of the estimator, and describe the complete data adaptive method for selecting tuning parameters. We illustrate the method with simulations and real data.
Joint work with Suhasini Subba Rao and Jonas Krampe.

Bootstrap convergence rates for the maximum of an increasing number of autocovariances and autocorrelations under strict stationarity

Alexander Braumann
TU Braunschweig

We consider maximum deviations of sample autocovariances and autocorrelations from their theoretical counterparts. The maximum is taken over a number of lags that increases with the number of observations. The asymptotic distribution of such statistics assuming strict stationary of the underlying time series is of Gumbel type. It is well known that speed of convergence of maximum statistics to the Gumbel distribution is rather slow. This also holds for the well-known autoregressive (AR) sieve bootstrap, which is asymptotically valid in this case but suffers from the same slow convergence rate (Braumann et al. 2021). We use Gaussian approximation for high-dimensional time series to show that for the class of strictly stationary processes a wild-type bootstrap is asymptotically valid and that the Kolmogorov distance between the bootstrap statistic and the statistic of interest vanishes at a polynomial convergence rate.
Joint work with Jens-Peter Kreiss and Marco Meyer.

Hybrid Autoregressive-Sieve Bootstrap

Linus Dowidat
TU Braunschweig

The classical autoregressive (AR) sieve bootstrap for sample autocovariances may fail because the limiting distribution depends on the fourth-order moment structure of the underlying process. We propose an AR-sieve bootstrap for general stationary time series that generates the innovation process by sampling from a normal distribution and show that the second-order part of the limiting distribution is preserved while the fourth-order term vanishes. In a second step, we combine this procedure with convolved subsampling, resulting in a consistent hybrid AR-sieve bootstrap.

Detecting moments of significant change in nonstationary time series of random geometric objects

William Knauth
Columbia University

In this article, we investigate the problem of detecting and estimating location of abrupt changes or anomalies in the marginal measure of Polish-valued stochastic processes that are naturally subject to changing dynamics. The approach, which consists of a sliding window of moving partial sum statistics, leverages weak invariance principles introduced in Van Delft 2026. To account for the complicated limiting behavior, which fundamentally differs under stationary and nonstationary regimes, a suitable wild dependent bootstrap fclt is derived under mild assumptions. We not only show that our method is consistent for estimating both the number and location of breaks, but we obtain explicit rates providing statistical guarantees for the researcher. Finally, we illustrate the flexibility and evaluate the performance of our approach in a variety of simulated environments, ranging from problems $\mathbb{R}^k$ to point clouds sampled from random sequences of shapes.

Statistical Inference for CEStGM

Jaeseon Lee
Texas A&M University

Conditionally exponential stationary graphical models (CEStGM) is a general and flexible framework for modelling mixed-type multivariate time series in which each univariate conditional distribution belongs to an exponential family. Under suitable compatibility conditions, these conditional specifications define a stationary process, and their coefficients encode a process-wide undirected conditional independence graph.
In this poster, we consider inferential aspects of CEStGM. In particular we focus on parameter estimation of CEStGM and conditional graph

estimation from data. Since a full joint likelihood is typically infeasible to construct, we propose a conditional partial likelihood estimator based on tractable univariate conditional distributions. The estimated parameters are then used to recover neighborhood sets in the process-wide conditional independence graph.
We also establish large sample properties of the proposed estimator. In particular, we show root-n consistency and derive the asymptotic normality. Since the estimator is based on a partial likelihood rather than the full likelihood, its limiting variance has a sandwich form, which needs to be estimated. We propose a hybrid Gibbs-sampler parametric bootstrap scheme to estimate the asymptotic variance and to conduct graphical inference. Further, we propose a simple diagnostic method based on sample correlations of the conditional Pearson residuals for checking for model adequacy. We illustrate our method with simulations.

Augmented Dynamic Regression

Yufei Li
King's Business School London

The recent work on regression modeling that permits general heterogeneity is extended to allow for lagged dependent variables. The purpose is to explore to what extent the generality of the setting, the simplicity of assumptions, and the ease of computation of standard errors can be preserved. Theoretical properties of regression estimation and inference are accompanied by Monte Carlo experiments and an empirical application.

Detecting Periodicity of a General Stationary Time Series via AR(2)-Model Fitting

Panagiotis Maouris
University of Cyprus, Nicosia

Estimating the periodicity of a stationary time series via fitting a second order stationary autoregressive (AR(2)) model has been initiated by the seminal paper of Yule (1927). We investigate properties of this procedure when applied to a general stationary processes possessing a spectral density with a dominant peak at some unknown frequency $\lambda_0 \in (0, \pi)$. For this, a general class of stationary processes is considered with spectral densities having an arbitrary sharp peak. It is shown within this class, that if the peak of the spectral density is strong enough (in a sense to be specified), then the AR(2) model, which best (in mean square sense) approximates the underlying process, correctly identifies the frequency λ_0. To investigate consistency properties of the corresponding AR(2) based estimator of λ_0, a near to pole asymptotic framework is adopted. Triangular arrays of stationary stochastic processes are considered that possess a spectral density the peak of which at frequency λ_0 becomes more pronounced as the sample size n of the observed time series increases to infinity. It is then shown that, depending on the rate at which the sharpness of the spectral density at λ_0 gets closer to that of a pole, the AR(2) based estimator achieves a rate of convergence which is larger than the parametric $n^{-1/2}$ rate and can be arbitrarily close to $n^{-2/3}$, the best rate that can be achieved by this estimator.
Joint work with Jens-Peter Kreiss and Efstathios Paparoditis.

Multiscale Dynamic Dependence Estimation over Networks

Cristian Varon
University of York

Many modern datasets comprise multivariate time series whose cross-dependencies are structured by an underlying network. In such settings, the information at each node is often encompassed in a nonstationary process, giving rise to complex interactions between temporal dynamics and network topology. This paper introduces Locally Stationary Wavelet processes on Networks (Net-LSW), a new framework for modelling multiscale, time-varying dependencies that explicitly incorporates the network structure. Unlike traditional multivariate approaches, the Net-LSW process encodes the graph directly in the covariance structure of the process's random increments. We establish a key theoretical result showing that the absence of an edge in the graph corresponds to a zero entry in the associated network localized precision matrix across all wavelet scales and times. For inference on the local cross-nodal (partial) dependence, we develop a novel subprocess-based estimation procedure and theoretically prove estimator consistency. Simulation studies further demonstrate that the proposed framework accurately recovers evolving dependence structures whilst respecting the underlying graph topology.
Joint work with Marina Knight at York and Matt Nunes at Bath.

A simple spectral goodness-of-fit test for general time series models

Haihan (Mark) Yu
University of Rhode Island

A simple spectral goodness-of-fit test for general time series models Abstract: We revisit a test statistic originally proposed by Milhøj (1981) for assessing autoregressive and moving average (ARMA) models with independent innovations, and modify it to obtain a goodness-of-fit test for general time series models. Under the null hypothesis, the modified statistic asymptotically removes the contribution of fourth-order cumulant structures, which are typically the main obstacle to extending linear time series inference tools to general nonlinear processes. This removal obviates the need to estimate high-order terms, resulting in a tuning-parameter-free test that is computationally simple to implement. Importantly, the elimination of fourth-order effects is induced by the test's form and does not depend on the specific underlying model. The proposed test is therefore widely applicable, and its asymptotic null distribution is established under ψ-weak dependence conditions that encompass both linear and nonlinear time series. Numerical studies demonstrate that the test not only achieves accurate type I error control but also attains power comparable to existing bootstrap-based and data-driven portmanteau tests.

List of Participants

Andreas Anastasiou University of Cyprus

Sumanta Basu Cornell University

Patrice Bertail University Nanterre, Paris

Noah Bolanos Texas A&M University

Alexander Braumann TU Braunschweig

Anna Bykhovskaya Duke University

Likai Chen Washington University St. Louis

Haeran Cho University Bristol

Rainer Dahlhaus University Heidelberg

Richard A. Davis Columbia University New York

Anne van Delft Columbia University New York

Paul Doukhan University Cergy-Pontoise, Paris

Linus Dowidat TU Braunschweig

Edward Eriksson Max-Planck Institute Leipzig

Christian Francq University Lille

Liudas Giraitis London School of Economics

Siegfried Hörmann Graz University of Technology

Carsten Jentsch TU Dortmund

Moritz Jirak University Vienna

William Knauth Columbia University New York

Georg Köstenberger University Vienna

Piotr Kokoszka Colorado State University

Jens-Peter Kreiss TU Braunschweig

Jaeseon Lee Texas A&M University

Anne Leucht University Bamberg

Yufei Li King's Business School London

Alexander Lindner University Ulm

Miles Lopez University of California Davis

Panagiotis Maouris University of Cyprus

George Michailidis University of California Los Angeles

Michael H. Neumann University Jena

Dan Nordman Iowa State University

Efstathios Paparoditis University of Cyprus

Dimitris N. Politis University of California San Diego

Mohsen Pourahmadi Texas A&M University

Daniel Rademacher Graz University of Technology

Gregory Rice University of Waterloo

Hanlin Shang Macquarie University

Suhasini Subba Rao Texas A&M University

Lionel Truquet University Rennes ENSAI

Cristian Varon University York

Martin Wendler Otto-von-Guericke-Universitẗ Magdeburg

Kejin Wu Loyola University Chicago

Han Xiao Rutgers University

Han Yan London School of Economics

Qiwei Yao London School of Economics

Haihan (Mark) Yu University of Rhode Island

Yunyi Zhang Chinese University of Hong Kong

Zhou Zhou University of Toronto

Pavlos Zoubouloglou University Münster