BSSASA 2026

Book of Abstracts

Berlin Summer School on Applied Stochastic Analysis

Berlin, July 28–31, 2026

ORGANISED AND SUPPORTED BY

Contents

Introduction & Information

It is our great pleasure to welcome you to the **Berlin Summer School on Applied Stochastic Analysis**.

Venue

The summer school will take place at TU Berlin (IMoS building) at Fasanenstraße 89, 10623 Berlin. The building can be accessed only from Mueller-Breslau-Straße, via a path between the HF and L buildings.

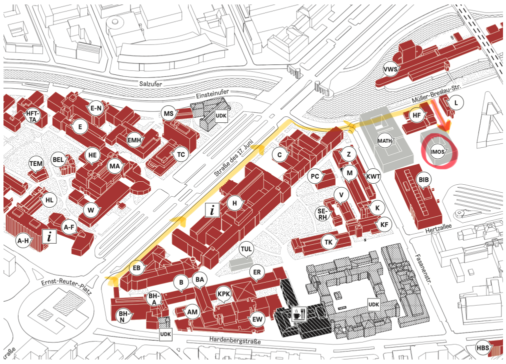

Mini-courses

The scientific programme will feature two mini-courses, each consisting of three 90-minute lectures, delivered by Josef Teichmann and Pierre Del Moral. Josef Teichmann will discuss recent developments in *Neural Jump ODEs and Generative Modeling*, while Pierre Del Moral's mini-course is titled *An Introduction to Schrödinger Bridges and Sinkhorn Semigroups*.

Contributed Talks

Contributed talks will be given slots of 30 minutes each consisting of 25 minutes for each talk and 5 minutes for questions and smooth transition between speakers. The abstracts are presented in the subsequent pages.

Recommendations for Lunch Breaks

For your lunch breaks you have several options:

There is a café on the ground floor of the University Library, Café Nero, approximately 5 minutes on foot from the venue. It offers freshly prepared lunch dishes, including pasta, soups, salads and vegetarian or vegan options, as well as smaller snacks and drinks. Please note that the café accepts cash only.

The restaurant Schleusenkrug features a beer garden and serves traditional German cuisine. It is located approximately 5 minutes on foot from the venue.

The Mensa TU Hardenbergstraße is approximately 10–12 minutes on foot from the venue. It offers several vegan and vegetarian dishes every day, as well as selected meat and fish dishes. Guests require a MensaCard, which can be purchased at the checkout or at a card machine in the Mensa.

The Mathe-Café is located in the Mathematics Building at Straße des 17. Juni 136, approximately 8–10 minutes on foot from the venue. It mainly serves Indian-inspired lunch dishes, including curries, chicken dishes, paneer, lentils, rice and samosas. Vegetarian options are usually available. The menu changes regularly and can be checked on the website.

There are numerous cafés and restaurants around Savignyplatz, approximately 12–15 minutes on foot from the venue. The area offers a broad selection of cuisines, including Italian, Asian, German, Mediterranean and vegetarian food. Prices and opening hours vary considerably, so participants may wish to check the individual restaurant before going there.

Dave B. – Taste Five is located at Otto-Suhr-Allee 6–16, approximately 15 minutes on foot from the venue. It is a food-court-style restaurant with several different food counters under one roof. The selection generally includes pizza, burgers, salads, vegetarian or vegan dishes, snacks and coffee.

Conference Dinner

The conference dinner will take place on Wednesday at 19:30 at Dicke Wirtin. Located near Savignyplatz, the restaurant is a traditional Old Berlin pub serving hearty regional cuisine. It is approximately 10–15 minutes on foot from the venue.

Website

Further information is available at `https://www.wias-berlin.de/workshops/IRTG2026/`.

Organisers

Peter Bank, Jannis Dause, Jakob Kellermann, Martin Nicolaus, Gemma Sedrakjan, Wilhelm Stannat, Roman Wüthrich.

Timetable

	Tuesday 28.7.	Wednesday 29.7.	Thursday 30.7.	Friday 31.7.
9:00		**Lecture 2** Del Moral	**Lecture 1** Teichmann	**Lecture 3** Teichmann
10:30	**Registration & Welcome**	Coffee break		
11:00	**Lecture 1** Del Moral	**Session 3** Haihan, Yang, Spille	**Session 4** Aqsha, Rauscher, Dattge	**Session 6** Duan, Xu, Chinta
12:30	**Lunch break**			
14:00	**Session 1** Cinque, Gökalp, Rahmani	**Lecture 3** Del Moral	**Lecture 2** Teichmann	**Session 7** Nero, Arndt
15:00				
15:30	Coffee break			
16:00	**Session 2** Abedi, Wunderlich, Likibi Pellat	**Summer School Photo & Social get-together** Activity flexible according to the weather	**Session 5** Zoubouloglou, Vencato, Alonso Martin	
19:30		**Conference dinner**		

Contributed Talks

Session 1

Estimating and Forecasting the Intraday Dynamics of Spot Betas

Paolo Cinque

Scuola Normale Superiore

We develop a complete inferential and forecasting framework for the intraday trajectory of the spot beta coefficient estimated from high-frequency, noisy, and asynchronously observed prices. First, we establish a joint central limit theorem for the Fourier spot variance and covariance estimators under microstructure noise, and derive, via the delta method, the asymptotic distribution of the spot beta estimator. The asymptotic variance admits a closed-form global minimizer in the tuning parameters, yielding a fully data-driven and feasible selection rule for the cutting frequencies. Second, we apply the feasible estimator to a panel of twenty NASDAQ-listed stocks over 2022–2024, documenting four stylized facts of the estimated intraday beta curves: a persistent U-shaped intraday pattern, occasional idiosyncratic decoupling episodes, strong serial dependence with an intraday seasonal structure, and a low intrinsic dimensionality captured by two Robust Dynamic Functional Principal Components (RDFPCA) accounting for roughly 95% of cross-day variability. Building on these findings, we produce one-day-ahead forecasts of the full intraday beta trajectory by combining RDFPCA score extraction with univariate time series models and machine learning models. Out-of-sample evaluation shows that score-based specifications, in particular an AR model on the leading scores, consistently outperform functional and machine-learning benchmarks, while remaining parsimonious and fully interpretable.

Parameter estimation in a fully coupled partially observed Ornstein-Uhlenbeck process

Hasan Mert Gökalp

Technische Universität Berlin

We study a two-dimensional Ornstein–Uhlenbeck system where only the first coordinate is observed, whereas the second coordinate remains hidden. We estimate the coupling parameter in the drift of the observed coordinate. Within the Ibragimov–Hasminskii framework, we prove

consistency, asymptotic normality, convergence of moments and asymptotic efficiency of the maximum likelihood estimator under stability and identifiability assumptions as the time horizon grows. We also apply the results to estimation of the reaction-parameter in partially observed linear stochastic reaction–diffusion equations. Finally, we briefly mention ongoing work on parameter estimation for SPDEs with spatio-temporal point process observations.

A Flexible SPDE/LGCP/Hawkes Framework with Application to the Epidemic-Type Aftershock Sequence Model (ETAS)

Sofiane Takieddine Rahmani
University of Potsdam

We propose a general framework for incorporating stochastic partial differential equation (SPDE) representations and Log-Gaussian Cox Processes (LGCPs) into Hawkes-type self-exciting point processes, with the Epidemic-Type Aftershock Sequence (ETAS) model for earthquake occurence (Ogata, 1998). In this formulation, the background intensity is modeled as a latent Gaussian random field represented through an SPDE, while the conditional intensity of the Hawkes process combines this structured background component with the triggering mechanism governing self-excitation.

Within a Bayesian framework, the proposed approach enables joint estimation of latent background fields and Hawkes triggering parameters while substantially reducing computational cost relative to Gaussian Process implementations based on dense covariance structures. Using ETAS as an example, the framework provides accurate inference of both background and triggering components and naturally accommodates large and high-resolution event catalogs. Validation tests have been carried out on both synthetic seismicity catalogs and real earthquake data, demonstrating that the methodology accurately recovers model parameters, reproduces spatial background seismicity patterns, and remains stable and computationally efficient under realistic conditions.

Ogata, Y. Space-time point-process models for earthquake occurrences. Annals of the Institute of Statistical Mathematics, 50(2):379-402,

---

Session 2

Fractional Sobolev paths on Wasserstein spaces and their energy-minimizing particle representations

Ehsan Abedi

Technische Universität Berlin

We study a generalization of the Monge–Kantorovich optimal transport problem. Given a prescribed family of time-dependent probability measures (μ_t), we aim to find, among all path-continuous stochastic processes whose one-dimensional time marginals coincide with (μ_t) (if there is any), a process that minimizes a given energy. After discussing existing results on Sobolev energy, we investigate fractional Sobolev energy. Given a path (μ_t) on a p-Wasserstein space with fractional Sobolev regularity $W^{\alpha,p}$, where $1/p < \alpha < 1$, we provide conditions under which we prove the existence of a process that minimizes the energy, and construct a process that realizes the regularity of (μ_t). From a technical perspective, our construction relies on a characterization of nonlinear Besov spaces established by Liu–Prömel–Teichmann (2020). If time permits, we will discuss some applications.

Weak convergence of Volterra-type stochastic integrals in Skorokhod space

Fabrice Wunderlich

University of Luxembourg

We study weak convergence of Volterra-type stochastic integrals whose integrands depend on both the integration variable and the terminal time. Such stochastic convolutions arise naturally in non-Markovian limit problems.

We obtain weak convergence results in both the J1 and M1 Skorokhod topologies under natural conditions on the semimartingale integrators, the integrands, and their interaction. The framework applies in particular to stochastic convolutions with singular kernels and is motivated by scaling limits arising in nearly unstable Hawkes models and rough volatility.

This is joint work with Giulia Livieri and Andreas Søjmark.

Rearranged Stochastic Heat Equation with Entropy gradient structure

Rhoss Likibi Pellat
Université Côte d'Azur

In this talk, we extend a previously introduced one-dimensional diffusion model on probability measures, defined via the rearranged stochastic heat equation, by penalizing the dynamics with an additional entropy-driven gradient-descent term. By means of a splitting argument, we prove that despite the opposite effects of rearrangement and entropy minimization, the resulting penalized stochastic heat equation is well defined. We study several properties of the associated dynamics and show, in particular, that solutions admit a density satisfying a corrected version of the Dean–Kawasaki equation. Moreover, smoothing properties established for the stochastic heat equation are shown to persist, which, together with the existence of a density, leads to regularization results for mean-field models depending on the pointwise value of the density.

This is based on a joint work with François Delarue.

Session 3

Existence, Uniqueness and Propagation of Chaos for Distribution-Dependent SDEs with Local Interactions

Yang Haihan

Nanjing University of Science and Technology

In this talk, I will introduce our joint results with Michael A. Kouritzin on nonlinear stochastic differential equations with local interactions. Particles exhibit local interactions through the drift coefficient b and the diffusion coefficients σ, α, where σ, b and α depend upon the particle distribution and on other particles within the interaction radius r. The main goals are to demonstrate the *sub-explosive* solution of DDSDEs, and by applying the modified Yamada-Watanabe theorem and the Euler-Maruyama (EM) methods, we establish the existence and uniqueness for this particle system with potentially discontinuous coefficients. Finally, we show a propagation of chaos result in Wasserstein distance for DDSDEs with local interactions.

Equilibrium Strategies for the N-agent Mean-Variance Investment Problem over a Random Horizon

Ying Yang

Southern University of Science and Technology

We study equilibrium feedback strategies for a family of dynamic mean-variance problems with competition among a large group of agents. We assume that the time horizon is random and each agent's risk aversion depends dynamically on the current wealth. We consider both the finite population game and the corresponding mean-field one. Each agent can invest in a risk-free asset and a specific individual stock, which is correlated with other stocks by a common noise. By applying stochastic control theory, we derive the extended Hamilton-Jacobi-Bellman (HJB) system of equations for both n-agent and mean-field games. Under an exponentially distributed random horizon, in each case, we explicitly obtain the equilibrium feedback strategies and the value function. Our results show that the agent's equilibrium feedback strategy depends not only on his/her current wealth but also on the wealth of other competitors. Moreover, when the risk aversion is state-independent and the risk-free interest rate is zero, the equilibrium strategies degenerate to constants, which is identical to the unique equilibrium obtained in [Lacker, D. and Zariphopoulou, T. (2019). Mean field and n-agent games for optimal investment under relative performance criteria. Mathematical Finance, 29(4):1003–1038] with exponential risk preferences; when the competition parameter goes to zero and the risk aversion equals some specific value, the equilibrium strategies coincide with the ones derived in [Landriault, D., Li, B., Li, D., and Young, V. R. (2018). Equilibrium strategies for the mean- variance investment problem over a random horizon. SIAM Journal on Financial Mathematics, 9(3):1046–1073].

A Novel Approach to Peng's Maximum Principle for McKean-Vlasov SDE

Johan Benedikt Spille

Technische Universität Berlin

We present a novel approach to the proof of Peng's maximum principle for McKean-Vlasov stochastic differential equations (SDE). The main step is the introduction of a third adjoint equation, a conditional McKean-Vlasov backward SDE, to accommodate the dualization of quadratic terms containing two independent copies of the first-order variational process. This is an intrinsic extension of the maximum principle from Peng for standard SDE and gives a conceptually consistent proof. Our approach will be useful in further extensions to the common noise setting and the infinite dimensional setting.

Session 4

Solving linear-quadratic stochastic control problems with signatures

Muhammad Alif Aqsha
University of Oxford

We study a signature-driven numerical scheme to solve multi-dimensional linear-quadratic (LQ) stochastic control problems. Using that linear signature functionals are dense in the natural class of admissible controls, we show that our approach turns the original LQ problem into a deterministic convex quadratic polynomial optimisation. To underpin a numerical approach based on truncated signatures, we prove that the problem's value function can be approximated by finite-dimensional polynomial approximations when the truncation levels are chosen sufficiently high. Remarkably, our numerical experiments show very decent accuracy already for small truncation levels. Key tools for our analysis are (i) the algebraic representation of controlled stochastic differential equations and the associated cost function as linear functionals of the path signatures of the driving noise, (ii) the convergence of the truncated linear functionals, and (iii) the density of signature controls.

Inverting a Measure from Its Expected Signature via Conditional Gradient Descent

Marco Rauscher
Technische Universität München

We address the inverse problem of reconstructing a probability measure on path space from truncated expected signature data. Interpreting signatures as terminal states of path-controlled controlled differential equations in a truncated tensor algebra, we seek a law whose expected signature matches a prescribed or empirical target. This gives rise to a convex moment-matching problem over probability measures, where the objective is to minimise the quadratic discrepancy between the target expected signature and the one generated by the candidate law, subject to suitable moment constraints.

Our approach is based on a conditional gradient, or Frank–Wolfe, scheme that constructs a minimising sequence of finitely supported measures. At each iteration, the algorithm adds a new atomic path measure by solving an optimal control problem driven by the current expected-signature residual. Equivalently, this step generates a path whose signature yields the steepest improvement in the moment-matching objective. An extreme-point analysis explains the intrinsic sparsity of the construction and ensures that all iterates remain atomic, while dual certificates provide first-order characterisations of the resulting support paths.

We further discuss an extension beyond truncated signatures using the signature kernel and its Goursat-type characterisation. The resulting kernel formulation allows evaluation and optimisation directly in the associated reproducing kernel Hilbert space, avoiding explicit tensor-coordinate

computations and providing a route towards exact inversion from full expected signatures.

When Does SINDy Converge?
Least-Squares Consistency and Challenges with LASSO

Martin Dattge

King's College London

Sparse Identification of Nonlinear Dynamics (SINDy) is a widely used framework for discovering governing equations of dynamical systems from data, yet rigorous convergence guarantees remain limited, especially when sparsity-inducing regularization is used. In this work, we first analyze the original SINDy formulation without regularization and show that, under suitable assumptions, the estimated vector field converges to the true underlying vector field generating the data. When adding LASSO regularization, our analysis suggests that only imposing sparsity on the coefficients may be insufficient and that the choice of the regularization parameter is crucial.

Session 5

The Fourth Moment and Breuer-Major Theorems for Hilbert space valued random variables.

Pavlos Zouboulоglou
University of Münster

In this work, we establish conditions ensuring convergence in distribution of a sequence admitting a Wiener-Itô chaos representation to a nondegenerate Gaussian measure on a separable Hilbert space. A key ingredient in our approach is a Stein-Malliavin bound formulated with respect to a distance that metrizes weak convergence of probability measures on separable Hilbert spaces.

The results are infinite-dimensional extensions of the classical real-valued Fourth-Moment Theorem of Nualart and Peccati [Ann. Probab. 33, 177-193 (2005)], building upon the work by Bourguin and Campese [Electron. J. Probab. 25, 1-30 (2020)]. As an application of the Fourth Moment Theorem, we develop a Breuer-Major theorem for random variables taking values in separable Hilbert spaces, and discuss several of its applications to functional data analysis.

Densities in rough stochastic filtering with jumps: A Malliavin Calculus approach.

David Vencato
University of Oslo

Recent advances in rough stochastic differential equations have opened new perspectives on stochastic filtering with highly irregular signals, notably through the framework of Bugini, Friz and Lê (2025). In parallel, Bugini, Coghi and Nilssen (2024) established Malliavin differentiability and a Hörmander theory in the geometric rough path setting.

Motivated by these developments, we study Malliavin differentiability for rough stochastic differential equations driven by general càdlàg rough paths, building on recent progress in rough analysis with jumps by Allan and Pieper (2026).

Our approach is based on the identification of a suitable Sobolev-type structure for stochastic controlled paths. It incorporates Malliavin derivatives in a way that is compatible with the controlled path structure and enables an intrinsic formulation of Malliavin differentiability in this setting.

We discuss how this framework can be applied to the study of densities in rough stochastic filtering with jumps.

This is ongoing joint work with Salvador Ortiz-Latorre, Jost Pieper and Frank Proske.

Statistical Inference for Homogenization Limits Driven by Wiener or Hermite Processes.

Pablo Ramses Alonso Martin

University of Warwick

We study the effective estimation of the diffusivity and Hurst parameter for the homogenized limit of a class of slow/fast systems. Depending on the system parameters, this limit solves a stochastic differential equation driven by either a Wiener process or a Hermite process. In the class of models we consider, the fast variable is a fractional Ornstein–Uhlenbeck process. We show that estimators constructed from the homogenized limit remain consistent when applied to appropriately subsampled data generated by the original slow/fast system.

A key tool in our analysis is the consistency of renormalized quadratic variations for a family of additive functionals of the fast process. Using Wiener chaos expansions, we obtain an (L^2)-orthogonal decomposition of these renormalized quadratic variations. This allows us to show that, under appropriate subsampling conditions, the consistency properties of the estimators are preserved even when the data is generated by the slow/fast system rather than the homogenized limit. We also show that, under stricter subsampling conditions, a non-central limit theorem is preserved in the case where the fluctuations of the estimator around the true value are non-Gaussian. As a direct consequence of convergence in (L^2), we obtain consistency of an estimator for the limiting self-similarity that does not require knowledge of the limiting diffusivity. Finally, we show that our results apply to a class of one-dimensional fluctuation models. Talk based on https://arxiv.org/abs/2604.27242.

Session 6

Inference-Time Alignment for Diffusion Models via Variationally Stable Doob's Matching

Chenguang Duan

RWTH Aachen University

Inference-time alignment for diffusion models aims to adapt a pre-trained reference diffusion model toward a target distribution without retraining the reference score network, thereby preserving the generative capacity of the reference model while enforcing desired properties at the inference time. A central mechanism for achieving such alignment is guidance, which modifies the sampling dynamics through an additional drift term. In this work, we introduce variationally stable Doob's matching, a novel framework for provable guidance estimation grounded in Doob's h-transform. Our approach formulates guidance as the gradient of logarithm of an underlying Doob's h-function and employs gradient-regularized regression to simultaneously estimate both the h-function and its gradient, resulting in a consistent estimator of the guidance. Theoretically, we establish non-asymptotic convergence rates for the estimated guidance. Moreover, we analyze the resulting controllable diffusion processes and prove non-asymptotic convergence guarantees for the generated distributions in the 2-Wasserstein distance. Finally, we show that variationally stable guidance estimators are adaptive to unknown low dimensionality, effectively mitigating the curse of dimensionality under low-dimensional subspace assumptions.

Gradient Flow Sampler-based Distributionally Robust Optimization

Zusen Xu

Weierstrass Institute

We propose a mathematically principled PDE gradient flow framework for distributionally robust optimization (DRO). Exploiting the recent advances in the intersection of Monte Carlo sampling and statistical optimal transport, we show that our theoretical framework can be implemented as practical algorithms for sampling from worst-case distributions and, consequently, DRO.

While numerous previous works have relied on dual reformulation techniques, we contribute a sound and complete gradient flow view based on SDEs or PDEs that can be used to construct new algorithms for general, potentially non-convex, losses.

As an application, we solve a class of Wasserstein and entropy-regularized DRO problems using the recently-discovered Wasserstein Fisher-Rao and Stein variational gradient flows. Notably, we also show some simple reductions of our framework recover exactly previously proposed popular DRO methods, and provide new insights into their theoretical limits and optimization dynamics of DRO. Numerical studies based on stochastic gradient descent on machine learning tasks provide

empirical backing for our theoretical findings.

Optimal MDP Repair

Venkata Harshavardhan Chinta
University of Liverpool

Markov Decision Processes (MDPs) are a standard model for stochastic systems, where a central objective is to maximise the probability of reaching a set of target states. In many scenarios, this probability is fundamentally limited by the underlying system dynamics rather than by suboptimal policies. We address this limitation by introducing *optimal MDP repair*, where the system can be improved via buyable action upgrades that modify transition behaviour — at a cost.

We consider two settings. In the *static* setting, upgrades incur a one-off cost and permanently extend the available actions. In the *dynamic* setting, upgrades are applied during execution and incur a cost per use, leading to accumulated costs over executions. For both settings, we study the budget-constrained problem and the cost-optimal problem. We show that both problems are NP-complete in both settings, and provide mixed-integer linear programming encodings that enable automated repair strategy synthesis using off-the-shelf optimisation solvers. We implement the proposed approach and evaluate it experimentally, demonstrating its feasibility and effectiveness.

Session 7

Probability on Iterated Graph Systems

Nero Z. Li
Imperial College London

Iterated graph systems are a natural extension of what physicists call hierarchical lattices; they are fractal graphs generated by recursive iteration. To show that this model is interesting, we will present two problems on iterated graph systems: Brownian motion and Bernoulli percolation.For the former, we prove that for any edge iterated graph system the walk dimension and the spectral dimension (a form of heat kernel estimate) admit exact solutions in terms of fractal dimensions, and also construct Brownian motion on the percolation cluster, which was left open by [Hambly & Kumagai, 2010]. For the latter, we answer a question on the universality class of Bernoulli percolation raised in 1986, and thereby give a complete classification of the universality classes. There will be no detailed proofs in this talk; we hope that everyone can find new problems of their own on iterated graph systems.

On Large Deviations for the Empirical Measure of Piecewise Deterministic Markov Processes

Cinja Arndt
University of Gothenburg

We consider piecewise deterministic Markov processes, arising in applications such as continuous-time MCMC. A large deviation principle for the empirical measure is one way of assessing the speed of convergence to the stationary distribution. Using the weak convergence approach, we derive Lyapunov-type conditions for the large deviations upper bound, and verify these for the Zig-Zag process and Bouncy Particle Sampler, two widely used sampling algorithms. For the lower bound, we require a sufficiently regular near minimizer of the variational problem associated with the Laplace principle, together with an associated near-optimal control, and establish these in the one-dimensional Zig-Zag setting.

Speaker Index