

Schedule

Mon, Jul 28**Session**

08:00–17:30	Conference Check-In (HH Lobby)
08:45–09:00	Opening Ceremony (HH Auditorium)
09:00–10:00	Plenary Talk by Rohan Sawhney, Nvidia Corporation, Monte Carlo Methods in Computer Graphics (HH Auditorium)
10:00–10:30	Coffee Break (HH Lobby)
10:30–12:30	Stochastic Computation and Complexity, Part I (HH Auditorium)
10:30–12:30	Domain Uncertainty Quantification (HH Ballroom)
10:30–12:30	Nested expectations: models and estimators, Part I (HH 002)
10:30–12:30	Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part I (WH Auditorium)
10:30–12:30	Technical Session - Markov Chain Monte Carlo (WH 115)
12:30–14:00	Lunch Break (MTCC Commons)
14:00–15:00	Plenary Talk by Christiane Lemieux, U of Waterloo, Golden ratio nets and sequences (HH Auditorium)
15:00–15:30	Coffee Break (HH Lobby)
15:30–17:30	Stochastic Computation and Complexity, Part II (HH Auditorium)
15:30–17:30	Recent advances in optimization under uncertainty (HH Ballroom)
15:30–17:30	Computational Methods for Low-discrepancy Sampling and Applications (HH 002)
15:30–17:30	Technical Session - Quasi-Monte Carlo, Part I (WH Auditorium)
15:30–17:30	Technical Session - PDEs and SDEs (WH 115)
17:30–19:30	Welcome Reception (HH Lobby)

Tue, Jul 29**Session**

08:30–17:30	Registration Desk Open (HH Lobby)
09:00–10:00	Plenary Talk by Peter Glynn, Stanford U, Combining Simulation and Linear Algebra: COSIMLA (HH Auditorium)
10:00–10:30	Coffee Break (HH Lobby)
10:30–12:30	Stochastic Computation and Complexity, Part III (HH Auditorium)
10:30–12:30	Next-generation optimal experimental design: theory, scalability, and real world impact: Part I (HH Ballroom)
10:30–12:30	Heavy-tailed Sampling (HH 002)
10:30–12:30	Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part I (WH Auditorium)
10:30–12:30	Technical Session - Bayesian Methods (WH 115)
12:30–14:00	Lunch Break (On your own)
14:00–15:00	Plenary Talk by Roshan Joseph, Georgia Institute of Technology, Sensitivity and Screening: From Monte Carlo to Experimental Design (HH Auditorium)
15:00–15:30	Coffee Break (HH Lobby)
15:30–17:30	Stochastic Computation and Complexity, Part IV (HH Auditorium)
15:30–17:30	Next-generation optimal experimental design: theory, scalability, and real world impact: Part II (HH Ballroom)
15:30–17:30	Advances in Rare Events Simulation (HH 002)
15:30–17:30	Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part II (WH Auditorium)
15:30–17:30	Technical Session - Quasi-Monte Carlo, Part II (WH 115)
18:00–20:00	Chicago White Sox vs. Philadelphia Phillies (must purchase tickets beforehand) (Meet in HH Lobby)

Wed, Jul 30	Session
08:30–16:30	Registration Desk Open (HH Lobby)
09:00–10:00	Plenary Talk by Veronika Ročková, U of Chicago, AI-Powered Bayesian Inference (HH Auditorium)
10:00–10:30	Coffee Break (HH Lobby)
10:30–12:30	Stochastic Computation and Complexity, Part V (HH Auditorium)
10:30–12:30	Statistical Design of Experiments (HH Ballroom)
10:30–12:30	Advances in Adaptive Hamiltonian Monte Carlo (HH 002)
10:30–12:30	Technical Session - Simulation (WH Auditorium)
10:30–12:30	Technical Session - Sampling (WH 115)
12:30–14:00	Lunch Break (On your own)
14:00–16:00	Stochastic Optimization (HH Auditorium)
14:00–16:00	Recent Progress on Algorithmic Discrepancy Theory and Applications (HH Ballroom)
14:00–16:00	Monte Carlo Applications in High-performance Computing, Computer Graphics, and Computational Science (HH 002)
14:00–16:00	Technical Session - Statistics (WH Auditorium)
16:00–16:30	Coffee Break (HH Lobby)
18:00–20:30	Conference Banquet (Bridgeport Art Center, 1200 W 35th Street)

Thu, Jul 31	Session
08:30–17:30	Registration Desk Open (HH Lobby)
09:00–10:00	Plenary Talk by Uros Seljak, UC Berkeley, Gradient-Based MCMC Sampling: Methods and Optimization Strategies (HH Auditorium)
10:00–10:30	Coffee Break (HH Lobby)
10:30–12:30	QMC and Applications Part I (HH Auditorium)
10:30–12:30	Analysis of Langevin and Related Sampling Algorithms, Part I (HH Ballroom)
10:30–12:30	Nested expectations: models and estimators, Part II (HH 002)
10:30–12:30	Technical Session - Finance (WH Auditorium)
10:30–12:30	Technical Session - ML & Optimization (WH 115)
12:30–14:00	Lunch Break (On your own)
14:00–15:00	Plenary Talk by Nicolas Chopin, Institut Polytechnique de Paris, Saddlepoint Monte Carlo and its application to exact ecological inference (HH Auditorium)
15:00–15:30	Coffee Break (HH Lobby)
15:30–17:30	QMC and Applications Part II (HH Auditorium)
15:30–17:30	Analysis of Langevin and Related Sampling Algorithms, Part II (HH Ballroom)
15:30–17:30	Recent Advances in Stochastic Gradient Descent (HH 002)
15:30–17:30	Technical Session - Sampling (WH Auditorium)
15:30–17:30	Technical Session - SDEs (WH 115)
18:30–20:30	Steering Committee Meeting (by invitation) (Mikami Izakaya & Ramen, 1400 S Michigan Ave)
19:00–21:00	Early Career Dinner (Time Out Market, 916 W Fulton Market)

Fri, Aug 1	Session
08:30–12:15	Registration Desk Open (HH Lobby)
09:00–11:00	Forward and Inverse Problems for Stochastic Reaction Networks (HH Auditorium)
09:00–11:00	Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part II (HH Ballroom)
09:00–11:00	Technical Session - Simulation (HH 002)
09:00–11:00	Technical Session - Sampling and Markov Chain Monte Carlo (WH Auditorium)
11:00–11:30	Coffee Break (HH Lobby)
11:30–12:30	Plenary Talk by Michaela Szölgyenyi, U of Klagenfurt, An optimal transport approach to quantifying model uncertainty of SDEs (HH Auditorium)
12:30–12:40	Closing Ceremony (HH Auditorium)

Mon, Jul 28, 2025 – Morning

08:00–17:30	Conference Check-In, HH Lobby				
08:45–09:00	Opening Ceremony by Fred Hickernell, Nicole Beebe, and Kevin Corlette, HH Auditorium				
09:00–10:00	Plenary Talk: Rohan Sawhney, Nvidia Corporation, Monte Carlo Methods in Computer Graphics , p. 43 Chair: <i>Michael Mascagni</i> , HH Auditorium				
10:00–10:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session Stochastic Computation and Complexity, Part I, p. 54 Chair: <i>Stefan Heinrich</i>	HH Ballroom Special Session Domain Uncertainty Quantification, p. 55 Chair: <i>André-Alexander Zepernick</i>	HH 002 Special Session Nested expectations: models and estimators, Part I, p. 56 Chair: <i>Arved Bartuska</i>	WH Auditorium Special Session Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part I, p. 57 Chair: <i>Mike Giles</i>	WH 115 Technical Session Markov Chain Monte Carlo Chair: <i>Philippe Gagnon</i>
10:30–11:00	<i>Andreas Neuenkirch</i> , A strong order 1.5 boundary-preserving discretization scheme for scalar SDEs defined in a domain, p. 91	<i>André-Alexander Zepernick</i> , Domain UQ for stationary and time-dependent PDEs using QMC, p. 94	<i>Abdul Lateef Haji Ali</i> , An Adaptive Sampling Algorithm for Level-set Approximation, p. 97	<i>Pieterjan Robbe</i> , Multilevel quasi-Monte Carlo without replications, p. 100	<i>Zhihao Wang</i> , Stereographic Multi-Try Metropolis Algorithms for Heavy-tailed Sampling, p. 180
11:00–11:30	<i>Christopher Rauhögger</i> , An adaptive Milstein-type method for strong approximation of systems of SDEs with a discontinuous drift coefficient, p. 92	<i>Carlos Jerez-Hanckes</i> , Domain Uncertainty Quantification for Electromagnetic Wave Scattering via First-Order Sparse Boundary Element Approximation, p. 95	<i>Vinh Hoang</i> , Posterior-Free A-Optimal Bayesian Design of Experiments via Conditional Expectation, p. 98	<i>Irina-Beatrice Haas</i> , A nested Multilevel Monte Carlo framework for efficient simulations on FPGAs, p. 100	<i>Ruben Seyer</i> , Creating rejection-free samplers by rebalancing skew-balanced jump processes, p. 181
11:30–12:00	<i>Verena Schwarz</i> , Strong order 1 adaptive approximation of jump-diffusion SDEs with discontinuous drift, p. 93	<i>Jürgen Dölz</i> , Quantifying uncertainty in spectral clusterings: expectations for perturbed and incomplete data, p. 96	<i>Vesa Kaarnioja</i> , QMC for Bayesian optimal experimental design with application to inverse problems governed by PDEs, p. 99	<i>Mike Giles</i> , CUDA implementation of MLMC on NVIDIA GPUs, p. 101	<i>Philippe Gagnon</i> , Theoretical guarantees for lifted samplers, p. 182
12:00–12:30	<i>Toni Karvonen</i> , Approximation in Hilbert spaces of the Gaussian and related analytic kernels, p. 93	<i>Harri Hakula</i> , Model Problems for PDEs on Uncertain Domains, p. 97		<i>Chung Ming Loi</i> , Scalable and User-friendly QMC Sampling with UMBridge, p. 102	

Mon, Jul 28, 2025 – Afternoon

12:30–14:00	Lunch Break, MTCC Commons				
14:00–15:00	Plenary Talk: <i>Christiane Lemieux, U of Waterloo, Golden ratio nets and sequences</i> , p. 44 Chair: <i>Nathan Kirk</i> , HH Auditorium				
15:00–15:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session Stochastic Computation and Complexity, Part II, p. 59 Chair: <i>Larisa Yaroslavtseva</i>	HH Ballroom Special Session Recent advances in optimization under uncertainty, p. 60 Chair: <i>Phillip A. Guth</i>	HH 002 Special Session Computational Methods for Low-discrepancy Sampling and Applications, p. 61 Chair: <i>Nathan Kirk</i>	WH Auditorium Technical Session Quasi-Monte Carlo, Part I Chair: <i>Peter Kritzer</i>	WH 115 Technical Session PDEs and SDEs Chair: <i>Håkon Hoel</i>
15:30–16:00	<i>Michael Gnewuch</i> , Optimality of deterministic and randomized QMC-cubatures on several scales of function spaces, p. 103	<i>Tapio Helin</i> , Stability of Expected Utility in Bayesian Optimal Experimental Design, p. 106	<i>François Clément</i> , Searching Permutations for Constructing Low-Discrepancy Point Sets and Investigating the Kritzinger Sequence, p. 109	<i>Christian Weiss</i> , Halton Sequences, Scrambling and the Inverse Star-Discrepancy, p. 192	<i>Leon Wilkosz</i> , Forward Propagation of Low Discrepancy Through McKean–Vlasov Dynamics: From QMC to MLQMC, p. 211
16:00–16:30	<i>Kateryna Pozharska</i> , Optimal designs for function discretization and construction of tight frames, p. 104	<i>Karina Koval</i> , Subspace accelerated measure transport methods for fast and scalable sequential experimental design, p. 107	<i>Nathan Kirk</i> , Minimizing the Stein Discrepancy, p. 110	<i>Sifan Liu</i> , Transport Quasi-Monte Carlo, p. 192	<i>Miguel Alvarez</i> , A New Approach for Unbiased Estimation of Parameters of Partially Observed Diffusions, p. 212
16:30–17:00	<i>Leszek Plaskota</i> , Complexity of approximating piecewise smooth functions in the presence of deterministic or random noise, p. 105	<i>Johannes Milz</i> , Randomized quasi-Monte Carlo methods for risk-averse stochastic optimization, p. 108	<i>Makram Chahine</i> , Improving Efficiency of Sampling-based Motion Planning via Message-Passing Monte Carlo, p. 111	<i>Ambrose Emmett-Iwaniw</i> , Using Normalizing Flows for Efficient Quasi-Random Sampling for Copulas, p. 193	<i>Håkon Hoel</i> , High-order adaptive methods for exit times of diffusion processes and reflected diffusions, p. 212
17:00–17:30	<i>Larysa Matiukha</i> , The Quality of Lattice Sequences, p. 105	<i>Arved Bartuska</i> , Efficient expected information gain estimators based on the randomized quasi-Monte Carlo method, p. 109	<i>Gregory Seljak</i> , An Empirical Evaluation of Robust Estimators for RQMC, p. 112	<i>Claude Hall</i> , Optimization of Kronecker Sequences, p. 194	<i>Thomas Cass</i> , Generative Modelling of Levy Area for High-Order SDE Simulation, p. 213
17:30–19:30	Welcome Reception, HH Lobby				

Tue, Jul 29, 2025 – Morning

08:30–17:30	Registration Desk Open, HH Lobby				
09:00–10:00	Plenary Talk: <i>Peter Glynn, Stanford U, Combining Simulation and Linear Algebra: COSIMLA</i> , p. 45 Chair: <i>Chang-Han Rhee</i> , HH Auditorium				
10:00–10:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session Stochastic Computation and Complexity, Part III, p. 63 Chair: <i>Leszek Plaskota</i>	HH Ballroom Special Session Next-generation optimal experimental design: theory, scalability, and real world impact: Part I, p. 64 Chair: <i>Alen Alexanderian</i>	HH 002 Special Session Heavy-tailed Sampling, p. 66 Chair: <i>Sebastiano Grazzi</i>	WH Auditorium Special Session Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part I, p. 68 Chair: <i>Sou-Cheng Choi</i>	WH 115 Technical Session Bayesian Methods Chair: <i>Hamza Ruzayqat</i>
10:30–11:00	<i>Jean-François Chassagneux</i> , Computing the stationary measure of McKean-Vlasov SDEs, p. 113	<i>Xun Huan</i> , Optimal Pilot Sampling for Multi-fidelity Monte Carlo Methods, p. 115	<i>Sebastiano Grazzi</i> , Parallel computations for Metropolis Markov chains based on Picard maps, p. 118	<i>Jonathan Weare</i> , Functional estimation of the marginal likelihood, p. 121	<i>Lorenzo Nagar</i> , Optimizing Generalized Hamiltonian Monte Carlo for Bayesian Inference applications, p. 183
11:00–11:30	<i>Noufel Frikha</i> , On the convergence of the Euler-Maruyama scheme for McKean-Vlasov SDEs, p. 114	<i>Adrien Corenflos</i> , A recursive Monte Carlo approach to optimal Bayesian experimental design, p. 116	<i>Federica Milinanni</i> , A large deviation principle for Metropolis-Hastings sampling, p. 119	<i>Nikhil Bansal</i> , Randomized QMC Methods via Combinatorial Discrepancy, p. 122	<i>Hamza Ruzayqat</i> , Bayesian Anomaly Detection in Variable-Order and Variable-Diffusivity Fractional Mediums, p. 185
11:30–12:00	<i>Sotirios Sabanis</i> , Wasserstein Convergence of Score-based Generative Models under Semiconvexity and Discontinuous Gradients, p. 114	<i>Ayoub Belhadji</i> , Weighted quantization using MMD: From mean field to mean shift via gradient flows, p. 117	<i>Xingyu Wang</i> , Sharp Characterization and Control of Global Dynamics of SGDs with Heavy Tails, p. 120	<i>Michael Mascagni</i> , The Walk on Spheres Monte Carlo Algorithm for Solving Partial Differential Equations, p. 123	<i>Arghya Datta</i> , Theoretical Guarantees of Mean Field Variational Inference for Bayesian Principal Component Analysis, p. 186
12:00–12:30				<i>Hwanwoo Kim</i> , Enhancing Gaussian Process Surrogates for Optimization and Posterior Approximation via Random Exploration, p. 124	<i>Jimmy Lederman</i> , Bayesian Analysis of Latent Underdispersion Using Discrete Order Statistics, p. 187

Tue, Jul 29, 2025 – Afternoon

12:30–14:00	Lunch Break, On your own				
14:00–15:00	Plenary Talk: <i>Roshan Joseph, Georgia Institute of Technology, Sensitivity and Screening: From Monte Carlo to Experimental Design</i> , p. 46 Chair: <i>Simon Mak</i> , HH Auditorium				
15:00–15:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session Stochastic Computation and Complexity, Part IV, p. 69 Chair: <i>Thomas Müller-Gronbach</i>	HH Ballroom Special Session Next-generation optimal experimental design: theory, scalability, and real world impact: Part II, p. 70 Chair: <i>Xun Huan</i>	HH 002 Special Session Advances in Rare Events Simulation, p. 72 Chair: <i>Shyam Mohan Subbiah Pillai</i>	WH Auditorium Special Session Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part II, p. 73 Chair: <i>Yuhan Ding</i>	WH 115 Technical Session Quasi-Monte Carlo, Part II Chair: <i>Christian Weiss</i>
15:30–16:00	<i>Larisa Yaroslavtseva</i> , Optimal strong approximation of SDEs with Hölder continuous drift coefficient, p. 124	<i>Alen Alexanderian</i> , Goal-Oriented Sensor Placement for Infinite-Dimensional Bayesian Inverse Problems, p. 127	<i>Nicola Branchini</i> , Revisiting self-normalized importance sampling: new methods and diagnostics, p. 130	<i>Takashi Goda</i> , Quasi-uniform quasi-Monte Carlo digital nets, p. 133	<i>Peter Kritzer</i> , Approximation using median lattice algorithms, p. 195
16:00–16:30	<i>Gunther Leobacher</i> , Tractability of L_2 -approximation and integration in weighted Hermite spaces of finite smoothness, p. 125	<i>Jacopo Iollo</i> , Diffusion-Based Bayesian Experimental Design: Advancing BED for Practical Applications, p. 128	<i>Bruno Tuffin</i> , Asymptotic robustness of smooth functions of rare-event estimators, p. 131	<i>Ziang Niu</i> , WITHDRAWN—Boosting the inference for generative models by (Quasi-)Monte Carlo resampling, p. 134	<i>Yang Liu</i> , Convergence Rates of Randomized Quasi-Monte Carlo Methods under Various Regularity Conditions, p. 195
16:30–17:00	<i>Alexander Steinicke</i> , Malliavin differentiation of Lipschitz SDEs and BSDEs and an Application to Quadratic Forward-Backward SDEs, p. 126	<i>Tommie Catanach</i> , Robust Bayesian Optimal Experimental Design under Model Misspecification, p. 129	<i>Eya Ben Amar</i> , Importance Sampling Methods with Stochastic Differential Equations for the Estimation of the Right Tail of the CCDF of the Fade Duration, p. 132	<i>Chenyang Zhong</i> , A hit-and-run approach for sampling and analyzing ranking models, p. 135	<i>Jakob Dilen</i> , Use of rank-1 lattices in the Fourier neural operator, p. 196
17:00–17:30	<i>Fred J. Hickernell</i> , A Unified Treatment of Tractability for Approximation Problems Defined on Hilbert Spaces, p. 126		<i>Shyam Mohan Subbiah Pillai</i> , Estimating rare event probabilities associated with McKean–Vlasov SDEs, p. 132	<i>Di Yu</i> , Interior-Point Frank-Wolfe (IPFW) for Linearly Constrained Functional Optimization Over Probability Spaces, p. 136	<i>Aadit Jain</i> , Investigating the Optimum RQMC Batch Size for Betting and Empirical Bernstein Confidence Intervals, p. 197
18:00–20:00	Chicago White Sox vs. Philadelphia Phillies (must purchase tickets beforehand), Meet in HH Lobby				

Wed, Jul 30, 2025 – Morning

08:30–16:30	Registration Desk Open, HH Lobby				
09:00–10:00	Plenary Talk: <i>Veronika Ročková, U of Chicago, AI-Powered Bayesian Inference</i> , p. 47 Chair: <i>Art B. Owen</i> , HH Auditorium				
10:00–10:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session Stochastic Computation and Complexity, Part V, p. 74 Chair: <i>Andreas Neuenkirch</i>	HH Ballroom Special Session Statistical Design of Experiments, p. 75 Chair: <i>Simon Mak</i>	HH 002 Special Session Advances in Adaptive Hamiltonian Monte Carlo, p. 76 Chair: <i>Art Owen</i>	WH Auditorium Technical Session Simulation Chair: <i>Toon Ingelaere</i>	WH 115 Technical Session Sampling Chair: <i>Nicola Branchini</i>
10:30–11:00	<i>Stefan Heinrich</i> , On the quantum complexity of parametric integration in Sobolev spaces, p. 137	<i>Simon Mak</i> , Respecting the boundaries: Space-filling designs for surrogate modeling with boundary information, p. 139	<i>Bob Carpenter</i> , GIST: Gibbs self-tuning for locally adapting Hamiltonian Monte Carlo, p. 143	<i>Philippe Blondeel</i> , Combining quasi-Monte Carlo with Stochastic Optimal Control for Trajectory Optimization of Autonomous Vehicles in Mine Counter Measure Simulations, p. 216	<i>Joonha Park</i> , Sampling from high-dimensional, multimodal distributions using automatically tuned, tempered Hamiltonian Monte Carlo, p. 198
11:00–11:30	<i>Bernd Käßemodel</i> , Quantum Integration in Tensor Product Besov Spaces, p. 138	<i>Andrews Boahen</i> , Active Learning for Nonlinear Calibration, p. 140	<i>Nawaf Bou-Rabee</i> , Acceleration of the No-U-Turn Sampler, p. 144	<i>Rino Persiani</i> , A Monte Carlo Approach to Designing a Novel Sample Holder for Enhanced UV-Vis Spectroscopy, p. 217	<i>Arne Bouillon</i> , Localized consensus-based sampling for non-Gaussian distributions, p. 199
11:30–12:00	<i>Nikolaos Makras</i> , Taming the Interacting Particle Langevin Algorithm – The Superlinear Case, p. 138	<i>Qian Xiao</i> , Optimal design of experiments with quantitative-sequence factors, p. 141	<i>Chirag Modi</i> , ATLAS: Adapting Trajectory Lengths and Step-Size for Hamiltonian Monte Carlo, p. 145	<i>Prasanth Shyamsundar</i> , ARCANE Reweighting: A technique to tackle the sign problem in the simulation of collider events in high-energy physics, p. 218	
12:00–12:30	<i>Iosif Lytras</i> , Sampling with Langevin Dynamics from non-smooth and non-logconcave potentials., p. 139	<i>Chaofan Huang</i> , Factor Importance Ranking and Selection using Total Indices, p. 142	<i>Trevor Campbell</i> , AutoStep: Locally adaptive involutive MCMC, p. 146	<i>Nicole Aretz</i> , Multifidelity and Surrogate Modeling Approaches for Uncertainty Quantification in Ice Sheet Simulations, p. 219	

Wed, Jul 30, 2025 – Afternoon

12:30–14:00	Lunch Break, On your own				
	HH Auditorium Special Session Stochastic Optimization, p. 78 Chair: <i>Shane Henderson</i>	HH Ballroom Special Session Recent Progress on Algorithmic Discrepancy Theory and Applications, p. 79 Chair: <i>Haotian Jiang</i>	HH 002 Special Session Monte Carlo Applications in High-performance Computing, Computer Graphics, and Computational Science, p. 80 Chair: <i>Michael Mascagni</i>	WH Auditorium Technical Session Statistics Chair: <i>Yiming Xu</i>	
14:00–14:30	<i>Raghu Bollapragada</i> , Monte Carlo Based Adaptive Sampling Approaches for Stochastic Optimization, p. 147	<i>Haotian Jiang</i> , Algorithmic Discrepancy Theory: An Overview, p. 149	<i>Arash Fahim</i> , Gaining efficiency in Monte Carlo policy gradient methods for stochastic optimal control, p. 152	<i>Kazeem Adeleke</i> , Empirical Statistical Comparative Analysis of SNP Heritability Estimators and Gradient Boosting Machines (GBM) Using Genetic Data from the UK Biobank, p. 220	
14:30–15:00	<i>Shane Henderson</i> , A New Convergence Analysis of Two Stochastic Frank-Wolfe Algorithms, p. 148	<i>Peng Zhang</i> , Improving the Design of Randomized Experiments via Discrepancy Theory, p. 150	<i>Sharanya Jayaraman</i> , Examining the Fault Tolerance of High-Performance Monte Carlo Applications through Simulation, p. 153	<i>Carles Domingo-Enrich</i> , Cheap permutation testing, p. 221	
15:00–15:30	<i>Akshita Gupta</i> , Stochastic Gradient with Testing Functionals, p. 149	<i>Aleksandar Nikolov</i> , Online Factorization for Online Discrepancy Minimization, p. 151	<i>Rohan Sawhney</i> , Building Monte Carlo “Renderers” for Physics, p. 154	<i>Christopher Draper</i> , Moving PCG beyond LCGs, p. 222	
15:30–16:00			<i>Silei Song</i> , WoS-NN: Collaborating Walk-on-Spheres with Machine Learning to Solve Elliptic PDEs, p. 155	<i>Yiming Xu</i> , Hybrid least squares for learning functions from highly noisy data, p. 222	
16:00–16:30	Coffee Break, HH Lobby				
18:00–20:30	Conference Banquet, Bridgeport Art Center, 1200 W 35th Street				

Thu, Jul 31, 2025 – Morning

08:30–17:30	Registration Desk Open, HH Lobby				
09:00–10:00	Plenary Talk: Uros Seljak, UC Berkeley, Gradient-Based MCMC Sampling: Methods and Optimization Strategies , p. 48 Chair: <i>Nicolas Chopin</i> , HH Auditorium				
10:00–10:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session QMC and Applications Part I, p. 81 Chair: <i>Michael Gnewuch</i>	HH Ballroom Special Session Analysis of Langevin and Related Sampling Algorithms, Part I, p. 82 Chair: <i>Xiaoou Cheng</i>	HH 002 Special Session Nested expectations: models and estimators, Part II, p. 83 Chair: <i>Abdul-Lateef Haji-Ali</i>	WH Auditorium Technical Session Finance Chair: <i>Aleksei Sorokin</i>	WH 115 Technical Session ML & Optimization Chair: <i>Frédéric Blondeel</i>
10:30–11:00	<i>Felix Bartel</i> , Exact discretization, tight frames and recovery via D-optimal designs, p. 156	<i>Krishnakumar Balasubramanian</i> , Finite-Particle Convergence Rates for Stein Variational Gradient Descent, p. 159	<i>Matteo Raviola</i> , Stochastic gradient with least-squares control variates, p. 162	<i>Vincent Zhang</i> , Characterizing Efficacy of Geometric Brownian Motion Expectation-based Simulations on Low-Volatility American Common Stocks, p. 203	<i>Frédéric Blondeel</i> , Learning cooling strategies in simulated annealing through binary interactions, p. 214
11:00–11:30	<i>Mou Cai</i> , L2-approximation: using randomized lattice algorithms and QMC hyperinterpolation, p. 157	<i>Lihan Wang</i> , Convergence rates of kinetic Langevin dynamics with weakly confining potentials, p. 160	<i>Philipp Guth</i> , A one-shot method for Bayesian optimal experimental design, p. 162	<i>Hao Quan</i> , Efficient Pricing for Variable Annuity via Simulation, p. 204	<i>Du Ouyang</i> , Accuracy of Discretely Sampled Stochastic Policies in Continuous-Time Reinforcement Learning, p. 215
11:30–12:00	<i>Zhijian He</i> , High-dimensional density estimation on unbounded domain, p. 158	<i>Xiaoou Cheng</i> , Delocalization of Bias in Unadjusted Hamiltonian Monte Carlo, p. 161	<i>Sara Pérez-Vieites</i> , Langevin-based strategies for nested particle filters, p. 163		<i>Yiqing Zhou</i> , Minimizing Functions with Sparse Samples: A Fast Interpolation Approach, p. 216
12:00–12:30	<i>Frances Y. Kuo</i> , Application of QMC to Oncology, p. 158				

Thu, Jul 31, 2025 – Afternoon

12:30–14:00	Lunch Break, On your own				
14:00–15:00	Plenary Talk: <i>Nicolas Chopin, Institut Polytechnique de Paris, Saddlepoint Monte Carlo and its application to exact ecological inference</i> , p. 50 Chair: <i>Bruno Tuffin</i> , HH Auditorium				
15:00–15:30	Coffee Break, HH Lobby				
	HH Auditorium Special Session QMC and Applications Part II, p. 84 Chair: <i>Takashi Goda</i>	HH Ballroom Special Session Analysis of Langevin and Related Sampling Algorithms, Part II, p. 85 Chair: <i>Yifan Chen</i>	HH 002 Special Session Recent Advances in Stochastic Gradient Descent, p. 86 Chair: <i>Jing Dong</i>	WH Auditorium Technical Session Sampling Chair: <i>Joonha Park</i>	WH 115 Technical Session SDEs Chair: <i>Fabio Zoccolan</i>
15:30–16:00	<i>Dirk Nuyens</i> , Approximation of multivariate periodic functions, p. 164	<i>Molei Tao</i> , Langevin-Based Sampling under Nonconvex Constraints, p. 166	<i>Jose Blanchet</i> , Inference for Stochastic Gradient Descent with Infinite Variance, p. 169	<i>Josephine Westermann</i> , Polynomial approximation for efficient transport-based sampling, p. 199	<i>Fabio Zoccolan</i> , Dynamical Low-Rank Approximation for SDEs: an interacting particle-system ROM, p. 208
16:00–16:30	<i>Art Owen</i> , Randomized QMC with one categorical variable, p. 164	<i>Yifan Chen</i> , Convergence of Unadjusted Langevin in High Dimensions: Delocalization of Bias, p. 167	<i>Chang-Han Rhee</i> , Exit-Time Analysis of Stochastic Gradient Descent via Kesten’s Recursion, p. 170	<i>Soumyadip Ghosh</i> , Fast Approximate Matrix Inversion via MCMC for Linear System Solvers, p. 200	<i>Adrien Richou</i> , A probabilistic Numerical method for semi-linear elliptic Partial Differential Equations, p. 209
16:30–17:00	<i>Zexin Pan</i> , QMC confidence intervals using quantiles of randomized nets, p. 165	<i>Fuzhong Zhou</i> , Entropy methods for the delocalization of bias in Langevin Monte Carlo, p. 168	<i>Jing Dong</i> , Stochastic Gradient Descent with Adaptive Data, p. 170	<i>Ally Kwan and Lijia Lin</i> , Investigating general L_2 discrepancies with Message-Passing Monte Carlo, p. 201	<i>Anke Wiese</i> , A Chen-Fliess series for stochastic differential equations driven by Lévy processes, p. 209
17:00–17:30	<i>Kosuke Suzuki</i> , Quasi-uniform quasi-Monte Carlo lattice point sets, p. 166	<i>Siddharth Mitra</i> , Convergence of Φ -Divergence and Φ -Mutual Information Along Langevin Markov Chains, p. 168		<i>Jimmy Nguyen and Anders Pride</i> , Discrepancy calculation and generating vector searches for Kronecker and extensible lattice sequences, p. 202	<i>Riccardo Saporiti</i> , Comparing Probabilistic Load Forecasters: Stochastic Differential Equations and Deep Learning, p. 210
18:30–20:30	Steering Committee Meeting (by invitation), Mikami Izakaya & Ramen, 1400 S Michigan Ave				
19:00–21:00	Early Career Dinner, Time Out Market, 916 W Fulton Market				

Fri, Aug 1, 2025

08:30–12:15	Registration Desk Open, HH Lobby				
	HH Auditorium Special Session Forward and Inverse Problems for Stochastic Reaction Networks, p. 87 Chair: <i>Sophia Münker</i>	HH Ballroom Special Session Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part II, p. 88 Chair: <i>Pieterjan Robbe</i>	HH 002 Technical Session Simulation Chair: <i>Nicole Aretz</i>	WH Auditorium Technical Session Sampling and Markov Chain Monte Carlo Chair: <i>Soumyadip Ghosh</i>	
09:00–09:30	<i>Zhou Fang</i> , Fixed-budget simulation method for growing cell populations, p. 171	<i>Niklas Baumgarten</i> , A High-performance Multi-level Monte Carlo Software for Full Field Estimates and Applications in Optimal Control, p. 175	<i>Yashveer Kumar</i> , Monte Carlo simulation approach to solve distributed order fractional mathematical model, p. 188	<i>Daniel Yukimura</i> , Quantitative results on sampling from quasi-stationary distributions, p. 205	
09:30–10:00	<i>Sophia Münker</i> , Dimensionality Reduction for Efficient Rare Event Estimation, p. 172	<i>Aleksei Sorokin</i> , Fast Gaussian Processes, p. 176	<i>Reuben Cohn-Gordon</i> , Gradient-based MCMC in high dimensions, p. 189	<i>Toon Ingelaere</i> , Multilevel simulation of ensemble Kalman methods: interactions across levels, p. 206	
10:00–10:30	<i>Maksim Chupin</i> , Filtered Markovian Projection: Dimensionality Reduction in Filtering for Stochastic Reaction Networks, p. 173	<i>Johannes Krotz</i> , Hybrid Monte Carlo methods for kinetic transport, p. 177	<i>Serena Fattori</i> , Benchmarking the Geant4-DNA 'UHDR' Example for Monte Carlo Simulation of pH Effects on Radiolytic Species Yields Using a Mesoscopic Approach, p. 190	<i>Annabelle Carrell</i> , Low-Rank Thinning, p. 207	
10:30–11:00	<i>Muruhan Rathinam</i> , State and parameter inference in stochastic reaction networks, p. 174	<i>Joseph Farmer</i> , Flow-Based Monte Carlo Transport Simulation, p. 178	<i>Chi-Ok Hwang</i> , First-passage-based Last-passage Algorithm for Charge Density on a Conducting Surface, p. 191		
11:00–11:30	Coffee Break, HH Lobby				
11:30–12:30	Plenary Talk: <i>Michaela Szölgényi, U of Klagenfurt, An optimal transport approach to quantifying model uncertainty of SDEs</i> , p. 52 Chair: <i>Gunther Leobacher</i> , HH Auditorium				
12:30–12:40	Closing Ceremony by Fred Hickernell, HH Auditorium				

Plenary Talks



Monte Carlo Methods in Computer Graphics

Rohan Sawhney

NVIDIA

rsawhney@nvidia.com

Monte Carlo methods have transformed computer graphics and scientific computing alike by providing a versatile framework for solving complex integrals and simulating physical processes. In rendering, Monte Carlo techniques such as path tracing have enabled physically based image synthesis by stochastically simulating light transport across scenes with intricate geometry and materials. Sophisticated algorithms—including Metropolis light transport and ReSTIR—extend this foundation, combining Markov Chain Monte Carlo and resampled importance sampling to achieve high-quality results with manageable computational budgets.

Beyond image generation, Monte Carlo approaches are increasingly applied to solve partial differential equations central to modeling the natural world. Notably, the walk on spheres (WoS) algorithm offers a grid-free alternative to conventional finite element or finite difference solvers. By reformulating PDEs like the Poisson equation as recursive integral equations, WoS leverages random sampling to estimate solutions without requiring volumetric meshing or the assembly of global linear systems. This approach exhibits several powerful advantages—trivial parallelism, the ability to evaluate the solution at arbitrary points, and favorable scaling on modern hardware such as GPUs.

This talk will survey the historical development, mathematical underpinnings, and practical innovations of Monte Carlo methods across these domains, illustrating how ideas from rendering and PDE simulation increasingly cross-pollinate. We will explore how the shared challenges of high-dimensional integration, variance reduction, and efficient sampling drive advances in algorithms that bridge computer graphics and computational physics, blurring the boundary between visualization and simulation. Finally, the talk will highlight emerging frontiers, including differentiable solvers and neural denoising.



Golden ratio nets and sequences

Christiane Lemieux

University of Waterloo

clemieux@uwaterloo.ca

Coauthor(s): Nathan Kirk and Jaspar Wiart

In this talk, we discuss nets and sequences constructed in an irrational base, focusing on the case of a base given by the golden ratio φ . We provide a complete framework to study equidistribution properties of nets in base φ , which among other things requires the introduction of a new concept of prime elementary intervals which differ from the standard definition used for integer bases. We define the one-dimensional van der Corput sequence in base φ and two-dimensional Hammersley point sets in base φ and we prove some properties for $(0, 1)$ -sequences and $(0, m, 2)$ -nets in base φ , respectively. This part of the talk is based on [1].

Building on this new framework, we propose an *interlaced Halton sequence* that makes use of integer *and* irrational-based van der Corput sequences and show empirically improved performance compared to the traditional Halton sequence [2]. In addition, we propose a scrambling algorithm for irrational-based digital sequences, which leverages dependence properties of scrambled digital nets [3].

- [1] N. Kirk, C. Lemieux and J. Wiart. Golden ratio nets and sequences. To appear in *Functiones et Approximatio*, 2025.
- [2] N. Kirk, C. Lemieux. An improved Halton sequence for implementation in quasi-Monte Carlo methods. *Proceedings of the 2024 Winter Simulation Conference*, 431–442, IEEE Press, Piscataway, NJ, 2024.
- [3] C. Lemieux and J. Wiart. On the distribution of scrambled $(0, m, s)$ -nets over unanchored boxes. In: *Monte Carlo and Quasi-Monte Carlo Methods 2020*, A. Keller (ed), Springer, 187-230, 2022.



Combining Simulation and Linear Algebra: COSIMLA

Peter W. Glynn

Stanford University

glynn@stanford.edu

Coauthor(s): Zeyu Zheng

In numerical computation for Markov chains and jump processes, matrix-based linear algebraic methods leverage fully the special structure of such models, allowing one to efficiently compute highly accurate solutions quickly. When the number of states is large or infinite, Monte Carlo simulation is an appealing alternative, typically allowing low-accuracy solutions to be computed efficiently. In this talk, we describe COSIMLA, COmbined SIMulations and Linear Algebra. This new class of algorithms combines the best of the two numerical approaches, using matrix methods to compute expectations and probabilities in the truncated core of the state space, while one uses Monte Carlo to simulate path excursions outside the truncation. As a result, one can now compute high-accuracy solutions for models with a very large state space. We show how the method applies to computing equilibrium quantities and various transient characteristics of Markov chains. These algorithms can typically be viewed as an application of conditional Monte Carlo. We also discuss how stratification can be conveniently applied in this setting to provide further variance reductions.



Sensitivity and Screening: From Monte Carlo to Experimental Design

Roshan Joseph

Georgia Institute of Technology, Atlanta

roshan@gatech.edu

Identifying the most important factors affecting the output of a system from a set of potentially important factors is an important problem in scientific investigations. If a computational model is available to predict the output, we can use global sensitivity analysis to quantify the importance of each factor. There are many Monte Carlo-based methods available to estimate global sensitivity indices. However, their computation can become costly if the model is computationally expensive. In such cases, carefully designed experiments can be used for screening the factors. In this talk, I will explain some of these techniques and the latest developments, including their applications in active learning. I will also briefly explain how to estimate the sensitivity indices from noisy data when we do not know or have access to the model that generated the data.

- [1] Xiao, Q., Joseph, V. R., and Ray, D. M. (2023). *Maximum One-Factor-At-A-Time Designs for Screening in Computer Experiments*. *Technometrics*, 65, 220-230.
- [2] Song, D. and Joseph, V. R. (2025). *Efficient Active Learning Strategies for Computer Experiments*. <https://arxiv.org/abs/2501.13841>.
- [3] Huang, C. and Joseph, V. R. (2025). *Factor Importance Ranking and Selection using Total Indices*. *Technometrics*, <https://doi.org/10.1080/00401706.2025.2483531>.



AI-Powered Bayesian Inference

Veronika Ročková

University of Chicago

veronika.rockova@chicagobooth.edu

Coauthor(s): Sean O'Hagan

The advent of Generative Artificial Intelligence (GAI) has heralded an inflection point that changed how society thinks about knowledge acquisition. While GAI cannot be fully trusted for decision-making, it may still provide valuable information that can be integrated into a decision pipeline. Rather than seeing the lack of certitude and inherent randomness of GAI as a problem, we view it as an opportunity. Indeed, variable answers to given prompts can be leveraged to construct a prior distribution which reflects assuredness of AI predictions. This prior distribution may be combined with tailored datasets for a fully Bayesian analysis with an AI-driven prior. In this paper, we explore such a possibility within a non-parametric Bayesian framework. The basic idea consists of assigning a Dirichlet process prior distribution on the data-generating distribution with AI generative model as its baseline. Hyperparameters of the prior can be tuned out-of-sample to assess the informativeness of the AI prior. Posterior simulation is achieved by computing a suitably randomized functional on an augmented data that consists of observed (labeled) data as well as fake data whose labels have been imputed using AI. This strategy can be parallelized and rapidly produces iid samples from the posterior by optimization as opposed to sampling from conditionals. Our method enables (predictive) inference and uncertainty quantification leveraging AI predictions in a coherent probabilistic manner.



Gradient-Based MCMC Sampling: Methods and Optimization Strategies

Uroš Seljak

UC Berkeley and Lawrence Berkeley National Laboratory

useljak@berkeley.edu

Coauthor(s): Reuben Cohn-Gordon, Jakob Robnik

Gradient-based Markov Chain Monte Carlo (MCMC) methods significantly outperform gradient-free alternatives in sampling efficiency, particularly in high-dimensional spaces where they have become the standard approach. These methods leverage gradient information to guide the sampling process more intelligently than random-walk approaches. Two fundamental approaches that dominate this field are 1) Hamiltonian Monte Carlo (HMC), which employs principles from classical mechanics, treating the sampling problem as simulating Hamiltonian dynamics on an extended phase space. This approach naturally incorporates momentum variables that help the sampler traverse the parameter space more efficiently than simple random walks. 2) Langevin Monte Carlo (LMC), which utilizes stochastic differential equations that incorporate both gradient information and controlled noise injection. Recent theoretical developments have produced microcanonical versions of both Hamiltonian and Langevin samplers (MCHMC and MCLMC). These variants demonstrate measurably superior sampling efficiency compared to their canonical predecessors.

In addition to the choice of the method, practitioners face numerous algorithmic choices that can significantly impact performance: 1) Metropolis Adjustment: The decision whether to include Metropolis-Hastings correction steps involves trading exact preservation of the target distribution against computational speed. 2) Preconditioning: Incorporating problem-specific geometric information through preconditioning matrices can dramatically improve convergence rates, particularly for ill-conditioned target distributions. 3) Hyperparameter Tuning: Critical parameters include step sizes, trajectory lengths for HMC, and damping coefficients for Langevin methods. Recently, well tuned black-box methods have been developed that approach optimal performance. 4) Parallelization Strategy: parallel sampling on a GPU or CPU cluster enables dramatically reduced wall clock time to reach the required target accuracy. 5) Numerical Integration: Higher-order integrators can improve accuracy at the cost of additional gradient evaluations per step.

The goal of this talk is to provide guidance to the optimal choice among these methods,

which depends on specific application requirements including computational budget, accuracy demands, and problem dimensionality. Understanding the theoretical trade-offs enables practitioners to select and configure samplers that best match their particular constraints and objectives.



Saddlepoint Monte Carlo and its Application to Exact Ecological Inference

Nicolas Chopin

ENSAE, Institut Polytechnique de Paris

nicolas.chopin@ensae.fr

Coauthor(s): Théo Voldoire, Guillaume Rateau, Robin J. Ryder

In ecological inference, one wishes to model individual items, but perform inference based only on aggregate data. For instance, in two-round elections, we are interested in the behaviour of individual voters, but only have access to aggregate vote numbers at each precinct. We develop an exact method for a large class of Ecological Inference Bayesian models, which scales to the large data setting. Our approach solves a more general problem: assuming X is a random vector and A a non-invertible matrix, one sometimes need to perform inference while only having access to samples of $Y = AX$. The corresponding likelihood is typically intractable. One may still be able to perform exact Bayesian inference using a pseudo-marginal sampler, but this requires an unbiased estimator of the intractable likelihood.

We propose saddlepoint Monte Carlo, a method for obtaining an unbiased estimate of the density of Y with very low variance, for any model belonging to an exponential family. Our method relies on importance sampling and characteristic functions, with insights brought by the standard saddlepoint approximation scheme with exponential tilting. We show that saddlepoint Monte Carlo makes it possible to perform exact inference on particularly challenging problems and datasets. We present a study of the carryover of votes between the two rounds of various French elections, using the finest available data (number of votes for each candidate in about 60,000 polling stations over most of the French territory).

We show that existing, popular approximate methods for ecological inference can lead to substantial bias; saddlepoint Monte Carlo is immune from this bias, and can handle ecological inference in the large data framework. We also present original results for the 2024 legislative elections on political centre-to-left and left-to-centre conversion rates when the far-right is present in the second round. Finally, we discuss other exciting applications for saddlepoint Monte Carlo in privacy and inverse problems, such as dealing with inference with empirical quantiles for continuous data.

[1] Voldoire T., Chopin N., Rateau G. and Ryder R.J. (2024). Monte Carlo and its

Application to Exact Ecological Inference, *arxiv 2410.18243*, <https://arxiv.org/abs/2410.18243>,



An optimal transport approach to quantifying model uncertainty of SDEs

Michaela Szölgényi

University of Klagenfurt

michaela.szoelgyenyi@aau.at

Coauthor(s): Benjamin A. Robinson

A fundamental question in stochastic modelling is that of quantifying the effects of model uncertainty. In this context it is of interest to compute a distance between different stochastic models. A reasonable choice of distance is a modification of the Wasserstein distance on the space of probability measures called the adapted Wasserstein distance, as it appears in bicausal optimal transport.

We solve constrained optimal transport problems in which the marginal laws are given by the laws of solutions of stochastic differential equations (SDEs). We consider SDEs with irregular coefficients, making only minimal regularity assumptions. Numerical methods are employed as a theoretical tool to bound the adapted Wasserstein distance. This opens the door for computing the adapted Wasserstein distance in a simple way. We show that this method can be applied to quantifying model uncertainty in stochastic optimisation problems.

Our approach successfully brings together optimal transport and numerical analysis of SDEs.

Special Sessions

Stochastic Computation and Complexity, Part I

Organizers:

Stefan Heinrich

RPTU Kaiserslautern-Landau

heinrich@informatik.uni-kl.de

Thomas Müller-Gronbach

University of Passau

Thomas.Mueller-Gronbach@uni-passau.de

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Session Description:

The session is devoted to algorithms and complexity for

— quadrature and strong approximation of SDEs and SPDEs, in particular under nonstandard assumptions,

— high and infinite dimensional integration and approximation, and

— stochastic optimization and neural networks, including connections to functional analysis and stochastic analysis.

Mon, Jul 28, 2025–Morning, 10:30 – 12:30, HH Auditorium

Andreas Neuenkirch

A Strong Order 1.5 Boundary-Preserving Discretization Scheme for Scalar SDEs Defined in a Domain p. [91](#)

Christopher Rauhögger

An Adaptive Milstein-Type Method for Strong Approximation of Systems of SDEs With a Discontinuous Drift Coefficient p. [92](#)

Verena Schwarz

Strong Order 1 Adaptive Approximation of Jump-Diffusion SDEs With Discontinuous Drift p. [93](#)

Toni Karvonen

Approximation in Hilbert Spaces of the Gaussian and Related Analytic Kernels p. [93](#)

Domain Uncertainty Quantification

Organizers:

André-Alexander Zepernick
Free University of Berlin
a.zepernick@fu-berlin.de

Philipp A. Guth
RICAM, Austrian Academy of Sciences
philipp.guth@ricam.oeaw.ac.at

Vesa Kaarnioja
Free University of Berlin
vesa.kaarnioja@fu-berlin.de

Session Description:

Uncertainty in computational measurement models poses significant challenges in engineering and applied mathematics, where inaccuracies in material properties or geometric domains can greatly impact outcomes. Geometric errors, such as manufacturing imperfections or improper modeling in applications like electronic design and tomography, can be the dominant error contributor. Some approaches to modeling domain uncertainty include homogenization, perturbation, and reference mapping techniques, which facilitate the analysis of uncertainty propagation within computational measurement models. This session brings together leading experts to present recent theoretical and computational advancements in the study of domain uncertainty quantification.

Mon, Jul 28, 2025–Morning, 10:30 – 12:30, HH Ballroom

André-Alexander Zepernick
Domain UQ for Stationary and Time-Dependent PDEs Using QMC p. 94

Carlos Jerez-Hanckes
Domain Uncertainty Quantification for Electromagnetic Wave Scattering via First-Order Sparse Boundary Element Approximation p. 95

Jürgen Dölz
Quantifying Uncertainty in Spectral Clusterings: Expectations for Perturbed and Incomplete Data p. 96

Harri Hakula
Model Problems for PDEs on Uncertain Domains p. 97

Nested Expectations: Models and Estimators, Part I

Organizers:

Arved Bartuska

King Abdullah University of Science and Technology/RWTH Aachen University

arved.bartuska@kaust.edu.sa

Abdul-Lateef Haji-Ali

Heriot-Watt University

a.hajiali@hw.ac.uk

Session Description:

Nested expectations arise in many applications, such as in engineering, mathematical finance, and medical decision-making. In addition to their nested structure, numerical estimations of such expectations are often complicated by singularities or discontinuities. Moreover, approximations when evaluating inner expectations using, for example, finite element or time-stepping schemes render traditional estimation methods such as double-loop Monte Carlo prohibitively expensive. This session will explore models and applications with this structure and methods for efficient estimation.

Mon, Jul 28, 2025–Morning, 10:30 – 12:30, HH 002

Abdul Lateef Haji Ali

An Adaptive Sampling Algorithm for Level-Set Approximation p. 97

Vinh Hoang

Posterior-Free A-Optimal Bayesian Design of Experiments via Conditional Expectation p. 98

Vesa Kaarnioja

QMC for Bayesian Optimal Experimental Design With Application to Inverse Problems Governed by PDEs p. 99

Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part I

Organizers:

Sou-Cheng T. Choi

Illinois Institute of Technology

schoi32@iit.edu

Pieterjan Robbe

Sandia National Laboratories

pmrobbe@sandia.gov

Mike Giles

University of Oxford

mike.giles@maths.ox.ac.uk

Session Description:

Monte Carlo (MC) or quasi-Monte Carlo (QMC) algorithms are widely used in various fields such as finance, physics, and engineering for their ability to handle high-dimensional integration problems. The development and maintenance of software for (quasi-)Monte Carlo ((Q)MC) algorithms can significantly enhance the accessibility and usability of these techniques. This special session aims to bring together experts from academia and industry to discuss recent advances in (Q)MC software, share best practices, and explore future directions, fostering collaboration among researchers and practitioners. Topics of interest for the session include:

- Novel hardware or architectural designs for open-source (Q)MC libraries.
- Best collaborative practices for developing and maintaining efficient and reliable (Q)MC software.
- Challenges and opportunities in integrating (Q)MC methods with machine learning and AI techniques.
- High-performance computing solutions for (Q)MC software.
- Adaptation of (Q)MC software to application fields such as finance, computer graphics, sensitivity analysis, Bayesian optimization, and uncertainty quantification.
- Innovative approaches to enhancing and extending existing (Q)MC tools.

Mon, Jul 28, 2025–Morning, 10:30 – 12:30, WH Auditorium

Pieterjan Robbe

Multilevel Quasi-Monte Carlo Without Replications p. [100](#)

Irina-Beatrice Haas

A Nested Multilevel Monte Carlo Framework for Efficient Simulations on FPGAs p. [100](#)

Mike Giles

CUDA Implementation of MLMC on NVIDIA GPUs p. [101](#)

Chung Ming Loi

Scalable and User-Friendly QMC Sampling With UMBridge p. [102](#)

Stochastic Computation and Complexity, Part II

Organizers:

Stefan Heinrich

RPTU Kaiserslautern-Landau

heinrich@informatik.uni-kl.de

Thomas Müller-Gronbach

University of Passau

Thomas.Mueller-Gronbach@uni-passau.de

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Session Description:

The session is devoted to algorithms and complexity for

— quadrature and strong approximation of SDEs and SPDEs, in particular under nonstandard assumptions,

— high and infinite dimensional integration and approximation, and

— stochastic optimization and neural networks, including connections to functional analysis and stochastic analysis.

Mon, Jul 28, 2025–Afternoon, 15:30 – 17:30, HH Auditorium

Michael Gnewuch

Optimality of Deterministic and Randomized QMC-Cubatures on Several Scales of Function Spaces p. [103](#)

Kateryna Pozharska

Optimal Designs for Function Discretization and Construction of Tight Frames p. [104](#)

Leszek Plaskota

Complexity of Approximating Piecewise Smooth Functions in the Presence of Deterministic or Random Noise p. [105](#)

Larysa Matiukha

The Quality of Lattice Sequences p. [105](#)

Recent Advances in Optimization Under Uncertainty

Organizers:

Philipp A. Guth

RICAM, Austrian Academy of Sciences

philipp.guth@ricam.oeaw.ac.at

Vesa Kaarnioja

Free University of Berlin

vesa.kaarnioja@fu-berlin.de

Session Description:

The quantification of uncertainties associated with large-scale optimization problems based on partial differential equation models typically involves a number of uncertainties. Some uncertain parameters may include, e.g., the material parameters, domain shape or sensor locations used to collect measurement data. The associated challenging high-dimensional integration problems can be solved efficiently using, e.g., multilevel Monte Carlo or quasi-Monte Carlo methods. This session aims to cover some recent developments in the computational and theoretical treatment of this actively developing field of research.

Mon, Jul 28, 2025–Afternoon, 15:30 – 17:30, HH Ballroom

Tapio Helin

Stability of Expected Utility in Bayesian Optimal Experimental Design p. 106

Karina Koval

Subspace Accelerated Measure Transport Methods for Fast and Scalable Sequential Experimental Design p. 107

Johannes Milz

Randomized Quasi-Monte Carlo Methods for Risk-Averse Stochastic Optimization p. 108

Arved Bartuska

Efficient Expected Information Gain Estimators Based on the Randomized Quasi-Monte Carlo Method p. 109

Computational Methods for Low-Discrepancy Sampling and Applications

Organizers:

Nathan Kirk

Illinois Institute of Technology

nkirk@iit.edu

François Clément

University of Washington

fclement@uw.edu

Session Description:

This session aims to showcase recent advancements in the optimization of sample point distributions [1, 2] and their applications. Some of the methods on display will range from deep learning methods to permutation optimization and greedy approaches [3], showcasing the usefulness of the L_2 -discrepancies in optimizing the L_∞ discrepancies. As a consequence of some of these improved low-discrepancy sets, an application will be shown in improved path planning in robotics [4]. Several other applications will be explored in the context of using the median over the mean of r RQMC estimates as proposed in several recent papers including [5].

- [1] T. K. Rusch, N. Kirk, M. Bronstein, C. Lemieux and D. Rus, *Message-Passing Monte Carlo: Generating low-discrepancy point sets via graph neural networks*, PNAS **121** (40) e2409913121 (2024)
- [2] F. Clément, C. Doerr, K. Klamroth, L. Paquete, *Transforming the Challenge of Constructing Low-Discrepancy Point Sets into a Permutation Selection Problem*, <https://arxiv.org/abs/2407.11533>.
- [3] F. Clément, *Outperforming the Best 1D Low-Discrepancy Constructions with a Greedy Algorithm*, <https://arxiv.org/abs/2406.18132>.
- [4] M. Chahine, T. K. Rusch, Z. J. Patterson and D. Rus, *Improving Efficiency of Sampling-based Motion Planning via Message-Passing Monte Carlo*, <https://arxiv.org/abs/2410.03909>
- [5] P. L'Ecuyer, M. K. Nayakama, A. B. Owen and B. Tuffin, *Confidence Intervals for Randomized Quasi-Monte Carlo Estimators*, Proceedings of the 2023 Winter Simulation Conference (2023)

Mon, Jul 28, 2025–Afternoon, 15:30 – 17:30, HH 002

François Clément

Searching Permutations for Constructing Low-Discrepancy Point Sets and Investigating the Kritzing Sequence p. 109

Nathan Kirk

Minimizing the Stein Discrepancy p. [110](#)

Makram Chahine

Improving Efficiency of Sampling-Based Motion Planning via Message-Passing Monte Carlo
p. [111](#)

Gregory Seljak

An Empirical Evaluation of Robust Estimators for RQMC p. [112](#)

Stochastic Computation and Complexity, Part III

Organizers:

Stefan Heinrich

RPTU Kaiserslautern-Landau

heinrich@informatik.uni-kl.de

Thomas Müller-Gronbach

University of Passau

Thomas.Mueller-Gronbach@uni-passau.de

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Session Description:

The session is devoted to algorithms and complexity for

— quadrature and strong approximation of SDEs and SPDEs, in particular under nonstandard assumptions,

— high and infinite dimensional integration and approximation, and

— stochastic optimization and neural networks, including connections to functional analysis and stochastic analysis.

Tue, Jul 29, 2025–Morning, 10:30 – 12:30, HH Auditorium

Jean-François Chassagneux

Computing the Stationary Measure of McKean-Vlasov SDEs p. [113](#)

Noufel Frikha

On the Convergence of the Euler-Maruyama Scheme for McKean-Vlasov SDEs p. [114](#)

Sotirios Sabanis

Wasserstein Convergence of Score-Based Generative Models Under Semiconvexity and Discontinuous Gradients p. [114](#)

Next-Generation Optimal Experimental Design: Theory, Scalability, and Real World Impact: Part I

Organizers:

Florence Forbes

Inria, France

florence.forbes@inria.fr

Xun Huan

University of Michigan, USA

xhuan@umich.edu

Youssef Marzouk

Massachusetts Institute of Technology, USA

ymarz@mit.edu

Session Description:

Optimal experimental design (OED) provides a mathematical framework for identifying candidate data or experimental configurations that are most useful for inference, prediction, or some other downstream goal. Though OED is hardly a new topic,¹ the need for advances in OED has never been greater than it is today. Myriad application areas have witnessed, on the one hand, an explosion in the volume of data that can be acquired, and on the other, the use of increasingly complex and computationally intensive models to interpret these data. Yet large volumes of data do not by default carry a commensurate amount of information. Moreover, we inevitably face constraints: on the costs of experimentation or data acquisition, on data storage and communication, and on the computational effort of statistical inference in complex models. OED directly addresses the associated trade-offs—e.g., between experimentation, measurement, and/or processing *costs* and the *quality* of subsequent and decision making. See *e.g.* [1] for a recent review of OED topics, which provides numerous other references. This session will highlight computational developments at the OED research frontier. Methods for stochastic simulation, high-dimensional approximation or integration, and stochastic optimization are central to modern OED and to the scaling of OED to large parameter spaces and complex statistical models. Modern machine learning methodologies—from neural network surrogates, to deep reinforcement learning in sequential OED, to modern generative models and transport methods for simulation-based inference—also play a catalyzing role in such OED approaches. Talks in this session will illuminate these emerging interactions and their role in realizing Bayesian, decision-theoretic, and information-theoretic formulations of OED for truly complex problems. Session speakers will also discuss ongoing work to develop theoretical guarantees for these new OED methodologies, and showcase applications to real-world problems ranging from sensor steering to seismology.

¹At the first session of the Indian Statistical Conference in 1938, R. Fisher supposedly said, “To call in the statistician after the experiment is done may be no more than asking him to perform a postmortem examination: he may be able to say what the experiment died of.”

- [1] X. Huan, J. Jagalur, Y. Marzouk (2024). Optimal experimental design: Formulations and computations, *Acta Numerica* **33**, 715–840.

Tue, Jul 29, 2025–Morning, 10:30 – 12:30, HH Ballroom

Xun Huan

Optimal Pilot Sampling for Multi-Fidelity Monte Carlo Methods p. [115](#)

Adrien Corenflos

A Recursive Monte Carlo Approach to Optimal Bayesian Experimental Design p. [116](#)

Ayoub Belhadji

Weighted Quantization Using MMD: From Mean Field to Mean Shift via Gradient Flows
p. [117](#)

Heavy-Tailed Sampling

Organizers:

Alex Shestopaloff

Queen Mary University of London, UK

a.shestopaloff@qmul.ac.uk

Jun Yang

University of Copenhagen, Denmark

jy@math.ku.dk

Session Description:

Heavy-tailed distributions frequently arise in modern statistics, machine learning, and applied sciences, yet their intricate properties pose significant challenges for computational inference. This session, Heavy-Tailed Sampling, brings together cutting-edge advances in Monte Carlo methods and stochastic optimization to address these challenges. The topics span theoretical breakthroughs and practical algorithms, showcasing how heavy-tailed phenomena influence algorithmic design and performance. Our session aims to inspire new methods and applications in the broader Monte Carlo community. The session will explore:

- Langevin Monte Carlo for Heavy-Tailed Distributions: A comprehensive complexity analysis of Langevin-based samplers for heavy-tailed targets using weighted Poincaré inequalities. The results reveal fundamental limits of mean-square analysis and include innovative techniques for gradient approximation.
- Stereographic MCMC: A novel class of samplers that map Euclidean spaces onto spheres to resolve mixing issues inherent to heavy-tailed targets. These methods, featuring uniform ergodicity and rapid convergence, capitalize on the "blessings of dimensionality" to enhance performance in high dimensions.
- Large Deviation Principles in MCMC: A groundbreaking application of large deviation theory to assess and improve MCMC algorithms. This approach extends to Metropolis-Hastings and related methods on general state spaces, providing new insights into empirical measure convergence and rate functions.
- Heavy-Tailed Phenomena in Stochastic Gradient Descent (SGD): An analysis of heavy-tailed noise in SGD and its impact on escaping sharp minima in deep learning. The session highlights a variant of SGD with gradient truncation, offering theoretical and empirical evidence of enhanced generalization through flatter minima.

Tue, Jul 29, 2025–Morning, 10:30 – 12:30, HH 002

Sebastiano Grazi

Parallel Computations for Metropolis Markov Chains Based on Picard Maps

p. 118

Federica Milinanni

A Large Deviation Principle for Metropolis-Hastings Sampling p. 119

Xingyu Wang

Sharp Characterization and Control of Global Dynamics of SGDs With Heavy Tails p. 120

Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part I

Organizers:

Sou-Cheng T. Choi

Illinois Institute of Technology

schoi32@iit.edu

Yuhan Ding

Illinois Institute of Technology

yding2@iit.edu

Session Description:

(Quasi-)Monte Carlo ((Q)MC) and Markov Chain Monte Carlo (MCMC) algorithms are fundamental tools in computational mathematics, with a wide range of applications spanning finance, physics, engineering, and more. These methods have proven invaluable in solving high-dimensional problems where traditional numerical techniques often fail, and continue to expand their reach into emerging fields such as artificial intelligence, climate modeling, precision medicine, and data science. Recent advances in the theoretical foundations of these methods, including convergence rates, complexity analysis, sampling techniques, error analysis, variance reduction, optimal stopping conditions, and ergodic properties, have significantly improved their accuracy, efficiency, and reliability. In addition, interdisciplinary applications are driving new developments, such as machine learning, Bayesian inference, stochastic optimization, and uncertainty quantification. This special session aims to bring together leading experts from academia and industry to share breakthroughs, foster interdisciplinary collaboration, and identify future research directions in the broad field of Monte Carlo methods. Participants will benefit from insights into cutting-edge research and practical applications of (Q)MC and MCMC methods, as well as opportunities to network with peers and thought leaders.

Tue, Jul 29, 2025–Morning, 10:30 – 12:30, WH Auditorium

Jonathan Weare

Functional Estimation of the Marginal Likelihood p. [121](#)

Nikhil Bansal

Randomized QMC Methods via Combinatorial Discrepancy p. [122](#)

Michael Mascagni

The Walk on Spheres Monte Carlo Algorithm for Solving Partial Differential Equations
p. [123](#)

Hwanwoo Kim

Enhancing Gaussian Process Surrogates for Optimization and Posterior Approximation via
Random Exploration p. [124](#)

Stochastic Computation and Complexity, Part IV

Organizers:

Stefan Heinrich

RPTU Kaiserslautern-Landau

heinrich@informatik.uni-kl.de

Thomas Müller-Gronbach

University of Passau

Thomas.Mueller-Gronbach@uni-passau.de

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Session Description:

The session is devoted to algorithms and complexity for

— quadrature and strong approximation of SDEs and SPDEs, in particular under nonstandard assumptions,

— high and infinite dimensional integration and approximation, and

— stochastic optimization and neural networks, including connections to functional analysis and stochastic analysis.

Tue, Jul 29, 2025–Afternoon, 15:30 – 17:30, HH Auditorium

Larisa Yaroslavtseva

Optimal Strong Approximation of SDEs With Hölder Continuous Drift Coefficient p. [124](#)

Gunther Leobacher

Tractability of L_2 -Approximation and Integration in Weighted Hermite Spaces of Finite Smoothness p. [125](#)

Alexander Steinicke

Malliavin Differentiation of Lipschitz SDEs and BSDEs and an Application to Quadratic Forward-Backward SDEs p. [126](#)

Fred J. Hickernell

A Unified Treatment of Tractability for Approximation Problems Defined on Hilbert Spaces p. [126](#)

Next-Generation Optimal Experimental Design: Theory, Scalability, and Real World Impact: Part II

Organizers:

Florence Forbes

Inria, France

florence.forbes@inria.fr

Xun Huan

University of Michigan, USA

xhuan@umich.edu

Youssef Marzouk

Massachusetts Institute of Technology, USA

ymarz@mit.edu

Session Description:

Optimal experimental design (OED) provides a mathematical framework for identifying candidate data or experimental configurations that are most useful for inference, prediction, or some other downstream goal. Though OED is hardly a new topic,² the need for advances in OED has never been greater than it is today. Myriad application areas have witnessed, on the one hand, an explosion in the volume of data that can be acquired, and on the other, the use of increasingly complex and computationally intensive models to interpret these data. Yet large volumes of data do not by default carry a commensurate amount of information. Moreover, we inevitably face constraints: on the costs of experimentation or data acquisition, on data storage and communication, and on the computational effort of statistical inference in complex models. OED directly addresses the associated trade-offs—e.g., between experimentation, measurement, and/or processing *costs* and the *quality* of subsequent and decision making. See *e.g.* [1] for a recent review of OED topics, which provides numerous other references. This session will highlight computational developments at the OED research frontier. Methods for stochastic simulation, high-dimensional approximation or integration, and stochastic optimization are central to modern OED and to the scaling of OED to large parameter spaces and complex statistical models. Modern machine learning methodologies—from neural network surrogates, to deep reinforcement learning in sequential OED, to modern generative models and transport methods for simulation-based inference—also play a catalyzing role in such OED approaches. Talks in this session will illuminate these emerging interactions and their role in realizing Bayesian, decision-theoretic, and information-theoretic formulations of OED for truly complex problems. Session speakers will also discuss ongoing work to develop theoretical guarantees for these new OED methodologies, and showcase applications to real-world problems ranging from sensor steering to seismology.

²At the first session of the Indian Statistical Conference in 1938, R. Fisher supposedly said, “To call in the statistician after the experiment is done may be no more than asking him to perform a postmortem examination: he may be able to say what the experiment died of.”

- [1] X. Huan, J. Jagalur, Y. Marzouk (2024). Optimal experimental design: Formulations and computations, *Acta Numerica* **33**, 715–840.

Tue, Jul 29, 2025–Afternoon, 15:30 – 17:30, HH Ballroom

Alen Alexanderian

Goal-Oriented Sensor Placement for Infinite-Dimensional Bayesian Inverse Problems p. 127

Jacopo Iollo

Diffusion-Based Bayesian Experimental Design: Advancing BED for Practical Applications
p. 128

Tommie Catanach

Robust Bayesian Optimal Experimental Design Under Model Misspecification p. 129

Advances in Rare Events Simulation

Organizers:

Nadhir Ben Rached

University of Leeds, United Kingdom

N.BenRached@leeds.ac.uk

Shyam Mohan Subbiah Pillai

RWTH Aachen University, Germany

subbiah@uq.rwth-aachen.de

Raúl Tempone

King Abdullah University of Science and Technology, Saudi Arabia

raul.tempone@kaust.edu.sa

Session Description:

Rare events are events with small probabilities, but their occurrences are critical in many real-life applications. The problem of estimating rare event probabilities is encountered in various engineering applications (finance, wireless communications, system reliability, biology, etc.). Naive Monte Carlo simulations are, in this case, substantially expensive. This session focuses on advances in methods belonging to the class of variance reduction techniques. These alternative methods deliver, when appropriately used, accurate estimates with a substantial amount of variance reduction compared to the naive Monte Carlo estimator.

Tue, Jul 29, 2025–Afternoon, 15:30 – 17:30, HH 002

Nicola Branchini

Revisiting Self-Normalized Importance Sampling: New Methods and Diagnostics p. 130

Bruno Tuffin

Asymptotic Robustness of Smooth Functions of Rare-Event Estimators p. 131

Eya Ben Amar

Importance Sampling Methods With Stochastic Differential Equations for the Estimation of the Right Tail of the CCDF of the Fade Duration p. 132

Shyam Mohan Subbiah Pillai

Estimating Rare Event Probabilities Associated With McKean--Vlasov SDEs p. 132

Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods, Part II

Organizers:

Sou-Cheng T. Choi

Illinois Institute of Technology

schoi32@iit.edu

Yuhan Ding

Illinois Institute of Technology

yding2@iit.edu

Session Description:

(Quasi-)Monte Carlo ((Q)MC) and Markov Chain Monte Carlo (MCMC) algorithms are fundamental tools in computational mathematics, with a wide range of applications spanning finance, physics, engineering, and more. These methods have proven invaluable in solving high-dimensional problems where traditional numerical techniques often fail, and continue to expand their reach into emerging fields such as artificial intelligence, climate modeling, precision medicine, and data science. Recent advances in the theoretical foundations of these methods, including convergence rates, complexity analysis, sampling techniques, error analysis, variance reduction, optimal stopping conditions, and ergodic properties, have significantly improved their accuracy, efficiency, and reliability. In addition, interdisciplinary applications are driving new developments, such as machine learning, Bayesian inference, stochastic optimization, and uncertainty quantification. This special session aims to bring together leading experts from academia and industry to share breakthroughs, foster interdisciplinary collaboration, and identify future research directions in the broad field of Monte Carlo methods. Participants will benefit from insights into cutting-edge research and practical applications of (Q)MC and MCMC methods, as well as opportunities to network with peers and thought leaders.

Tue, Jul 29, 2025—Afternoon, 15:30 – 17:30, WH Auditorium

Takashi Goda

Quasi-Uniform Quasi-Monte Carlo Digital Nets p. [133](#)

Ziang Niu

WITHDRAWN---Boosting the Inference for Generative Models by (Quasi-)Monte Carlo Resampling p. [134](#)

Chenyang Zhong

A Hit-And-Run Approach for Sampling and Analyzing Ranking Models p. [135](#)

Di Yu

Interior-Point Frank-Wolfe (IPFW) for Linearly Constrained Functional Optimization Over Probability Spaces p. [136](#)

Stochastic Computation and Complexity, Part V

Organizers:

Stefan Heinrich

RPTU Kaiserslautern-Landau

heinrich@informatik.uni-kl.de

Thomas Müller-Gronbach

University of Passau

Thomas.Mueller-Gronbach@uni-passau.de

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Session Description:

The session is devoted to algorithms and complexity for

— quadrature and strong approximation of SDEs and SPDEs, in particular under nonstandard assumptions,

— high and infinite dimensional integration and approximation, and

— stochastic optimization and neural networks, including connections to functional analysis and stochastic analysis.

Wed, Jul 30, 2025–Morning, 10:30 – 12:30, HH Auditorium

Stefan Heinrich

On the Quantum Complexity of Parametric Integration in Sobolev Spaces p. [137](#)

Bernd Käßemodel

Quantum Integration in Tensor Product Besov Spaces p. [138](#)

Nikolaos Makras

Taming the Interacting Particle Langevin Algorithm -- The Superlinear Case p. [138](#)

Iosif Lytras

Sampling With Langevin Dynamics From Non-Smooth and Non-Logconcave Potentials.
p. [139](#)

Statistical Design of Experiments

Organizers:

Lulu Kang

University of Massachusetts Amherst

lulukang@umass.edu

Chunfang Devon Lin

Queen's University

devon.lin@queensu.ca

Session Description:

This session explores innovative methodologies for optimizing experimental design and factor analysis in complex, high-dimensional, and resource-constrained settings. The first talk introduces QuIP, a novel framework for designing experiments with qualitative factors using integer programming and Gaussian process models, demonstrating its effectiveness in path planning and rover trajectory optimization. The second talk addresses the challenge of cost-efficient predictive computing by proposing a multi-fidelity emulator design inspired by Multilevel Monte Carlo methods, which ensures predictive accuracy while minimizing computational costs under a tight budget. The third talk shifts focus to experiments involving both quantitative and sequence factors, presenting a new class of optimal quantitative-sequence (QS) designs that are flexible, space-filling, and asymptotically orthogonal, making them ideal for high-dimensional applications in medical science and bio-engineering. Finally, the fourth talk introduces FIRST, a model-free framework for factor importance ranking and selection using total Sobol' indices, offering a computationally efficient and consistent approach to identifying key factors in regression and classification tasks. Together, these talks highlight cutting-edge advancements in experimental design, optimization, and factor analysis, with broad applications across scientific and engineering disciplines.

Wed, Jul 30, 2025–Morning, 10:30 – 12:30, HH Ballroom

Simon Mak

Respecting the Boundaries: Space-Filling Designs for Surrogate Modeling With Boundary Information p. [139](#)

Andrews Boahen

Active Learning for Nonlinear Calibration p. [140](#)

Qian Xiao

Optimal Design of Experiments With Quantitative-Sequence Factors p. [141](#)

Chaofan Huang

Factor Importance Ranking and Selection Using Total Indices p. [142](#)

Advances in Adaptive Hamiltonian Monte Carlo

Organizer:

Art B. Owen

Stanford University

owen@stanford.edu

Session Description:

Hamiltonian Monte Carlo (HMC) is one of the most effective tools for high-dimensional Markov chain Monte Carlo. It is the default algorithm used in probabilistic programming languages for Bayesian computation, including `Stan` [2], `PyMC` and `NumPyro` (Python), and `Turing.jl` (Julia). While HMC handles some of the most difficult MCMC problems, it does so through the use of several tuning parameters, which can be challenging to set. Significant progress came from the no-U-turn sampler (NUTS) [3] and apogee-to-apogee path sampler [7], both of which dynamically adapt path lengths. More recent progress includes delayed rejection HMC [5], which locally adapts step sizes, and Gibbs self tuning (GIST) [1], which treats tuning parameters as random variables to maintain detailed balance. Chirag Modi's ATLAS [5] leverages local Hessians, delayed rejection, and GIST. AutoStep MCMC [4] adapts step sizes locally to match the variable geometry of target distributions. Bob Carpenter will provide an introduction to HMC, NUTS, and GIST for non-specialists. Nawaf Bou-Rabee will elaborate on GIST. Chirag Modi will discuss delayed rejection and ATLAS. Trevor Campbell will present AutoStep and related methods.

- [1] Bou-Rabee, B. Carpenter, and M. Marsden. GIST: Gibbs self-tuning for locally adaptive HMC. [arXiv:2404.15253](https://arxiv.org/abs/2404.15253), 2024.
- [2] Carpenter, B. et al. 2017. Stan: A probabilistic programming language. *J. Stat. Soft.*, 76
- [3] Hoffman, M. D. and Gelman, A. 2014. The no-U-turn sampler: Adaptively setting path lengths in HMC. *J. Mach. Learn. Res.*, 15(1).
- [4] Liu, T., Campbell, T., et al. 2024. AutoStep: Locally adaptive involutive MCMC. [arXiv:2410.18929](https://arxiv.org/abs/2410.18929), 2024.
- [5] Modi, C. 2024. ATLAS: Adapting trajectory lengths and step-size for HMC. [arXiv:2410.21587](https://arxiv.org/abs/2410.21587)
- [6] Modi, C., Barnett, A. and Carpenter, B. 2024. Delayed rejection Hamiltonian Monte Carlo for sampling multiscale distributions. *Bayesian Analysis*, 19(3), 2024.
- [7] Sherlock, C., Urbas, S. and Ludkin, M. 2023. The apogee to apogee path sampler. *JCGS*, 32(4).

Wed, Jul 30, 2025–Morning, 10:30 – 12:30, HH 002

Bob Carpenter

GIST: Gibbs Self-Tuning for Locally Adapting Hamiltonian Monte Carlo p. [143](#)

Nawaf Bou-Rabee

Acceleration of the No-U-Turn Sampler p. [144](#)

Chirag Modi

ATLAS: Adapting Trajectory Lengths and Step-Size for Hamiltonian Monte Carlo p. [145](#)

Trevor Campbell

AutoStep: Locally Adaptive Involutive MCMC p. [146](#)

Stochastic Optimization

Organizer:

Shane G. Henderson

Cornell University

sgh9@cornell.edu

Session Description:

In many applications, one wishes to solve an optimization problem $\min_{x \in X} f(x)$, where $f(\cdot)$ and/or its derivatives can only be evaluated through noisy estimates obtained using Monte Carlo simulation. Such problems are ubiquitous in machine learning, and also arise in stochastic simulation applications. This session will consist of 3 talks in the area.

Wed, Jul 30, 2025–Afternoon, 14:00 – 16:00, HH Auditorium

Raghu Bollapragada

Monte Carlo Based Adaptive Sampling Approaches for Stochastic Optimization p. 147

Shane Henderson

A New Convergence Analysis of Two Stochastic Frank-Wolfe Algorithms p. 148

Akshita Gupta

Stochastic Gradient With Testing Functionals p. 149

Recent Progress on Algorithmic Discrepancy Theory and Applications

Organizer:

Haotian Jiang

University of Chicago

jhtdavid@uchicago.edu

Session Description:

Discrepancy theory studies the irregularities of distributions. Typical questions studied in discrepancy theory include: “What is the most uniform way of distributing n points in the unit square, and how big must the irregularity be?”, “What is the best way to divide a set of n objects into two groups that are as ‘similar’ as possible?” These questions have been studied since the 1930s and progress on them has found extensive applications in many areas of mathematics, computer science, statistics, finance, etc. The past decade has seen tremendous progress in designing efficient algorithms for discrepancy questions. These developments have led to many surprising applications in areas such as differential privacy, graph sparsification, approximation algorithms and rounding, kernel density estimation, randomized controlled trials, and quasi-Monte Carlo methods. The goal of this special session is to present several exciting recent progress in this direction, and to facilitate cross-fertilization across different areas. A few related recent papers in this direction are listed below.

- [1] Harshaw, Christopher, Fredrik Sävje, Daniel A. Spielman, and Peng Zhang (2024). *Balancing covariates in randomized experiments with the Gram-Schmidt walk design*. Journal of the American Statistical Association 119, no. 548 (2024): 2934-2946.
- [2] Bansal, Nikhil, and Haotian Jiang (2025). *Quasi-Monte Carlo Beyond Hardy-Krause*. In Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2051-2075. Society for Industrial and Applied Mathematics, 2025.
- [3] Aistleitner, Christoph, Dmitriy Bilyk, and Aleksandar Nikolov (2016). *Tusnády’s problem, the transference principle, and non-uniform QMC sampling*. In Monte Carlo and Quasi-Monte Carlo Methods: MCQMC 2016, Stanford, CA, August 14-19 12, pp. 169-180. Springer International Publishing, 2018.

Wed, Jul 30, 2025–Afternoon, 14:00 – 16:00, HH Ballroom

Haotian Jiang

Algorithmic Discrepancy Theory: An Overview p. [149](#)

Peng Zhang

Improving the Design of Randomized Experiments via Discrepancy Theory p. [150](#)

Aleksandar Nikolov

Online Factorization for Online Discrepancy Minimization p. [151](#)

Monte Carlo Applications in High-Performance Computing, Computer Graphics, and Computational Science

Organizer:

Michael Mascagni

Florida State University and the National Institute of Standards and Technology

mascagni@fsu.edu

Session Description:

Monte Carlo methods are useful for solving problems in a variety of areas. We have four talks organized that span several areas. First, we consider how Monte Carlo methods can provide fault tolerance to large computations via work on simulating soft and hard faults in Monte Carlo computation on a state-of-the-art computer. Next, we consider using Monte Carlo to create a fast and efficient computer graphics renderer. Next we consider two talks on applications of Monte Carlo to the solution of partial differential equations. One of these talks deals specifically with equations that arise in financial computing.

Wed, Jul 30, 2025–Afternoon, 14:00 – 16:00, HH 002

Arash Fahim

Gaining Efficiency in Monte Carlo Policy Gradient Methods for Stochastic Optimal Control
p. [152](#)

Sharanya Jayaraman

Examining the Fault Tolerance of High-Performance Monte Carlo Applications Through Simulation p. [153](#)

Rohan Sawhney

Building Monte Carlo “Renderers” for Physics p. [154](#)

Silei Song

WoS-NN: Collaborating Walk-On-Spheres With Machine Learning to Solve Elliptic PDEs
p. [155](#)

QMC and Applications Part I

Organizers:

Michael Gnewuch

University of Osnabrück

michael.gnewuch@uni-osnabrueck.de

Takashi Goda

The University of Tokyo

goda@frcer.t.u-tokyo.ac.jp

Peter Kritzer

Austrian Academy of Sciences

peter.kritzer@oeaw.ac.at

Session Description:

Quasi-Monte Carlo (QMC) methods have been widely studied as an effective tool for high-dimensional integration and have found applications in various fields, including computational finance, computer graphics, data compression, partial differential equations with random coefficients, and optimization. Despite their success, ongoing theoretical developments and the expansion of application areas continue to drive this research field forward. This special session is devoted to showcasing recent advances in the theory of QMC methods and their applications.

Thu, Jul 31, 2025–Morning, 10:30 – 12:30, HH Auditorium

Felix Bartel

Exact Discretization, Tight Frames and Recovery via D-Optimal Designs p. [156](#)

Mou Cai

L2-Approximation: Using Randomized Lattice Algorithms and QMC Hyperinterpolation
p. [157](#)

Zhijian He

High-Dimensional Density Estimation on Unbounded Domain p. [158](#)

Frances Y. Kuo

Application of QMC to Oncology p. [158](#)

Analysis of Langevin and Related Sampling Algorithms, Part I

Organizers:

Yifan Chen

Courant Institute of Mathematical Sciences, New York University

yifan.chen@nyu.edu

Xiaoou Cheng

Courant Institute of Mathematical Sciences, New York University

chengxo@nyu.edu

Jonathan Weare

Courant Institute of Mathematical Sciences, New York University

weare@nyu.edu

Session Description:

Many Markov Chain Monte Carlo (MCMC) samplers are based on stochastic dynamics. Langevin dynamics serves as a fundamental basis for a vast family of extensions, such as unadjusted Langevin algorithms, kinetic/underdamped Langevin algorithms, Hamiltonian Monte Carlo, and the No-U-Turn Sampler (NUTS). The gradient flow structure of Langevin dynamics also motivates the development of a large class of novel algorithms such as Stein variational gradient descent, birth-death process approaches, and those based on Fisher-Rao gradient flows. These methods have become ubiquitous across various fields, including molecular dynamics, Bayesian statistics, and machine learning. Recent years have seen significant theoretical advances in analyzing such methods, particularly in high-dimensional settings and non-convex cases. This special session aims to bring together researchers from different communities (probability, statistics, scientific computing, theoretical computer science, machine learning, etc.) working on analysis of sampling dynamics of Langevin and beyond to present recent progress, discuss challenges, and share ideas.

Thu, Jul 31, 2025–Morning, 10:30 – 12:30, HH Ballroom

Krishnakumar Balasubramanian

Finite-Particle Convergence Rates for Stein Variational Gradient Descent p. 159

Lihan Wang

Convergence Rates of Kinetic Langevin Dynamics With Weakly Confining Potentials p. 160

Xiaoou Cheng

Delocalization of Bias in Unadjusted Hamiltonian Monte Carlo p. 161

Nested Expectations: Models and Estimators, Part II

Organizers:

Arved Bartuska

King Abdullah University of Science and Technology/RWTH Aachen University

arved.bartuska@kaust.edu.sa

Abdul-Lateef Haji-Ali

Heriot-Watt University

a.hajiali@hw.ac.uk

Session Description:

Nested expectations arise in many applications, such as in engineering, mathematical finance, and medical decision-making. In addition to their nested structure, numerical estimations of such expectations are often complicated by singularities or discontinuities. Moreover, approximations when evaluating inner expectations using, for example, finite element or time-stepping schemes render traditional estimation methods such as double-loop Monte Carlo prohibitively expensive. This session will explore models and applications with this structure and methods for efficient estimation.

Thu, Jul 31, 2025–Morning, 10:30 – 12:30, HH 002

Matteo Raviola

Stochastic Gradient With Least-Squares Control Variates p. 162

Philipp Guth

A One-Shot Method for Bayesian Optimal Experimental Design p. 162

Sara Pérez-Vieites

Langevin-Based Strategies for Nested Particle Filters p. 163

QMC and Applications Part II

Organizers:

Michael Gnewuch

University of Osnabrück

michael.gnewuch@uni-osnabrueck.de

Takashi Goda

The University of Tokyo

goda@frcer.t.u-tokyo.ac.jp

Peter Kritzer

Austrian Academy of Sciences

peter.kritzer@oeaw.ac.at

Session Description:

Quasi-Monte Carlo (QMC) methods have been widely studied as an effective tool for high-dimensional integration and have found applications in various fields, including computational finance, computer graphics, data compression, partial differential equations with random coefficients, and optimization. Despite their success, ongoing theoretical developments and the expansion of application areas continue to drive this research field forward. This special session is devoted to showcasing recent advances in the theory of QMC methods and their applications.

Thu, Jul 31, 2025–Afternoon, 15:30 – 17:30, HH Auditorium

Dirk Nuyens

Approximation of Multivariate Periodic Functions p. 164

Art Owen

Randomized QMC With One Categorical Variable p. 164

Zexin Pan

QMC Confidence Intervals Using Quantiles of Randomized Nets p. 165

Kosuke Suzuki

Quasi-Uniform Quasi-Monte Carlo Lattice Point Sets p. 166

Analysis of Langevin and Related Sampling Algorithms, Part II

Organizers:

Yifan Chen

Courant Institute of Mathematical Sciences, New York University

yifan.chen@nyu.edu

Xiaoou Cheng

Courant Institute of Mathematical Sciences, New York University

chengxo@nyu.edu

Jonathan Weare

Courant Institute of Mathematical Sciences, New York University

weare@nyu.edu

Session Description:

Many Markov Chain Monte Carlo (MCMC) samplers are based on stochastic dynamics. Langevin dynamics serves as a fundamental basis for a vast family of extensions, such as unadjusted Langevin algorithms, kinetic/underdamped Langevin algorithms, Hamiltonian Monte Carlo, and the No-U-Turn Sampler (NUTS). The gradient flow structure of Langevin dynamics also motivates the development of a large class of novel algorithms such as Stein variational gradient descent, birth-death process approaches, and those based on Fisher-Rao gradient flows. These methods have become ubiquitous across various fields, including molecular dynamics, Bayesian statistics, and machine learning. Recent years have seen significant theoretical advances in analyzing such methods, particularly in high-dimensional settings and non-convex cases. This special session aims to bring together researchers from different communities (probability, statistics, scientific computing, theoretical computer science, machine learning, etc.) working on analysis of sampling dynamics of Langevin and beyond to present recent progress, discuss challenges, and share ideas.

Thu, Jul 31, 2025–Afternoon, 15:30 – 17:30, HH Ballroom

Molei Tao

Langevin-Based Sampling Under Nonconvex Constraints p. 166

Yifan Chen

Convergence of Unadjusted Langevin in High Dimensions: Delocalization of Bias p. 167

Fuzhong Zhou

Entropy Methods for the Delocalization of Bias in Langevin Monte Carlo p. 168

Siddharth Mitra

Convergence of Φ -Divergence and Φ -Mutual Information Along Langevin Markov Chains
p. 168

Recent Advances in Stochastic Gradient Descent

Organizer:

Jing Dong

Columbia University

jing.dong@gsb.columbia.edu

Session Description:

Stochastic Gradient Descent (SGD) is a cornerstone optimization method in machine learning, renowned for its efficiency in handling large-scale data. Its iterative approach enables the processing of extensive datasets by updating model parameters using randomly selected data subsets, thereby reducing computational costs. Despite its widespread adoption, traditional SGD faces challenges such as convergence to sharp minima, and sensitivity to data distribution shifts. Addressing these challenges is crucial for enhancing model generalization, robustness, and overall performance in diverse applications. This session aims to delve into recent developments that address these challenges in SGD, presenting innovative methodologies and theoretical insights to enhance its effectiveness in complex learning scenarios. The session will have three to four speakers. Currently, the confirmed speakers are Jose Blanchet (Stanford University), Chang-Han Rhee (Northwestern University), and Jing Dong (Columbia University). Each will present their recent works on stochastic gradient descent, ranging from SGD and heavy-tailed phenomenon to SGD with adaptively generated data. Collectively, these talks will shed light on cutting-edge advancements in SGD methodologies, providing both theoretical frameworks and practical strategies to enhance optimization in complex, real-world applications.

Thu, Jul 31, 2025–Afternoon, 15:30 – 17:30, HH 002

Jose Blanchet

Inference for Stochastic Gradient Descent With Infinite Variance p. 169

Chang-Han Rhee

Exit-Time Analysis of Stochastic Gradient Descent via Kesten’s Recursion p. 170

Jing Dong

Stochastic Gradient Descent With Adaptive Data p. 170

Forward and Inverse Problems for Stochastic Reaction Networks

Organizers:

Sophia Münker

RWTH Aachen University

muenker@uq.rwth-aachen.de

Chiheb Ben Hammouda

Utrecht University

c.benhammouda@uu.nl

Raúl Tempone

RWTH Aachen University

tempone@uq.rwth-aachen.de

Session Description:

This session aims to bring together experts working on stochastic reaction networks and pure jump processes for modeling stochastic biological and chemical systems. The session is about recent advances in Monte Carlo methods, variance and dimension reduction techniques that are relevant for tackling forward and inverse problems.

Fri, Aug 1, 2025–Morning, 09:00 – 11:00, HH Auditorium

Zhou Fang

Fixed-Budget Simulation Method for Growing Cell Populations p. 171

Sophia Münker

Dimensionality Reduction for Efficient Rare Event Estimation p. 172

Maksim Chupin

Filtered Markovian Projection: Dimensionality Reduction in Filtering for Stochastic Reaction Networks p. 173

Muruhan Rathinam

State and Parameter Inference in Stochastic Reaction Networks p. 174

Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part II

Organizers:

Sou-Cheng T. Choi

Illinois Institute of Technology

schoi32@iit.edu

Pieterjan Robbe

Sandia National Laboratories

pmrobbe@sandia.gov

Mike Giles

University of Oxford

mike.giles@maths.ox.ac.uk

Session Description:

Monte Carlo (MC) or quasi-Monte Carlo (QMC) algorithms are widely used in various fields such as finance, physics, and engineering for their ability to handle high-dimensional integration problems. The development and maintenance of software for (quasi-)Monte Carlo ((Q)MC) algorithms can significantly enhance the accessibility and usability of these techniques. This special session aims to bring together experts from academia and industry to discuss recent advances in (Q)MC software, share best practices, and explore future directions, fostering collaboration among researchers and practitioners. Topics of interest for the session include:

- Novel hardware or architectural designs for open-source (Q)MC libraries.
- Best collaborative practices for developing and maintaining efficient and reliable (Q)MC software.
- Challenges and opportunities in integrating (Q)MC methods with machine learning and AI techniques.
- High-performance computing solutions for (Q)MC software.
- Adaptation of (Q)MC software to application fields such as finance, computer graphics, sensitivity analysis, Bayesian optimization, and uncertainty quantification.
- Innovative approaches to enhancing and extending existing (Q)MC tools.

Fri, Aug 1, 2025–Morning, 09:00 – 11:00, HH Ballroom

Niklas Baumgarten

A High-Performance Multi-Level Monte Carlo Software for Full Field Estimates and Applications in Optimal Control p. 175

Aleksei Sorokin

Fast Gaussian Processes p. 176

Johannes Krotz

Hybrid Monte Carlo Methods for Kinetic Transport p. 177

Joseph Farmer

Flow-Based Monte Carlo Transport Simulation p. 178

Abstracts

A strong order 1.5 boundary-preserving discretization scheme for scalar SDEs defined in a domain

Andreas Neuenkirch

University of Mannheim

neuenkirch@uni-mannheim.de

Coauthor(s): Ruishu Liu, Xiaojie Wang

Special session: Stochastic Computation and Complexity [p.54](#)

We study the strong approximation of scalar SDEs, which take values in a domain and have non-Lipschitz coefficients. By combining a Lamperti-type transformation with a semi-implicit discretization approach and a taming strategy, we construct a domain-preserving scheme that strongly converges under weak assumptions. Moreover, we show that this scheme has a strong convergence order 1.5 under additional assumptions on the coefficients of the SDE. In our scheme, the domain preservation is a consequence of the semi-implicit discretization approach, while the taming strategy allows controlling terms of the scheme that admit singularities but are required to obtain the desired order.

Our general convergence results are applied to various SDEs from applications, with sub-linearly or super-linearly growing and non-globally Lipschitz coefficients.

- [1] Ruishu Liu, Andreas Neuenkirch and Xiaojie Wang (2024+). A strong order 1.5 boundary-preserving discretization scheme for scalar SDEs defined in a domain. *Mathematics of Computation*. doi:10.1090/mcom/4014 (to appear, online first)

Christopher Rauhögger*University of Passau*

christopher.rauhogger@uni-passau.de

StochasticComputationandComplexity

We consider d -dimensional systems of SDEs with a discontinuous drift coefficient. More precisely, we assume that there exists a C^5 -hypersurface $\Theta \subseteq \mathbb{R}^d$ such that the drift coefficient is intrinsic Lipschitz continuous on $\mathbb{R}^d \setminus \Theta$ and has intrinsic Lipschitz continuous derivative on $\mathbb{R}^d \setminus \Theta$. Furthermore, the diffusion coefficient is C^1 on \mathbb{R}^d and commutative with a bounded derivative that is intrinsic Lipschitz continuous on $\mathbb{R}^d \setminus \Theta$.

It was proven in [1] for $d = 1$ and more recently in [2] for general $d \in \mathbb{N}$ that in this setting a transformed Milstein scheme achieves an L_p -error rate of order at least $3/4-$ in terms of the number of evaluations of the driving Brownian motion. Furthermore, it was proven in [3] that for $d = 1$ in the same setting an adaptive Milstein-type scheme achieves an L_p -error rate of order at least 1 in terms of the average number of evaluations of the driving Brownian motion.

In this talk, we present a generalisation of the result from [3] to higher dimensions. More precisely, we introduce an adaptive transformed Milstein scheme which can be used for the approximation of solutions of d -dimensional systems of SDEs at the final time point in this setting and prove that this scheme achieves an L_p -error rate of order at least 1 in terms of the average number of evaluations of the driving Brownian motion.

- [1] Müller-Gronbach, Thomas & Yaroslavtseva, Larisa. (2022). *A strong order 3/4 method for SDEs with discontinuous drift coefficient*. IMA Journal of Numerical Analysis. 42. 229-259
- [2] Rauhögger, Christopher. (2025+). *Milstein-type methods for strong approximation of systems of SDEs with a discontinuous drift coefficient*. In preparation
- [3] Yaroslavtseva, Larisa. (2022). *An adaptive strong order 1 method for SDEs with discontinuous drift coefficient*. Journal of Mathematical Analysis and Applications. 513. 2. Paper Number 126180, 29

Strong order 1 adaptive approximation of jump-diffusion SDEs with discontinuous drift

Verena Schwarz
 University of Klagenfurt
verena.schwarz@aau.at

Coauthor(s): Stochastic Computation and Complexity

In this talk, we present an adaptive approximation scheme for jump-diffusion SDEs with discontinuous drift and (possibly) degenerate diffusion. The scheme is a transformation-based doubly-adaptive quasi-Milstein scheme, which is doubly-adaptive in the sense that it is jump-adapted, i.e. all jump times of the Poisson noise are grid points, and it includes an adaptive stepsize strategy to account for the discontinuities of the drift. It is proven to have strong convergence rate 1 in L^p for $p \in [1, \infty)$ with respect to the average computational cost for these SDEs. To obtain our result, we prove that under slightly stronger assumptions which are still weaker than those in existing literature, a related doubly-adaptive quasi-Milstein scheme has a convergence order 1.

Approximation in Hilbert spaces of the Gaussian and related analytic kernels

Toni Karvonen
 Lappeenranta–Lahti University of Technology LUT
toni.karvonen@lut.fi

Coauthor(s): Yuya Suzuki

We consider linear approximation based on function evaluations in reproducing kernel Hilbert spaces of certain analytic weighted power series kernels on the interval $[-1, 1]$. This class contains the popular Gaussian kernel $K(x, y) = \exp(-\frac{1}{2}\varepsilon^2(x - y)^2)$. We derive almost matching upper and lower bounds on the worst-case error. When applied to the Gaussian kernel, our results state that, up to a sub-exponential factor, the n th minimal error decays as $(\varepsilon/2)^n (n!)^{-1/2}$. The proofs are based on weighted polynomial interpolation and classical polynomial coefficient estimates that we use to bound the Hilbert space norm of a weighted polynomial fooling function.

Domain UQ for stationary and time-dependent PDEs using QMC*André-Alexander Zepernick*

Free University of Berlin

a.zepernick@fu-berlin.de

Coauthor(s): Ana Djurdjevac, Vesa Kaarnioja, Claudia Schillings

Special session: Domain Uncertainty Quantification p.55

The problem of modelling processes with partial differential equations posed on random domains arises in various applications like biology or engineering. We study uncertainty quantification for partial differential equations subject to domain uncertainty. For the random domain parameterization, we adopt an approach, which was also examined by Chernov and Lê [1,2] as well as Harbrecht, Schmidlin, and Schwab [3], where one assumes the input random field to be Gevrey regular. This approach has the advantage of being substantially more general than models which assume a particular parametric representation of the input random field such as a Karhunen–Loève series expansion. As model problems we consider both the Poisson equation as well as the heat equation and design randomly shifted lattice quasi-Monte Carlo (QMC) cubature rules for the computation of response statistics subject to domain uncertainty. The QMC rules obtained in [4] exhibit dimension-independent, faster-than-Monte Carlo cubature convergence rates. Our theoretical results are illustrated by numerical examples.

- [1] Chernov, Alexey, & Lê, Tùng (2024). Analytic and Gevrey class regularity for parametric elliptic eigenvalue problems and applications. *SIAM Journal on Numerical Analysis*, **62**(4), 1874–1900.
- [2] Chernov, A., & Lê, Tùng (2024). Analytic and Gevrey class regularity for parametric semilinear reaction-diffusion problems and applications in uncertainty quantification. *Computers & Mathematics with Applications*, **164**, 116–130.
- [3] Harbrecht, Helmut, Schmidlin, Marc, & Schwab, Christoph (2024). The Gevrey class implicit mapping theorem with application to UQ of semilinear elliptic PDEs. *Mathematical Models and Methods in Applied Sciences.*, **34**(5), 881–917.
- [4] Djurdjevac, Ana, Kaarnioja, Vesa, Schillings, Claudia & Zepernick, André-Alexander (2025). Uncertainty quantification for stationary and time-dependent PDEs subject to Gevrey regular random domain deformations. Preprint, *arXiv:2502.12345 [math.NA]*.

Domain Uncertainty Quantification for Electromagnetic Wave Scattering via First-Order Sparse Boundary Element Approximation

Carlos Jerez-Hanckes

INRIA Chile

carlos.jerez@inria.cl

Coauthor(s): Paul Escapil-Inchauspé

Special session: Domain uncertainty quantification p.55

Quantifying the effects on electromagnetic waves scattered by objects of uncertain shape is key for robust design, particularly in high-precision applications. Assuming small random perturbations departing from a nominal domain, the first-order sparse boundary (FOSB) element method has been proven to directly compute statistical moments with poly-logarithmic complexity [1,2] for a prescribed accuracy, without resorting to computationally intense Monte Carlo (MC) simulations. However, implementing FOSB is not straightforward as the lack of compelling computational results for EM scattering attests [3]. In this work, we present a first full 3D implementation of FOSB for shape-related uncertainty quantification (UQ) in EM scattering [4]. In doing so, we address several implementation issues such as ill-conditioning and large computational and memory requirements and present a comprehensive, state-of-the-art, easy-to-use, open-source computational framework to directly apply this technique when dealing with complex objects. Exhaustive numerical experiments confirm our claims and demonstrate the technique’s applicability and provide pathways for further improvement.

- [1] Jerez-Hanckes, Schwab (2017). *Electromagnetic Wave Scattering by Random Surfaces: Uncertainty Quantification via Sparse Tensor Boundary Elements*, IMA Journal of Numerical Analysis **37**(3), 1175–1210.
- [2] Hiptmair, Jerez-Hanckes, Schwab (2013). *Sparse Tensor Edge Elements*, BIT Numerical Mathematics **53**, 925–943.
- [3] Escapil-Inchauspé, Jerez-Hanckes (2020). *Helmholtz Scattering by Random Domains: First-Order Sparse Boundary Elements Approximation*, SIAM Journal of Scientific Computing **42**(5), A2561–A2592.
- [4] Escapil-Inchauspé, Jerez-Hanckes (2024). *Shape Uncertainty Quantification for Electromagnetic Wave Scattering via First-Order Sparse Boundary Element Approximation*, IEEE Transactions in Antennas & Propagation **72**(8):6627–6637.

Quantifying uncertainty in spectral clusterings: expectations for perturbed and incomplete data

Jürgen Dölz

University of Bonn

doelz@ins.uni-bonn.de

Coauthor(s): Jolanda Weygandt

Special session: Domain Uncertainty Quantification [p.55](#)

Spectral clustering is a popular unsupervised learning technique which partitions a set of unlabelled data into disjoint clusters. However, the data under consideration are often experimental data, implying that the data is subject to measurement errors and measurements may even be lost or invalid. These uncertainties in the input data induce corresponding uncertainties in the resulting clusters. In this talk, we model the uncertainties as random, implying that the clusters need to be considered random as well. We further discuss a mathematical framework based on random set theory for the computational approximation of statistically expected clusterings.

Model Problems for PDEs on Uncertain Domains

Harri Hakula

Aalto University

Harri.Hakula@aalto.fi

Coauthor(s): Domain Uncertainty Quantification

Partial differential equation related uncertainty quantification has become one of the topical research areas in applied mathematics and, in particular, engineering. Stochastic finite element methods are applied both in source and eigenvalue problems. Remarkably, computational function theory provides a rich set of invariants and identities that can be applied in designing model problems where the domain is random or uncertain. In this talk the focus is on conformal capacity in a simple, yet general case where the sides of a quadrilateral are assumed to be random and parameterised with a suitable Karhunen-Loève expansion [1]. Lattice quasi-Monte Carlo (QMC) cubature rules are used for computing the expected value of the solution to the resulting Poisson problem subject to domain uncertainty.

High-order finite element methods (*hp*-FEM) are used in the deterministic problems. The special features related to modelling random domains in *hp*-context are discussed. Convergence properties of the lattice QMC quadratures are presented. The talk concentrates on numerical experiments demonstrating the theoretical error estimates. The new results on the associated Steklov eigenvalue problem are also covered.

- [1] Hakula, H., Harbrecht, H., Kaarnioja, V., Kuo, F. Y., & Sloan, I. H. (2024). Uncertainty quantification for random domains using periodic random variables. *Numerische Mathematik*, 156(2), 273–317. <https://doi.org/10.1007/s00211-023-01392-6>

An Adaptive Sampling Algorithm for Level-set Approximation

Abdul-Latefe Haji-Ali

Heriot-Watt University

a.hajiali@hw.ac.uk

Coauthor(s): Matteo Croci and Ian CJ Powell

Let $D \subset\subset \mathbb{R}^d$ be a d -dimensional domain with compact closure. We consider the approxi-

mation of the $d - 1$ dimensional zero level-set $\mathcal{L}_0 := \{x \in \overline{D} : f(x) = 0\}$ where the Lipschitz function f is either accessible directly or when $f(x) = \mathbb{E}[\tilde{f}(x)]$ for all $x \in \overline{D}$. Given an approximation scheme with a priori error bounds and L^p bounds on the \tilde{f} -approximation error, we propose a grid-based adaptive sampling scheme which produces an approximation to \mathcal{L}_0 with expected cost-complexity reduction of $\varepsilon^{\left(\frac{p+1}{\alpha p}\right)}$ compared to a non-adaptive scheme, where α is the known convergence rate of an interpolation scheme. We provide the numerical analysis and experiments to show that these rates hold in practice.

Mon, Jul 28 11:00–11:30

Posterior-Free A-Optimal Bayesian Design of Experiments via Conditional Expectation

Vinh Hoang

RWTH-Aachen University

hoang@uq.rwth-aachen.de

Coauthor(s): Luis Espath, Sebastian Krumscheid, Raúl Tempone

We propose a novel approach for solving the A-optimal Bayesian design of experiments that does not require sampling or approximating the posterior distribution. In this setting, the objective function is the expected conditional variance (ECV). Our method estimates the ECV by leveraging conditional expectation, which we approximate using its orthogonal projection property. We derive an asymptotic error bound for this estimator and validate it through numerical experiments. The method is particularly efficient when the design parameter space is continuous. In such scenarios, the conditional expectation can be approximated non-locally using tools such as neural networks. To reduce the number of evaluations of the measurement model, we incorporate transfer learning and data augmentation. Numerical results show that our method significantly reduces model evaluations compared to standard importance sampling-based techniques. Code available at: <https://github.com/vinh-tr-hoang/DOEviaPACE>.

- [1] Hoang, V., Espath, L., Krumscheid, S., & Tempone, R. (2025). Scalable method for Bayesian experimental design without integrating over posterior distribution. *SIAM ASA Journal on Uncertainty Quantification*, 13(1), 114-139. <https://doi.org/10.1137/23M1603364>

QMC for Bayesian optimal experimental design with application to inverse problems governed by PDEs

Vesa Kaarnioja

Free University of Berlin

vesa.kaarnioja@fu-berlin.de

Coauthor(s): Claudia Schillings

Special session: Nested expectations: models and estimators, Part I p.56

The goal in Bayesian optimal experimental design (OED) is to maximize the expected information gain for the reconstruction of unknown quantities in an experiment by optimizing the placement of measurements. The objective function in the resulting optimization problem involves a multivariate double integral over the high-dimensional parameter and data domains. For the efficient approximation of these integrals, we consider a sparse tensor product combination of quasi-Monte Carlo (QMC) cubature rules over the parameter and data domains. For the parameterization of the unknown quantities, we consider a model recently studied by Chernov and Lê [1,2] as well as Harbrecht, Schmidlin, and Schwab [3] in which the input random field is assumed to belong to a Gevrey class. The Gevrey class contains functions that are infinitely many times continuously differentiable with a growth condition on the higher-order partial derivatives, but which are not analytic in general. Using the techniques developed in [4], we investigate efficient Bayesian OED for inverse problems governed by partial differential equations (PDEs).

- [1] Chernov, Alexey, & Lê, Tùng (2024). Analytic and Gevrey class regularity for parametric elliptic eigenvalue problems and applications. *SIAM Journal on Numerical Analysis*, **62**(4), 1874–1900.
- [2] Chernov, Alexey, & Lê, Tùng (2024). Analytic and Gevrey class regularity for parametric semilinear reaction-diffusion problems and applications in uncertainty quantification. *Computers & Mathematics with Applications*, **164**, 116–130.
- [3] Harbrecht, Helmut, Schmidlin, Marc, & Schwab, Christoph (2024). The Gevrey class implicit mapping theorem with applications to UQ of semilinear elliptic PDEs. *Mathematical Models and Methods in Applied Sciences*, **34**(5), 881–917.
- [4] Kaarnioja, Vesa, & Schillings, Claudia (2024). Quasi-Monte Carlo for Bayesian design of experiment problems governed by parametric PDEs. Preprint, *arXiv:2405.03529 [math.NA]*.

Multilevel quasi-Monte Carlo without replications*Pieterjan Robbe*

Sandia National Laboratories

pmrobbe@sandia.gov

Coauthor(s): Aleksei Sorokin, Gianluca Geraci, Fred J. Hickernell, Mike Eldred

Special session: Hardware or Software for (Quasi-)Monte Carlo Algorithms p.57

In this talk, we explore a novel approach to multilevel quasi-Monte Carlo (MLQMC) sampling that eliminates the need for stochastic replications. Our approach for estimating the level-wise variances is based on the Bayesian cubature framework introduced in [1]. Empirical results from a series of numerical experiments illustrate the effectiveness of our method in various applications. We discuss the integration of our new method in Dakota, Sandia's flagship UQ software package.

- [1] Jagadeeswaran, R., & Hickernell, F. J. (2019). Fast automatic Bayesian cubature using lattice sampling. *Statistics and Computing*, **29**(6), 1215–1229.

A nested Multilevel Monte Carlo framework for efficient simulations on FPGAs*Irina-Beatrice Haas*

Mathematical Institute, University of Oxford

irina-beatrice.haas@maths.ox.ac.uk

Coauthor(s): Michael B. Giles

Special session: Hardware or Software for (Quasi-)Monte Carlo Algorithm p.57

Multilevel Monte Carlo (MLMC) is a computational method that reduces the cost of Monte Carlo simulations by combining SDE approximations with multiple resolutions. A further avenue for significantly reducing cost and improving power efficiency of MLMC, notably for financial option pricing, is the use of low-precision calculations on configurable hardware devices such as Field-Programmable Gate Arrays (FPGAs). With this goal in mind, in this talk, we propose a new MLMC framework that exploits approximate random variables and

fixed-point operations with optimised precision to compute most SDE paths with a lower cost.

For the generation of random Normal increments, we discuss several methods based on the approximation of the inverse normal CDF (see e.g., [3]), and we argue that these methods could be implemented on FPGAs using small Look-Up-Tables to generate random numbers more efficiently than on CPUs or GPUs.

To set the bit-width of variables in the path generation we propose a rounding error model and optimise the precision of all variables on each Monte Carlo level. This optimisation stage is independent of the desired overall accuracy, and can therefore be performed off-line.

With these two key improvements, our proposed framework [2] offers higher computational savings than the existing mixed-precision MLMC framework [1].

- [1] Brugger, C., de Schryver, C., Wehn, N., Omland, S., Hefter, M., Ritter, K., Kostiuk, A., & Korn, R. (2014). Mixed precision multilevel Monte Carlo on hybrid computing systems. In Proceedings of the IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr) (pp. 215–222). IEEE. <https://doi.org/10.1109/CIFEr.2014.6924076>
- [2] Haas, I. B., & Giles, M. B. (2025). A nested MLMC framework for efficient simulations on FPGAs. Accepted to appear in Monte Carlo Methods and Applications. arXiv preprint arXiv:2502.07123. <https://arxiv.org/abs/2502.07123>
- [3] Giles, M.B., & Sheridan-Methven, O. (2023). Approximating Inverse Cumulative Distribution Functions to Produce Approximate Random Variables. ACM Transactions on Mathematical Software, 49(3), no 26. <https://doi.org/10.1145/3604935>

Mon, Jul 28 11:30–12:00

CUDA implementation of MLMC on NVIDIA GPUs

Mike Giles

University of Oxford

mike.giles@maths.ox.ac.uk

Coauthor(s): Hardware or Software for (Quasi-)Monte Carlo Algorithms

This talk will discuss the implementation of Multilevel Monte Carlo on NVIDIA GPUs. It will focus on some of the tricks needed for best performance, such as on-the-fly generation of random numbers within CUDA kernels, and latency-hiding through GPU computations in anticipation of requests from the CPU host process which manages the MLMC optimisation.

It will also discuss the opportunities and challenges in exploiting mixed-precision computing, using nested MLMC to perform most calculations at half precision (fp16) and just a few at single precision (fp32).

- [1] Giles, M.B. & Sheridan-Methven, O. (2022) Analysis of nested multilevel Monte Carlo using approximate Normal random variables. *SIAM/ASA Journal on Uncertainty Quantification*, **10**(1), 200–226.

Mon, Jul 28 12:00–12:30

Scalable and User-friendly QMC Sampling with UMBridge

Chung Ming Loi

Durham University

chung.m.loi@durham.ac.uk

Coauthor(s): Anne Reinarz

Special session: Hardware and Software for Quasi-Monte Carlo Methods p.57

Uncertainty quantification (UQ) plays a crucial role in geoscience: Bayesian inference determines model parameters, such as the permeability and porosity of the sub-surface, that are typically impossible to determine accurately from observations. In practice, it is crucial to study the uncertainty in the inferred parameters to correctly quantify risk and make decisions. Despite its scientific value, performing UQ for an application is often a lengthy process due to a need for interdisciplinary expertise in both UQ and advanced simulation codes. In this talk, we will look at improving the workflow and computational efficiency of quasi-Monte Carlo (i.e., sampling/ensemble based) approaches to UQ applications. We introduce UM-Bridge [2], a universal software interface that facilitates integration of complex simulation models with an entire range of leading UQ packages. By separating concerns between simulation and UQ, UM-Bridge allows rapid development of cutting-edge applications. The newly implemented load balancing framework in UM-Bridge further enables scaling workloads to High Performance Computing clusters.

- [1] Graham, I. G., Kuo, F. Y., Nuyens, D., Scheichl, R., and Sloan, I. H. (2011). Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications. *Journal of Computational Physics*, 230(10), 3668-3694.
- [2] L. Seelinger, V. Cheng-Seelinger, A. Davis, M. Parno, and A. Reinarz, “UM-Bridge: Uncertainty quantification and modeling bridge,” *Journal of Open Source Software*, vol. 8, no. 83, p. 4748, 2023. [Online]. Available: <https://doi.org/10.21105/joss.04748>

Optimality of deterministic and randomized QMC-cubatures on several scales of function spaces

Michael Gnewuch

University of Osnabrueck

mgnewuch@uos.de

Coauthor(s): Josef Dick, Lev Markhasin, Winfried Sickel, Yannick Meiners

Special session: Stochastic Computation and Complexity p.59

We study the integration problem over the s -dimensional unit cube on four scales of Banach spaces of integrands. First, we consider Haar wavelet spaces $H_{p,q,\alpha}$, $1 \leq p, q \leq \infty$, $\alpha > 1/p$, consisting of functions whose Haar wavelet coefficients exhibit a certain decay behavior measured by the parameters p, q , and, most importantly, α . We study the worst-case error of a deterministic cubature rule over the norm unit ball and provide upper bounds for quasi-Monte Carlo (QMC) cubature rules based on arbitrary (t, m, s) -nets as well as matching lower error bounds for arbitrary cubature rules. These results show that using arbitrary (t, m, s) -nets as integration nodes yields the best possible rate of convergence. In the Hilbert space setting $p = 2 = q$ it was earlier shown by Heinrich, Hickernell and Yue [2] that scrambled (t, m, s) -nets yield optimal convergence rates in the randomized setting, where the randomized worst-case error is considered.

We establish several suitable function space embeddings that allow to transfer the deterministic and randomized upper error bounds on Haar wavelet spaces to certain spaces of fractional smoothness $1/p < \alpha \leq 1$ and to Sobolev and Besov spaces of dominating mixed smoothness $1/p < a \leq 1$. Known lower bounds for Sobolev and Besov spaces of dominating mixed smoothness show that (deterministic or suitably randomized) (t, m, s) -nets yield optimal convergence rates also on the corresponding scales of spaces.

The talk is based on the preprint [1] and the master thesis of my student Yannick Meiners.

- [1] M. Gnewuch, J. Dick, L. Markhasin, W. Sickel, QMC integration based on arbitrary (t, m, s) -nets yields optimal convergence rates on several scales of function spaces, preprint 2024, arXiv:2409.12879.
- [2] S. Heinrich, F. J. Hickernell and R. X. Yue, Optimal quadrature for Haar wavelet spaces, *Math. Comput.*, 73 (2004), 259–277.

Optimal designs for function discretization and construction of tight frames*Kateryna Pozharska*

Institute of Mathematics of NAS of Ukraine and Chemnitz University of Technology

pozharska.k@gmail.com

Coauthor(s): Felix Bartel, Lutz Kämmerer, Martin Schäfer, Tino Ullrich

Special session: Stochastic Computation and Complexity p.59

In the talk we will present a direct and constructive approach for the construction of tight frames and exact Marcinkiewicz-Zygmund inequalities in the Lebesgue space [1]. It is based on a similar procedure of maximization of the determinant of a certain Gramian matrix with respect to points and weights, already used in [2] for discretization problem for the uniform norm, and results in a discrete measure with at most $n^2 + 1$ atoms, which accurately subsamples the L_2 -norm of complex-valued functions contained in a given n -dimensional subspace.

This approach can as well be used for the reconstruction of functions from general RKHS in L_p where one only has access to the most important eigenfunctions. The general results apply to the d -sphere or multivariate trigonometric polynomials on \mathbb{T}^d spectrally supported on arbitrary finite index sets $I \subset \mathbb{Z}^d$. Numerical experiments indicate the sharpness of this result.

- [1] Bartel, Felix, & Kämmerer, Lutz, & Pozharska, Kateryna, & Schäfer, Martin, & Ullrich, Tino (2024). *Exact discretization, tight frames and recovery via D -optimal designs*. arXiv:2412.02489.
- [2] Krieg, David, & Pozharska, Kateryna, & Ullrich, Mario & Ullrich, Tino (2024). *Sampling projections in the uniform norm*. arXiv:2401.02220.

Complexity of approximating piecewise smooth functions in the presence of deterministic or random noise

Leszek Plaskota

University of Warsaw

L.Plaskota@mimuw.edu.pl

Coauthor(s): Stochastic Computation and Complexity

Consider the smoothness class of 1-periodic functions $f : \mathbb{R} \rightarrow \mathbb{R}$ for which

$$|f^{(r)}(x) - f^{(r)}(y)| \leq |x - y|^\rho, \quad x, y \in \mathbb{R},$$

where $r \in \{0, 1, 2, \dots\}$ and $0 < \rho \leq 1$. It is well known that the optimal worst-case error of L^p -approximation ($1 \leq p \leq \infty$) of such functions that can be achieved from n exact evaluations of f is proportional to $e_n = n^{-(r+\rho)}$. Less obvious is what happens when the functions are piecewise smooth only with unknown break points. Even less obvious is the situation when the function values are additionally corrupted by some noise, i.e., when evaluating the value of f at x we obtain $y = f(x) + \xi$ where $|\xi| \leq \delta$ (deterministic noise) or ξ is a zero-mean random variable of variance σ^2 (random noise). In this talk, we construct an algorithm which despite the presence of noise and break points achieves the worst case L^p -error still proportional to e_n provided the noise level δ or σ is of the same order e_n (except the case of $p = \infty$ and random noise where we have an additional logarithmic factor in the error). The algorithm uses divided differences and special adaptive extrapolation technique to locate the break points and approximate in their neighborhoods.

The Quality of Lattice Sequences

Larysa Matiukha

Illinois Institute of Technology

lmatiukha@hawk.iit.edu

Coauthor(s): Yuhan Ding, Fred J. Hickernell

Lattices are a popular choice of nodes for approximating multidimensional integrals by a sample mean, typically using sample sizes of the form $n = b^m$ for some prime base b and non-negative integer m . However, a computational time budget or hardware failure may prevent us from choosing the preferred number of samples. In this talk, we present an upper

bound on the figure of merit P_α for extensible lattice sequences with arbitrary n , derived in Banach space setting. We show that while the error decays relatively slowly for general n , it improves significantly when n is a small integer multiple of a power of the base, reaching the optimal decay of $\mathcal{O}(n^{-\alpha})$. Additionally, we investigated a related figure of merit $P_{\alpha,2}$ in a Hilbert space, allowing easier computation for arbitrary n . Numerical results for $P_{1,2}$ using both equal and optimally chosen sample weights show a decay of $\mathcal{O}(n^{-1})$ in both cases, while optimal weights lead to a generally smaller and non-increasing error.

Mon, Jul 28 15:30–16:00

Stability of Expected Utility in Bayesian Optimal Experimental Design

Tapio Helin

LUT University

tapio.helin@lut.fi

Coauthor(s): Recent advances in optimization under uncertainty

We explore the stability properties of utility functions in Bayesian optimal experimental design, introducing novel utilities inspired by the analysis. We establish a general framework to study the behavior of expected utility under perturbations in infinite-dimensional setting and prove its convergence properties. Our results provide theoretical guarantees for the robustness of Bayesian design criteria and offer insights into their practical applicability in complex experimental settings.

Karina Koval*Heidelberg University*

karina.koval@iwr.uni-heidelberg.de

TiangangCui, RolandHerzog, RobertScheichl

Coauthor(s): Recent advances in optimization under uncertainty

In this talk, we present a novel approach for sequential optimal experimental design (sOED) for Bayesian inverse problems governed by expensive forward models with large-dimensional unknown parameters. Our goal is to design experiments that maximize the expected information gain (EIG) from prior to posterior, a task that is computationally challenging. This task becomes more complex in sOED, where we must approximate the incremental expected information gain (iEIG) multiple times in distinct stages, often dealing with intractable prior and posterior distributions. To address this, we propose a derivative-based upper bound for iEIG that guides experimental design and enables parameter dimension reduction through likelihood-informed subspaces. By combining this approach with transport map surrogates for the sequence of posteriors, we develop a unified framework for parameter dimension-reduced sOED. We demonstrate the effectiveness of our approach with numerical examples inspired by groundwater flow and photoacoustic imaging.

Randomized quasi-Monte Carlo methods for risk-averse stochastic optimization*Johannes Milz*

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

johannes.milz@isye.gatech.edu

Coauthor(s): Olena Melnikov

Special session: Recent advances in optimization under uncertainty p.60

We establish epigraphical and uniform laws of large numbers for sample-based approximations of law invariant risk functionals. These sample-based approximation schemes include Monte Carlo (MC) and certain randomized quasi-Monte Carlo integration (RQMC) methods, such as scrambled net integration. Our results can be applied to the approximation of risk-averse stochastic programs and risk-averse stochastic variational inequalities. Our numerical simulations empirically demonstrate that RQMC approaches based on scrambled Sobol' sequences can yield smaller bias and root mean square error than MC methods for risk-averse optimization.

Efficient expected information gain estimators based on the randomized quasi-Monte Carlo method

Arved Bartuska

King Abdullah University of Science and Technology/RWTH Aachen University

arved.bartuska@kaust.edu.sa

Coauthor(s): André Gustavo Carlon, Luis Espath, Sebastian Krumscheid, Raúl Tempone

Efficient estimation of the expected information gain (EIG) of an experiment allows for design optimization in a Bayesian setting. This task faces computational challenges, particularly when the experiment model requires numerical discretization schemes. We demonstrate various methods to make such estimations feasible, combining quasi-Monte Carlo (QMC), randomized QMC (rQMC), and multilevel methods.

Analytical error bounds are made possible by Owen's [1] and He et al.'s [2] work on singular integrands combined with a truncation scheme of the observation noise present in experiment models. Applications from Bayesian experimental design demonstrate the improved convergence behavior of the proposed methods compared to traditional Monte Carlo-based estimators.

- [1] Owen, Art B. (2006). *Halton sequences avoid the origin*. SIAM Review, 48:487–503.
- [2] He, Zhijian, & Zheng, Zhan, & Wang, Xiaoqun. (2023). *On the error rate of importance sampling with randomized quasi-Monte Carlo*. SIAM Journal on Numerical Analysis, 61(2):10.1137/22M1510121.

Searching Permutations for Constructing Low-Discrepancy Point Sets and Investigating the Kritzing Sequence

François Clément

Department of Mathematics, University of Washington

fclement@uw.edu

Coauthor(s): Carola Doerr, Kathrin Klamroth, Luís Paquete

This talk focuses on two different approaches for the construction of low-discrepancy sets that are quite different from traditional approaches, yet yield excellent empirical results in

two dimensions. The first of these, which will be the main focus of my talk, is based on selecting the relative position of the different points we wish to place, before using non-linear programming methods to obtain a point set with extremely low star discrepancy. In [1], we showed that this method consistently outperformed all other existing techniques. It is however subject to computational limits: finding good permutation choices, with or without optimization, is the next key step in improving our understanding of low-discrepancy structures.

In the second part of the talk, I will quickly highlight some extended numerical experiments on the sequence introduced by Kritzinger in [2], showing that despite the lack of theoretical results proving that it is a low-discrepancy sequence, it performs at least as well as known sequences in one dimension, despite being constructed greedily.

- [1] F. Clément, C. Doerr, K. Klamroth, L. Paquette (2024). *Transforming the Challenge of Constructing Low-Discrepancy Point Sets into a Permutation Selection Problem.*, to appear, arxiv: <https://arxiv.org/abs/2407.11533>.
- [2] R. Kritzinger (2022). *Uniformly Distributed Sequence generated by a greedy minimization of the L_2 discrepancy.*, Moscow Journal of Combinatorics and Number Theory, **11**(2), 215–236.

Mon, Jul 28 16:00–16:30

Minimizing the Stein Discrepancy

Nathan Kirk

Illinois Institute of Technology

nkirk@iit.edu

Coauthor(s): T. K. Rusch, D. Rus, J. Zech, F. Clément, D. Chen, C. Doerr, L. Paquette

Special session: Computational Methods for Low-discrepancy Sampling and Applications
[p.61](#)

Approximating a probability distribution using a discrete set of points is a fundamental task in modern scientific computation, with applications in uncertainty quantification. Points whose empirical distribution is close to the true distribution are called low-discrepancy. We discuss recent advances in this area, including the use of Stein discrepancies and various optimization techniques. In particular, we introduce Stein-Message-Passing Monte Carlo (Stein-MPMC) [1], an extension of the original Message-Passing Monte Carlo model and the first machine-learning algorithm for generating low-discrepancy (space-filling) point sets. Additionally, we present a generalized Subset Selection algorithm [2], a much simpler yet

highly effective optimization method.

- [1] N. Kirk, T. K. Rusch, J. Zech, D. Rus, *Low Stein Discrepancy via Message-Passing Monte Carlo*, Preprint (2025)
- [2] D. Chen, F. Clément, C. Doerr, N. Kirk, L. Paquette, *Optimizing Kernel Discrepancies via Subset Selection*, Preprint (2025)

Mon, Jul 28 16:30–17:00

Improving Efficiency of Sampling-based Motion Planning via Message-Passing Monte Carlo

Makram Chahine

CSAIL, Massachusetts Institute of Technology

chahine@mit.edu

Coauthor(s): T. Konstantin Rusch, Zach J. Patterson, and Daniela Rus

Sampling-based motion planning methods, while effective in high-dimensional spaces, often suffer from inefficiencies due to irregular sampling distributions, leading to suboptimal exploration of the configuration space. In this talk, we present an approach that enhances the efficiency of these methods by utilizing low-discrepancy distributions generated through Message-Passing Monte Carlo (MPMC). MPMC leverages Graph Neural Networks (GNNs) to generate point sets that uniformly cover the space, with uniformity assessed using the \mathcal{L}_p -discrepancy measure, which quantifies the irregularity of sample distributions. By improving the uniformity of the point sets, our approach significantly reduces computational overhead and the number of samples required for solving motion planning problems. Experimental results demonstrate that our method outperforms traditional sampling techniques in terms of planning efficiency.

Gregory Seljak*Université de Montréal*

gregory.de.salaberry.seljak@umontreal.ca

PierreL'Ecuyer, ChristianeLemieux

Coauthor(s): Computational Methods for Low-discrepancy Sampling and Applications

Randomized quasi-Monte Carlo (RQMC) traditionally takes a low-discrepancy (QMC) point set, makes r independent randomizations of it to obtain r replicates of an unbiased RQMC estimator, then computes the average and variance of these r estimates to obtain a final estimate and perhaps a confidence interval [4]. Some methods construct the points by optimizing refined figures-of-merit adapted to the considered integrand and apply minimal randomization such as a random (digital) shift. Other methods randomize the parameters of the QMC point sets more extensively (e.g., the generating vectors or matrices). The second kind of method is easier to apply because it requires much less knowledge of the integrand, but bad parameter values may be drawn once in a while, leading to (rare) RQMC replicates having a large conditional variance that produce bad outliers. To reduce the impact of such outliers, one approach studied recently is to use the median of the r replicates instead of the mean as a final estimator [2, 3, 5, 6]. Other types of robust estimators could also be used in place of the median [1]. In this talk, we report extensive experiments that compare the mean square errors and convergence of various estimators (the mean, the median, and other robust estimators) defined in terms of r RQMC replicates. We also discuss the computation of confidence intervals for the mean when using such estimators.

- [1] E. Gobet, M. Lerasle, and D. Métivier. Accelerated convergence of error quantiles using robust randomized quasi Monte Carlo methods. <https://hal.science/hal-03631879>, 2024.
- [2] T. Goda and P. L'Ecuyer. Construction-free median quasi-Monte Carlo rules for function spaces with unspecified smoothness and general weights. *SIAM Journal on Scientific Computing*, 44(4):A2765–A2788, 2022.
- [3] T. Goda, K. Suzuki, and M. Matsumoto. A universal median quasi-Monte Carlo integration. *SIAM Journal on Numerical Analysis*, 62(1):533–566, 2024.
- [4] P. L'Ecuyer, M. Nakayama, A. B. Owen, and B. Tuffin. Confidence intervals for randomized quasi-Monte Carlo estimators. In *Proceedings of the 2023 Winter Simulation Conference*, pages 445–456. IEEE Press, 2023.
- [5] Z. Pan and A. B. Owen. Super-polynomial accuracy of one dimensional randomized nets using the median of means. *Mathematics of Computation*, 92(340):805–837, 2023.

- [6] Z. Pan and A. B. Owen. Super-polynomial accuracy of multidimensional randomized nets using the median of means. *Mathematics of Computation*, 93(349):2265–2289, 2024.

Tue, Jul 29 10:30–11:00

Computing the stationary measure of McKean-Vlasov SDEs

Jean-François CHASSAGNEUX

ENSAE-CREST & IP Paris

chassagneux@ensae.fr

Coauthor(s): Gilles Pagès

Special session: Stochastic Computation and Complexity [p.63](#)

Under some confluence assumption, it is known that the stationary distribution of a McKean-Vlasov SDE is the limit of the empirical measure of its associated self-interacting diffusion. Our numerical method consists of introducing the Euler scheme with decreasing step size of this self-interacting diffusion and seeing its empirical measure as the approximation of the stationary distribution of the original McKean-Vlasov SDEs. This simple approach is successful (under some reasonable assumptions) as we are able to prove convergence with a rate for the Wasserstein distance between the two measures both in the L2 and almost sure sense. In this talk, I will first explain the rationale behind this approach and I will discuss the various convergence results we have obtained so far.

On the convergence of the Euler-Maruyama scheme for McKean-Vlasov SDEs*Noufel Frikha*

Université Paris 1 Panthéon-Sorbonne

noufel.frikha@univ-paris1.fr

Coauthor(s): Clément Rey, Xuanye Song

Relying on the backward Kolmogorov PDE stated on the Wasserstein space, we obtain several new results concerning the approximation error of some non-linear diffusion process in the sense of McKean-Vlasov by the corresponding Euler-Maruyama discretization scheme of its system of interacting particles. We notably present explicit error estimates, at the level of the trajectories, at the level of the semigroup (stated on the Wasserstein space) and at the level of the densities. Some Gaussian density estimates of the transition density and its first-order derivative for the Euler-Maruyama scheme are also established. This presentation is based on joint works with Clément Rey (Ecole Polytechnique) and Xuanye Song (Université Paris Cité).

- [1] Frikha, N. & Song, X. (2025). *On the convergence of the Euler-Maruyama scheme for McKean-Vlasov SDEs*, arXiv:2503.22226.
- [2] Frikha, N. & Rey, C. (2025). *On the weak convergence of the Euler-Maruyama scheme for McKean-Vlasov SDEs: expansion of the densities*.

Wasserstein Convergence of Score-based Generative Models under Semiconvexity and Discontinuous Gradients*Sotirios Sabanis*

University of Edinburgh & National Technical University of Athens & Archimedes/Athena Research Centre

s.sabanis@ed.ac.uk

Coauthor(s): Stefano Bruno

Special session: Stochastic Computation and Complexity [p.63](#)

Score-based Generative Models (SGMs) approximate a data distribution by perturbing it with Gaussian noise and subsequently denoising it via a learned reverse diffusion process.

These models excel at modeling complex data distributions and generating diverse samples, achieving state-of-the-art performance across domains such as computer vision, audio generation, reinforcement learning, and computational biology. Despite their empirical success, existing Wasserstein-2 convergence analysis typically assumes strong regularity conditions—such as smoothness or strict log-concavity of the data distribution—that are rarely satisfied in practice.

In this work, we establish the first non-asymptotic Wasserstein-2 convergence guarantees for SGMs targeting semiconvex distributions with potentially discontinuous gradients. Our upper bounds are explicit and sharp in key parameters, achieving optimal dependence of $O(\sqrt{d})$ on the data dimension d and convergence rate of order one. The framework accommodates a wide class of practically relevant distributions, including symmetric modified half-normal distributions, Gaussian mixtures, double-well potentials, and elastic net potentials. By leveraging semiconvexity without requiring smoothness assumptions on the potential such as differentiability, our results substantially broaden the theoretical foundations of SGMs, bridging the gap between empirical success and rigorous guarantees in non-smooth, complex data regimes.

- [1] Bruno, Stefano & Sabanis, Sotirios (2025). *Wasserstein Convergence of Score-based Generative Models under Semiconvexity and Discontinuous Gradients*. ArXiv.

Tue, Jul 29 10:30–11:00

Optimal Pilot Sampling for Multi-fidelity Monte Carlo Methods

Xun Huan

University of Michigan, Department of Mechanical Engineering

xhuan@umich.edu

Coauthor(s): Thomas Coons, Aniket Jivani

Bayesian optimal experimental design (OED) aims to maximize an expected utility, often chosen to be the expected information gain (EIG), over a given design space. Estimating EIG typically relies on Monte Carlo methods, which require repeated evaluations of a computational model simulating the experimental process. However, when the model is expensive to evaluate, standard Monte Carlo becomes impractical.

Multi-fidelity variants of Monte Carlo, such as Approximate Control Variate (ACV) estimators, can significantly expedite such estimations by leveraging an ensemble of low-fidelity models that approximate the high-fidelity model with varying degrees of accuracy and cost. To apply these techniques in an error-optimal manner, the covariance matrix across model

outputs must be estimated from independent pilot model evaluations. This step incurs a significant but often overlooked computational cost. Furthermore, the optimal allocation of computational resources between covariance estimation and ACV estimation remains an open problem. Existing approaches fail to accommodate optimal estimators and may not be accurate with small pilot sample sizes.

In this work, we introduce a novel framework for dynamically allocating resources between these two tasks. Our method employs Bayesian inference to quantify uncertainty in the covariance matrix and derives an adaptive expected loss metric to determine when to terminate pilot sampling. We demonstrate and analyze our framework through a benchmark nonlinear OED problem.

Tue, Jul 29 11:00–11:30

A recursive Monte Carlo approach to optimal Bayesian experimental design

Adrien Corenflos

Department of Statistics, University of Warwick

adrien.corenflos@warwick.ac.uk

Coauthor(s): Hany Abdulsamad, Sahel Iqbal, Sara Pérez-Vieites, Simo Särkkä

Special session: Next-generation optimal experimental design: theory, scalability, and real world impact: Part I [p.64](#)

Bayesian experimental design is concerned with designing experiments that maximize information on a latent parameter of interest. This can be formally understood as minimizing the *expected* entropy over the parameter, given the input, the expectation being taken over the data. Solving this problem is intractable *as is*, and several surrogate loss functions have been proposed to learn policies that, when deployed in nature, help inference by sampling more informative data. A drawback of most of these, however, is that they introduce a substantial amount of bias, or otherwise exhibit a high variance. In this talk, we will introduce another surrogate formulation of optimal Bayesian design as a risk-sensitive policy optimization, compatible with non-exchangeable models. Under this formulation, minimizing the entropy of the posterior can be understood as sampling from a posterior distribution over the (random) designs. We will then discuss two nested sequential Monte Carlo algorithms [1,2] to infer these optimal designs, and discuss how to embed them within a particle Markov chain Monte Carlo framework to perform gradient-based policy learning. We will discuss the respective advantages and drawbacks of both algorithms as well as those of alternative methods.

[1] Iqbal, S., Corenflos, A., Särkkä, S., & Abdulsamad, H.(2024). *Nesting Particle Filters*

for *Experimental Design in Dynamical Systems*. In Forty-first International Conference on Machine Learning (ICML).

- [2] Iqbal, S., Abdulsamad, H., Pérez-Vieites, S., Särkka, S., & Corenflos, A. (2024). *Recursive nested filtering for efficient amortized Bayesian experimental design*. In NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty.

Tue, Jul 29 11:30–12:00

Ayoub Belhadji

Massachusetts Institute of Technology
abelhadj@mit.edu

DanielSharp, YoussefMarzouk

Coauthor(s): Next-generation optimal experimental design: theory, scalability, and real world impact: Part I

Approximating a probability distribution using a set of particles is a fundamental problem in machine learning and statistics, with applications including quantization and optimal design. Formally, we seek a finite weighted mixture of Dirac measures that best approximates the target distribution. While much existing work relies on the Wasserstein distance to quantify approximation errors, maximum mean discrepancy (MMD) has received comparatively less attention, especially when allowing for variable particle weights. We study the quantization problem from the perspective of minimizing MMD via gradient flow in the Wasserstein–Fisher–Rao (WFR) geometry. This gradient flow yields an ODE system from which we further derive a fixed-point algorithm called *mean shift interacting particles* (MSIP). We show that MSIP extends the non-interacting mean shift algorithm, widely used for identifying modes in kernel density estimates. Moreover, we show that MSIP can be interpreted as preconditioned gradient descent, and that it acts as a relaxation of Lloyd’s algorithm for clustering. Our numerical experiments demonstrate that MSIP and the WFR ODEs outperform other algorithms for quantization of multi-modal and high-dimensional targets.

Parallel computations for Metropolis Markov chains based on Picard maps*Sebastiano Grazzi*

Department of Decision Sciences and BIDSa, Bocconi University, Via Roentgen, 1, 20136, Milan, Italy

sebastiano.grazzi@unibocconi.it

Coauthor(s): Giacomo Zanella

Special session: Heavy-tailed Sampling p.66

We develop parallel algorithms for simulating zeroth-order Metropolis Markov chains based on the Picard map. For Random Walk Metropolis Markov chains targeting log-concave distributions π on \mathbb{R}^d , our algorithm generates samples close to π in $\mathcal{O}(\sqrt{d})$ iterations with $\mathcal{O}(\sqrt{d})$ parallel processors, therefore speeding-up the convergence of the corresponding sequential implementation by a factor \sqrt{d} . Furthermore, a modification of our algorithm generates samples from an approximate measure $\tilde{\pi}_\epsilon$ in $\mathcal{O}(1)$ iterations and $\mathcal{O}(d)$ parallel processors. In this talk, I will present the methodology, the analysis and numerical simulations. Our algorithms are straightforward to implement and may constitute a useful tool for practitioners seeking to sample from a prescribed distribution π using only point-wise evaluations proportional to π .

A large deviation principle for Metropolis-Hastings sampling

Federica Milinanni

KTH Royal Institute of Technology

fedmil@kth.se

Coauthor(s): Pierre Nyquist

Special session: Heavy-tailed Sampling [p.66](#)

Sampling algorithms from the class of Markov chain Monte Carlo (MCMC) methods are widely used across scientific disciplines. Good performance measures are essential to analyse these methods, to compare different MCMC algorithms, and to tune parameters within a given method. Common tools that are used for analysing convergence properties of MCMC algorithms are, e.g., mixing times, spectral gap and functional inequalities (e.g., Poincaré, log-Sobolev). A further, rather novel, approach consists of the use of large deviations theory to study the convergence of empirical measures of MCMC chains. At the heart of large deviations theory is the large deviation principle, which allows us to describe the rate of convergence of the empirical measures through a so-called rate function.

In this talk, we will consider Markov chains generated via MCMC methods of Metropolis-Hastings type for sampling from a target distribution on a Polish space. We will state a large deviation principle for the corresponding empirical measure, show examples of algorithms from this class for which the theorem applies, and illustrate how the result can be used to tune algorithms' parameters.

Sharp Characterization and Control of Global Dynamics of SGDs with Heavy Tails

Xingyu Wang

University of Amsterdam

x.wang4@uva.nl

Coauthor(s): Chang-Han Rhee, Sewoong Oh

Special session: Heavy tailed sampling [p.66](#)

The empirical success of deep learning is often attributed to the mysterious ability of stochastic gradient descents (SGDs) to avoid sharp local minima in the loss landscape, as sharp minima are believed to lead to poor generalization. To unravel this mystery and potentially further enhance such capability of SGDs, it is imperative to go beyond the traditional local convergence analysis and obtain a comprehensive understanding of SGDs' global dynamics within complex non-convex loss landscapes. In this talk, we characterize the global dynamics of SGDs through the heavy-tailed large deviations and local stability framework. This framework systematically characterizes the rare events in heavy-tailed dynamical systems; building on this, we characterize intricate phase transitions in the first exit times, which leads to the heavy-tailed counterparts of the classical Freidlin-Wentzell and Eyring-Kramers theories. Moreover, applying this framework to SGD, we reveal a fascinating phenomenon in deep learning: by injecting and then truncating heavy-tailed noises during the training phase, SGD can almost completely avoid sharp minima and hence achieve better generalization performance for the test data.

Functional estimation of the marginal likelihood*Jonathan Weare*

Courant Institute of Mathematical Sciences, New York University

weare@nyu.edu

Coauthor(s): Omiros Papaspiliopoulos, Timothée Stumpf-Fétizon

We propose a framework for computing, optimizing and integrating with respect to a smooth marginal likelihood in statistical models that involve high-dimensional parameters/latent variables and continuous low-dimensional hyperparameters. The method requires samples from the posterior distribution of the parameters for different values of the hyperparameters on a simulation grid and returns inference on the marginal likelihood defined everywhere on its domain, and on its functionals. We explain the relationship between the method and many of the methods that have been used in this context, including sequential Monte Carlo, Gibbs sampling, Monte Carlo maximum likelihood, and umbrella sampling. We establish the consistency of the proposed estimators as the sampling effort increases, both when the simulation grid is kept fixed and when it becomes dense in the domain. We showcase the approach on Gaussian process regression and classification and crossed effect models.

Randomized QMC Methods via Combinatorial Discrepancy

Nikhil Bansal

University of Michigan

bansaln@umich.edu

Coauthor(s): Haotian Jiang

Special session: Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods
[p.68](#)

A folklore result in geometric discrepancy theory, called the Transference Principle, allows one to use methods from combinatorial discrepancy to show the existence of good QMC point sets with low star discrepancy. Unfortunately, this connection was not of much use until recently, as most results in combinatorial discrepancy were non-algorithmic. Recently however, there has been a major revolution in algorithmic combinatorial discrepancy [1].

In this talk, I will explain how recent innovations in combinatorial discrepancy can be used to give an algorithm to produce randomized QMC point sets that achieve substantially better error than given by the classical Koksma-Hlawka inequality. Moreover, the algorithm only requires random samples, as opposed to carefully chosen points, and also optimally combines the best features of both MC and QMC methods.

The talk does not require any background on combinatorial discrepancy and will be a gentle introduction to the area. These results appear in the paper [2].

- [1] Bansal, Nikhil (2022). *Discrepancy and Related Algorithms*. Survey article accompanying an invited talk at International Congress of Mathematicians (ICM) 2022.
- [2] Bansal, Nikhil and Jiang, Haotian (2025). *Quasi-Monte Carlo Beyond Hardy-Krause*. Proceedings of Symposium on Discrete Algorithms (SODA). Invited to Journal of the ACM (JACM).

The Walk on Spheres Monte Carlo Algorithm for Solving Partial Differential Equations

Michael Mascagni

Florida State University and NIST, USA

mascagni@fsu.edu

Coauthor(s): Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods

The stochastic representation of the solutions of partial differential equations (PDEs) and integral equations (IEs) has been known for decades. These representations have led to a variety of Monte Carlo methods for the numerical solution of both PDEs and IEs, among them is a method called the Walk on Spheres (WoS) algorithm. WoS has been quite successful in solving linear elliptic and parabolic PDEs with Dirichlet boundary conditions in a very cost-effective way. This numerical analysis work was taken up by the computer graphics community, who are experts in Monte Carlo through their use of it in image rendering via ray tracing. Their interest stems from the fact that WoS and ray tracing both permit the representation of the geometry of the problem in a very compact and computational efficient way via Bounding Volume Hierarchy. This has led to the expansion of the scope and efficiency of WoS in a wide variety of ways which will be presented here.

Enhancing Gaussian Process Surrogates for Optimization and Posterior Approximation via Random Exploration

Hwanwoo Kim

Duke University

hwanwoo.kim@duke.edu

Coauthor(s): Daniel Sanz-Alonso

Special session: Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods
p.68

In this talk, we propose novel noise-free Bayesian optimization strategies that rely on a random exploration step to enhance the accuracy of Gaussian process surrogate models. The new algorithms retain the ease of implementation of the classical GP-UCB algorithm, but the additional random exploration step accelerates their convergence, nearly achieving the optimal convergence rate. Furthermore, to facilitate Bayesian inference with an intractable likelihood, we propose to utilize the optimization iterates as design points to build a Gaussian process surrogate model for the unnormalized log-posterior density. We show that the Hellinger distance between the true and the approximate posterior distributions decays at a near-optimal rate. We demonstrate the effectiveness of our algorithms in benchmarking non-convex problems test functions for optimization, and in a black-box engineering design problem. We also showcase the effectiveness of our posterior approximation approach in Bayesian inference for parameters of dynamical systems.

Optimal strong approximation of SDEs with Hölder continuous drift coefficient

Larisa Yaroslavtseva

University of Graz

larisa.yaroslavtseva@uni-graz.at

Coauthor(s): Simon Ellinger and Thomas Müller-Gronbach

We study strong approximation of the solution of a scalar stochastic differential equation (SDE)

$$\begin{aligned} dX_t &= \mu(X_t) dt + dW_t, \quad t \in [0, 1], \\ X_0 &= x_0 \end{aligned} \tag{6.1}$$

at the final time point 1 in the case that the drift coefficient μ is α -Hölder continuous with $\alpha \in (0, 1]$. Recently, it was shown in [1] that for such SDEs the equidistant Euler approximation achieves an L^p -error rate of at least $(1 + \alpha)/2$, up to an arbitrary small ε , in terms of the number of evaluations of the driving Brownian motion W . In this talk we present a matching lower error bound. More precisely, we show that the L^p -error rate $(1 + \alpha)/2$ can not be improved in general by no numerical method based on finitely many evaluations of W at fixed time points. For the proof of this result we choose μ to be the Weierstrass function and we employ the coupling of noise technique recently introduced in [2].

- [1] Butkovsky, O., Dareiotis, K., & Gerencsér, M. (2021). Approximation of SDEs: a stochastic sewing approach. *Probab. Theory Related Fields*, **181**(4), 975–1034
- [2] Müller-Gronbach, T. & Yaroslavtseva, L. (2023). Sharp lower error bounds for strong approximation of SDEs with a discontinuous drift coefficient by coupling of noise. *Ann. Appl. Probab.* **33**, 902–935.

Tue, Jul 29 16:00–16:30

Tractability of L_2 -approximation and integration in weighted Hermite spaces of finite smoothness

Gunther Leobacher

University of Graz

gunther.leobacher@uni-graz.at

Coauthor(s): Adrian Ebert, Friedrich Pillichshammer

Special session: Stochastic Computation and Complexity p.69

We consider integration and L_2 -approximation for functions over \mathbb{R}^s from weighted Hermite spaces as introduced in [1]. We study tractability of the integration and L_2 -approximation problem for the different Hermite spaces, which describes the growth rate of the information complexity when the error threshold ε tends to 0 and the problem dimension s grows to infinity. Our main results are characterizations of tractability in terms of the involved weights, which model the importance of the successive coordinate directions for functions from the weighted Hermite spaces.

- [1] Ch. Irrgeher and G. Leobacher. High-dimensional integration on the \mathbb{R}^d , weighted Hermite spaces, and orthogonal transforms. *J. Complexity* 31: 174–205, 2015.

Tue, Jul 29 16:30–17:00

Malliavin differentiation of Lipschitz SDEs and BSDEs and an Application to Quadratic Forward-Backward SDEs

Alexander Steinicke

University of Leoben

alexander.steinicke@unileoben.ac.at

Coauthor(s): Hannah Geiss, Céline Labart, Adrien Richou

Special session: Stochastic Computation and Complexity p.69

Geiss and Zhou [1] showed that SDEs and BSDEs with Lipschitz generators admit Malliavin differentiability in the Brownian setting. We extend and apply this result in the Lévy case and present a differentiation formula for coefficients that are Lipschitz in the solution variable with respect to the Skorohod metric. The obtained formula then allows us to show the existence and uniqueness of solutions to a class of quadratic and superquadratic forward-backward SDE systems.

- [1] Geiss, S. and Zhou, X. (2024). *Coupling of Stochastic Differential Equations on the Wiener Space*. <https://arxiv.org/pdf/2412.10836>.

Tue, Jul 29 17:00–17:30

A Unified Treatment of Tractability for Approximation Problems Defined on Hilbert Spaces

Fred J. Hickernell

Illinois Institute of Technology

hickernell@iit.edu

Coauthor(s): Onyekachi Emenike and Peter Kritzer

The presenter and co-authors showed in [2] how to unify many of the equivalent conditions for various notions of tractability for approximation problems defined on Hilbert spaces. This talk describes those results.

Let $\{\mathcal{F}_d\}_{d \in \mathbb{N}}$ and $\{\mathcal{G}_d\}_{d \in \mathbb{N}}$ be sequences of Hilbert spaces, and let $\{\text{SOL}_d : \mathcal{F}_d \rightarrow \mathcal{G}_d\}_{d \in \mathbb{N}}$ be a sequence of linear solution operators with positive singular values $\lambda_{1,d} \geq \lambda_{2,d} \geq \dots$. For

input functions lying in \mathcal{B}_d , the *unit ball* in \mathcal{F}_d , and when arbitrary linear functionals are available, the information complexity of this problem is known to be $\text{scomp}(\varepsilon, d) = \min\{n \in \mathbb{N}_0 : \lambda_{n+1,d} \leq \varepsilon\}$ (see [3]). Different notions of tractability are defined in terms of how quickly $\text{comp}(\varepsilon, d)$ increases as ε^{-1} and/or d tend to infinity.

Equivalence results on tractability replace a condition on ordered singular values by a more accessible condition in terms of weighted sums of powers of singular values. The theorems in [2] unify these arguments by providing proofs that cover a variety of special cases.

The work in [1] presents a setting where the input functions lie not in the unit ball, but in a cone, \mathcal{C}_d , of tame functions. We speculate on what the conditions on the singular values might be equivalent to various notions of tractability for such classes of input functions.

1. Y. Ding, F. J. Hickernell, P. Kritzer, and S. Mak, *Adaptive approximation for multivariate linear problems with inputs lying in a cone*, Multivariate Algorithms and Information-Based Complexity (F. J. Hickernell and P. Kritzer, eds.), DeGruyter, Berlin/Boston, 2020, pp. 109–145.
2. O. Emenike, F. J. Hickernell, and P. Kritzer, *A unified treatment of tractability for approximation problems defined on Hilbert spaces*, J. Complexity **84** (2024), 101856.
3. E. Novak and H. Woźniakowski, *Tractability of multivariate problems Volume I: Linear information*, EMS Tracts in Mathematics, no. 6, European Mathematical Society, Zürich, 2008.

Tue, Jul 29 15:30–16:00

Goal-Oriented Sensor Placement for Infinite-Dimensional Bayesian Inverse Problems

Alen Alexanderian

North Carolina State University

alexanderian@ncsu.edu

Coauthor(s): Next-generation optimal experimental design: theory, scalability, and real world impact

We consider optimal experimental design (OED) for infinite-dimensional Bayesian inverse problems governed by partial differential equations with infinite-dimensional inversion parameters. Specifically, we focus on the case where we seek sensor placements that minimize the uncertainty in a prediction or goal functional. To address this, we propose a goal-oriented OED (gOED) approach that uses a quadratic approximation of the parameter-to-prediction mapping to obtain a measure of posterior uncertainty in the prediction quantity of interest

(QoI). We focus on linear inverse problems in which the prediction is a nonlinear functional of the inversion parameters. We seek to find sensor placements that result in minimized posterior variance of the prediction QoI. In this context, and under the assumption of Gaussian prior and noise models, we derive a closed-form expression for the gOED criterion. We also discuss efficient and accurate computational approaches for computing the gOED objective and its optimization. We illustrate the proposed approach in model inverse problems governed by an advection-diffusion equation.

Tue, Jul 29 16:00–16:30

Bayesian Experimental Design (BED) is a powerful tool to reduce the cost of running a sequence of experiments. When based on the Expected Information Gain (EIG), design optimization corresponds to the maximization of some intractable expected contrast between prior and posterior distributions. Scaling this maximization to high dimensional and complex settings has been an issue due to BED inherent computational complexity. In this work, we introduce a pooled posterior distribution with cost-effective sampling properties and provide a tractable access to the EIG contrast maximization via a new EIG gradient expression. Diffusion-based samplers are used to compute the dynamics of the pooled posterior and ideas from bi-level optimization are leveraged to derive an efficient joint sampling-optimization loop. The resulting efficiency gain allows to extend BOED to the well-tested generative capabilities of diffusion models. By incorporating generative models into the BOED framework, we expand its scope and its use in scenarios that were previously impractical. Numerical experiments and comparison with state-of-the-art methods show the potential of the approach. As a practical application, we showcase how our method accelerates Magnetic Resonance Imaging (MRI) acquisition times while preserving image quality. This presentation will also detail how Diffuse, a new modifiable Python package for diffusion models facilitate composability and research in diffusion models through its simple and intuitive API, allowing researchers to easily integrate and experiment with various model components.

Robust Bayesian Optimal Experimental Design under Model Misspecification

Tommie A. Catanach

Sandia National Laboratories

tacatan@sandia.gov

Coauthor(s): Next-generation optimal experimental design: theory, scalability, and real world impact: Part II

Bayesian Optimal Experimental Design (BOED) has become a powerful tool for improving uncertainty quantification by strategically guiding data collection. However, the reliability of BOED depends critically on the validity of its underlying assumptions and the possibility of model discrepancy. In practice, the chosen data acquisition strategy may inadvertently reinforce prior assumptions—overlooking data that could challenge them—or rely on low-fidelity models whose error is not well characterized, leading to biased inferences. These biases can be particularly severe because BOED often targets extreme parameter regions as the most “informative,” potentially magnifying the impact of model error.

In this talk, we present a new information criterion, Expected Generalized Information Gain (EGIG)[1], that explicitly accounts for model discrepancy in BOED. EGIG augments standard Expected Information Gain by balancing the trade-off between experiment performance (i.e., how much information is gained) and robustness (i.e., how susceptible the design is to model misspecification). Concretely, EGIG measures how poorly inference under an incorrect model might perform, compared to a more appropriate model for the experiment. We will discuss the theoretical underpinnings of EGIG, as well as nested Monte Carlo algorithms for incorporating it into BOED for nonlinear inference problems. These methods handle both quantifiable discrepancies (e.g., low-fidelity vs. high-fidelity models) and unknown discrepancies represented by a distribution of potential errors, thereby enhancing the robustness and reliability of BOED in real-world settings.

SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

- [1] Catanach, T. A., & Das, N. (2023). *Metrics for Bayesian optimal experiment design under model misspecification*. In 2023 62nd IEEE Conference on Decision and Control (CDC) (pp. 7707-7714). IEEE.

Revisiting self-normalized importance sampling: new methods and diagnostics

Nicola Branchini

University of Edinburgh

n.branchini@sms.ed.ac.uk

Coauthor(s): Víctor Elvira

Importance sampling (IS) can often be implemented only with normalized weights, yielding the popular self-normalized IS (SNIS) estimator. However, proposal distributions are often learned and evaluated using criteria designed for the unnormalized IS (UIS) estimator.

In this talk, we aim to present a unified perspective on recent methodological advances in understanding and improving SNIS. We propose and compare two new frameworks for adaptive importance sampling (AIS) methods tailored to SNIS. Our first framework exploits the view of SNIS as a ratio of two UIS estimators, coupling two separate AIS samplers in a joint distribution selected to minimize asymptotic variance. Our second framework instead proposes the first MCMC-driven AIS sampler directly targeting the (often overlooked) optimal SNIS proposal.

We also establish a close connection between the optimal SNIS proposal and so-called subtractive mixture models (SMMs), where negative coefficients are possible - motivating the study of the properties of the first IS estimators using SMMs.

Finally, we propose new Monte Carlo diagnostics specifically for SNIS. They extend existing diagnostics for numerator and denominator by incorporating their statistical dependence, drawing on different notions of tail dependence from multivariate extreme value theory.

- [1] Branchini, N., & Elvira, V. (2024). Generalizing self-normalized importance sampling with couplings. arXiv preprint arXiv:2406.19974.
- [2] Branchini, N., & Elvira, V. (2025). Towards adaptive self-normalized importance samplers. In submission at Statistical Signal Processing Workshop (SSP), 2025.
- [3] **Zellinger, L. & Branchini, N. (equal contribution)**, Elvira, V., & Vergari, A. Scalable expectation estimation with subtractive mixture models. In submission at Frontiers in Probabilistic Inference: Learning meets Sampling (workshop at ICLR 2025).
- [4] Branchini, N., & Elvira, V. The role of tail dependence in estimating posterior expectations. In NeurIPS 2024 Workshop on Bayesian Decision-making and Uncertainty.

Asymptotic robustness of smooth functions of rare-event estimators

Bruno Tuffin

Inria

bruno.tuffin@inria.fr

Coauthor(s): Marvin K. Nakayama

Special session: Advances in Rare Events Simulation p.72

In many rare-event simulation problems, an estimand is expressed as a smooth function of several quantities, each estimated by simulation but not necessarily all of their estimators are critically influenced by the rarity of the event of interest. An example arises in the estimation of the mean time to failure of a regenerative system, usually expressed as the ratio of two quantities to be estimated, the denominator being the only one entailing a rare event in a highly reliable context.

In general, there has been to our knowledge no work investigating the efficiency of estimating $\alpha = g(\boldsymbol{\theta})$ for some known smooth function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ and where $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_d) \in \mathbb{R}^d$ is a vector of unknown parameters, for some $d \geq 1$, each estimated by simulation.

We will provide during the talk conditions under which having efficient estimators of each individual mean leads to an efficient estimator of the function of the means. Our conditions are described for several asymptotic robustness properties: logarithmic efficiency, bounded relative error and vanishing relative error. We illustrate this setting through several examples, and numerical results complement the theory.

- [1] Nakayama, Marvin K. and Bruno Tuffin. Efficiency of Estimating Functions of Means in Rare-Event Contexts. In the *Proceedings of the 2023 Winter Simulation Conference*, San Antonio, TX, USA, December 2023.

Importance Sampling Methods with Stochastic Differential Equations for the Estimation of the Right Tail of the CCDF of the Fade Duration

Eya Ben Amar

King Abdullah University of Science and Technology

eya.benamar@kaust.edu.sa

Coauthor(s): Nadhir Ben Rached, Raúl Tempone, and Mohamed-Slim Alouini

In this work, we explore using stochastic differential equations (SDEs) to study the performance of wireless communication systems. Particularly, we investigate the fade duration metric, representing the time during which the signal remains below a specified threshold within a fixed time interval. We estimate the complementary cumulative distribution function (CCDF) of the fade duration using Monte Carlo simulations. We propose an optimal importance sampling (IS) estimator that provides accurate estimates of the CCDF tail, corresponding to rare event probabilities, at a significantly reduced computational cost. Additionally, we introduce a novel multilevel Monte Carlo method combined with IS and discuss its efficiency in further reducing the computational cost of the IS estimator.

Estimating rare event probabilities associated with McKean–Vlasov SDEs

Shyam Mohan Subbiah Pillai

RWTH Aachen University, Germany

subbiah@uq.rwth-aachen.de

Coauthor(s): Nadhir Ben Rached, Abdul-Lateef Haji-Ali, Raúl Tempone

Special session: Advances in Rare Events Simulation [p.72](#)

McKean–Vlasov stochastic differential equations (MV-SDEs) arise as the mean-field limits of stochastic interacting particle systems, with applications in pedestrian dynamics, collective animal behavior, and financial market modelling. This work develops an efficient method for estimating rare event probabilities associated with MV-SDEs by combining multilevel Monte Carlo (MC) with importance sampling (IS). To apply a measure change for IS, we first reformulate the MV-SDE as a standard SDE by conditioning on its law, leading to the decoupled MV-SDE. We then formulate the problem of finding the optimal IS measure change as a stochastic optimal control problem that minimizes the variance of the MC esti-

mator. The resulting partial differential equation is solved numerically to obtain the optimal IS measure change. Building on this IS scheme and the decoupling approach, we introduce a double loop Monte Carlo (DLMC) estimator. To further improve computational efficiency, we extend DLMC to a multilevel setting, reducing its computational complexity. To enhance variance convergence in the level differences for the discontinuous indicator function, we propose two key techniques: (1) numerical smoothing via one-dimensional integration over a carefully chosen variable and (2) an antithetic sampler to increase correlation between fine and coarse SDE paths. By integrating IS with efficient multilevel sampling, we develop the multilevel double loop Monte Carlo (MLDLMC) estimator. We demonstrate its effectiveness on the Kuramoto model from statistical physics, showing a reduction in computational complexity from $\mathcal{O}(\text{TOL}_r^{-4})$ using DLMC to $\mathcal{O}(\text{TOL}_r^{-3})$ using MLDLMC with IS, for estimating rare event probabilities up to a prescribed relative error tolerance TOL_r . This talk is primarily based on [1,2].

- [1] Ben Rached, N., Haji-Ali, A. L., Subbiah Pillai, S. M., & Tempone, R. (2024). Double-loop importance sampling for McKean–Vlasov stochastic differential equation. *Statistics and Computing*, 34(6), 197.
- [2] Ben Rached, N., Haji-Ali, A. L., Subbiah Pillai, S. M., & Tempone, R. (2025). Multilevel importance sampling for rare events associated with the McKean–Vlasov equation. *Statistics and Computing*, 35(1), 1.

Tue, Jul 29 15:30–16:00

Quasi-uniform quasi-Monte Carlo digital nets

Takashi Goda

Graduate School of Engineering, The University of Tokyo

goda@frcer.t.u-tokyo.ac.jp

Coauthor(s): Josef Dick, Kosuke Suzuki

Special session: Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods
[p.73](#)

We investigate the quasi-uniformity properties of digital nets, a class of quasi-Monte Carlo point sets. Quasi-uniformity is a space-filling property that plays a crucial role in applications such as designs of computer experiments and radial basis function approximation. However, it remains open whether common low-discrepancy digital nets satisfy quasi-uniformity.

In this talk, we introduce the concept of *well-separated* point sets as a tool for constructing quasi-uniform low-discrepancy digital nets. We establish an algebraic criterion to determine

whether a given digital net is well-separated and use this criterion to construct an explicit example of a two-dimensional digital net that is both low-discrepancy and quasi-uniform. Furthermore, we present counterexamples of low-discrepancy digital nets that fail to achieve quasi-uniformity, highlighting the limitations of existing constructions.

- [1] T. Goda, The Sobol' sequence is not quasi-uniform in dimension 2. *Proc. Amer. Math. Soc.*, 152(8):3209–3213, 2024.
- [2] J. Dick, T. Goda, & K. Suzuki, On the quasi-uniformity properties of quasi-Monte Carlo digital nets and sequences. *arXiv preprint arXiv:2501.18226*, 2025.

Tue, Jul 29 16:00–16:30

TALK WITHDRAWN—Boosting the inference for generative models by (Quasi-)Monte Carlo resampling

Ziang Niu

University of Pennsylvania

ziangniu@wharton.upenn.edu

Coauthor(s): Bhaswar B. Bhattacharya, François-Xavier Briol, Anirban Chatterjee, Johanna Meier.

Special session: Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods
p.73

In the era of generative models, statistical inference based on classical likelihood for such models has faced a challenge. This is due to highly nontrivial model structures and thus computing the likelihood functions is almost impossible. Often time, information on these generative models can only be obtained by sampling from the models but the necessity of sampling can further pose a tradeoff between computational burden and statistical accuracy. In this talk, we propose a framework for statistical inference for generative models leveraging the techniques from (Quasi-)Monte Carlo. Despite the unavoidable balance of statistical accuracy and computation, both computational and statistical performances can be boosted by employing (Quasi-)Monte Carlo techniques. The presentation will be based on two papers.

- [1] A. Chatterjee, Z. Niu & B. Bhattacharya (2024). *A kernel-based conditional two-sample test using nearest neighbors (with applications to calibration, regression curves, and simulation-based inference)*. Preprint.
- [2] Z. Niu, J. Meier & F-X. Briol (2023). *Discrepancy-based inference for intractable generative models using Quasi-Monte Carlo*. *Electronic Journal of Statistics*.

A hit-and-run approach for sampling and analyzing ranking models

Chenyang Zhong

Department of Statistics, Columbia University

cz2755@columbia.edu

Coauthor(s): Frontiers in (Quasi-)Monte Carlo and Markov Chain Monte Carlo Methods

The analysis of ranking data has gained much recent interest across various applications, including recommender systems, market research, and electoral studies. This talk focuses on the Mallows permutation model, a probabilistic model for ranking data introduced by C. L. Mallows. The Mallows model specifies a family of non-uniform probability distributions on permutations and is characterized by a distance metric on permutations. We focus on two popular choices: the L^1 (Spearman's footrule) and L^2 (Spearman's rank correlation) distances. Despite their widespread use in statistics and machine learning, Mallows models with these metrics present significant computational challenges due to the intractability of their normalizing constants.

Hit and run algorithms form a broad class of MCMC algorithms, including Swendsen-Wang and data augmentation. In this talk, I will first explain how to sample from Mallows models using hit and run algorithms. For both models, we establish $O(\log n)$ mixing time upper bounds, which provide the first theoretical guarantees for efficient sampling and enable computationally feasible Monte Carlo maximum likelihood estimation. Then, I will also discuss how the hit and run algorithms can be utilized to prove theorems about probabilistic properties of the Mallows models.

Interior-Point Frank-Wolfe (IPFW) for Linearly Constrained Functional Optimization Over Probability Spaces

Di Yu

Purdue University

pasupath@purdue.edu

Coauthor(s): Raghu Pasupathy, S. G. Henderson

Many challenging problems in statistics and applied probability can be posed as that of minimizing a smooth functional over a linearly constrained space of probability measures supported on a compact subset of the Euclidean space. Examples include the classical experimental design problem, the P-means problem, the problem of moments, and certain problems involving service systems. In this talk, we propose interior-point Frank-Wolfe (IPFW) type first-order recursions that handle linear functional constraints while operating on probability spaces. A barrier objective as in classical interior point is first posed, and the resulting central path is solved imprecisely using the Frank-Wolfe (FW) recursion. Importantly, and as in the unconstrained context, the FW subproblem can be solved in “closed form,” as a point mass distribution concentrating on the minimum of the barrier’s *influence function*. The resulting iteration is a simply specified update involving point-mass distributions. We provide convergence and complexity calculations, and some discussion of the derivative-free context.

On the quantum complexity of parametric integration in Sobolev spaces

Stefan Heinrich

heinrich@informatik.uni-kl.de

Stochasticcomputationandcomplexity

We consider the following problem of parametric integration in Sobolev spaces. We seek to approximate

$$S : W_p^r(D) \rightarrow L_q(D_1), \quad (Sf)(s) = \int_{D_2} f(s, t) dt \quad (s \in D_1),$$

where

$$D = [0, 1]^d = D_1 \times D_2, \quad D_1 = [0, 1]^{d_1}, \quad D_2 = [0, 1]^{d_2}, \\ 1 \leq p, q \leq \infty, \quad d, d_1, d_2, r \in \mathbf{N}, \quad d = d_1 + d_2, \quad \frac{r}{d_1} > \left(\frac{1}{p} - \frac{1}{q} \right)_+.$$

We study the complexity of this problem in the quantum setting of Information-Based Complexity [1]. Under the assumption that $W_p^r(D)$ is embedded into $C(D)$ (embedding condition) the case $p = q$ was solved by Wiegand [2]. Here we treat the case $p = q$ without embedding condition and the general case $p \neq q$ with or without the embedding condition. We also compare the rates with those in the (classical) randomized setting [3].

- [1] Heinrich, Stefan (2002). Quantum summation with an application to integration. *Journal of Complexity* 18, 1–50.
- [2] Wiegand, Carsten (2006). *Optimal Monte Carlo and Quantum Algorithms for Parametric Integration*. Shaker Verlag.
- [3] Heinrich, Stefan (2024). Randomized complexity of parametric integration and the role of adaptation II. Sobolev spaces. *Journal of Complexity* 82, 101823.

Quantum Integration in Tensor Product Besov Spaces*Bernd Käßemodel*

Faculty of Mathematics, Technische Universität Chemnitz

bernd.kaessemodel@mathematik.tu-chemnitz.de

Coauthor(s): Tino Ullrich

Special session: Stochastic Computation and Complexity p.74

We begin with a brief introduction to the basic concepts of quantum computing and quantum information-based complexity for multivariate integration and approximation problems in various smoothness classes. We then discuss characterizations of functions in tensor product Besov spaces (mixed smoothness) using the tensorized Faber-Cieselski basis with coefficients based on mixed iterated differences. Relying on such a decomposition we develop a quantum algorithm to establish bounds for the worst case quantum integration error for this function class.

Taming the Interacting Particle Langevin Algorithm – The Superlinear Case*Nikolaos Makras*

University of Edinburgh, School of Mathematics

N.Makras@sms.ed.ac.uk

Coauthor(s): Tim Johnston, Sotirios Sabanis

Recent advances in stochastic optimization have yielded the interacting particle Langevin algorithm (IPLA), which leverages the notion of interacting particle systems (IPS) to efficiently sample from approximate posterior densities. This becomes particularly crucial in relation to the framework of Expectation-Maximization (EM), where the E-step is computationally challenging or even intractable. Although prior research has focused on scenarios involving convex cases with gradients of log densities that grow at most linearly, our work extends this framework to include polynomial growth. Taming techniques are employed to produce an explicit discretization scheme that yields a new class of stable, under such non-linearities, algorithms which are called tamed interacting particle Langevin algorithms (tIPLA). We obtain non-asymptotic convergence error estimates in Wasserstein-2 distance for the new class under the best known rate.

- [1] Tim Johnston, Nikolaos Makras and Sotirios Sabanis, *Taming the Interacting Particle Langevin Algorithm – The Superlinear Case* (2024). Preprint: [arxiv:2403.19587](https://arxiv.org/abs/2403.19587)

Wed, Jul 30 12:00–12:30

Archimedes/Athena Research Centre and National Technical University of Athens

i.lytras@athenarc.gr

Tim Johnston, Nikos Makras, Sotirios Sabanis

In this article, we study the problem of sampling from distributions whose densities are not necessarily smooth nor log-concave. We propose a simple Langevin-based algorithm that does not rely on popular but computationally challenging techniques, such as the Moreau Yosida envelope or Gaussian smoothing. We derive non-asymptotic guarantees for the convergence of the algorithm to the target distribution in Wasserstein distances. Non asymptotic bounds are also provided for the performance of the algorithm as an optimizer, specifically for the solution of associated excess risk optimization problems.

Possible extensions to potentials with log-gradients that grow super-linearly may also be discussed.

This is based on the joint work in [1].

- [1] Johnston, T., Lytras, I., Makras, N., Sabanis, S. (2025). The Performance Of The Unadjusted Langevin Algorithm Without Smoothness Assumptions. arXiv preprint arXiv:2502.03458.

Wed, Jul 30 10:30–11:00

Respecting the boundaries: Space-filling designs for surrogate modeling with boundary information

Simon Mak

Duke University

sm769@duke.edu

Coauthor(s): Yen-Chun Liu

Gaussian process (GP) surrogate models are widely used for emulating expensive computer simulators, and have led to important advances in science and engineering. One challenge with fitting such surrogates is the costly generation of training data, which can require thousands of CPU hours per run. Recent promising work has investigated the integration of known boundary information for surrogate modeling, which can greatly reduce its required training sample size and thus computational cost. There is, however, little work exploring the important question of how such experiments should be designed given boundary information. We propose here a new class of space-filling designs, called boundary maximin designs, for effective GP surrogates with boundary information. Our designs rely on a new space-filling criterion derived from the asymptotic D-optimal designs of the boundary GPs of Vernon et al. (2019) and Ding et al. (2019), which can incorporate a broad class of known boundaries, including axis-parallel and/or perpendicular boundaries. To account for effect sparsity given many input parameters, we further propose a new boundary maximum projection design that jointly factors in boundary information and ensures good projective properties. Numerical experiments and an application in particle physics demonstrate improved surrogate performance with the proposed boundary maximin designs over the state-of-the-art.

Wed, Jul 30 11:00–11:30

Active Learning for Nonlinear Calibration

Andrews Boahen

Department of Statistics and Probability, Michigan State University (USA)

boahenan@msu.edu

Coauthor(s): Junoh Heo, Chih-Li Sung

Combining computer models with real observations is a valuable way to analyze complex systems. One of the key challenges in model calibration for complex physical systems is accounting for uncertainty around the model's form. The Kennedy O'Hagan's (KOH) framework provides a way to account for model form uncertainty when the computer model outputs are linearly related to the real observations. In this paper, we extend the KOH model by proposing a nonlinear wrapping function of the design inputs and the computer model outputs. We develop the statistical inference related to our model and propose active learning approaches to simultaneously collect both physical and computer model data by taking advantage of the closed-forms of the posterior mean and variance of our model under some specified kernel matrices. We show that under some regularity conditions, when $n \rightarrow \infty$, the estimate of the calibration parameter $\hat{\theta}_n$ converges to the true calibration parameter θ^* and derive uncertainty quantification results for $\hat{\theta}_n$.

Optimal design of experiments with quantitative-sequence factors*Qian Xiao*

Shanghai Jiao Tong University

qian.xiao@sjtu.edu.cn

Coauthor(s): Yaping Wang, Sixu Liu

A new type of experiments with joint considerations of quantitative and sequence factors are recently drawing much attention in medical science, bio-engineering and many other disciplines. The input spaces of such experiments are semi-discrete and often very large. Thus, efficient and economic experimental designs are required. Based on the transformations and aggregations of good lattice point sets, we construct a new class of optimal quantitative-sequence (QS) designs which are marginally coupled, pair-balanced, space-filling and asymptotically orthogonal. The proposed QS designs have certain flexibility in run and factor sizes, and are especially appealing for high-dimensional cases.

Factor Importance Ranking and Selection using Total Indices

Chaofan Huang

Work done during C. Huang's Ph.D. studies at Georgia Institute of Technology

10billhuang01@gmail.com

Coauthor(s): V. Roshan Joseph

Factor importance measures the impact of each feature on output prediction accuracy. In this paper, we focus on the *intrinsic importance* proposed by Williamson et al. (2023), which defines the importance of a factor as the reduction in predictive potential when that factor is removed. To bypass the modeling step required by the existing estimator, we present the equivalence between predictiveness potential and total Sobol' indices from global sensitivity analysis, and introduce a novel model-free consistent estimator that can be directly computed from noisy data. Integrating with forward selection and backward elimination gives rise to **FIRST**, Factor Importance Ranking and Selection using Total (Sobol') indices. Extensive simulations are provided to demonstrate the effectiveness of **FIRST** on regression and binary classification problems, and a clear advantage over the state-of-the-art methods.

GIST: Gibbs self-tuning for locally adapting Hamiltonian Monte Carlo*Bob Carpenter*

Flatiron Institute

bcarpenter@flatironinstitute.org

I will introduce the Gibbs self-tuning (GIST) framework for automatically tuning Metropolis algorithms and apply it to locally adapting the number of steps, step size, and mass matrix in Hamiltonian Monte Carlo (HMC). After resampling momentum, each iteration of GIST Gibbs samples tuning parameters based on the current position and momentum. Like with HMC, the Metropolis step for the Hamiltonian proposal is adjusted for the Gibbs sampling.

I will demonstrate how randomized Hamiltonian Monte Carlo (HMC), multinomial HMC, the No-U-Turn Sampler (NUTS), and the Apogee-to-Apogee Sampler can be cast as GIST samplers that adapt the number of steps. I will then introduce an approach to tuning step size that can be naturally combined with NUTS and demonstrate its effectiveness empirically on both multivariate normal and multiscale distributions with varying curvature.

I will sketch an approach to local mass matrix adaptation, demonstrate how it works for log concave distributions, then outline the problems remaining for efficiency and generalizability to other distributions.

Acceleration of the No-U-Turn Sampler

Nawaf Bou-Rabee

Rutgers University

nawaf.bourabee@rutgers.edu

The No-U-Turn Sampler (NUTS) is a state-of-the-art Markov chain Monte Carlo method widely used in Bayesian computation [7,8], yet its theoretical properties remain poorly understood. In this talk, I will present the first rigorous mixing time bounds for NUTS on a class of high-dimensional Gaussian targets with both high- and low-variance directions. Our analysis uncovers a striking phase transition: when initialized from the region of concentration, NUTS exhibits ballistic mixing in regimes dominated by high-variance directions, and diffusive mixing otherwise [1]. These results provide the first theoretical evidence that the U-turn mechanism can recover the acceleration of critically damped randomized Hamiltonian Monte Carlo [4,5,6]. Our analysis combines a sharp concentration result for the U-turn condition [1,2] with a recent coupling framework for localized mixing [3].

- [1] Bou-Rabee, Nawaf, & Oberdörster, Stefan (2025). *Acceleration of the No-U-turn Sampler*. Forthcoming.
- [2] Bou-Rabee, Nawaf, & Oberdörster, Stefan (2024). *Mixing of the No-U-Turn Sampler and the Geometry of Gaussian Concentration*. arXiv:2410.06978 [math.PR]
- [3] Bou-Rabee, Nawaf, & Oberdörster, Stefan (2024). *Mixing of Metropolis-Adjusted Markov Chains via Couplings: The High Acceptance Regime*. Electronic Journal of Probability. Vol. 29, No. 89, pages 1-27.
- [4] Lu, Jianfeng & Wang, Lihan (2022). *On explicit L^2 -convergence rate estimate for piecewise deterministic Markov processes in MCMC algorithms*. Annals of Applied Probability, Vol. 32, No. 2, pages 1333-1361.
- [5] Eberle, Andreas & Lörler, Francis (2024). *Non-reversible lifts of reversible diffusion processes and relaxation times*. Probability Theory and Related Fields.
- [6] Durmus, Alain; Gruffaz, Samuel; Kailas, Miika; Saksman, Eero, & Vihola, Matti (2023). *On the convergence of dynamic implementations of Hamiltonian Monte Carlo and no U-turn samplers*. arXiv:2307.03460 [stat.CO]
- [7] Hoffman, Matthew & Gelman, Andrew (2014). *The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo*. Journal of Machine Learning Research. Vol. 15, No. 1, pages 1593-1623.
- [8] Carpenter, Bob; Gelman, Andrew; Hoffman, Matthew; Lee, Daniel; Goodrich, Ben; Betancourt, Michael; Brubaker, Marcus; Guo, Jiqiang; Li, Peter, and Riddell, Allen (2016). *Stan: A Probabilistic Programming Language*. Journal of Statistical Software, Vol. 76, No. 1, pages 1-32.

ATLAS: Adapting Trajectory Lengths and Step-Size for Hamiltonian Monte Carlo

Chirag Modi

New York University

modichirag@nyu.edu

Hamiltonian Monte-Carlo (HMC) and its auto-tuned variant, the No U-Turn Sampler (NUTS)

can struggle to accurately sample distributions with complex geometries, e.g., varying curvature, due to their constant step size for leapfrog integration and fixed mass matrix. In this

talk, I will present a strategy to locally adapt the step size parameter of HMC at every iteration by evaluating a low-rank approximation of the local Hessian and estimating its largest

eigenvalue. I will then combine it with a strategy to similarly adapt the trajectory length by monitoring the no U-turn condition, resulting in an adaptive sampler, ATLAS: adapting trajectory length and step-size. I will further use a delayed rejection framework for making multiple proposals that improve the computational efficiency of ATLAS, and develop an

approach for automatically tuning its hyperparameters during warmup. Finally, I will compare ATLAS with NUTS on a suite of synthetic and real world examples, and show that

i) unlike NUTS, ATLAS is able to accurately sample difficult distributions with complex geometries, ii) it is computationally competitive to NUTS for simpler distributions, and iii) it is more robust to the tuning of hyperparameters.

AutoStep: Locally adaptive involutive MCMC*Trevor Campbell*

University of British Columbia

trevor@stat.ubc.ca

Coauthor(s): Tiange Liu, Nikola Surjanovic, Miguel Biron-Lattes, Alexandre Bouchard-Côté

Special session: Advances in Adaptive Hamiltonian Monte Carlo [p.76](#)

Many common Markov chain Monte Carlo (MCMC) kernels can be formulated using a deterministic involutive proposal with a step size parameter. Selecting an appropriate step size is often a challenging task in practice; and for complex multiscale targets, there may not be one choice of step size that works well globally. In this talk, I'll address this problem with a novel class of involutive MCMC methods—AutoStep MCMC—that selects an appropriate step size at each iteration adapted to the local geometry of the target distribution. I'll present theoretical results guaranteeing that under mild conditions AutoStep MCMC is π -invariant, irreducible, and aperiodic, and provides bounds on expected energy jump distance and cost per iteration. The talk will conclude with empirical results examining the robustness and efficacy of the proposed step size selection procedure.

Monte Carlo Based Adaptive Sampling Approaches for Stochastic Optimization*Raghu Bollapragada*

The University of Texas at Austin

raghu.bollapragada@utexas.edu

Coauthor(s): Shagun Gupta

Special session: Stochastic Optimization p.78

Stochastic optimization problems arise in a wide range of applications, from acoustic/geophysical inversion to deep learning. The scale, computational cost, and difficulty of these models make classical optimization techniques impractical. To address these challenges, in this talk, we propose new stochastic optimization methods using Monte Carlo-based adaptive sampling approaches. These approaches adaptively control the accuracy in the stochastic approximations at each iteration of the optimization algorithm by controlling the sample sizes used in these approximations to achieve efficiency and scalability. Furthermore, these approaches are well-suited for distributed computing implementations. We show that these approaches achieve optimal theoretical global convergence and complexity results for strongly convex, general convex, and non-convex problems and illustrate our algorithm's performance on machine learning models.

A New Convergence Analysis of Two Stochastic Frank-Wolfe Algorithms

Shane G. Henderson

Cornell University

sgh9@cornell.edu

Coauthor(s): Natthawut Boonsiriphatthanajaroen

Special session: Stochastic Optimization p.78

We study the convergence properties of the original and away-step Frank-Wolfe algorithms for linearly constrained stochastic optimization assuming the availability of unbiased objective function gradient estimates. The objective function is not restricted to a finite summation form, like in previous analyses tailored to machine-learning applications. To enable the use of concentration inequalities we assume either a uniform bound on the variance of gradient estimates or uniformly sub-Gaussian tails on gradient estimates. With one of these regularity assumptions along with sufficient sampling, we can ensure sufficiently accurate gradient estimates. We then use a Lyapunov argument to obtain the desired complexity bounds, relying on existing geometrical results for polytopes.

Stochastic Gradient with Testing Functionals

Akshita Gupta

Purdue University

gupta417@purdue.edu

Coauthor(s): R. Bollapragada, R. Pasupathy, A. Yip

A folklore stochastic gradient (SG) algorithm works as follows. Execute SG with a *fixed step* $\eta > 0$ until the iterates approach “stationarity,” so that no further gains are possible without a reduction in step size. Now re-execute SG starting from the final iterate of the previous execution, but with the smaller fixed step $\eta \leftarrow \beta_0 \eta, \beta_0 < 1$ until the iterates again approach “stationarity,” and so on. This simple restart scheme, originally conceived by G. Pflug in 1983, has been rediscovered and commented on by several prominent authors over the decades, with some reporting strikingly favorable numerical results, robustness, and a self-correcting nature. And yet, no complete analysis of such algorithms exists to date, probably due to the need to handle stopping times, and the need to rigorously detect stationarity. In this work, we first unify adaptive fixed-step stochastic gradient methods through *testing functionals* — essentially stopping time policies that use algorithmic history to decide stopping. We characterize general conditions on testing functionals to guarantee algorithm consistency and optimality, and then identify a specific testing functional that is based on a statistic entering a fixed stopping region, whose optimal size is intimately tied to SG’s fixed step size. The proposed testing functional framework also provides a path to analyze the many adaptive SG heuristics that have emerged over the years.

Algorithmic Discrepancy Theory: An Overview

Haotian Jiang

University of Chicago

jhtdavid@uchicago.edu

Recent Progress on Algorithmic Discrepancy Theory and Applications

Combinatorial discrepancy theory studies the following question: given a universe of elements $U = \{1, \dots, n\}$ and a collection $\mathcal{S} = \{S_1, \dots, S_m\}$ of subsets of U , how well can we partition U into two pieces, so that all sets in \mathcal{S} are split as evenly as possible. The study of this question has found extensive applications in many areas of mathematics, computer science, statistics, finance, etc.

The past decade has seen tremendous progress in designing efficient algorithms for this problem. These developments have led to many surprising applications in areas such as differential privacy, graph sparsification, approximation algorithms and rounding, kernel density estimation, randomized controlled trials, and quasi-Monte Carlo methods.

In this talk, I will briefly survey some recent algorithmic developments in this area.

Wed, Jul 30 14:30–15:00

Improving the Design of Randomized Experiments via Discrepancy Theory

Peng Zhang

Rutgers University

pz149@rutgers.edu

Coauthor(s): Recent Progress on Algorithmic Discrepancy Theory and Applications

Randomized controlled trials (RCTs) or A/B tests are the “gold standard” for estimating the causal effects of new treatments. In a trial, we want to randomly assign experimental units into two groups so that certain unit-specific pre-treatment variables, called covariates, are balanced across different groups. Balancing covariates improves causal effect estimates if covariates correlate with treatment outcomes. Simultaneously, we want our assignment of the units to be robust or sufficiently random such that our estimate is not bad if covariates do not correlate with treatment outcomes. We will show a close connection between the design of RCTs and discrepancy theory and how recent advances in algorithmic discrepancy theory could improve the design of RCTs.

Online Factorization for Online Discrepancy Minimization

Aleksandar Nikolov

University of Toronto

anikolov@cs.toronto.edu

Coauthor(s): Haohua Tang

A recent line of work, initiated by a paper of Bansal and Spencer [2], has made remarkable progress in developing online algorithms for combinatorial discrepancy minimization. In the online discrepancy minimization model, the algorithm receives a sequence of elements of some set system, or, more generally, a sequence of vectors, to color. The new element/vector must be colored immediately upon being received. We now have online algorithms that nearly match some important (offline) discrepancy bounds when the sequence of elements/vectors is stochastic or oblivious: examples include Spencer's theorem [2], and Banaszczyk's theorem [1,3]. In this work, we explore if the factorization method for discrepancy minimization [4] can also be adapted to the offline setting. We introduce a model of online factorization of matrices, and, as a case study, show how to factor online the incidence matrix of a collection of n halfspaces in d dimensions. As a result, we obtain an online algorithm to color an oblivious sequence of T points in d dimensions, so that the discrepancy with respect to a pre-specified collection of n halfspaces is on the order of $n^{\frac{1}{2} - \frac{1}{2d}}$ up to logarithmic terms.

- [1] Ryan Alweiss, Yang P. Liu, Mehtaab Sawhney: Discrepancy minimization via a self-balancing walk. STOC 2021: 14-20
- [2] Nikhil Bansal, Joel H. Spencer: On-line balancing of random inputs. Random Struct. Algorithms 57(4): 879-891 (2020)
- [3] Janardhan Kulkarni, Victor Reis, Thomas Rothvoss: Optimal Online Discrepancy Minimization. STOC 2024: 1832-1840
- [4] Jiri Matousek, Aleksandar Nikolov, Kunal Talwar, Factorization Norms and Hereditary Discrepancy, IMRN (3), 751-780, 2020.

Gaining efficiency in Monte Carlo policy gradient methods for stochastic optimal control

Arash Fahim

Florida State University

afahim@fsu.edu

Coauthor(s): Md. Arafatur Rahman, Citi Group, NYU

In this paper, we propose an efficient implementation of a deep policy gradient method (PGM) for stochastic optimal control problems in continuous time. The proposed method has the ability to distribute the allocation of computational resources, i.e., the number of Monte Carlo sample paths of a controlled diffusion process and complexity of the neural network architecture, more efficiently and improves performance for certain continuous-time problems that require a fine time discretization to achieve a desired accuracy. At each step of the method, we train a policy, modeled by a neural network, for a discretized optimal control problem in a different time scale. The first step has the coarsest time discretization. As we proceed to further steps, the time grid renders finer and a new policy is trained on the finer time scale. We provide a theoretical result on efficiency gained by this method and conclude the paper by numerical experiments on a linear-quadratic stochastic optimal control problem.

Examining the Fault Tolerance of High-Performance Monte Carlo Applications through Simulation

Sharanya Jayaraman

Department of Computer Science, Florida State University

sjayaraman@fsu.edu

Coauthor(s): Monte Carlo Applications in High-performance Computing, Computer Graphics, and Computational Science

Monte Carlo methods are naturally fault-tolerant due to their stochastic nature and are highly scalable, making them well-suited for high-performance computing environments^[1]. In this work, we use a custom-built, cloud-based exascale simulator to investigate the resilience of Monte Carlo and multi-level Monte Carlo methods under various hardware and software fault scenarios. We examine the behavior of the calculations and the propagation of inaccuracies in the presence and absence of common mitigation strategies, such as checkpointing^[2]. Our simulations offer insights into how faults propagate through computations, enabling us to evaluate techniques for constraining and mitigating the impact of unpredictable failures. The results may be used as guidance for deploying Monte Carlo-based applications on exascale systems, thereby enhancing their reliability.

References

- [1] Pauli, Stefan, Arbenz, Peter, & Schwab, Christoph, (2015). *Intrinsic fault tolerance of multilevel Monte Carlo methods*, Journal of Parallel and Distributed Computing **84**, 24-36.
- [2] Chang, C., Deringer, V.L., Katti, K.S. et al. (2023) *Simulations in the era of exascale computing*, Nat Rev Mater **8**, 309–313

Building Monte Carlo “Renderers” for Physics

Rohan Sawhney

Nvidia Corporation

rsawhney@nvidia.com

Coauthor(s): Monte Carlo Applications in High-performance Computing, Computer Graphics, and Computational Science

Accurately analyzing large amounts of geometric data is critical for many scientific and engineering applications. Techniques based on partial differential equations (PDEs) provide powerful tools for analyzing physical systems, but conventional solvers are not yet at a stage where they “just work” on problems of real-world complexity. A constant challenge is spatial discretization, which divides the domain into a high-quality volumetric mesh for PDE-based analysis. Unfortunately, this approach does not scale well to modern computer architectures, and as such, there remains a large divide between our ability to *visualize* and *analyze* the natural world.

This talk provides a broad overview of grid-free Monte Carlo methods for solving PDEs based on the walk on spheres (WoS) algorithm. WoS makes a radical departure from conventional solvers by reformulating fundamental PDEs like the Poisson equation as recursive integrals that can be solved using the Monte Carlo method. Since these integrals closely resemble those in light transport, one can leverage deep knowledge from Monte Carlo rendering to build new scalable algorithms with vastly different numerical tradeoffs, such as trivial parallelism, the ability to evaluate the solution at any point in the domain without solving a global system of equations, and avoiding volumetric meshing altogether.

WoS-NN: Collaborating Walk-on-Spheres with Machine Learning to Solve Elliptic PDEs*Silei Song*

Department of Computer Science, Florida State University

ss19cu@fsu.edu

Coauthor(s): Michael Mascagni, Arash Fahim

Special session: Monte Carlo Applications in High-performance Computing, Computer Graphics, and Computational Science p.80

Solving elliptic partial differential equations (PDEs) is a fundamental step in various scientific and engineering studies. As a classic stochastic solver, the Walk on Spheres (WoS) method is a well-established and efficient algorithm that provides accurate local estimates for PDEs. However, limited by the curse of dimensionality, WoS may not offer sufficiently precise global estimations, which becomes more serious in high-dimensional scenarios. Recent developments in machine learning offer promising strategies to address this limitation. By integrating machine learning techniques with WoS and space discretization approaches, we developed a novel stochastic solver, WoS-NN. This new method solves elliptic problems with Dirichlet boundary conditions, facilitating precise and rapid global solutions and gradient approximations. A typical experimental result demonstrated that the proposed WoS-NN method provides accurate field estimations, reducing 76.32% errors while using only 8% of path samples compared to the conventional WoS method, which saves abundant computational time and resource consumption. WoS-NN can also be utilized as a fast and effective gradient estimator based on established implementations of the original WoS method. This new method reduced the impacts of the curse of dimensionality and can be widely applied to areas like geometry processing, bio-molecular modeling, financial mathematics, etc.

Exact discretization, tight frames and recovery via D -optimal designs*Felix Bartel*

University of New South Wales

f.bartel@unsw.edu.au

Coauthor(s): Lutz Kämmerer, Kateryna Pozharska, Martin Schäfer, and Tino Ullrich

Special session: QMC and Applications Part I or II [p.81](#)

D -optimal designs originate in statistics literature as an approach for optimal experimental designs. In numerical analysis points and weights resulting from maximal determinants turned out to be useful for quadrature and interpolation. Also recently, two of the present authors and coauthors investigated a connection to the discretization problem for the uniform norm. Here we use this approach of maximizing the determinant of a certain Gramian matrix with respect to points and weights for the construction of tight frames and exact Marcinkiewicz-Zygmund inequalities in L_2 . We present a direct and constructive approach resulting in a discrete measure with at most $N \leq n^2 + 1$ atoms, which discretely and accurately subsamples the L_2 -norm of complex-valued functions contained in a given n -dimensional subspace. This approach can as well be used for the reconstruction of functions from general RKHS in L_2 where one only has access to the most important eigenfunctions. We verifiably and deterministically construct points and weights for a weighted least squares recovery procedure and pay in the rate of convergence compared to earlier optimal, however probabilistic approaches. The general results apply to the d -sphere or multivariate trigonometric polynomials on T_d spectrally supported on arbitrary finite index sets $I \subset \mathbb{Z}^d$. They can be discretized using at most $|I|^2 - |I| + 1$ points and weights. Numerical experiments indicate the sharpness of this result. As a negative result we prove that, in general, it is not possible to control the number of points in a reconstructing lattice rule only in the cardinality $|I|$ without additional condition on the structure of I . We support our findings with numerical experiments.

L_2 -approximation: using randomized lattice algorithms and QMC hyperinterpolation*Mou Cai*

Graduate School of Engineering, The University of Tokyo

caimoumou@g.ecc.u-tokyo.ac.jp

Coauthor(s): Congpei An, Takashi Goda, Yoshihito Kazashi

Abstract

We propose a randomized lattice algorithm for approximating multivariate periodic functions over the d -dimensional unit cube from the weighted Korobov space with mixed smoothness $\alpha > 1/2$ and product weights $\gamma_1, \gamma_2, \dots \in [0, 1]$. This randomization involves drawing the number of points for function evaluations randomly, and selecting a good generating vector for rank-1 lattice points using the randomized component-by-component algorithm. We prove that our randomized algorithm achieves a worst-case root mean squared L_2 -approximation error of order $M^{-\alpha/2-1/8+\varepsilon}$ for an arbitrarily small $\varepsilon > 0$, where M denotes the maximum number of function evaluations, and that the error bound is independent of the dimension d if the weights satisfy $\sum_{j=1}^{\infty} \gamma_j^{1/\alpha} < \infty$. Our upper bound converges faster than a lower bound on the worst-case L_2 -approximation error for deterministic rank-1 lattice-based approximation proved by Byrenheid, Kämmerer, Ullrich, and Volkmer (2017). We also show a lower error bound of order $M^{-\alpha/2-1/2}$ for our randomized algorithm. Finally, we present a generalization of hyperinterpolation over the unit cube named Quasi-Monte Carlo (QMC) hyperinterpolation. This new approximation scheme can be integrated with the Lasso technique to enhance sparsity and denoising capabilities.

High-dimensional density estimation on unbounded domain*Zhijian He*

South China University of Technology

hezhijian@scut.edu.cn

Coauthor(s): Ziyang Ye, Haoyuan Tan, Xiaoqun Wang

Special session: QMC and Applications Parts I and II p.81

This talk will present a kernel-based method to approximate probability density functions of unbounded random variables taking values in high-dimensional spaces. Building upon the framework of Kazashi and Nobile [1], our estimator is a linear combination of kernel functions whose coefficients are determined a linear equation. We first transform the unbounded sample domain into a hypercube and then use rank-1 lattice points as the interpolation nodes. We establish a rigorous error analysis for the mean integrated squared error (MISE) under an exponential decay condition. Under a suitable smoothness assumption, our method attains an MISE rate approaching $O(N^{-1})$ for N independent identically distributed observations. Numerical experiments validate our theoretical findings and demonstrate the superior performance of the proposed estimator compared to state-of-the-art alternatives.

- [1] Kazashi, Yoshihito, & Nobile, Fabio. (2023). Density estimation in RKHS with application to Korobov spaces in high dimensions, *SIAM Journal on Numerical Analysis*, **61**(2), 1080-1102.

Application of QMC to Oncology*Frances Y. Kuo*

UNSW Sydney, Australia

f.kuo@unsw.edu.au

Coauthor(s): Alexander D. Gilbert, Dirk Nuyens, Graham Pash, Ian H. Sloan, Karen E. Willkox

Special session: QMC and Applications Part I p.81

Tumor models largely focus on two key characteristics: infiltration of the tumor into the

surrounding healthy tissue modelled by diffusion, and proliferation of the existing tumor modelled by logistic growth in tumor cellularity. Together with terms to model chemotherapy and radiotherapy treatments, these give rise to a nonlinear parabolic reaction-diffusion PDE with random diffusion and proliferation coefficients. We show that QMC methods can be successful in computing quantities of interest arising from this tough application in oncology.

Thu, Jul 31 10:30–11:00

Finite-Particle Convergence Rates for Stein Variational Gradient Descent

Krishnakumar Balasubramanian

University of California, Davis

kbala@ucdavis.edu

We provide finite-particle convergence rates for the Stein Variational Gradient Descent (SVGD) algorithm in the Kernelized Stein Discrepancy (KSD) and Wasserstein-2 metrics. Our key insight is that the time derivative of the relative entropy between the joint density of N particle locations and the N -fold product target measure, starting from a regular initial distribution, splits into a dominant ‘negative part’ proportional to N times the expected KSD^2 and a smaller ‘positive part’. This observation leads to KSD rates of order $1/\sqrt{N}$, in both continuous and discrete time, providing a near optimal (in the sense of matching the corresponding i.i.d. rates) double exponential improvement over the recent result by [1]. Under mild assumptions on the kernel and potential, these bounds also grow polynomially in the dimension d . By adding a bilinear component to the kernel, the above approach is used to further obtain Wasserstein-2 convergence in continuous time. For the case of ‘bilinear + Matérn’ kernels, we derive Wasserstein-2 rates that exhibit a curse-of-dimensionality similar to the i.i.d. setting. We also obtain marginal convergence and long-time propagation of chaos results for the time-averaged particle laws.

- [1] Shi, J., & Mackey, L. (2023). A finite-particle convergence rate for Stein variational gradient descent. *Advances in Neural Information Processing Systems*, 36, 26831-26844.

Convergence rates of kinetic Langevin dynamics with weakly confining potentials

Lihan Wang

Carnegie Mellon University

lihanw@andrew.cmu.edu

Coauthor(s): Giovanni Brigati, Gabriel Stoltz, Andi Q. Wang

We discuss the long-time convergence behavior of kinetic Langevin dynamics, and show how the growth of the potentials in space and velocity impact the convergence rates of the dynamics, via weighted and weak Poincaré inequalities. The analysis is inspired by the Armstrong-Mourrat variational framework for hypocoercivity, which combines a weighted Poincaré-Lions inequality in time-augmented state space and an L^2 energy estimate.

Delocalization of Bias in Unadjusted Hamiltonian Monte Carlo

Xiaou Cheng

New York University

chengxo@nyu.edu

Coauthor(s): Yifan Chen, Jonathan Niles-Weed, Jonathan Weare

Special session: Analysis of Langevin and Related Sampling Algorithms, Part I p.82

Hamiltonian Monte Carlo is a commonly used algorithm to sample high dimensional probability distributions. However, for strongly log-concave distributions, existing analyses of the unadjusted algorithm show that the number of iterations follows a power law in terms of the dimension d , to ensure convergence within a desired error in the W_2 metric. Also, because of the large bias, Hamiltonian Monte Carlo is often Metropolized to remove the bias effectively. [1] suggests that for unadjusted Langevin algorithm, similar power law dimension scaling of convergence and bias in the W_2 metric can be misleading. There, for strongly log-concave distributions with certain sparse interactions, the *marginal* distribution of a small number of K variables can be well-approximated in the W_2 metric, with a small number of iterations proportional to K up to *logarithmic* terms in d . A novel W_{2,ℓ^∞} metric is used in analysis. We show that this *delocalization of bias* effect also exists in unadjusted Hamiltonian Monte Carlo with the leapfrog integrator, which suggests that Metropolization may not be necessary in this situation. A key observation is that the propagator of the leapfrog integrator is closely related to Chebyshev polynomials.

- [1] Chen, Y., Cheng, X., Niles-Weed, J., & Weare, J. (2024). Convergence of Unadjusted Langevin in High Dimensions: Delocalization of Bias. *arXiv preprint arxiv:2408.13115*.

Stochastic gradient with least-squares control variates*Matteo Raviola*

École polytechnique fédérale de Lausanne

matteo.raviola@epfl.ch

Coauthor(s): Fabio Nobile, Nathan Schaeffer

The stochastic gradient (SG) method is a widely used approach for solving stochastic optimization problems, but its convergence is typically slow. Existing variance reduction techniques, such as SAGA [1], improve convergence by leveraging stored gradient information; however, they are restricted to settings where the objective functional is a finite sum, and their performance degrades when the number of terms in the sum is large. In this work, we propose a novel approach which also works when the objective is given by an expectation over random variables with a continuous probability distribution. Our method constructs a control variate by fitting a linear model to past gradient evaluations using weighted discrete least-squares, effectively reducing variance while preserving computational efficiency. We establish theoretical sublinear convergence guarantees and demonstrate the method's effectiveness through numerical experiments on random PDE-constrained optimization.

- [1] Defazio, A., Bach, F., & Lacoste-Julien, S. (2014). *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*. *Advances in neural information processing systems*, 27.

A one-shot method for Bayesian optimal experimental design*Philipp A. Guth*

RICAM, Austrian Academy of Sciences

philipp.guth@ricam.oeaw.ac.at

Coauthor(s): Robert Gruhlke, Claudia Schillings

Special session: Nested expectations: models and estimators, Part II p.83

Bayesian optimal experimental design (BOED) problems often involve nested integrals, making their direct computation challenging. To address this, a one-shot optimization approach is proposed, which decouples the design parameters from the forward model during the op-

timization process. In addition, the solution of the forward model can be replaced by a surrogate that is trained during the one-shot optimization. This allows for the generation of computationally inexpensive samples. Efficient sampling strategies are particularly important in BOED, as they reduce the high computational cost of nested integration, ultimately making the optimization more tractable.

Thu, Jul 31 11:30–12:00

Langevin-based strategies for nested particle filters

Sara Pérez-Vieites

Aalto University

sara.perezvieites@aalto.fi

Coauthor(s): Nicola Branchini, Víctor Elvira and Joaquín Míguez

Many problems in some of the most active fields of science require to estimate parameters and predict the evolution of complex dynamical systems using sequentially collected data. The nested particle filter (NPF) framework stands out since it is the only fully recursive probabilistic method for Bayesian inference. That is, it computes the joint posterior distribution of the parameters and states while maintaining a computational complexity of $\mathcal{O}(T)$, which makes it particularly suitable for long observation sequences.

A key strategy to keep particle diversity in the parameter space, given the static nature of the parameters, is jittering. The parameter space is explored by perturbing a subset of particles with arbitrary variance or applying a controlled variance to all particles. As the perturbations are controlled, it ensures convergence to the true posterior distribution while keeping the full framework recursive. However, this is not an efficient exploration strategy, particularly for problems with a higher dimension in the parameter space.

To address this limitation, we propose a Langevin-based methodology within the NPF framework. A challenge is that the required score function is intractable. We propose to approximate the score with an accurate method that is provably stable over time, and to explore strategies to reduce its computational cost while retaining accuracy. This approach significantly improves the scalability of NPF in the parameter dimension, while still ensuring asymptotic convergence to the true posterior, as well as maintaining computational feasibility.

Approximation of multivariate periodic functions

Dirk Nuyens

KU Leuven, Belgium

dirk.nuyens@kuleuven.be

Coauthor(s): Laurence Wilkes

Special session: QMC and Applications p.84

We study approximation of multivariate periodic functions using n function values. We use generated sets as the sample points, first introduced in [1]. We prove existence and convergence of the almost optimal L_2 error and obtain similar bounds as for the least squares algorithm from [2] which uses unstructured points.

- [1] Kämmerer. Reconstructing multivariate trigonometric polynomials by sampling along generated sets. In Monte Carlo and Quasi-Monte Carlo Methods 2012 (Dick, Kuo, Peters, Sloan), pages 439–454, 2013.
- [2] Krieg, Ullrich. Function values are enough for L_2 approximation. *Foundations of Computational Mathematics*, 21(4):1141–1151, 2021.

Randomized QMC with one categorical variable

Art B. Owen

Stanford University

owen@stanford.edu

Coauthor(s): Valerie Ho, Zexin Pan

Randomized quasi-Monte Carlo (RQMC) methods benefit from smoothness in the integrand. In some applications, one of the input variables takes only a modest finite number $L \geq 2$ of values and can be considered a categorical variable. Such a variable introduces a discontinuity in the integrand. It is an RQMC-friendly discontinuity because the discontinuity is axis parallel when each level of the categorical variable corresponds to a single sub-interval of $[0, 1]$. A naturally occurring use case has the categorical variable correspond to a component

in a distribution that is a mixture of L other distributions. Within the mixture setting, mixture importance sampling is a prominent example.

If category ℓ has probability α_ℓ in the motivating model then it would naturally get an interval of width α_ℓ . If the RQMC method uses scrambled Sobol' points then we can make the problem even more RQMC-friendly by instead giving level ℓ an interval of width $\beta_\ell = 2^{-\kappa_\ell}$ for some integers $\kappa_\ell \geq 1$ and then incorporating sampling ratios α_ℓ/β_ℓ . These β_ℓ are negative powers of 2 that sum to 1. The number of ways to choose β_1, \dots, β_L as a function of L is given by sequence A002572 in the online encyclopedia of integer sequences.

Under usual assumptions on the convergence rates for RQMC estimates, we find that the asymptotically optimal β_ℓ are more nearly equal than the original α_ℓ are. We find some rules for making that allocation under uncertainty about the appropriate RQMC rate.

The mixing problem was studied by [1] where they use instead L different QMC samples instead of one sample in which one of the variables yields ℓ . The first QMC variable was used by [2] to allocate points to different processors in parallel computing.

- [1] Cui, T., J. Dick, and F. Pillichshammer (2023). Quasi-Monte Carlo methods for mixture distributions and approximated distributions via piecewise linear interpolation. Technical report, arXiv:2304.14786
- [2] Keller, A. and L. Grunschloß (2012). Parallel quasi-Monte Carlo integration by partitioning low-discrepancy sequences. In Monte Carlo and Quasi-Monte Carlo Methods 2010, pp. 487–498. Springer.

Thu, Jul 31 16:30–17:00

QMC confidence intervals using quantiles of randomized nets

Zexin Pan

Johann Radon Institute for Computational and Applied Mathematics

zexin.pan@oeaw.ac.at

The median of linearly scrambled digital net estimates has been shown to converge to the target integral value at nearly the optimal rate across various function spaces [1]. In this talk, we explore how quantiles of these estimates can be used to construct confidence intervals for the target integral. In particular, we demonstrate that as the sample size increases, the error distribution becomes increasingly symmetric, and we quantify the rate of this symmetrization for a class of smooth integrands.

- [1] Pan, Zexin. (2025). Automatic optimal-rate convergence of randomized nets using

median-of-means. *Mathematics of Computation*, to appear.

Thu, Jul 31 17:00–17:30

Quasi-uniform quasi-Monte Carlo lattice point sets

Kosuke Suzuki

Yamagata University

kosuke-suzuki@sci.kj.yamagata-u.ac.jp

Coauthor(s): Josef Dick, Takashi Goda, Gerhard Larcher, Friedrich Pillichshammer

Special session: QMC and Applications p.84

The discrepancy of a point set quantifies how well the points are distributed, with low-discrepancy point sets demonstrating exceptional uniform distribution properties. Such sets are integral to quasi-Monte Carlo methods, which approximate integrals over the unit cube for integrands of bounded variation. In contrast, quasi-uniform point sets are characterized by optimal separation and covering radii, making them well-suited for applications such as radial basis function approximation. This paper explores the quasi-uniformity properties of quasi-Monte Carlo point sets constructed from lattices. Specifically, we analyze rank-1 lattice point sets, Fibonacci lattice point sets, Frolov point sets, and $(n\alpha)$ -sequences, providing insights into their potential for use in applications that require both low-discrepancy and quasi-uniform distribution. As an example, we show that the $(n\alpha)$ -sequence with $\alpha_j = 2^{j/(d+1)}$ for $j \in \{1, 2, \dots, d\}$ is quasi-uniform and has low-discrepancy.

Thu, Jul 31 15:30–16:00

Molei Tao

Georgia Tech

mtao@gatech.edu

[Kijung Jeon](#), [Michael Muehlebach](#)

Coauthor(s): Analysis of Langevin and Related Sampling Algorithms

Given an unnormalized density ρ and a constraint set Σ in \mathbb{R}^n , we aim at sampling from

a constrained distribution $Z^{-1}\rho(x)I_{\Sigma}(x)dx$. While the case when Σ is convex has been extensively studied, *no* convexity is needed in this talk, in which case quantitative results are scarce. Our method admits multiple interpretations, but this talk will focus on a Langevin perspective, where overdamped Langevin dynamics is first modified, and then discretized so that a sampling algorithm can be constructed. The quantitative convergence of the continuous dynamics will be detailed, but if time permits, the performance of the time-discretization (i.e., the actual sampler) will also be discussed.

Thu, Jul 31 16:00–16:30

Convergence of Unadjusted Langevin in High Dimensions: Delocalization of Bias

Yifan Chen

UCLA, Courant Institute

yifanc96@gmail.com

Coauthor(s): Xiaoou Cheng, Jonathan Niles-Weed, Jonathan Weare

The unadjusted Langevin algorithm is commonly used to sample probability distributions in extremely high-dimensional settings. However, existing analyses of the algorithm for strongly log-concave distributions suggest that, as the dimension d of the problem increases, the number of iterations required to ensure convergence within a desired error in the W_2 metric scales in proportion to d or \sqrt{d} . In this work, we argue that, the behavior for a *small number* of variables can be significantly better: a number of iterations proportional to K , up to logarithmic terms in d , often suffices for the algorithm to converge to within a desired W_2 error for all K -marginals. We refer to this effect as *delocalization of bias*. We show that the delocalization effect does not hold universally and prove its validity for Gaussian distributions and strongly log-concave distributions with certain sparse/local interactions. Our analysis relies on a novel W_{2,ℓ^∞} metric to measure convergence. A key technical challenge we address is the lack of a one-step contraction property in this metric. Finally, we use asymptotic arguments to explore potential generalizations of the delocalization effect beyond the Gaussian and sparse/local interactions setting.

Entropy methods for the delocalization of bias in Langevin Monte Carlo*Fuzhong Zhou*

Columbia University, IEOR department

fz2329@columbia.edu

Coauthor(s): Daniel Lacker

Analysis of Langevin and Related Sampling Algorithms, Part II.

The unadjusted Langevin algorithm is widely used for sampling from complex high-dimensional distributions. It is well known to be biased, with the bias as measured in squared Wasserstein-2 distance scaling linearly with the dimension. However, the remarkable recent work [1] has revealed a *delocalization of bias* effect: For a class of distributions with *sparse interactions*, the bias between lower-dimensional marginals scales only with the lower dimension, not the full dimension. In this work, we strengthen the results of [1] in the sparse interaction regime by removing a logarithmic factor, measuring distance in KL-divergence, and relaxing the strong log-concavity assumption. In addition, we expand the scope of the delocalization phenomenon by showing that it holds for a different class of distributions, with *weak interactions*. Our proofs are based on a hierarchical analysis of the KL-divergence between marginals.

- [1] Y. Chen, X. Cheng, J. Niles-Weed, and J. Weare (2024). *Convergence of unadjusted Langevin in high dimensions: Delocalization of bias*. arXiv preprint arXiv:2408.13115.

Convergence of Φ -Divergence and Φ -Mutual Information Along Langevin Markov Chains*Siddharth Mitra*

Yale University

siddharth.mitra@yale.edu

Coauthor(s): Jiaming Liang, Andre Wibisono

The mixing time of a Markov chain determines when the marginal law of the Markov chain is close to the stationary distribution and can be studied in many statistical divergences such as KL divergence and chi-squared divergence, all the way to families of divergences such as Φ -divergences. However, the mixing time does not determine the dependency between samples along the Markov chain, which can be measured in terms of their mutual information,

chi-squared mutual information, or more generally their Φ -mutual information. In this talk, we study the mixing time of Langevin Markov chains in Φ -divergence and also study the Φ -mutual information between the iterates. The Markov chains we focus on are the Langevin Dynamics in continuous-time, and the Unadjusted Langevin Algorithm and Proximal Sampler in discrete-time and we show that for these Markov chains, the Φ -divergence and the Φ -mutual information decreases exponentially fast. Our proof technique is based on showing the Strong Data Processing Inequalities (SDPIs) hold along the Markov chains. To prove fast mixing of the Markov chains, we show the SDPIs hold for the stationary distribution. In contrast, to prove the contraction of Φ -mutual information, we need to show the SDPIs hold along the entire trajectories of the Markov chains; we prove this when the iterates along the Markov chains satisfy the corresponding Φ -Sobolev inequality.

Thu, Jul 31 15:30–16:00

Inference for Stochastic Gradient Descent with Infinite Variance

Jose Blanchet

Stanford University

jose.blanchet@stanford.edu

Coauthor(s): Peter Glynn, Aleks Mijatovic, Wenhao Yang

Special session: Recent Advances in Stochastic Gradient Descent p.86

Stochastic gradient descent (SGD) with infinite variance gradients arises, perhaps surprisingly, quite often in applications. Even in settings involving “finite variance” in theory, infinite variance models appear to provide a better statistical fit over spatial and temporal scales of interest in applied settings. Motivated by this, we investigate a general methodology that enables the development of valid confidence regions for SGD with infinite variance. Along the way, we also obtain key results and properties for SGD with infinite variance, for example, asymptotic limits, optimal convergence rates, etc., which are counterparts of celebrated results known only in the finite variance case.

Exit-Time Analysis for Kesten’s Recursion in Stochastic Gradient Descent*Chang-Han Rhee*

Northwestern University

chang-han.rhee@northwestern.edu

Coauthor(s): Jeeho Ryu, Insuk Seo

Stochastic gradient descent has been a classical subject in operations research and stochastic simulation literature. On the other hand, Kesten’s recursion has been studied extensively in probability theory and related fields. Recently, Kesten’s recursion has garnered renewed interest as a model for stochastic gradient descent (SGD) with a quadratic objective function and the emergence of heavy-tailed dynamics in machine learning. In particular, due to the connection between the heavy-tailed behaviors of SGD and the generalization performance of the neural network it trains, the emergence and characterization of heavy tails in SGD have been revisited. Unlike the classical contexts, these developments call for analysis of its asymptotic behavior under both negative and positive Lyapunov exponents. In this talk, I’ll discuss the exit times of Kesten’s stochastic recurrence equation in both cases. Depending on the sign of the Lyapunov exponent, the exit time scales either polynomially or logarithmically as the radius of the exit boundary increases.

- [1] Rhee, Chang-Han, Jeeho Ryu, & Insuk Seo (2025+). Exit time analysis for Kesten’s stochastic recurrence equations. arXiv:2503.05219.

Stochastic Gradient Descent with Adaptive Data*Jing Dong*

Columbia University

jing.dong@gsb.columbia.edu

Coauthor(s): Ethan Che, Xin Tong

Special session: Recent Advances in Stochastic Gradient Descent p.86

Stochastic gradient descent (SGD) is a powerful optimization technique that is particularly useful in online learning scenarios. Its convergence analysis is relatively well understood under the assumption that the data samples are independent and identically distributed (iid).

However, applying SGD to policy optimization problems in operations research involves a distinct challenge: the policy changes the environment and thereby affects the data used to update the policy. The adaptively generated data stream involves samples that are non-stationary, no longer independent from each other, and affected by previous decisions. The influence of previous decisions on the data generated introduces bias in the gradient estimate, which presents a potential source of instability for online learning not present in the iid case. In this paper, we introduce simple criteria for the adaptively generated data stream to guarantee the convergence of SGD. We show that the convergence speed of SGD with adaptive data is largely similar to the classical iid setting, as long as the mixing time of the policy-induced dynamics is factored in. Our Lyapunov-function analysis allows one to translate existing stability analysis of stochastic systems studied in operations research into convergence rates for SGD, and we demonstrate this for queueing and inventory management problems. We also showcase how our result can be applied to study the sample complexity of an actor-critic policy gradient algorithm.

Fri, Aug 1 09:00–09:30

Fixed-budget simulation method for growing cell populations

Zhou Fang

Academy of Mathematics and Systems Science, Chinese Academy of Sciences

zhfang@amss.ac.cn

Coauthor(s): Shaoqing Chen, Zheng Hu, Da Zhou

Special session: Forward and Inverse Problems for Stochastic Reaction Networks [p.87](#)

Investigating the dynamics of growing cell populations is crucial for unraveling key biological mechanisms in living organisms, with many important applications in therapeutics and biochemical engineering. Classical agent-based simulation algorithms are often inefficient for these systems because they track each individual cell, making them impractical for fast (or even exponentially) growing cell populations. To address this challenge, we introduce a novel stochastic simulation approach based on a Feynman-Kac-like representation of the population dynamics. This method, named the Feynman-Kac-inspired Gillespie's Stochastic Simulation Algorithm (FKG-SSA), always employs a fixed number of independently simulated cells for Monte Carlo computation of the system, resulting in a constant computational complexity regardless of the population size. Furthermore, we theoretically show the statistical consistency of the proposed method, indicating its accuracy and reliability. Finally, a couple of biologically relevant numerical examples are presented to illustrate the approach. Overall, the proposed FKG-SSA effectively addresses the challenge of simulating growing cell populations, providing a solid foundation for better analysis of these systems.

Dimensionality Reduction for Efficient Rare Event Estimation

Sophia Münker

RWTH Aachen University

muenker@uq.rwth-aachen.de

Coauthor(s): Chiheb Ben Hammouda, Nadhir Ben Rached, Raúl Tempone

Special session: Forward and Inverse Problems for Stochastic Reaction Networks [p.87](#)

A Stochastic Reaction Network (SRN) is a continuous-time, discrete-space Markov chain that models the random interaction of d species through reactions, commonly applied in biochemical systems. We are interested in efficiently estimating rare event probabilities, where we consider path-dependent observables. Therefore, we present an importance sampling (IS) method based on the discrete Tau-Leap (TL) scheme to enhance the performance of Monte Carlo (MC) estimators. The primary challenge in IS is selecting an appropriate change of probability measure to significantly reduce variance, which often requires deep insights into the underlying problem. To address this, we propose a generic approach to obtain an efficient path-dependent measure change, based on an original connection between finding optimal IS parameters and solving a variance minimization problem using a stochastic optimal control (SOC) formulation [1]. The optimal IS parameters can be derived by solving a Hamilton-Jacobi-Bellman equation.

To address the curse of dimensionality, we propose the Markovian Projection (MP) technique to reduce the SRN to a lower-dimensional SRN (called MP-SRN) while preserving the marginal distribution of the original high-dimensional system. When solving the resulting SOC problem numerically to derive the variance reducing IS parameters, we derive the parameter for a reduced-dimensional model. These IS parameters can be applied to the full-dimensional SRN in the forward run. Analysis and numerical experiments demonstrate that our IS strategies substantially reduce the variance of the MC estimator, leading to lower computational complexity in the rare event regime compared to standard MC methods.

At the end of the talk, we give a small outlook on a multilevel-IS scheme to further improve the efficiency of the estimator.

- [1] Ben Hammouda, C., Ben Rached, N., Tempone, R., & Wiechert, S. (2024). Automated importance sampling via optimal control for stochastic reaction networks: A Markovian projection-based approach. *Journal of Computational and Applied Mathematics*, 446, 115853.

Filtered Markovian Projection: Dimensionality Reduction in Filtering for Stochastic Reaction Networks

Maksim Chupin

King Abdullah University of Science and Technology (KAUST)

maksim.chupin@kaust.edu.sa

Coauthor(s): Chiheb Ben Hammouda, Sophia Munker, Raúl Tempone

Special session: Forward and Inverse Problems for Stochastic Reaction Networks p.87

Stochastic reaction networks (SRNs) model stochastic effects for various applications, including intracellular chemical or biological processes and epidemiology. A key challenge in practical problems modeled by SRNs is that only a few state variables can be dynamically observed. Given the measurement trajectories, one can estimate the conditional probability distribution of unobserved (hidden) state variables by solving a system filtering equation. The current numerical methods, such as the Filtered Finite State Projection [1], are hindered by the curse of dimensionality, significantly affecting their computational performance. To overcome this, we propose to use a dimensionality reduction technique based on the Markovian projection (MP), initially introduced for forward problems [2]. In this work, we explore how to adapt the existing MP approach to the filtering problem and introduce a novel version of the MP, the Filtered MP, that guarantees the consistency of the resulting estimator [3]. The novel method employs a reduced-variance particle filter for estimating the jump intensities of the projected model and solves the filtering equations in a low-dimensional space, improving computational efficiency over existing methods.

- [1] D’Ambrosio, E., Fang, Z., Gupta, A., Kumar, S., & Khammash, M. (2022). Filtered finite state projection method for the analysis and estimation of stochastic biochemical reaction networks. *bioRxiv*, 2022-10.
- [2] Hammouda, C. B., Rached, N. B., Tempone, R., & Wiechert, S. (2024). Automated importance sampling via optimal control for stochastic reaction networks: A Markovian projection-based approach. *Journal of Computational and Applied Mathematics*, 446, 115853.
- [3] Hammouda, C. B., Chupin, M., Munker, S., & Tempone, R. (2025). Filtered Markovian Projection: Dimensionality Reduction in Filtering for Stochastic Reaction Networks. *arXiv preprint arXiv:2502.07918*.

State and parameter inference in stochastic reaction networks

Muruhan Rathinam

University of Maryland Baltimore County

muhan@umbc.edu

Coauthor(s): Mingkai Yu

Continuous time Markov chain models are widely used to model intracellular chemical reactions networks that arise in systems and synthetic biology. In this talk, we address the problem of inference of state and parameters of such systems from partial observations. We present details of recent particle filtering methods that are applicable to two different scenarios: one in which the observations are made continuously in time and the other in which the observations are made in discrete snapshots of time. We provide the theoretical justification as well as numerical results to illustrate these methods.

- [1] Rathinam, Muruhan & Yu, Mingkai (2021). *State and parameter estimation from exact partial state observation in stochastic reaction networks*. The Journal of Chemical Physics. 154(3).
- [2] Rathinam, Muruhan & Yu, Mingkai (2023). Stochastic Filtering of Reaction Networks Partially Observed in Time Snapshots. Journal of Computational Physics. Volume 515, 15 October 2024, 113265.

A High-performance Multi-level Monte Carlo Software for Full Field Estimates and Applications in Optimal Control

Niklas Baumgarten

Heidelberg University

niklas.baumgarten@uni-heidelberg.de

Coauthor(s): Robert Scheichl, Robert Kutri, David Schneiderhan, Sebastian Krumscheid, Christian Wieners, Daniele Corallo

Special session: Hardware or Software for (Quasi-)Monte Carlo Algorithms, Part II [p.88](#)

We introduce a high-performance software framework for multi-level Monte Carlo (MLMC) and multi-level stochastic gradient descent (MLSGD) methods, designed for efficient uncertainty quantification and optimal control in high-dimensional PDE applications. The software operates under strict memory and CPU-time constraints, requiring no prior knowledge of problem regularity or memory demands.

Built on the budgeted MLMC method with a sparse multi-index update algorithm, it enables full spatial-domain estimates at the same computational cost as single-point evaluations while maintaining memory usage comparable to the deterministic problem. Additionally, the framework integrates an MLSGD method for optimal control, achieving faster convergence at reduced computational costs compared to standard stochastic gradient descent methods.

We validate the software on applications such as stochastic PDE-based sampling, subsurface flow, mass transport, acoustic wave equations, and high-dimensional control problems.

Fast Gaussian Processes

Aleksei G Sorokin

Illinois Institute of Technology, Department of Applied Mathematics.

Sandia National Laboratories.

asorokin@hawk.iit.edu

Special session: Hardware and Software for (Quasi-)Monte Carlo Algorithms, Part 2 p.88

Gaussian process regression (GPR) on N data points typically costs $\mathcal{O}(N^3)$ as one must compute the inverse and determinant of a dense, unstructured Gram matrix. Here we present `fastgps`¹, a Python software which performs GPR at only $\mathcal{O}(N \log N)$ cost by forcing nice structure into the Gram matrices. Specifically, when one controls the design of experiments and is willing to use special kernel forms, pairing certain low discrepancy (LD) sequences with shift invariant kernels yields Gram matrices diagonalizable by fast transforms. Two such classes are:

1. Pairing *lattice points* with *shift invariant (SI) kernels* gives circulant Gram matrices diagonalizable by the fast Fourier transform (FFT).
2. Pairing *digital nets* with *digitally shift invariant (DSI) kernels* gives Gram matrices diagonalizable by the fast Walsh Hadamard transform (FWHT).

`fastgps` supports a number of features which will be discussed.

1. **Kernel hyperparameter optimization** of marginal log likelihood (MLL), cross validation (CV), or generalized cross validation (GCV) loss.
2. **Fast Bayesian cubature** for uncertainty quantification in Quasi-Monte Carlo.
3. **Fast multi-task GPR** with support for different sample sizes for each task. This is useful for multi-fidelity simulations and Multilevel Monte Carlo (MLMC).
4. **Efficient variance projections** for non-greedy Bayesian optimization in MLMC.
5. **Derivative informed GPR** for simulations coupled with automatic differentiation.
6. **Batched GPR** for simultaneously modeling vector-output simulations.
7. **GPU support** enabled by the PyTorch² stack .
8. **Flexible LD sequences and SI/DSI kernels** from QMCPy³.

¹<https://alegresor.github.io/fastgps/>

²<https://pytorch.org/>

³<https://qmcssoftware.github.io/QMCSSoftware/>

Hybrid Monte Carlo methods for kinetic transport

Johannes Krotz

University of Notre Dame

jkrotz@nd.edu

Coauthor(s): Hardware or Software for (Quasi-)Monte Carlo Algorithm

We present a hybrid method for time-dependent particle transport problems that combines Monte Carlo (MC) estimation with deterministic solutions based on discrete ordinates. For spatial discretizations, the MC algorithm computes a piecewise constant solution and the discrete ordinates use bilinear discontinuous finite elements. From the hybridization of the problem, the resulting problem solved by Monte Carlo is scattering free, resulting in a simple, efficient solution procedure. Between time steps, we use a projection approach to “relabel” collided particles as uncollided particles. From a series of standard 2-D Cartesian test problems we observe that our hybrid method has improved accuracy and reduction in computational complexity of approximately an order of magnitude relative to standard discrete ordinates solutions.

Flow-Based Monte Carlo Transport Simulation

Joseph Farmer

University of Notre Dame

jfarmer@nd.edu

Coauthor(s): Prof. Ryan McClarren, Aidan Murray, Johannes Krotz

Special session: Hardware or Software for Quasi-Monte Carlo Methods p.88

Monte Carlo (MC) particle transport methods are the gold standard for high-fidelity simulations but can be computationally prohibitive, especially in optically thick media or problems requiring a large number of particle histories. The performance bottleneck often lies in the sequential, stochastic random walk, where each particle undergoes numerous collision events that are difficult to vectorize and lead to control flow divergence on parallel architectures like GPUs.

Our approach replaces the intra-cell random walk with a surrogate model trained via statistical learning. This surrogate effectively learns the transport kernel, providing a direct map from a particle's input state to its output state. This is accomplished by training a neural network to model the velocity field of a continuous probability path that transforms a simple noise distribution into the complex, multi-modal distribution of particle output states observed in MC simulations. This hybrid approach retains the Monte Carlo foundation of stochastic particle sampling and history tracking within the transport computation, but replaces the collision-by-collision simulation with a single forward pass through the neural network that samples from the learned distributions.

We demonstrate the efficacy of this approach by embedding our learned model into a multi-cell transport solver and validating it against canonical benchmarks, including the heterogeneous Reed's problem [2]. The results show that our method accurately reproduces the scalar flux profiles of both high-fidelity MC simulations and known analytic solutions, with comparable statistical error characteristics. This accuracy is achieved with computational speed-ups, particularly in optically thick regions where the cost of traditional MC is highest. The performance gains stem from decoupling the simulation cost from the material cross-sections, offering a scalable path for accelerating transport in challenging regimes.

The presentation will cover the mathematical formulation of the transport kernel learning problem, the architecture which captures the output distributions, the once-and-for-all training strategy, and computational results in both 1D and 2D geometries.

- [1] Y. Lipman, R. T. Q. Chen, et al. (2022). Flow Matching for Generative Modeling. *arXiv preprint arXiv:2210.02747*.
- [2] W. H. Reed. (1971). New Difference Schemes for the Neutron Transport Equation. *Nuclear Science and Engineering*, 46(2):309-314.

Contributed Talks

Mon, Jul 28 10:30–11:00

Stereographic Multi-Try Metropolis Algorithms for Heavy-tailed Sampling

Zhihao Wang

University of Copenhagen

zw@math.ku.dk

Coauthor(s): Jun Yang

We introduce a novel family of gradient-free Markov chain Monte Carlo (MCMC) algorithms that integrate the principles of multi-try Metropolis (MTM) and stereographic MCMC, designed specifically for efficient sampling of heavy-tailed distributions. Through scaling analysis and extensive simulations, we demonstrate that the proposed stereographic multi-try Metropolis (SMTM) algorithm outperforms both traditional Euclidean MTM and existing stereographic random-walk Metropolis. Furthermore, the SMTM algorithm has the potential to benefit from modern hardware, such as GPUs, allowing for improved performance through parallel implementation.

Creating rejection-free samplers by rebalancing skew-balanced jump processes*Ruben Seyer*

Chalmers University of Technology and University of Gothenburg

rubense@chalmers.se

Coauthor(s): Erik Jansson, Moritz Schauer, Akash Sharma

Markov chain sampling methods form the backbone of modern computational statistics. However, many popular methods are prone to random walk behaviour, i.e. diffusion-like exploration of the sample space, leading to slow mixing that requires intricate tuning to alleviate. Non-reversible samplers can resolve some of these issues. We introduce a device that turns jump processes that satisfy a skew-detailed balance condition for a reference measure into a process that samples a target measure that is absolutely continuous with respect to the reference measure. This sampler is rejection-free, non-reversible and time-continuous. As an example, we apply the device to Hamiltonian dynamics discretized by the leapfrog integrator, resulting in a rejection-free non-reversible time-continuous version of Hamiltonian Monte Carlo (HMC). We prove the geometric ergodicity of the resulting sampler, and demonstrate its increased robustness to hyperparameter tuning compared to HMC through numerical examples. This comes at a computational cost at worst double that of HMC, in practice lower than other popular non-reversible samplers such as the Bouncy Particle Sampler.

Theoretical guarantees for lifted samplers

Philippe Gagnon

Université de Montréal

philippe.gagnon.3@umontreal.ca

Coauthor(s): Florian Maire

The work I would like to present is about a particular class of Markov chain Monte Carlo (MCMC) methods which use non-reversible Markov chains commonly referred to as *lifted Markov chains*; the methods are commonly referred to as *lifted samplers*. The methods are not particularly new (they date back at least to Horowitz [1]), but they have recently been the subject of significant research work motivated by a general belief that, in statistical applications, they lead to more efficient estimators than their reversible counterparts which correspond to Metropolis–Hastings (MH) algorithms (see, e.g., Andrieu and Livingstone [2]). It was thus somewhat surprising to observe that it is not always the case in some recent work (see, e.g., Gagnon and Maire [3]). One can thus wonder what degree of inefficiency these chains may exhibit in worst-case scenarios. This is an important question given that lifted samplers are popular in practice, a consequence of the fact that they are often as easy to implement on a computer as their MH counterparts and often have the same computational complexity.

The main contribution of our work is to provide an answer to this question under arguably the most general framework. We proceed by leveraging the seminal work of Tierney [4] to define a lifted version of a generalized MH algorithm. Virtually any (reversible) MCMC method can be seen as a special case of this generalized MH algorithm, ranging from the traditional MH algorithm of Hastings [5] to the reversible jump algorithm of Green [6]. Our main theoretical result allows for a comparison between the generalized MH algorithm and its lifted version in terms of the variance of produced estimators. It essentially guarantees that the variance of estimators produced by the lifted version cannot be more than twice that of estimators produced by the generalized MH algorithm. This result indicates that, while there is potentially a lot to gain from lifting a Markov chain, there is not much to lose. We also show that our result is optimal, in the sense that it is not possible to improve on the factor 2 without additional assumptions. The definition of the lifted version of the generalized MH algorithm allows us to understand how a lifted sampler can be constructed under such a general framework, which adds a methodological contribution to our theoretical contribution.

The efficiency of Markov chains is traditionally assessed by studying the characteristics of their Markov transition operators. To establish our theoretical result, we needed to connect the efficiency of two significantly different operators (those of the MH and lifted algorithms): in addition to not being defined on the same domain, one is self-adjoint while the other is not, which further complicates the analysis. One of our main achievements was to identify a

specific auxiliary operator which acts as a bridge and allows to connect the efficiency of the two aforementioned operators. This auxiliary operator is compared to the MH one through a Peskun ordering [7] established via a careful analysis of the Markov kernels, yielding sharp bounds. The connections between the MH and lifted algorithms is completed by comparing the auxiliary operator with the lifted one using a result in Andrieu and Livingstone [2].

- [1] Horowitz, A. M. (1991) A generalized guided Monte Carlo algorithm. *Phys. Lett. B*, **268**, 247–252.
- [2] Andrieu, C. and Livingstone, S. (2021) Peskun–Tierney ordering for Markovian Monte Carlo: Beyond the reversible scenario. *Ann. Statist.*, **49**, 1958–1981.
- [3] Gagnon, P. and Maire, F. (2024) An asymptotic Peskun ordering and its application to lifted samplers. *Bernoulli*, **30**, 2301–2325.
- [4] Tierney, L. (1998) A note on Metropolis–Hastings kernels for general state spaces. *Ann. Appl. Probab.*, **8**, 1–9.
- [5] Hastings, W. K. (1970) Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, **57**, 97–109.
- [6] Green, P. J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711–732.
- [7] Peskun, P. (1973) Optimum Monte-Carlo sampling using Markov chains. *Biometrika*, **60**, 607–612

Tue, Jul 29 10:30–11:00

Optimizing Generalized Hamiltonian Monte Carlo for Bayesian Inference applications

Lorenzo Nagar

BCAM - Basque Center for Applied Mathematics, Bilbao, Spain

lnagar@bcamath.org

Coauthor(s): Leonardo Gavira Balmacz, Hristo Inouzhe Valdes, Martín Parga Pazos, María Xosé Rodríguez-Álvarez, Jesús María Sanz-Serna, and Elena Akhmatskaya

In contrast to the widely adopted Hamiltonian/Hybrid Monte Carlo (HMC) [1], the Generalized Hamiltonian Monte Carlo (GHMC) algorithm [2, 3] leverages the irreversibility of its generated Markov chains, resulting in faster convergence to equilibrium and reduced asymptotic variance [4, 5].

Despite its theoretically predicted advantages, GHMC can be highly sensitive to the choice of a numerical integrator for the Hamiltonian equations and requires careful tuning of sim-

ulation parameters, such as the integration step size, the trajectory length, and the amount of random noise in momentum refreshment.

In this talk, we present a novel approach for finding optimal (in terms of sampling performance and accuracy) settings for a GHMC simulation. For an arbitrary simulated system, our methodology identifies a system-specific integration scheme that maximizes a conservation of energy for harmonic forces, along with appropriate randomization intervals for the simulation parameters, without incurring additional computational cost.

Numerical experiments on well-established statistical models exhibit, with the help of the state-of-the-art performance metrics, significant gains in GHMC sampling efficiency when optimally tuned hyperparameters are chosen instead of heuristic or recommended ones. Comparative performance of GHMC and HMC with optimal settings is also discussed.

Additionally, we apply our methodology to three real-world case studies:

- Patient resistance to endocrine therapy in breast cancer;
- Influenza A (H1N1) epidemics outbreak;
- Modeling of cell-cell adhesion dynamics.

- [1] Neal, Radford M. (2011). *MCMC Using Hamiltonian Dynamics*. Handbook of Markov Chain Monte Carlo (Eds: Brooks, S., Gelman, A., Jones, G., Meng, X.), Vol. 2 pp. 113–162. Chapman and Hall/CRC, New York.
- [2] Horowitz, Alan M. (1991). *A generalized guided Monte Carlo algorithm*. Physics Letters B 268.2: 247–252.
- [3] Kennedy, Anthony D. & Pendleton, Brian J. (2001). *Cost of the generalised hybrid Monte Carlo algorithm for free field theory*. Nuclear Physics B 607.3: 456–510.
- [4] Ottobre, Michela (2016). *Markov Chain Monte Carlo and Irreversibility*. Reports on Mathematical Physics 77: 267–292.
- [5] Duncan, Andrew B., Lelièvre, Tony & Pavliotis, Grigorios A. (2016). *Variance Reduction Using Nonreversible Langevin Samplers*. Journal of Statistical Physics 163: 457–491.

Bayesian Anomaly Detection in Variable-Order and Variable-Diffusivity Fractional Mediums

Hamza Ruzayqat

King Abdullah University of Science and Technology

Hamza.Ruzayqat@kaust.edu.sa

Coauthor(s): Omar Knio and George Turkiyyah

Fractional diffusion equations (FDEs) are powerful tools for modeling anomalous diffusion in complex systems, such as fractured media and biological processes, where nonlocal dynamics and spatial heterogeneity are prominent. These equations provide a more accurate representation of such systems compared to classical models but pose significant computational challenges, particularly for spatially varying diffusivity and fractional orders. In this talk, I will present a Bayesian inverse problem for FDEs in a 2-dimensional bounded domain with an anomaly of unknown geometric and physical properties, where the latter are the diffusivity and fractional order fields. To tackle the computational burden of solving dense and ill-conditioned systems, we employ an advanced finite-element scheme incorporating low-rank matrix representations and hierarchical matrices. For parameter estimation, we implement two surrogate-based approaches using polynomial chaos expansions: one constructs a 7-dimensional surrogate for simultaneous inference of geometrical and physical parameters, while the other leverages solution singularities to separately infer geometric features, then constructing a 2-dimensional surrogate to learn the physical parameters and hence reducing the computational cost immensely. These surrogates are used inside a Markov chain Monte Carlo algorithm to infer the unknown parameters.

Theoretical Guarantees of Mean Field Variational Inference for Bayesian Principal Component Analysis

Arghya Datta

Université de Montréal

arghyadatta8@gmail.com

Coauthor(s): Philippe Gagnon, Florian Maire

In this talk, we will investigate mean field variational inference for Bayesian principal component analysis (BPCA). Despite the wide usage of mean field variational inference for the BPCA model, there exists remarkably little theoretical justification. I will talk about new results on the convergence guarantees of the iterative coordinate ascent variational inference (CAVI) algorithm for the BPCA model. In particular, we will show that under reasonable technical assumptions on the initialization, CAVI converges exponentially fast to a local optimum. An interesting connection between the CAVI algorithm for the BPCA model and power iteration, which is a popular iterative numerical algorithm for finding singular vectors of a given matrix, will also be discussed.

Bayesian Analysis of Latent Underdispersion Using Discrete Order Statistics

Jimmy Lederman

Department of Statistics, University of Chicago

jlederman@uchicago.edu

Coauthor(s): Aaron Schein

Researchers routinely analyze count data using models based on a Poisson likelihood, for which there exist many analytically convenient and computationally efficient strategies for posterior inference. A limitation of such models however is the equidispersion constraint of the Poisson distribution. This restriction prevents the model’s likelihood, and by extension its posterior predictive distribution, from concentrating around its mode. As a result, these models are parametrically bound to produce probabilistic predictions with high uncertainty, even in cases where low uncertainty is supported by the data. While count data often exhibits overdispersion *marginally*, such data may nevertheless be consistent with a likelihood that is underdispersed *conditionally*, given parameters and latent variables. Detecting conditional underdispersion, however, requires one to fit the “right” model and thus the ability to build, fit, and critique a variety of different models with underdispersed likelihoods. Towards this end, we introduce a novel family of models for conditionally underdispersed count data whose likelihoods are based on order statistics of Poisson random variables. More specifically, we assume that each observed count coincides with the j^{th} order-statistic of D latent i.i.d. Poisson random variables, where j and D are user-defined hyperparameters. To perform efficient MCMC-based posterior inference in this family of models, we derive a data-augmentation strategy which samples the other $D-1$ latent variables from their exact conditional, given the observed (j, D) -order statistic. By relying on the explicit construction of a Poisson order statistic, this data augmentation strategy can be modularly combined with the many existing inference strategies for Poisson-based models. We generalize this approach beyond the Poisson to any non-negative discrete parent distribution and, in particular, show that models based on negative binomial order statistics can flexibly capture both conditional under and overdispersion. To illustrate our approach empirically, we build and fit models to three real count data sets of flight times, COVID-19 cases counts, and RNA-sequence data, and we demonstrate how models with underdispersed likelihoods can leverage latent structure to make more precise probabilistic predictions. Although the possibility of conditional underdispersion is often overlooked in practice, we argue that this is at least in part due to the lack of tools for modeling underdispersion in settings where complex latent structure is present.

Monte Carlo simulation approach to solve distributed order fractional mathematical model

Yashveer Kumar

INESC-ID, Rua Alves Redol 9, Lisbon, Portugal 1000-029

yashveerkumar.rs.mat18@inesc-id.pt

Juan A. Acebron, INESC-ID, Department of Mathematics, Carlos III University of Madrid, Spain.

Jose Monteiro, INESC-ID, IST, Universidade de Lisboa, Portugal.

This work presents a novel approach to solving time-distributed order fractional nonhomogeneous differential equations using the Monte Carlo simulation method. Fractional differential equations with distributed orders are critical for modeling complex systems with memory and hereditary effects, such as viscoelastic materials, anomalous diffusion, and biological processes. The inclusion of time-distributed orders introduces additional challenges in analytical and numerical solutions, especially in the presence of nonhomogeneous terms.

The proposed Monte Carlo method reformulates the distributed order fractional equation into an equivalent integral representation. By simulating random processes and utilizing probabilistic interpretations of fractional operators, the solution is computed as an average over numerous realizations. The flexibility of Monte Carlo simulations makes them particularly well-suited for addressing the inherent complexity of distributed order systems.

Numerical experiments validate the efficiency and effectiveness of the Monte Carlo approach, illustrating its capability to handle various distributed order kernels and nonhomogeneous terms. This method offers a robust, scalable, and versatile framework for solving fractional differential equations, paving the way for broader applications in science and engineering.

- [1] Niederreiter, Harald (1992). *Random number generation and quasi-Monte Carlo methods*. Society for Industrial and Applied Mathematics (SIAM).
- [2] Guidotti, Nicolas L. & Acebron, Juan A. & Monterio, Jose (2024). *A stochastic method for solving time-fractional differential equations*. *Computers and Mathematics with Applications*, 159, 240-253.
- [3] Kumar, Yashveer & Singh, Vineet Kumar (2021). *Computational approach based on wavelets for financial mathematical model governed by distributed order time-fractional partial differential equation*. *Mathematics and Computers in Simulation*, 190, 531-569.

Gradient-based MCMC in high dimensions

Reuben Cohn-Gordon

University of California, Berkeley

reubenharry@gmail.com

Coauthor(s): Jakob Robnik, Uroš Seljak

Sampling from distributions over \mathbb{R}^d for d larger than 10^4 arises as a computational challenge in many of the physical sciences, including particle physics [1], condensed matter physics [2], cosmology [3] and chemistry [4], as well as in Bayesian statistics and machine learning [5]. Commonly used gradient-based variants of Markov Chain Monte Carlo such as Hamiltonian Monte Carlo (HMC) [6] and in particular the No U-Turn Sampler [7], are designed for differentiable multivariate densities, but struggle in very high dimensions. We propose a general purpose approach to the gradient-based high-dimensional regime, based on two insights. First, in high-dimensional cases where limited asymptotic bias is acceptable, Markov Chain algorithms without Metropolis-Hastings (MH) adjustment are more statistically efficient; we provide theoretical and numerical evidence for this claim and show how to choose a step size to limit the incurred bias to an acceptable level. Second, in the case that MH adjustment is required, we show that a particular 4th-order integrator [8] drastically improves the statistical efficiency of HMC and related algorithms in high dimensions.

- [1] Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D. (1987). Hybrid Monte Carlo. *Physics Letters B*, 195(2), 216-222.
- [2] Lunts, P., Albergo, M. S., & Lindsey, M. (2023). Non-Hertz-Millis scaling of the antiferromagnetic quantum critical metal via scalable Hybrid Monte Carlo. *Nature Communications*, 14(1), 2547.
- [3] Lewis, A., & Bridle, S. (2002). Cosmological parameters from CMB and other data: A Monte Carlo approach. *Physical Review D*, 66(10), 103511.
- [4] Tuckerman, M. E. (2023). *Statistical mechanics: theory and molecular simulation*. Oxford University Press.
- [5] Cobb, A. D., & Jalaian, B. (2021, December). Scaling Hamiltonian Monte Carlo inference for Bayesian neural networks with symmetric splitting. In *Uncertainty in Artificial Intelligence* (pp. 675-685). PMLR.
- [6] Betancourt, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. arXiv preprint arXiv:1701.02434.
- [7] Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15(1), 1593-1623.
- [8] Takaishi, T., & De Forcrand, P. (2006). Testing and tuning symplectic integrators for the hybrid Monte Carlo algorithm in lattice QCD. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 73(3), 036706.

Benchmarking the Geant4-DNA 'UHDR' Example for Monte Carlo Simulation of pH Effects on Radiolytic Species Yields Using a Mesoscopic Approach

Serena Fattori

Istituto Nazionale di Fisica Nucleare (INFN), Laboratori Nazionali del Sud (LNS), Catania, Italy

serena.fattori@lns.infn.it

Coauthor(s): Hoang Ngoc Tran, Anh Le Tuan, Fateme Farokhi, Giuseppe Antonio Pablo Cirrone, Sebastien Incerti

Background and Aims

FLASH radiotherapy is an innovative cancer treatment technique that delivers high radiation doses in an extremely short time (≥ 40 Gy/s), inducing the so-called FLASH effect—characterized by the sparing of healthy tissue while maintaining effective tumor control. However, the mechanisms underlying the FLASH effect remain unclear, and ongoing research aims to elucidate them. One approach to investigating this phenomenon is through Monte Carlo simulations of particle transport and the resulting radiolysis in aqueous media, enabling comparisons between FLASH and conventional irradiation.

Methods

To provide a useful tool for investigating the effects of FLASH irradiation, the Geant4-DNA example "UHDR" was introduced in the beta release 11.2.0 of Geant4 (June 2023). This example incorporates a newly developed radiolysis chemical stage based on the diffusion-reaction master equation (RDME), a mesoscopic method that bridges microscopic particle-level interactions and macroscopic chemical kinetics. This approach allows the extension of the simulation time to minutes post-irradiation, enabling the validation of equilibrium processes that may play a crucial role on long time scales. In this context, the impact of pH on radiolytic species yields towards equilibrium is particularly important. For the first time in Geant4-DNA, the UHDR example allows taking into account the effect of different pH values on water radiolysis.

Results

This study aims to benchmark the capability of the UHDR example to accurately reproduce the effect of pH on radiolytic species yields. Preliminary results are currently under analysis for 1 MeV electron and 300 MeV proton irradiation in the conventional modality, with comparisons against literature data.

Conclusions

The ability to simulate the impact of pH on water radiolysis represents a significant advancement in studying the evolution of radiolytic species toward equilibrium. This improvement could provide valuable insights into potential differences in chemical evolution under FLASH irradiation compared to conventional irradiation.

Fri, Aug 1 10:30–11:00

Chi-Ok Hwang

Gwangju Institute of Science and Technology, Gwangju 61005, Republic of Korea
chwang@gist.ac.kr

[JinseongSon](#), [MaximilianoIslasSolis](#), [TsoggerelTsogbadrakh](#)

According to probabilistic potential theory, first- and last-passage algorithms have been developed. Usually the first-passage algorithms with an enclosing sphere are used for overall charge distribution on a closed conducting object and last-passage algorithms for charge density at a specific point on the conducting object. The first- and last-passage algorithms are inherently connected. In this paper, we combine the first- and last-passage algorithms. We develop an algorithm for computing charge density at a specific point on the conducting object via the overall charge density distribution on a conducting object which is the simulation result of the first-passage algorithm with an enclosing sphere. We demonstrate the algorithm for charge density on a sphere and on the unit cube held at unit potential. The results show good agreement with theoretical or other simulation ones.

Acknowledgments: This work was supported by the GIST Research Institute (GRI) in 2024.

- [1] J. Son, J. Im, and C.-O. Hwang, *Appl. Math. Comput.* submitted (2021).
- [2] H. Jang, U. Yu, Y. Chung, and C.-O. Hwang, *Adv. Theory Simul.* 3(8) (2020).
- [3] H. Jang, J. Given, U. Yu, and C.-O. Hwang, *Adv. Theory Simul.*
<https://onlinelibrary.wiley.com/doi/full/10.1002/adts.20200026> (2021).
- [4] C.-O. Hwang and T. Won, *J. Korean Phys. Soc.* 47, S464 (2005).
- [5] J. A. Given, C.-O. Hwang, and M. Mascagni, *Phys. Rev. E* 66, 056704 (2002).

Halton Sequences, Scrambling and the Inverse Star-Discrepancy

Christian Weiß

Ruhr West University of Applied Sciences

christian.weiss@hs-ruhrwest.de

Halton sequences are classical examples of multi-dimensional low-discrepancy sequences. Braaten and Weller discovered that scrambling strongly reduces their empirical star-discrepancy. A similar approach may be applied to certain multi-parameter subsequences of Halton sequences. Indeed, results from p-adic analysis guarantee that these subsequences still have the theoretical low-discrepancy property while scrambling has strong effects on the empirical star-discrepancy. By optimizing the parameters of these subsequences known empiric bounds for the inverse star-discrepancy can be improved.

- [1] E. Braaten, and G. Weller. (1979). An improved low-discrepancy sequence for multi-dimensional quasi-Monte Carlo integration. *Journal of Computational Physics*, 33(2): 249–258.
- [2] C. Weiß. (2024). Scrambled Halton Subsequences and Inverse Star-Discrepancy, arXiv: 2411.10363.

Transport Quasi-Monte Carlo

Sifan Liu

Center for Computational Mathematics, Flatiron Institute

sliu@flatironinstitute.org

Quasi-Monte Carlo (QMC) is a powerful method for evaluating high-dimensional integrals. However, its use is typically limited to distributions where direct sampling is straightforward, such as the uniform distribution on the unit hypercube or the Gaussian distribution. For general target distributions with potentially unnormalized densities, leveraging the low-discrepancy property of QMC to improve accuracy remains challenging. We propose training a transport map to push forward the uniform distribution on the unit hypercube to approximate the target distribution. Inspired by normalizing flows, the transport map is constructed as a composition of simple, invertible transformations. To ensure that RQMC achieves its superior error rate, the transport map must satisfy specific regularity conditions. We introduce a flexible parametrization for the transport map that not only meets these conditions

but is also expressive enough to model complex distributions. Our theoretical analysis establishes that the proposed transport QMC estimator achieves faster convergence rates than standard Monte Carlo, under mild and easily verifiable growth conditions on the integrand. Numerical experiments confirm the theoretical results, demonstrating the effectiveness of the proposed method in Bayesian inference tasks.

Mon, Jul 28 16:30–17:00

Using Normalizing Flows for Efficient Quasi-Random Sampling for Copulas

Ambrose Emmett-Iwaniw

University of Waterloo Department of Actuarial Science and Statistics

arsemmettiwaniw@uwaterloo.ca

Coauthor(s): Christiane Lemieux

In finance and risk management, copulas are used to model the dependence between stock prices and insurance losses to compute expectations of interest. Generally, Monte Carlo (MC) sampling is used to generate copula samples to approximate expectations. To reduce the variance of the approximation, we can use quasi-Monte Carlo (QMC) sampling to generate copula samples. This paper examines a new method to generate quasi-random samples from copulas requiring fewer training resources than previous methods such as the generative moment matching networks (GMMN) model [1]. Traditional methods that do not use generative models often rely on conditional distribution methods (CDM) to generate quasi-random samples from specific copulas [2]. CDM is limited to only a few parametric copulas (Gumbel has no efficient CDM to sample quasi-random samples) in low dimensions [2]. Here, we propose using a powerful and simple generative model called Normalizing Flows (NFs) to generate quasi-random samples for any copula, including cases where we only have data available. NFs are a type of explicit generative model that relies on transforming a simple density, such as a normal density, through efficient invertible transformations that rely on the change of variables formula into a density that models complex data that facilitates easy sampling and efficient inverting of samples from complex data to normal data and vice versa. The benefit of these NFs for copula modelling is that their training is efficient in terms of runtime, allowing for larger batch sizes compared to the GMMN model [1]. Also, it is sample-efficient; it only needs samples from the copula and not samples from the normal as the GMMN model [1] required. Once the NF model is trained, we can efficiently invert the model to take as input quasi-random samples to generate quasi-random copula samples. Through many different simulations and applications, we show our approach allows us to leverage the benefit of QMC in a variety of real-world settings involving dependent data.

- [1] Hofert, M., Prasad, A., and Zhu, M. (2021). Quasi-random sampling for multivariate distributions via generative neural networks. *Journal of Computational and Graphical Statistics*, 30(3):647–670.
- [2] Cambou, M., Hofert, M., and Lemieux, C. (2017). Quasi-random numbers for copula models. *Statistics and Computing*, 27:1307–1329.

Mon, Jul 28 17:00–17:30

Optimization of Kronecker Sequences

Claude D. Hall Jr.

Illinois Institute of Technology

cdhjr2000@gmail.com, chall14@iit.edu

Coauthor(s): Dr. Fred Hickernell, Dr. Nathan Kirk

Our goal is to optimize Kronecker sequences [1,2] for Quasi-Monte Carlo integration. A shifted Kronecker sequence can be defined as $\{x_i = i\boldsymbol{\alpha} + \boldsymbol{\Delta} \pmod{1} : i = 0, 1, 2, 3, \dots\} \subset [0, 1]^d$, where $\boldsymbol{\alpha} \in [0, 1]^d$ is our generating vector and $\boldsymbol{\Delta} \in [0, 1]^d$ is the shift. There is no preferred sample size for Kronecker sequences, whereas the similarly defined lattice sequences do have preferred sample sizes.

For discrepancies defined in terms of certain families of periodic reproducing kernels, when you replace any one of the elements α_j by $1 - \alpha_j$, the discrepancy stays unchanged. This allows one to shrink the search space for good generating vectors. Performing a component by component search for $\boldsymbol{\alpha}$, we find that discrepancy for the Kronecker sequence converges nearly like $O(n^{-1})$, especially for smaller dimensions. This talk presents the results of our component-by-component search algorithm and the performance of Kronecker sequences as measured by discrepancy and test cases.

- [1] G. Larcher, *On the distribution of s-dimensional Kronecker-sequences*, *Acta Arith.* (1988), 335–347.
- [2] H. Niederreiter, *Quasi-Monte Carlo methods and pseudo-random numbers*, *Bull. Amer. Math. Soc.* 84 (1978), 957–1041.

Approximation using median lattice algorithms

Peter Kritzer

Austrian Academy of Sciences

peter.kritzer@oeaw.ac.at

Coauthor(s): Takashi Goda, Zexin Pan

We consider L_2 -approximation of functions in a weighted Korobov space. We present a median algorithm, which is related to median integration rules, that have recently gained a lot of interest in the theory of quasi-Monte Carlo methods. Indeed, we use lattice rules as the underlying integration rules to approximate Fourier coefficients. As we will show, we can obtain a convergence rate that is arbitrarily close to optimal in terms of the number of evaluations needed of the function to be approximated.

Convergence Rates of Randomized Quasi-Monte Carlo Methods under Various Regularity Conditions

Yang Liu

King Abdullah University of Science and Technology

yang.liu.3@kaust.edu.sa

In this work, we analyze the convergence rate of randomized quasi-Monte Carlo (RQMC) methods under Owen’s boundary growth condition (Owen, 2006) via spectral analysis. We examine the RQMC estimator variance for two commonly studied sequences—the lattice rule and the Sobol’ sequence—using the Fourier transform and Walsh–Fourier transform, respectively. Under certain regularity conditions, our results reveal that the asymptotic convergence rate of the RQMC estimator’s variance closely aligns with the exponent specified in Owen’s boundary growth condition for both sequence types. We also provide an analysis for certain discontinuous integrands.

In addition, we investigate the L^p integrability of weak mixed first-order derivatives of the integrand and study the convergence rates of scrambled digital nets. We demonstrate that the generalized Vitali variation with parameter $\alpha \in [\frac{1}{2}, 1]$ from Dick and Pillichshammer (2010) is bounded above by the L^p norm of the weak mixed first-order derivative, where $p = \frac{2}{3-2\alpha}$. Consequently, when the weak mixed first-order derivative belongs to L^p for $1 \leq p \leq 2$, the variance of the scrambled digital nets estimator converges at a rate of $\mathcal{O}\left(N^{-4+\frac{2}{p}} \log^{s-1} N\right)$. Together, these results provide a comprehensive theoretical framework for understanding the

convergence behavior of RQMC methods and scrambled digital nets under various regularity assumptions.

- [1] Liu, Y. (2024). Randomized quasi-Monte Carlo and Owen’s boundary growth condition: A spectral analysis. arXiv preprint arXiv:2405.05181.
- [2] Liu, Y. (2025). Integrability of weak mixed first-order derivatives and convergence rates of scrambled digital nets. arXiv preprint arXiv:2502.02266.

Tue, Jul 29 16:30–17:00

Use of rank-1 lattices in the Fourier neural operator

Jakob Dilen

Department of Computer Science, KU Leuven

jakob.dilen@student.kuleuven.be

Coauthor(s): Frances Y. Kuo, Dirk Nuyens

The “Fourier neural operator” [2] is a variant of the “neural operator”. Its defining characteristic, compared to the regular neural operator, is that it transforms the input to the Fourier domain at the start of each layer. This transformation uses the d -dimensional FFT on a regular grid in d dimensions. We describe how to do this more efficiently using rank-1 lattice points, which allow for a one-dimensional FFT algorithm, see, e.g., [1].

- [1] F. Y. Kuo, G. Migliorati, F. Nobile, and D. Nuyens. *Function integration, reconstruction and approximation using rank-1 lattices*. Math. Comp., 90, April 2021.
- [2] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. *Fourier neural operator for parametric partial differential equations*. International Conference on Learning Representations, 2021.

Investigating the Optimum RQMC Batch Size for Betting and Empirical Bernstein Confidence Intervals

Aadit Jain

Rancho Bernardo High School

aaditdjain@gmail.com

Coauthor(s): Fred J. Hickernell, Art B. Owen, Aleksei G. Sorokin

The Betting [1] and Empirical Bernstein (EB) [1,2] confidence intervals (CIs) are finite sample (non-asymptotic) and require IID samples. Since both are non-asymptotic, they are much wider than confidence intervals based on the Central Limit Theorem (CLT) due to the stronger coverage property they provide. To apply these finite sample CIs to randomized quasi-Monte Carlo (RQMC), we take R independent replications of n RQMC points, averaging the n function evaluations within each replication. Given a fixed budget $N = nR$, we investigate the optimal n that minimizes the CI widths for both methods.

Using the code from [1], we ran simulations on various integrands (smooth, rough, one-dimensional, multi-dimensional) and ridge functions. Interestingly, the optimal n was quite small compared to N , often just 1 (plain IID), 2, or 4 when $N = 2^{10}$. Moreover, the optimal n appeared to grow quite slowly as N increased. Notably, both CI methods applied to RQMC outperformed plain IID when the optimal n is greater than 1.

This experimental trend aligns with our analysis of Bennett’s inequality for EB [2], which suggests that the optimum n is $O(N^{1/(2\theta+1)})$ for $\theta > 1/2$. Specifically, for $\theta = 3/2$, which occurs for smoother integrands, we obtain $n = O(N^{1/4})$. For $\theta = 1$, which corresponds to a typical Koksma-Hlawka rate, we get $n = O(N^{1/3})$. The ratio of RQMC EB CI widths to plain IID EB CI widths is $\Theta(N^{(1-2\theta)/(4\theta+2)})$. For $\theta = 1$, we get a ratio of $\Theta(N^{-1/6})$, while for $\theta = 3/2$, we get a more favorable width ratio of $\Theta(N^{-1/4})$.

On the other hand, CLT based CIs using RQMC are only asymptotically valid. The value of R could be any reasonable number that isn’t too small, and remains constant as the total sample size, N , increases. This means that $n = \Theta(N)$, which takes full advantage of the power of QMC. It is also important to note that both Betting and EB require the random variables to be bounded between 0 and 1, unlike CLT based CIs.

- [1] I. Waubdy-Smith and A. Ramdas. Estimating means of bounded random variables by betting. *J. Roy. Statist. Soc. B*, 86:1–27, 2024.
- [2] A. Maurer and M. Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, pages 1–9, 2009.

Joonha Park

University of Kansas, Department of Mathematics
j.park@ku.edu

Hamiltonian Monte Carlo (HMC) is widely used for sampling from high-dimensional target distributions with probability density known up to proportionality. While HMC possesses favorable dimension scaling properties, it encounters challenges when applied to strongly multimodal distributions. Traditional tempering methods, commonly used to address multimodality, can be difficult to tune, particularly in high dimensions. In this study, we propose a method that combines a tempering strategy with Hamiltonian Monte Carlo, enabling efficient sampling from high-dimensional, strongly multimodal distributions. Our approach involves proposing candidate states for the constructed Markov chain by simulating Hamiltonian dynamics with time-varying mass, thereby searching for isolated modes at unknown locations. Moreover, we develop an automatic tuning strategy for our method, resulting in an automatically-tuned, tempered Hamiltonian Monte Carlo (ATHMC). Unlike simulated tempering or parallel tempering methods, ATHMC provides a distinctive advantage in scenarios where the target distribution changes at each iteration, such as in the Gibbs sampler. We numerically show that our method scales better with increasing dimensions than an adaptive parallel tempering method and demonstrate its efficacy for a variety of target distributions, including mixtures of log-polynomial densities and Bayesian posterior distributions for a sensor network self-localization problem.

Localized consensus-based sampling for non-Gaussian distributions*Arne Bouillon*

KU Leuven, Leuven, Belgium

arne.bouillon@kuleuven.be

Coauthor(s): Alexander Bodard, Panagiotis Patrinos, Dirk Nuyens, and Giovanni Samaey

Drawing samples distributed according to a given unnormalized probability density function is a common task in Bayesian inverse problems. Algorithms based on an ensemble of interacting *particles*, moving in parameter space, are gaining in popularity for these problems since they are often parallelizable, derivative-free, and affine-invariant. However, most are only accurate for near-Gaussian target distributions; an example is the consensus-based sampling (CBS) method [1]. We propose a novel way to derive CBS from ensemble-preconditioned Langevin diffusions by first approximating the target potential by its anisotropic Moreau envelope, then approximating the proximal operator by a weighted mean, and finally assuming that the initial and target distributions are Gaussian. We adapt these approximations with non-Gaussian distributions in mind and arrive at a new interacting-particle method for sampling, which we call *localized consensus-based sampling*. Numerical tests illustrate that localized CBS compares favorably to alternative methods in terms of affine-invariance and performance on non-Gaussian distributions.

- [1] Carrillo, J. A., Hoffmann, F., Stuart, A. M., & Vaes, U. (2022). Consensus-based sampling. *Studies in Applied Mathematics*, **148**(3), 1069–1140.

Polynomial approximation for efficient transport-based sampling*Josephine Westermann*

Heidelberg University

josephine.westermann@uni-heidelberg.de

Sampling from non-trivial probability distributions is a fundamental challenge in uncertainty quantification and inverse problems, particularly when dealing with high-dimensional domains, costly-to-evaluate or unnormalized density functions, and non-trivial support structures. Measure transport via polynomial density surrogates [1] provides a systematic and constructive solution to this problem by reformulating it as a convex optimization task with

deterministic error bounds. This approach is particularly efficient for smooth problems but can become computationally demanding when dealing with highly concentrated posterior distributions, as constructing the density surrogate requires many evaluations in regions where the density is nearly zero—leading to an inefficient allocation of computational resources. In this talk, we explore how the computational cost in such cases can be further reduced by first approximating the potential function with a polynomial. This additional approximation step shifts the focus of expensive evaluations toward capturing the underlying system more effectively. The surrogate-based posterior can then be evaluated cheaply and approximated with high accuracy, enabling efficient transport-based sampling. We discuss the implications of this strategy and examine its potential to enhance performance in demanding inference problems.

- [1] Westermann, J., & Zech, J. (2025). *Measure transport via polynomial density surrogates*. *Foundations of Data Science*. <https://doi.org/10.3934/fods.2025001>

Thu, Jul 31 16:00–16:30

Fast Approximate Matrix Inversion via MCMC for Linear System Solvers

Soumyadip Ghosh

IBM Research

ghoshs@us.ibm.com

Coauthor(s): Vassil Alexandrov, Lior Horesh, Vasilieos Kalantzis, Anton Lebedev, WonKyung Lee, Yingdong Lu, Tomasz Nowicki, Shashanka Ubaru, Olha Yaman

A key prerequisite of modern iterative solvers of linear algebraic equations $Ax = b$ is the fast computation of a pre-conditioner matrix P that gives a good approximation to the (generalized) inverse of A such that the set of equations obtained by pre-multiplying with P , $PAx = Pb$, is solved quickly. We study the classical Ulan-von Neumann MCMC algorithm that was designed based on the Neumann infinite series representation of the inverse of a non-singular matrix. The parameters of the MCMC algorithm determine the overall time to solve $Ax = b$, which is a metric that reflects both the time to compute the MCMC preconditioner P and its quality as a preconditioner to solve the linear equations. Our main focus is on how the MCMC parameters should be tuned to speed up computations in applications that require repeated calls to the solver with varying matrices A , a common scenario for instance in numerical approximations of physical phenomena. We present a model that relates key features of matrices A with good choices of MCMC algorithm parameters that lead to a fast overall time to find a solution to $Ax = b$. A computationally efficient approach based on Bayesian experimental design is described to learn and update this model while minimizing

the number of runs of the expensive solver in application settings that solve of linear system over well-defined sets of A matrices. We present numerical experiments to illustrate the efficacy of this approach. In another contribution, we present a new MCMC algorithm which we term as *regenerative Ulam-von Neumann* algorithm. It exploits a regenerative structure present in the Neumann series that underlies the original algorithm and improves on it by producing an unbiased estimator of the matrix inverse. A rigorous analysis of performance of the algorithm is provided. This includes the variance of the estimator, which allows one to estimate the time taken to obtain solutions of a desired quality. Finally, numerical experiments verify the qualitative effectiveness of the proposed scheme.

Thu, Jul 31 16:30–17:00

Investigating general $L2$ discrepancies with Message-Passing Monte Carlo

Ally Kwan and Lijia Lin

Foothill College and Johns Hopkins University

ally.kwan@yale.edu, llin72@jh.edu

Coauthor(s): Nathan Kirk

Finding low-discrepancy (LD) points are essential to facilitate the use of quasi-Monte Carlo methods, which achieve faster convergence rate compared to traditional Monte Carlo methods. The Message-Passing Monte Carlo (MPMC)¹ method utilizes graph neural networks to generate LD point sets, where, in the original model, the learning objective is set to the classical $L2$ star discrepancy. In this talk, we employ MPMC as an optimization tool to investigate the properties of LD point sets optimized with respect to other $L2$ discrepancy measures known in the literature and show evidence that the star discrepancy should not always be the standard measure to use in optimization pipelines. Further, we demonstrate the new MPMC module in the QMCPy Python package that is available for use with functionality such as choosing the discrepancy learning objective, specifying weights for high-dimensional problems, randomizations, and pretrained points for particular N and d combinations to facilitate use of the MPMC method.

- [1] Rusch, T. Konstantin and Kirk, Nathan and Bronstein, Michael M. and Lemieux, Christiane and Rus, Daniela (2024). *Message-Passing Monte Carlo: Generating low-discrepancy point sets via graph neural networks*. Proceedings of the National Academy of Sciences (PNAS).

Discrepancy calculation and generating vector searches for Kronecker and extensible lattice sequences*Jimmy Nguyen and Anders Pride*

University of California, Irvine and Illinois Institute of Technology

jimmyhn7@uci.edu, apride@hawk.illinoistech.edu

We investigate extensible lattice sequences and Kronecker sequences for arbitrary sample sizes n , which are used for Quasi-Monte Carlo sampling. We give methods for efficiently computing the discrepancy and the weighted sum of squared discrepancy of these points, where the discrepancy is defined in terms of a symmetric, positive definite, periodic kernel of product form. We use these measures to perform component-by-component searches for generating vectors and compare these to popular choices. Additionally, we implement various stopping criteria for Quasi-Monte Carlo cubature and are adding our contributions to the Python QMCPy open source library.

Characterizing Efficacy of Geometric Brownian Motion Expectation-based Simulations on Low-Volatility American Common Stocks

Vincent X Zhang

University of Southern California

vxzhang@usc.edu

Coauthor(s): Heather Choi

In this manuscript, daily, monthly, and annual geometric Brownian motion forecasts are obtained and tested for reliability upon 21 stock symbols within NASDAQ of varying volatilities and drifts. Biweekly, monthly, biannual, and annual rolling windows were used as a preliminary filtering scheme to remove unreliable stock symbols, and then accuracy was further evaluated on stocks with higher accuracies in the first screening. Annual and 10-year windows were used to estimate the drift and diffusion component and then applied to obtain one-period-ahead geometric Brownian motion stock values and associated probabilities. Further building off of these one-period-ahead values, expected values for 1-252 periods were estimated. Expected values of each stock were estimated by totaling up the product of the stock value and its associated probabilities, and tested over multiple rolling-windows for reliability. The results indicate that geometric Brownian-simulated expected index values estimated using one thousand simulations can be slightly reliable if catered and re-optimized to specific stock characterizations, but only for a daily window, and even then only slightly preferable to flipping a coin. Expected values estimated with less than 100 simulations were thrown out, seen as unreliable.

- [1] Sinha, A. (2024). Daily and Weekly Geometric Brownian Motion Stock Index Forecasts. *Journal of Risk and Financial Management*, 17(10), 434.
<https://doi.org/10.3390/jrfm17100434>
- [2] Abbas, Anas, and Mohammed Alhagyan. 2023. Forecasting exchange rate of SAR/CNY by incorporating memory and stochastic volatility into GBM model. *Advances and Applications in Statistics* 86: 65–78
- [3] Demirel, Ugur, Handan Cam, and Ramazan Unlu. 2021. Predicting Stock Prices Using Machine Learning Methods and Deep Learning Algorithms: The Sample of the Istanbul Stock Exchange. *Gazi University Journal of Science* 34: 63–82
- [4]] Tie, Jingzhi, Hanqin Zhang, and Qing Zhang. 2018. An Optimal Strategy for Pairs Trading Under Geometric Brownian Motions. *Journal of Optimization Theory and Applications* 179: 654–75
- [5] Ramos, André Lubene, Douglas Batista Mazzinghy, Viviane da Silva Borges Barbosa, Michel Melo Oliveira, and Gilberto Rodrigues da Silva. 2019. Evaluation of an Iron Ore Price Forecast Using a Geometric Brownian Motion Model. *Revista Escola de Minas* 72: 9–15
- [6] Nordin, Norazman, Norizarina Ishak, Nurfadhлина Abdul Halim, Siti Raihana Hamzah,

and Ahmad Fadly Nurullah Rasadee. 2024. A Geometric Brownian Motion of ASEAN-5 Stock Indexes. In *AI and Business, and Innovation Research: Understanding the Potential and Risks of AI for Modern Enterprises*. Edited by Bahaaeddin Alareeni and Islam Elgedawy. New York: Springer, pp. 779–86

- [7] W Farida Agustini et al. 2018 *J. Phys.: Conf. Ser.* 974 012047
- [8] Liang, W., Li, Z., Chen, W. (2025). Enhancing Financial Market Predictions: Causality-Driven Feature Selection. In: Sheng, Q.Z., et al. *Advanced Data Mining and Applications. ADMA 2024. Lecture Notes in Computer Science()*, vol 15387. Springer, Singapore. https://doi.org/10.1007/978-981-96-0811-9_11
- [9] Prasad, Krishna, Bhuvana Prabhu, Lionel Pereira, Nandan Prabhu, and Pavithra S. 2022. Effectiveness of geometric Brownian motion method in predicting stock prices: Evidence from India. *Asian Journal of Accounting and Governance* 18: 121–34
- [10] Shafii, Nor Hayati, Nur Ezzati, Dayana Mohd, Rohana Alias, Nur Fatimah Fauzi, and Mathematical Sciences. 2019. Fuzzy Time Series and Geometric Brownian Motion in Forecasting Stock Prices in Bursa Malaysia. *Jurnal Intelek* 14: 240–50
- [11] Zhu, Song-Ping, and Xin-Jiang He. 2018. A New Closed-Form Formula for Pricing European Options under a Skew Brownian Motion. *The European Journal of Finance* 24: 1063–74

Thu, Jul 31 11:00–11:30

Efficient Pricing for Variable Annuity via Simulation

Hao Quan

University of Waterloo

h5quan@uwaterloo.ca

Coauthor(s): Ben (Mingbin) Feng

Variable Annuities (VAs) are insurance products that offer policyholders exposure to financial market upside potential while safeguarding against downside risk through optional riders, such as Guaranteed Minimum Death Benefits (GMDBs), Guaranteed Minimum Accumulation Benefits (GMABs), and Guaranteed Minimum Withdrawal Benefits (GMWBs). These riders, tailored to policyholders' needs, introduce a complex risk profile combining mortality and financial uncertainties, rendering VA pricing and fee determination computationally challenging. Due to this complexity, Monte Carlo simulation is often the only practical approach for valuing these contracts.

In this study, we address the problem of setting fair management fees for VA rider combinations using the equivalence principle, which balances the expected present value of premiums

and benefits. We formulate fee determination as a stochastic root-finding problem, expressed as

$$E[V(\varphi)] - P = 0,$$

where $E[V(\varphi)]$ denotes the expected present value of VA benefits under fee structure φ , and P represents the premium. The VA benefit $V(\varphi)$ reflects the evolution of the contract's shadow account value and various benefit guarantees over the contract's lifetime. As a result, estimating its expected value is computationally challenging. Moreover, solving the root-finding problem requires estimating the gradient of this expectation. To solve this, we employ stochastic gradient estimation techniques, such as finite differences and infinitesimal perturbation analysis (IPA). We analyze the theoretical properties and computational performance of the proposed root-finding algorithms, offering insights into their efficacy for VA pricing. Our results show that gradient estimation techniques have a significant impact on the efficiency and accuracy of estimating fair fees for various rider combinations.

Fri, Aug 1 09:00–09:30

Quantitative results on sampling from quasi-stationary distributions

Daniel Yukimura

IMPA, Rio de Janeiro, Brazil

yukimura@impa.br

Coauthor(s): Roberto I. Oliveira

We study the rate of convergence of Sequential Monte Carlo (SMC) methods for approximating the quasi-stationary distribution (QSD) of Markov processes. For processes with killing or absorption, the QSD appears as a stable behavior observed before extinction, or as the limiting distribution of the process conditioned on not being absorbed. We give quantitative lower and upper bounds for the particle filter approximation of these distributions. For the lower bound, we show that fast mixing is not enough to guarantee that simulation methods can converge in few steps. In the upper bound, we show that SMC with adaptive resampling has a rate depending on the number of steps, the mixing time, and in how fast the processed gets killed. Our result seems to be the first to have a quantitative dependency of this form that is valid for discrete time Markov chains in general state spaces. Our techniques and concentration results for bounding the approximations of SMC with adaptive resampling are also novel, and we believe they might be applicable in other scenarios that can benefit from the lower variance obtained due to an adaptive approach.

Multilevel simulation of ensemble Kalman methods: interactions across levels*Toon Ingelaere*

KU Leuven

toon.ingelaere@kuleuven.be

Coauthor(s): Giovanni Samaey

To solve problems in domains such as filtering, optimization, and posterior sampling, ensemble Kalman methods have recently received much attention. These parallelizable and often gradient-free algorithms use an ensemble of particles that evolve in time, based on a combination of well-chosen dynamics and interaction between the particles. For computationally expensive dynamics, the cost of attaining a high accuracy quickly becomes prohibitive. To improve the asymptotic cost-to-error relation, different multilevel Monte Carlo techniques have been proposed. These methods simulate multiple differently sized ensembles at different resolutions, corresponding to different accuracies and costs. While particles within one of these ensembles do interact with each other, a key question is whether and how particles should interact across ensembles and levels. In this talk, we will outline and compare the most common approaches to such multilevel ensemble interactions.

Low-Rank Thinning

Annabelle Michael Carrell

University of Cambridge

ac2411@cam.ac.uk

Coauthor(s): Albert Gong, Abhishek Shetty, Raaz Dwivedi, Lester Mackey

The goal in thinning is to summarize a dataset using a small set of representative points. Remarkably, sub-Gaussian thinning algorithms like Kernel Halving and Compress can match the quality of uniform subsampling while substantially reducing the number of summary points. However, existing guarantees cover only a restricted range of distributions and kernel-based quality measures and suffer from pessimistic dimension dependence. To address these deficiencies, we introduce a new low-rank analysis of sub-Gaussian thinning that applies to any distribution and any kernel, guaranteeing high-quality compression whenever the kernel or data matrix is approximately low-rank. To demonstrate the broad applicability of the techniques, we design practical sub-Gaussian thinning approaches that improve upon the best known guarantees for approximating attention in transformers, accelerating stochastic gradient training through reordering, and distinguishing distributions in near-linear time.

Dynamical Low-Rank Approximation for SDEs: an interacting particle-system ROM

Fabio Zoccolan

École Polytechnique Fédérale de Lausanne

fabio.zoccolan@epfl.ch

Coauthor(s): Dr. Yoshihito Kazashi, Prof. Fabio Nobile

The Dynamical Low-Rank Approximation (DLRA) technique is a time-dependent reduced-order model (ROM) known for its significant advantages in terms of computational time and accuracy. Its appeal in uncertainty quantification is due to the fact that its solution is composed of time-dependent deterministic and stochastic bases, allowing the approximation to better track the dynamics of the studied system. In the context of stochastic differential equations (SDEs) a rigorous mathematical setting was presented in [1], using the so-called Dynamically Orthogonal (DO) framework. The well-posedness of this setting is nontrivial due to the coupled nature of the DO system: for instance, the deterministic basis depends on all stochastic basis paths, and the equations involve the inversion of a Gramian matrix. When coming to stochastic discretization through a Monte-Carlo procedure, these features imply to deal with an interacting noisy particle dynamics. We proposed two fully discretized schemes based on the Monte-Carlo method, investigating their errors and analyzing possible issues arisen by the discretization of the Gramian inverse [2]. Theoretical results will be supported by numerical simulations.

- [1] Yoshihito Kazashi, Fabio Nobile, and Fabio Zoccolan. *Dynamical low-rank approximation for stochastic differential equations*. Mathematics of Computation (2024).
- [2] Yoshihito Kazashi, Fabio Nobile, and Fabio Zoccolan. *Numerical Methods for Dynamical low-rank approximation of stochastic differential equations, Part I & II*, in preparation (2025).

A probabilistic Numerical method for semi-linear elliptic Partial Differential Equations

Adrien Richou

adrien.richou@math.u-bordeaux.fr

EmmanuelGobet, CharuShardul, LukaszSzpruch

In this presentation, we study the numerical approximation of a class of Backward Stochastic Differential Equations (BSDEs) in an infinite horizon setting that provide a probabilistic representation for semi-linear elliptic Partial Differential Equations. In particular, we are also able to treat some ergodic BSDEs that are related to elliptic PDEs or ergodic type. In order to build our numerical scheme, we put forward a new representation of the PDE solution by using a classical probabilistic representation of the gradient. Then, based on this representation, we propose a fully implementable numerical scheme using a Picard iteration procedure, a grid space discretization and a Monte-Carlo approximation. We obtain an upper bound for the numerical error and we also provide some numerical experiments that show the efficiency of this approach for small dimensions. Some numerical experiments also show that it is possible to efficiently handle larger dimensions by replacing grid-based spatial discretization with neural networks. This presentation is based on [1] for the non-ergodic framework and [2] for results concerning the ergodic case.

- [1] Gobet, Emmanuel, & Richou, Adrien, & Shardul, Charu (2025). Numerical approximation of Markovian BSDEs in infinite horizon and elliptic PDEs. Draft.
- [2] Gobet, Emmanuel, & Richou, Adrien, & Szpruch, Lukasz (2025). Numerical approximation of ergodic BSDEs using non linear Feynman-Kac formulas, Preprint arXiv:2407.09034

A Chen-Fliess series for stochastic differential equations driven by Lévy processes

Anke Wiese

Heriot-Watt University, UK

A.Wiese@hw.ac.uk

Coauthor(s): Kurusch Ebrahimi-Fard, Frédéric Patras

Stochastic differential equations driven by Lévy processes have become established as models

to describe the evolution of random variables such as financial and economic variables and more recently of climate variables, when the stochastic system shows jump discontinuities. In this talk, we derive a novel series representation of the flowmap of such stochastic differential equations in terms of commutators of vector fields with stochastic coefficients, in other words a Chen–Strichartz formula. We provide an explicit expression for the components in this series. Our results extend previous results for deterministic and continuous stochastic differential equations.

Thu, Jul 31 17:00–17:30

Comparing Probabilistic Load Forecasters: Stochastic Differential Equations and Deep Learning

Riccardo Saporiti

EPFL, Lausanne, Switzerland

riccardo.saporiti@epfl.ch

Coauthor(s): Fabio Nobile, Celia García-Pareja

Generating probabilistic predictions for the electricity-load profile is the foundation of efficient use of renewable energy and diminishing carbon footprint.

In this talk, we consider the problem of creating probabilistic forecasts of the day-ahead electricity consumption profile of an agglomerate of buildings in the city of Lausanne (Switzerland) in the absence of an externally provided prediction function.

We propose a nonparametric, data-driven, approach based on Itô’ Stochastic Differential Equations (SDEs) [1]. Our work is novel in that the mean function of the SDE is expanded on a Fourier periodic basis, capturing intra-day and intra-week periodic features. Using a derivative tracking term, we impose the trajectories of the process to revert toward the mean. To model high-volatility levels associated with more uncertain electricity consumption regimes, we employ a square-root type diffusion coefficients.

Maximum-Likelihood estimation is used to infer the parameters of the model coherently with the available observations of the time history. We show that the maximization problem is well posed and that it admits at least one solution over the feasible domain.

We compare the probabilistic predictions generated by the SDE with Deep Learning based probabilistic forecaster. On the one hand, we introduce a Deep Learning forecaster based on Long short-term memory (LSTM) recurrent neural networks trained by minimizing the quantile loss function. This approach allows the generation of confidence intervals by sampling from the one-step-ahead univariate cumulative density function (CDF) associated with the electricity consumption of the future time instant. On the other hand, inspired by [2], we consider Multivariate Quantile Function Forecasters that, based on Normalizing Flows,

learn the multivariate cumulative density function of the day-ahead electricity consumption. Metrics such as Continuous ranked probability score and Prediction interval coverage percentage are used to assess the quality of the forecasts.

We show that SDEs generate reliable and interpretable predictions while presenting the most parsimonious and computationally efficient structure among the three models.

- [1] Riccardo Saporiti, Fabio Nobile, Celia García-Pareja. *Probabilistic Forecast of the Day-Ahead electricity consumption profile with Stochastic Differential Equations: a comparison with Deep Learning models*. In preparation.
- [2] Kelvin Kan, Francois-Xavier Aubet, Tim Januschowski, Youngsuk Park, Konstantinos Benidis, Lars Ruthotto, and Jan Gasthaus (2022). *Multivariate Quantile Function Forecaster*. Proceedings of The 25th International Conference on Artificial Intelligence and Statistics, PMLR 151:10603-10621, 2022.

Mon, Jul 28 15:30–16:00

Forward Propagation of Low Discrepancy Through McKean–Vlasov Dynamics: From QMC to MLQMC

Leon Wilkosz

King Abdullah University of Science and Technology (KAUST)

leon.wilkosz@kaust.edu.sa

Coauthor(s): Nadhir Ben Rached, Abdul-Lateef Haji-Ali, Raul Tempone

This work develops a particle system addressing the approximation of McKean–Vlasov stochastic differential equations (SDEs). The novelty of the approach lies in involving low-discrepancy sequences nontrivially in the construction of a particle system with coupled noise and initial conditions. Weak convergence for SDEs with additive noise is proven. A numerical study demonstrates that the novel approach presented here doubles the respective convergence rates for weak and strong approximation of the mean-field limit, compared with the standard particle system. These rates are proven in the simplified setting of a mean-field ordinary differential equation in terms of appropriate bounds involving the star discrepancy for low-discrepancy sequences with a group structure, such as Rank-1 lattice points. This construction nontrivially provides an antithetic multilevel quasi-Monte Carlo estimator. An asymptotic error analysis reveals that the proposed approach outperforms methods based on the classic particle system with independent initial conditions and noise.

A New Approach for Unbiased Estimation of Parameters of Partially Observed Diffusions

Miguel Alvarez

King Abdullah University of Science and Technology (KAUST)

miguelangel.alvarezballesteros@kaust.edu.sa

Coauthor(s): Ajay Jasra

In this talk, we consider the estimation of static parameters for a partially observed diffusion process with discrete-time observations over a fixed time interval. We develop particle filtering methods using time-discretization schemes, and we employ particle Markov chain Monte Carlo methods to estimate the smoothing distribution. In particular, we use backward sampling to address the issue of sample degeneracy. We use the score function and stochastic gradient ascent methods to maximize the likelihood of the observations. Parameter estimation in the diffusion term is possible by introducing bridge processes and the corresponding bridge-guiding proposals. To achieve unbiasedness, we adopt the Rhee and Glynn approach [1], in which the sources of bias are the number of stochastic gradient ascent steps and the time-discretization. Finally, we display numerically the method applying it to two systems.

- [1] Rhee, C. H., & Glynn, H. (2015). Unbiased estimation with square root convergence for SDE models, *Op. Res.*, **63**, 1026–1043.

High-order adaptive methods for exit times of diffusion processes and reflected diffusions

Håkon Hoel

University of Oslo

haakonah@math.uio.no

The Feynman–Kac formula connects domain-exit and boundary-reflection properties of stochastic differential equations (SDEs) and parabolic partial differential equations. The SDE viewpoint is particularly interesting for numerical methods, as it can be used with Monte Carlo methods to overcome the curse of dimensionality when solving high-dimensional PDE. This

however hinges on having an efficient numerical method for simulating the exit times of SDEs. Since exit times of diffusion processes are very sensitive to perturbations in initial conditions, it is challenging to construct such numerical methods.

This talk presents a high-order method with adaptive time-stepping for strong approximations of exit times. The method employs a high-order Itô–Taylor scheme for simulating SDE paths and carefully decreases the step size in the numerical integration as the diffusion process approaches the domain’s boundary. These techniques complement each other well: adaptive time-stepping improves the accuracy of the exit time by reducing the overshoot out of the domain, and high-order schemes improve the state approximation of the diffusion process, which is useful feedback to control the step size. We will also consider an ongoing extension of the numerical method to reflected diffusions.

Mon, Jul 28 17:00–17:30

Generative Modelling of Lévy Area for High-Order SDE Simulation

Thomas Cass

Imperial College London

thomas.cass@imperial.ac.uk

Coauthor(s): Andraž Jelinčič, Jiajie Tao, William F. Turner, James Foster, Hao Ni

Achieving high-order strong convergence in stochastic differential equation (SDE) simulations hinges on accurately sampling Lévy areas—iterated integrals of Brownian motion that are notoriously difficult to simulate, especially in higher dimensions. In this talk, we introduce LévyGAN, a novel generative model that leverages deep learning to produce high-fidelity samples of Lévy area conditioned on Brownian increments. Our architecture combines a GNN-inspired generator with a characteristic-function-based discriminator. A key innovation, the Bridge-flipping operation, guarantees exact matching of all odd moments, while our Chen-training method eliminates the need for costly training data. We demonstrate state-of-the-art performance in 4D Brownian motion and show practical gains in simulating the log-Heston model using multilevel Monte Carlo. This work opens new direction for efficient, high-order SDE simulation via generative modeling.

- [1] Jelinčič, A., Tao, J., Turner, W. F., Cass, T., Foster, J., & Ni, H. (2023). Generative Modelling of Lévy Area for High Order SDE Simulation. *accepted, to appear in SIAM Journal on Mathematics of Data Science*.

Learning cooling strategies in simulated annealing through binary interactions*Frédéric Blondeel*

KULeuven & UniFe

frederic.blondeel@kuleuven.be

Coauthor(s): Lorenzo Pareschi, Giovanni Samaey

Global optimization is amongst the hardest to solve problems. This is because finding the global minimum can usually only be guaranteed to be found in infinite time. Therefore one usually relies on meta-heuristic algorithms to guide the search and improve chances of successfully identifying the minimum. One particular family of algorithms is simulated annealing (SA). This family of algorithms is inspired by real-world metallurgy and is based on the Metropolis-Hastings algorithm. It works by randomly sampling N particles on the search space, each with a given temperature T . The movement of these particles is analogous to a Brownian random walk with step size proportional to temperature. The temperature is gradually cooled down over the course of the simulation according to some predefined schedule. Due to the Metropolis-Hastings-like acceptance-rejection rule, these particles can jump out of local minima and are expected to move towards the lowest energy state, i.e., the global minimum. It is the design of these cooling schedules for SA that we wish to improve. This is because they directly impact the efficiency of the optimization tool. Typically cooling schedules are inverse logarithmic or geometric decays in time. Here, we consider a collective SA dynamic where particles *interact* to learn the optimal temperature cooling strategy. This is inspired by the well-known particle-swapping technique known as parallel tempering (PT). To this aim we introduce a Boltzmann-type description where particles (partially) exchange their temperatures, therefore slowly cooling down the overall mean temperature. In order to simulate the dynamic we use a direct simulation Monte Carlo (DSMC) algorithm known as Nanbu-Babovsky. We show on various test functions (Ackley, Rastrigin, etc.) that this novel approach outperforms the standard SA with logarithmic and geometric annealing schedules.

Accuracy of Discretely Sampled Stochastic Policies in Continuous-Time Reinforcement Learning

Du Ouyang

Department of Mathematical Sciences, Tsinghua University

oyd21@mails.tsinghua.edu.cn

Coauthor(s): Yanwei Jia, Yufei Zhang

How to execute a stochastic policy in continuous-time environments is crucial for real-time operations and decision making. We show that, by sampling actions from a stochastic policy at a fixed time grid and then executing a piecewise constant control process, the controlled state process converges to the corresponding aggregated dynamics in the weak sense as the grid size shrinks to zero and obtain the convergence rate. Specifically, under sufficiently regular conditions on the coefficients, the optimal convergence rate of $O(|\mathcal{G}|)$ is achieved with respect to the time grid \mathcal{G} . For less regular coefficients, a convergence rate is established that varies according to the degree of regularity of the coefficients. Additionally, we also derive large deviation bounds for the weak error. Beyond weak error convergence, strong convergence results are proved with a convergence order of $O(|\mathcal{G}|^{1/2})$ in cases where volatility is uncontrolled. Furthermore, we provide a counterexample to demonstrate that no strong convergence occurs when volatility is controlled. Based on these results, we analyze the bias and variance of the policy evaluation and policy gradient estimators in various algorithms for continuous-time reinforcement learning caused by discrete sampling.

Yiqing Zhou*KU Leuven*

yiqing.zhou@kuleuven.be

KarstenNaert,DirkNuyens

Minimizing a function with limited sample points is challenging when function evaluations are costly. We propose a fast interpolation-based approach using the Fast Fourier Transform (FFT) to estimate the minimum more efficiently. By interpolating from a sparse set of samples, our method achieves high accuracy with significantly fewer function evaluations. Preliminary results demonstrate its effectiveness for smooth periodic functions.

Combining quasi-Monte Carlo with Stochastic Optimal Control for Trajectory Optimization of Autonomous Vehicles in Mine Counter Measure Simulations*Philippe Blondeel*

Belgian Royal Military Academy

Philippe.blondeel@mil.be

Coauthor(s): Filip Van Utterbeeck, Ben Lauwens

Modelling and simulating mine countermeasures (MCM) search missions performed by autonomous vehicles is a challenging endeavour. The goal of these simulations typically consists of calculating trajectories for autonomous vehicles in a designated zone such that the coverage (residual risk) of the zone is below a certain user-defined threshold. We have chosen to model and implement the MCM problem as a stochastic optimal control problem, see [1]. Mathematically, the MCM problem is defined as minimizing the total mission time needed to survey a designated zone Ω for a given residual risk of not detecting sea mines in the user-chosen square domain, i.e.,

$$\min T_f, \tag{6.2}$$

subjected to

$$\mathbb{E}[q(T_F)] := \int_{\Omega} e^{-\int_0^{T_F} \gamma(\mathbf{x}(\tau), \boldsymbol{\omega}) d\tau} \phi(\boldsymbol{\omega}) d\boldsymbol{\omega} \leq \text{Residual Risk}. \tag{6.3}$$

The output of our stochastic optimal control implementation consists of an optimal trajectory in the square domain for the autonomous vehicle. As shown in Eq. (6.3), the residual risk is mathematically represented as an expected value integral. In [2], we presented a novel

relaxation strategy for the computation of the residual MCM risk, used in our stochastic optimal control formulation. This novel relaxation strategy ensures that the residual risk obtained at the end of the optimisation run is below the maximally allowed user-requested residual risk. This was however not the case with our initial ‘naive’ implementation of the MCM problem. Our proposed relaxation strategy ensures that the user requested risk is satisfied by sequentially solving the stochastic optimal control problem with an ever increasing size of the domain. We combine this strategy with a quasi-Monte Carlo sampling scheme based on a Rank-1 Lattice rule for the computation of the expected value integral. We observe a speedup up to a factor two in terms of total computational cost in favour of quasi-Monte Carlo when compared to standard Monte Carlo.

- [1] Blondeel, P., Van Utterbeeck, F., Lauwens, B. (2024). Modeling sand ripples in mine countermeasure simulations by means of stochastic optimal control. In: The 19th European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS, Lisbon, Portugal (2024).
- [2] Blondeel, P., Van Utterbeeck, F., Lauwens, B. (2025). Application of quasi-Monte Carlo in Mine Countermeasure Simulations with a Stochastic Optimal Control Framework, *arXiv preprint*

Wed, Jul 30 11:00–11:30

A Monte Carlo Approach to Designing a Novel Sample Holder for Enhanced UV-Vis Spectroscopy

R. Persiani

INFN Section of Catania, Via S. Sofia 64, Catania, 95123, Italy

rino.persiani@ct.infn.it

Coauthor(s): A. Agugliaro, S. Albergo, R. De Angelis, I. Di Bari, A. Sciuto, A. Tricomi

UV-Vis spectroscopy is one of the most widely used techniques for identifying and quantifying substances in water and other solvents due to its speed and reliability. Its applications span diverse scientific fields, including chemistry, biochemistry, and medicine, as well as industrial sectors such as the pharmaceutical and food industries. Moreover, it plays a crucial role in environmental monitoring, particularly in assessing water quality. While alternative methods such as Raman and Ring Down spectroscopy have emerged, the core design of UV-Vis spectrometers has remained largely unchanged since their inception. Typically, these instruments employ a light source, a monochromator, a standard 1 cm cuvette as the sample holder, and one or more photosensors.

In this work, we present a novel approach that leverages Monte Carlo simulation to opti-

mize the design of the sample holder for enhanced UV-Vis spectroscopy. In particular, our setup adopts a pulsed light source and a Silicon Photomultiplier (SiPM) with single-photon counting capability. Using the optical transport package available in the Geant4 toolkit, we characterized and optimized the new design. The innovative holder, crafted in PTFE for its high UV reflectivity, resembles an integrating sphere, which increases the photon path length in the solution and thereby enhances absorbance in the presence of absorbing substances. We also present a comparison between experimental data and Monte Carlo predictions for validation. With this new sample holder, the spectrophotometer exhibits enhanced detection sensitivity, especially at low concentrations.

Wed, Jul 30 11:30–12:00

ARCANE Reweighting: A technique to tackle the sign problem in the simulation of collider events in high-energy physics

Prasanth Shyamsundar

Fermi National Accelerator Laboratory

prasanth@fnal.gov

Negatively weighted events, which appear in the Monte Carlo (MC) simulation of particle collisions, significantly increase the computational resource requirements of current and future collider experiments in high-energy physics. This work introduces an MC technique called ARCANE reweighting for reducing or eliminating negatively weighted events. The technique works by redistributing (via an additive reweighting) the contributions of different pathways within the simulator that lead to the same final event. The technique is exact and does not introduce any biases in the distributions of physical observables. ARCANE reweighting can be thought of as a variant of the parametrized control variates technique, with the added nuance that redistribution is performed using a deferred additive reweighting. The technique is demonstrated for the simulation of a specific collision process, namely $e^+e^- \rightarrow q\bar{q} + 1 \text{ jet}$. The technique can be extended to several other collision processes of interest as well. This talk is based on the Refs [1] and [2].

- [1] Shyamsundar, Prasanth (2025). *ARCANE Reweighting: A Monte Carlo Technique to Tackle the Negative Weights Problem in Collider Event Generation*. arXiv:2502.08052 [hep-ph].
- [2] Shyamsundar, Prasanth (2025). *A Demonstration of ARCANE Reweighting: Reducing the Sign Problem in the MC@NLO Generation of $e^+e^- \rightarrow q\bar{q} + 1 \text{ jet}$ Events*. arXiv:2502.08052 [hep-ph].

Multifidelity and Surrogate Modeling Approaches for Uncertainty Quantification in Ice Sheet Simulations

Nicole Aretz

Oden Institute for Computational Engineering and Sciences, University of Texas at Austin
nicole.aretz@austin.utexas.edu

Coauthor(s): Max Gunzburger, Mathieu Morlighem, Karen Willcox

Our work [1] uses multifidelity and surrogate modeling to achieve computationally tractable uncertainty quantification (UQ) for large-scale ice sheet simulations. UQ is of utmost importance to enable judicious policy decisions combating climate change. However, high-fidelity ice sheet models are typically too expensive computationally to permit Monte Carlo sampling. To reduce the computational cost while achieving the same target accuracy, we use multifidelity estimators to shift the computational burden onto less expensive surrogate models derived from coarser discretizations and approximated physics. In this talk, we compare three estimators — Multifidelity [2] and Multilevel [3] Monte Carlo, and the Best Linear Unbiased Estimator [4] — and present results for the expected ice mass loss of the Greenland ice sheet.

- [1] Aretz, N., Gunzburger, M., Morlighem, M., & Willcox, K. (2025). Multifidelity uncertainty quantification for ice sheet simulations. *Computational Geosciences*, 29(1), 1-22.
- [2] Peherstorfer, B., Willcox, K., & Gunzburger, M. (2016). Optimal model management for multifidelity Monte Carlo estimation. *SIAM Journal on Scientific Computing*, 38(5), A3163-A3194.
- [3] Giles, M. B. (2015). Multilevel Monte Carlo methods. *Acta Numerica*, 24, 259-328.
- [4] Schaden, D., & Ullmann, E. (2020). On multilevel best linear unbiased estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2), 601-635.

Empirical Statistical Comparative Analysis of SNP Heritability Estimators and Gradient Boosting Machines (GBM) Using Genetic Data from the UK Biobank

Kazeem Adeleke^{*1}

University of the West of England, UK

adedayo.adeleke@uwe.ac.uk

Coauthor(s): Peter Ogunyinka², Emmanuel Ologunleko³ and Dawud Agunbiade⁴

This study addresses the methodological challenges in estimating genetic heritability by comparing traditional statistical approaches with advanced machine learning techniques. We evaluated three distinct methods: sibling regression, LD-score regression, and Gradient Boosting Machines (GBMs), using both simulated datasets and real-world data from the UK Biobank. Our methodology involved generating simulated genotypes following Mendelian inheritance patterns and creating corresponding phenotypes incorporating family-specific genetic effect sizes. We conducted Genome-Wide Association Studies (GWAS) on firstborn children from each family and performed comprehensive heritability analyses using all three methods. Results demonstrated that while sibling regression effectively captured within-family genetic similarities and LD-score regression accounted for population-wide linkage disequilibrium patterns, GBMs showed superior capability in predicting phenotypes by capturing complex genetic interactions. The integration of GBMs with traditional methods revealed enhanced predictive power and provided new insights into the genetic architecture of complex traits. Our findings emphasize the value of combining conventional statistical approaches with machine learning techniques for more robust heritability estimation in large-scale UK Biobank studies.

Cheap permutation testing

Carles Domingo-Enrich

Microsoft Research New England

carlesd@microsoft.com

Coauthor(s): Raaz Dwivedi, Lester Mackey

Permutation tests are a popular choice for distinguishing distributions and testing independence, due to their exact, finite-sample control of false positives and their minimax optimality when paired with U-statistics. However, standard permutation tests are also expensive, requiring a test statistic to be computed hundreds or thousands of times to detect a separation between distributions. In this work, we offer a simple approach to accelerate testing: group your datapoints into bins and permute only those bins. For U and V-statistics, we prove that these cheap permutation tests have two remarkable properties. First, by storing appropriate sufficient statistics, a cheap test can be run in time comparable to evaluating a single test statistic. Second, cheap permutation power closely approximates standard permutation power. As a result, cheap tests inherit the exact false positive control and minimax optimality of standard permutation tests while running in a fraction of the time. We complement these findings with improved power guarantees for standard permutation testing and experiments demonstrating the benefits of cheap permutations over standard maximum mean discrepancy (MMD), Hilbert-Schmidt independence criterion (HSIC), random Fourier feature, Wilcoxon-Mann-Whitney, cross-MMD, and cross-HSIC tests.

Moving PCG beyond LCGs

Christopher Draper

Florida State University

chd16@fsu.edu

Coauthor(s): Michael Mascagni

PCG is a set of generators released by Melissa E. O’Neill in 2014 [1]. The original technical report outlined a number of lightweight scrambling techniques. Each scrambling technique offered some improvement to the quality of the linear congruential generators they were designed for. However the real strength of the scrambling techniques was that they could easily be combined in different combinations to offer much stronger improvements. The PCG technical report concludes with the creation of the PCG library, a popular PRNG library that implements a number of generators described in the technical report. Starting from the observation that the PCG work was narrowly focused on implementing their scrambling techniques for specific linear congruential generators, we explore the PCG scrambling techniques and their potential application for being applied to other PRNGs. We show the steps taken to generalize the PCG scrambling techniques to work with any arbitrary number of bits and parameter values. Then test the PCG scrambling techniques across different linear congruential generators and then test the PCG scrambling techniques across a number of different PRNGs.

- [1] Melissa E. O’Neill. 2014. PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation. Technical Report HMC-CS-2014-0905. Harvey Mudd College, Claremont, CA.

Hybrid least squares for learning functions from highly noisy data

Yiming Xu

University of Kentucky

yiming.xu@uky.edu

Coauthor(s): Ben Adcock, Bernhard Hientzsch, Akil Narayan

Motivated by the request for efficient estimation of conditional expectations, we consider a least-squares function approximation problem with heavily polluted data. In such scenarios,

existing methods based on the small noise assumption become suboptimal. We propose a hybrid approach that combines Christoffel sampling with optimal experimental design to address this issue. The proposed algorithm adheres to appropriate optimality criteria for both sample points generation and function evaluation, leading to improved computational efficiency and sample complexity. We also extend the algorithm to convex-constrained settings with similar theoretical guarantees. Moreover, when the target function is defined as the expectation of a random field, we introduce adaptive random subspaces to approximate the target function and establish results concerning its approximation capacity. Our findings are corroborated through numerical studies on synthetic data and a more challenging stochastic simulation problem in computational finance.