Workshop on Neural dynamical systems for time-series data

Book of abstracts

Contents

Preface 3
- Scientific committee 3
- Sponsoring institutions 3

Mini-courses 4
- Paul Hager 4
- Cris Salvi 4

Wednesday, 23.04.2025. 5
- Fabian Harang 5
- Marco Rauscher 5
- Lingyi Yang 6
- Nina Drobac 6
- Tomás Carrondo 7
- Julian Pachschwöll 7
- Ric Chan 8
- Peter Friz 8

Thursday, 24.04.2025. 9
- Sebastian Riedel 9
- Giovanni Ballarin 9
- Maud Lemercier 10
- Naomi Chmielewski 10
- Ya-Ping Hsieh 11
- Liana Akobian 11
- Esteban Hernandez-Vargas 11
- Benjamin Walker 12
- Christian Bayer 12

Friday, 25.04.2025. 13
- Amaury Durand 13
- Fernando Moreno-Pino 13
- Hossein Mohammadi 14

Preface

This event will explore cutting-edge topics, including neural and controlled differential equations, reservoir computing, signature and signature-kernel methods, and more. The workshop emphasizes applications in modeling and optimization across diverse fields such as finance, biology, and medicine.

Jointly organized by QUARIMAFI (University of Vienna's QUARIMAFI group), the Wolfgang Pauli Institute (WPI) and the LPSM (Laboratoire de Probabilités Statistique & Modélisation) this workshop provides a platform for researchers to share innovative ideas and establish new collaborations. To encourage in-depth discussions, a dedicated time for interaction will be scheduled on each day of the event. In addition to keynote presentations, PhD students and early-career researchers are invited to contribute talks showcasing their recent work.

The workshop will consist of three full days of keynote talks, contributed presentations, two mini-courses and structured discussions. The two mini-courses will be integrated into the main workshop schedule and are directed to Phd- and advanced Master students and roughly correspond to 1ECTS. A background in probability theory, applied stochastic processes, and statistics is recommended.

Scientific committee

Linus Bleistein
Inria Montpellier & Inria Paris, France

Nina Drobac
Sorbonne Université, France

Paul Hager
University of Vienna, Austria

Christa Cuchiero
University of Vienna, Austria

Adeline Fermanian
Califrais, France

Olivier Wintenberger
Sorbonne Université, Paris

Sponsoring institutions

LPSM Laboratoire de probabilités statistique & modélisation

wpi Wolfgang Pauli Institute

oskar morgenstern doctoral school

Mini-courses

Paul Hager

University of Vienna, Austria

Dynamic and Distributional Features on Path Space

Cris Salvi

Imperial College London, United Kingdom

Signature Methods in Finance

Wednesday, 23.04.2025.

Fabian Harang

BI Norwegian Business School, Norway

Signatures with memory - exploring the Volterra signature

In this talk, we discuss the Volterra signature, an extension of the classical rough path signature. This concept was originally proposed in H. and Tindel (2021) for Volterra-type processes driven by potentially singular kernels. Motivated by Volterra equations and their Picard iterations, we define the signature using a convolution-based algebraic structure. We explore its analytic and algebraic properties, focusing on universality in functional approximation and injectivity of the mapping from the signature to its driving signal. We also discuss the role of Chen's relation, and computational aspects of this signature over "simple Volterra paths".

Marco Rauscher

Technical University Munich, Germany

A Fourier Inversion Formula for the Truncated Signature Group

In this talk, we develop the unitary irreducible representations for the group of truncated signatures. We describe the concept of induced representations for locally compact, nilpotent Lie groups and characterize all equivalence classes of these representations. Finally, we exploit our characterization to determine the characters of the unitary irreducible representations to formulate and simplify the Fourier Inversion formula in this context.

Lingyi Yang

University of Oxford, United Kingdom

Detecting fake data through anomaly detection

In the last couple of years, we have seen an astronomical increase in the popularity of generative deep learning, for example in the form of ChatGPT and DALL-E. At our fingertips, we can generate essays, art, music, and much more.
With new technologies, come new challenges. These tools can be maliciously used for misinformation and plagiarism, and therefore detection tools need to keep up with the evolution of generative AI. We frame the problem of detecting fake/generated data as an anomaly detection problem. Anomaly detection aims to identify whether new data points significantly deviate from its training corpus and certain transformations may aid this.
We present SigMahaKNN, a pipeline designed for scoring anomalous streams based on clean corpus anomaly detection. The conformance scores of new samples are derived from a combination of path signatures with the Mahalanobis distance. Such an approach preserves desirable invariances, namely to affine transformations of the data and appending metadata. Our pipeline is versatile and can be used with a wide array of non-stationary, multi-modal tick data with complex missingness patterns. We showcase its effectiveness in detecting time series obtained from deep generative models for voice cloning.

Nina Drobac

Sorbonne Université, France

Signed, Sealed, Predicted: Time Series Forecasting with Signatures

Recent challenges in electricity load forecasting, such as the COVID-19 pandemic and the energy crisis, have highlighted the need for adaptive predictive models. In this talk, I will introduce a time series forecasting framework based on the signature - a non-parametric feature set for sequential data that effectively captures temporal dynamics and interaction patterns. I will first present theoretical guarantees that motivate the use of linear regression models on signatures calculated over a sliding window. To bridge the gap between theory and practice, I will showcase an online algorithm that exploits the algebraic structure of the signature space for more efficient computation. The presentation concludes with experimental results, demonstrating the framework's potential in forecasting electricity load.

Tomás Carrondo

University of Vienna, Austria

Well-posedness and approximation properties of signature CDEs

In this talk, I will introduce signature controlled differential equations (Sig-CDEs) as a natural and universal class within the framework of path-dependent controlled differential equations. A central open problem is the absence of a general existence and uniqueness theory, even in the simpler case of bounded variation drivers. I will present two equivalent formulations of Sig-CDEs and provide sufficient conditions for existence and uniqueness in each. Additionally, I will establish a stability result showing that any solution of a path-dependent CDE can be approximated, under mild assumptions and in a suitable sense, by a solution of a Sig-CDE. The aim is to offer new insights into the foundational theory of Sig-CDEs and their role in modeling path-dependent dynamics. This talk is based on recent and ongoing joint work with Christa Cuchiero, Paul Hager, and Fabian Harang.

Julian Pachschwöll

University of Vienna, Austria

Solving High-Dimensional Riccati Equations in Signature Volatility Models

We study signature-based volatility models where the volatility is given as a linear combination of signature terms of an underlying primary process, specified as a multivariate time-extended Ornstein-Uhlenbeck process. Using the affine framework introduced by Cuchiero et al., we view the log price enhanced with the signature of the primary process, as an finite-dimensional affine process. Under certain non-trivial assumptions, this allows us to express the characteristic function of the log price as the solution to an infinite-dimensional Riccati ODE. Truncating this system provides a practical method for approximating the characteristic function, enabling option pricing via Fourier methods. This talk focuses on the numerical solution of the truncated Riccati system, highlighting implementation challenges, particularly the need for fast and efficient evaluation of the ODE function. We discuss these aspects in detail and present numerical results demonstrating the performance and limitations of the approach.

Ric Chan

Singapore Management University, Singapore

Signature-Based Clustering of Financial Time Series Using Weighted Signature Kernels

This study introduces a novel methodology for clustering financial time series through the application of weighted signature kernels — a technique that blends stochastic analysis with modern machine learning. Leveraging the signature transform to effectively capture the sequential dynamics inherent in financial data, our approach computes a weighted signature kernel matrix that quantifies the similarity between multivariate time series. We validate our method on US equity data (NVDA, AAPL, MSFT, and GOOGL) spanning from 2022 to 2024, where log returns are standardized and fed into a PDE-based solver to obtain the signature kernel.
The clustering pipeline integrates k-means with silhouette score optimization to determine the optimal number of clusters, thereby revealing distinct market regimes and anomaly patterns. In addition, we extend the analysis by incorporating an optimal measure framework and evaluating hyperbolic development kernels to quantify the alignment with Wiener measure. This dual analysis not only refines the clustering outcomes but also provides a robust error metric that compares the optimal discrete measure against theoretical benchmarks. Supplementary quality assessments, including metrics such as the Calinski-Harabasz and Davies-Bouldin scores, further support the efficacy of the proposed method.
Overall, the methodology demonstrates that weighted signature kernels are a powerful tool for unveiling complex structural similarities in financial time series, offering promising applications in market segmentation, portfolio optimization, and anomaly detection.

Peter Friz

Technische Universität Berlin & Weierstrass Institute, Germany

On Expected Signature Kernels

The expected signature kernel arises in statistical learning tasks as a similarity measure of probability measures on path space. Computing this kernel for known classes of stochastic processes is an important problem that, in particular, can help reduce computational costs. Building on the representation of the expected signature of inhomogeneous Lévy processes as the development of a smooth path in the extended tensor algebra [F.-H.-Tapia, Forum of Mathematics: Sigma (2022), "Unified signature cumulants and generalized Magnus expansions"], we extend the arguments developed for smooth rough paths in [Lemercier-Lyons (2024), "A high-order solver for signature kernels"] to derive a PDE system for the expected signature of inhomogeneous Lévy processes. As a specific example, we demonstrate that the expected signature kernel of Gaussian martingales satisfies a Goursat PDE. (Joint work with P. Hager)

Thursday, 24.04.2025.

Sebastian Riedel

FernUni Hagen, Germany

Stochastic control with signatures

We present a new approach to study stochastic optimal control problems using the signature, an object originated from rough paths theory. We will show how to solve the optimal stopping problem and furthermore study optimal control of stochastic differential equations with the signature. This is joint work with Peter Bank, Christian Bayer, Paul Hager, Tobias Nauen and John Schoenmakers.

Giovanni Ballarin

University of St. Gallen, Switzerland

From Many Models, One: Macroeconomic Forecasting with Reservoir Ensembles

Model combination is a powerful approach to achieve superior performance with a set of models than by just selecting any single one. We study both theoretically and empirically the effectiveness of ensembles of Multi-Frequency Echo State Networks (MF-ESNs), which have been shown to achieve state-of-the-art macroeconomic time series forecasting results (Ballarin et al., 2024). Hedge and Follow-the-Leader schemes are discussed, and their online learning guarantees are extended to the case of dependent data. In applications, our proposed Ensemble Echo State Networks show significantly improved predictive performance compared to individual MF-ESN models.

Maud Lemercier

University of Oxford, United Kingdom

High order solvers for signature kernels

Signature kernels are at the core of several machine learning algorithms for analysing multivariate time series. The kernels of bounded variation paths, such as piecewise linear interpolations of time series data, are typically computed by solving a linear hyperbolic second-order PDE. However, this approach becomes considerably less practical for highly oscillatory inputs, due to significant time and memory complexities. To mitigate this issue, I will introduce a high order method which involves replacing the original PDE, which has rapidly varying coefficients, with a system of coupled equations with piecewise constant coefficients. These coefficients are derived from the first few terms of the log-signatures of the input paths and can be computed efficiently using existing Python libraries.

Naomi Chmielewski

EDF Lab & CentraleSupélec, France

Quantum Reservoir Computing and Risk Bounds

Quantum Reservoir Computing (QRC) is a machine learning paradigm that proposes to harness the dynamics of open quantum systems for time series forecasting tasks. QRC is a realisation of physical reservoir computing on a quantum computer. Compared to reservoir computing on classical systems, QRC is governed by quantum dynamics that are not classically simulable, even for small system sizes. This suggests that a relatively small quantum system might be adapted to the forecasting of highly non-linear time series. The classical time series data is injected into the quantum system at discrete time intervals, and the system is left to evolve according to its natural dynamics. After all data has been injected, the quantum system is measured, and a simple linear or polynomial regression is performed to fit the forecasting task. While different realisations of QRC have been proven to be universal for the approx- imation of fading memory maps, the generalisation error of such quantum reservoirs had not been studied previously. In our work we find an upper bound on the generalisation error for the universal QRC classes mentioned above that scales as sqrt(log(m)/m) with the number m of training samples. The result suggests a convergence speed of the generalisation error with the number of training samples. The generalisation error also contains a factor that is exponential in the number of qubits (i.e. the system size), which comes from the polynomial regression of the measured states. This is unfavourable but can be mitigated by using linear regression; however it is not clear how this change might impact the universality of the model, as the polynomial regression is a key part in the proof of universality. This remains an open question.

Ya-Ping Hsieh

ETH Zürich, Switzerland

Schrödinger Bridge Framework for Modeling Snapshot Data

The Schrödinger Bridge (SB) framework provides a principled approach to reconstructing dynamical processes from snapshot data, with deep connections to optimal transport and stochastic optimal control. In this talk, we present a novel training algorithm that establishes rigorous guarantees for learning SBs. Our method leverages classical optimization techniques, specifically mirror descent and stochastic approximation, to ensure efficient and theoretically grounded training.

Liana Akobian

University of Vienna, Austria

Uncovering Neural Control: A Dynamical Systems Approach to Disentangling Intrinsic and Controlled Neural Dynamics

Understanding how neurons interact to produce behavior is a key challenge in neuroscience. The dynamics of these interacting neurons define the computations that underlie the processing of sensory information, decision making, and the generation of motor output. Recent advances in dynamical system modeling have formalized observed neural activity as the temporal evolution of states within a neural state space governed by dynamical laws. While many existing models assume autonomous evolution, they may not sufficiently capture external perturbations that influence neural computation. In this work, we introduce a controlled decomposed linear dynamical system (cdLDS). In an unsupervised way, this algorithm learns unknown inputs that modulate neural state transitions, extending prior work using autonomous dynamical system models (dLDS). We apply cdLDS to whole brain activity data from C. elegans and demonstrate its ability to separate intrinsic neural dynamics from control signals. This decomposition provides insights into how external perturbations shape neural computation, offering a principled framework for understanding the impact of control mechanisms on neural dynamics. By bridging neuroscience and mathematical modeling, our work contributes to broader applications in biological systems and dynamical modeling.

Esteban Hernandez-Vargas

Univerisity of Idaho, United States of America

Hybrid Neural Differential Equations to model Unknown Mechanisms and States

Efforts to model complex systems increasingly face challenges from ambiguous relationships within the model, such as through partially unknown mechanisms or unmodelled intermediate states. Hybrid neural differential equations are a recent modeling framework that has been previously shown to enable the identification and prediction of complex phenomena, especially in the context of partially unknown mechanisms. We extend the application of hybrid neural differential equations to enable the incorporation of theorized but unmodelled states within differential equation models. We find that beyond their capability to incorporate partially unknown mechanisms, hybrid neural differential equations provide an effective method to include knowledge of unmeasured states into differential equation models.

Benjamin Walker

University of Oxford, United Kingdom

Linear Neural Controlled Differential Equations

Controlled differential equations (CDEs) describe the relationship between a control path and the evolution of a solution path. Neural CDEs (NCDEs) extend this concept by parameterising the CDE's vector field with neural networks, treating time series as observations from a control path, and interpreting the solution as a continuously evolving hidden state. Their robustness to irregular sampling makes NCDEs highly effective for real-world data modelling.
This talk highlights Linear Neural Controlled Differential Equations (LNCDEs), where the vector field is linear in the hidden state. LNCDEs combine the expressive power of non-linear recurrent neural networks with the computational parallelism of structured state-space models. However, their cubic computational cost in hidden dimension limits their scalability.
We introduce three novel architectures—sparse, Walsh–Hadamard, and block-diagonal LNCDEs—collectively called Structured Linear Controlled Differential Equations (SLiCEs). We theoretically show that SLiCEs maintain the expressiveness of dense LNCDEs while significantly reducing computational complexity. Empirical benchmarks on state-tracking tasks confirm their practical efficiency and scalability.

Christian Bayer

Weierstrass Institute, Germany

Pricing American options under rough volatility

Rough volatility models are an important class of stock price models, which are widely recognised for allowing excellent fits to market prices of options. However, the roughness of the volatility dynamics, and, even more so, the lack of Markov property lead to considerable numerical challenges, especially regarding path-dependent options. We introduce a range of efficient numerical methods for pricing of American options under rough volatility based on path signatures. After providing theoretical analysis of the methods, we verify their accuracy using numerical examples. (Based on joint works with P. Hager, L. Pelizzari, S. Riedel, J. Schoenmakers, and J. J. Zhu.)

Friday, 25.04.2025.

Amaury Durand

EDF R&D OSIRIS, France

Power comparison of sequential testing by betting procedures

We derive power guarantees of some sequential tests for bounded mean under general alternatives. We focus on testing procedures using nonnegative supermartingales which are anytime valid and consider alternatives which coincide asymptotically with the null (e.g. vanishing mean) while still allowing to reject in finite time. Introducing variance constraints, we show that the alternative can be broaden while keeping power guarantees for certain second-order testing procedures. We also compare different test procedures in multidimensional setting using characteristics of the rejection times. Finally, we extend our analysis to other functionals as well as testing and comparing forecasters. Our results are illustrated with numerical simulations including bounded mean testing and comparison of forecasters.

Fernando Moreno-Pino

University of Oxford, United Kingdom

Rough Transformers: Lightweight and Continuous Time Series Modelling through Signature Patching

Time-series data in real-world settings typically exhibit long-range dependencies and are observed at non-uniform intervals. In these settings, traditional sequence-based recurrent models struggle. To overcome this, researchers often replace recurrent architectures with Neural ODE-based models to account for irregularly sampled data and use Transformer-based architectures to account for long-range dependencies. Despite the success of these two approaches, both incur very high computational costs for input sequences of even moderate length. To address this challenge, we introduce the Rough Transformer, a variation of the Transformer model that operates on continuous-time representations of input sequences and incurs significantly lower computational costs. In particular, we propose multi-view signature attention, which uses path signatures to augment vanilla attention and to capture both local and global (multi-scale) dependencies in the input data, while remaining robust to changes in the sequence length and sampling frequency and yielding improved spatial processing. We find that, on a variety of time-series-related tasks, Rough Transformers consistently outperform their vanilla attention counterparts while obtaining the representational benefits of Neural ODE-based models, all at a fraction of the computational time and memory resources.

Hossein Mohammadi

University of Exeter, United Kingdom

Emulating Complex Dynamical Simulators with Random Fourier Features

A Gaussian process (GP)-based methodology is proposed to emulate complex dynamical computer models (or simulators). The method relies on emulating the numerical flow map of the system over an initial (short) time step, where the flow map is a function that describes the evolution of the system from an initial condition to a subsequent value at the next time step. This yields a probabilistic distribution over the entire flow map function, with each draw offering an approximation to the flow map. The model output time series is then predicted (under the Markov assumption) by drawing a sample from the emulated flow map (i.e., its posterior distribution) and using it to iterate from the initial condition ahead in time. Repeating this procedure with multiple such draws creates a distribution over the time series. The mean and variance of this distribution at a specific time point serve as the model output prediction and the associated uncertainty, respectively. However, drawing a GP posterior sample that represents the underlying function across its entire domain is computationally infeasible, given the infinite-dimensional nature of this object. To overcome this limitation, one can generate such a sample in an approximate manner using random Fourier features (RFF). RFF is an efficient technique for approximating the kernel and generating GP samples, offering both computational efficiency and theoretical guarantees. The proposed method is applied to emulate several dynamic nonlinear simulators including the well-known Lorenz and van der Pol models. The results suggest that our approach has a promising predictive performance and the associated uncertainty can capture the dynamics of the system appropriately.