北京大学 PEKING UNIVERSITY

北京大学统计科学中心 Center for Statistical Science, Peking University

贝叶斯分析与人工智能前沿研讨会

WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

TIME

September 10-12, 2025

VENUE

111 Zhihua Building, Peking University

Organized by

Center for Statistical Science, Peking University

Co-organized by

School of Mathematical Sciences, Peking University
Guanghua School of Management, Peking University
Beijing International Center for Mathematical Research

ORGANIZING COMMITTEE

CO-CHAIRS

Cheng Zhang　Feng Li

MEMBERS

Fang Yao　Zhihua Zhang　Zhengyu Huang　Ke Deng

会议议程 MEETING AGENDA

10-Sep

Time	Speaker	Title	Host
9:00-9:20		Registration	
9:20-9:30		Opening Remarks: Fang Yao	Cheng Zhang
9:30-10:30	Plenary Talk Jun S. Liu	Variational Approximation for Statistical Modeling	Zhihua Zhang
10:30-11:00		Tea Break	
11:00-12:00	Plenary Talk Marc A. Suchard	Simple approximations for scalable Bayesian computing in phylogenetics	Yijuan Hu
12:00-13:30		Lunch[1]	
13:30-14:10	Yian Ma (online)	Reverse diffusion Monte Carlo	Jiajie Zhu
14:10-14:50	Yingzhen Li	Variational Uncertainty Decomposition for In-Context Learning	
14:50-15:10		Tea Break	
15:10-15:50	Minh-Ngoc Tran	Intrinsic natural gradient: definition, computation and applications	Feng Li
15:50-16:30	Matias Quiroz	When Bayes gets hard: Inference in doubly intractable models	
16:30-18:00		Brainstorming Discussion	

11-Sep

Time	Speaker	Title	Host
8:30-9:10	Jiajie Zhu	Gradient Flows in the Joint and Marginal Entropy-Regularized Transport Geometry for Sampling and Inference	Cheng Zhang
9:10-9:50	Chi Zhang (online)	Bayesian tip dating approach	
9:50-10:10		Tea Break	
10:10-10:50	Kamélia Daudel	Learning with Importance Weighted Variational Inference	Minh-Ngoc Tran
10:50-11:30	Alexandre Bouchard Côté (online)	Computing Normalization Constants on GPUs	

12-Sep

Time	Speaker	Title	Host
9:00-9:40	Xin Tong	Diffusion models for high dimensional distributions	Zhengyu Huang
9:40-10:20	Yuling Jiao	Provable Diffusion Posterior Sampling for Bayesian Inversion	
10:20-11:00	Harrison Bo Hua Zhu	Recurrent Memory for Online Interdomain Gaussian Processes	Yingzhen Li
11:00-12:00		Tea Break/Discussion	

1=Chinese Restaurant at Shao Yuan, Floor 1 of Shao Yuan Building 7(Buffet)

All talks in Dingshisun Lecture Hall (丁石孙教室), Floor 1 of Zhihua Building (智华楼), Peking University

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

大会报告 PLENARY TALK

Jun S. Liu

Professor

Tsinghua university

Title: Variational Approximation for Statistical Modeling

Abstract: We use three statistical modeling examples to illustrate how to use variational methods to handle statistical estimation and inference tasks involving complex latent structures. In some settings, we can show that variational approximations can also yield good approximations to the marginal posterior distributions of the targeted parameters. The three examples are: linear regression with group variables; measurement error model analysis, and multivariate stochastic volatility models.

Bio: Dr. Jun Liu is Chair Professor at Tsinghua University and Professor Emeritus at Harvard University. He received his BS in mathematics from Peking University (1985) and Ph.D. in statistics from the University of Chicago (1991). He held faculty positions at Stanford University (1994–2004) and Harvard University (2000–2025). Liu has received the COPSS Presidents' Award (2002), Morningside Gold Medal in Applied Mathematics (2010), Pao-Lu Hsu Award (2017), and Jerome Sacks Award (2017). He has been a Medallion Lecturer (IMS, 2002), Bernoulli Lecturer (2004), and elected Fellow of IMS (2004), ASA (2005), and ISCB (2022). He was elected to the U.S. National Academy of Sciences in 2025.Liu's research spans Bayesian statistics, Monte Carlo methods, statistical machine learning, stochastic processes, and computational biology. He introduced statistical missing data formulations and Gibbs sampling strategies for biological sequence analyses, pioneered sequential Monte Carlo (SMC) and novel MCMC techniques, and contributed to high-dimensional Bayesian modeling and sufficient dimension reduction. He has co-authored over 300 publications with more than 90,000 citations, served as co-editor of JASA (2011–2014), and mentored 40 PhD students and 32 postdoctoral fellows.Liu's current interests include Bayesian methodologies, bioinformatics, statistical machine learning, Monte Carlo methods, and applications of artificial intelligence.

Marc A. Suchard

Professor

University of California Los Angeles

Title: Simple approximations for scalable Bayesian computing in phylogenetics

Abstract: Bayesian computation remains onerous at scale for inference under many discrete-valued stochastic process-based models, while these statistical models remain ubiquitous across biology and public health. In this talk, I will explore how one can construct reliable approximations for continuous-time Markov chain (CTMC) models. CTMCs underpin the most popular models for learning about how rapidly evolving pathogens change over time and space to give rise to human infection, and the dimensionality of these problems are daunting. These approximations enable the introduction of novel random-effects and non-parameteric CTMC models that capture biological realism previously missing. Applied to the analysis of yellow fever and Ebola viruses, these models remove bias in inference of the factors driving transmission, while the statistical machinery is over an order of magnitude more time efficient than conventional approaches.

Bio: Dr. Suchard is Professor in the Departments of Biostatistics and of Computational Medicine at the University of California, Los Angeles (UCLA). His research focuses on large-scale inference of stochastic processes in genomics and observational healthcare. He received his BA in biophyics from the University of California, Berkeley and MD and PhD in biomathematics from UCLA. He has authored over 350 peer-reviewed publications, is the 2003 recipient of the Savage Award, an international prize for the best Bayesian dissertation, received the 2006 and 2011 Mitchell Prizes, the premier award in applied Bayesian statistics, a 2008 Guggenheim Fellowship to promote Bayesian approaches in medicine and Research Gifts from Microsoft Corporation and Google to further statistical computing. He is also the recipient of the 2013 Committee of Presidents of Statistical Societies (COPSS) Presidents' Award for outstanding contributions to the statistics profession by a person aged 40 or under and the 2021 Jerome Sacks Award for Outstanding Cross-Disciplinary Research. He holds fellowships in the American Statistical Association, the Institute of Mathematical Statistics and the International Society for Bayesian Analysis.

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

邀请报告 INVITED TALK

Yian Ma

Assistant Professor

University of California, San Diego（UCSD）

Title: Reverse diffusion Monte Carlo

Abstract: I will introduce a novel Monte Carlo sampling approach that uses the reverse diffusion process. In particular, the intermediary updates—the score functions—can be explicitly estimated to arbitrary accuracy, leading to an unbiased Bayesian inference algorithm. I will then discuss how to use this idea to improve sampling in the diffusion models via reverse transition kernels.

Bio: Yian Ma is an assistant professor at the Halıcığolu Data Science Institute, UC San Diego, where he serves as the vice chair in charge of the graduate programs. Prior to UCSD, he spent a year as a visiting faculty at Google Research. Before that, he was a post-doctoral fellow at UC Berkeley, hosted by Mike Jordan. Yian completed his Ph.D. at University of Washington. His current research primarily revolves around scalable inference methods for credible machine learning, with application to time series data and sequential decision-making tasks. He has received the Facebook research award, the Stein fellowship, and the best paper awards at the Neurips and ICML workshops.

Yingzhen Li

Associate Professor

Imperial College London

Title: Variational Uncertainty Decomposition for In-Context Learning

Abstract: As large language models (LLMs) gain popularity in conducting prediction tasks in-context, understanding the sources of uncertainty in in-context learning becomes essential to ensuring reliability. The recent hypothesis of in-context learning performing predictive Bayesian inference opens the avenue for Bayesian uncertainty estimation, particularly for de- composing uncertainty into epistemic uncertainty due to lack of in-context data and aleatoric uncertainty inherent in the in-context prediction task. However, the decomposition idea re- mains under-explored due to the intractability of the latent parameter posterior from the underlying Bayesian model. In this work, we introduce a variational uncertainty decomposi- tion framework for in-context learning without explicitly sampling from the latent parameter posterior, by optimising auxiliary inputs as probes to obtain an upper bound to the aleatoric uncertainty of an LLM's in-context learning procedure. Through experiments on synthetic and real-world tasks, we show quantitatively and qualitatively that the decomposed uncer- tainties obtained from our method exhibit desirable properties of epistemic and aleatoric uncertainty.

Bio: Yingzhen Li is an Associate Professor in Machine Learning at the Department of Computing, Imperial College London, UK. Before that she was a senior researcher at Microsoft Research Cambridge, and previously she has interned at Disney Research. She received her PhD in engineering from the University of Cambridge, UK. Yingzhen is passionate about building reliable machine learning systems, and her approach combines both Bayesian statistics and deep learning. She has worked extensively on approximate inference methods with applications to Bayesian deep learning and deep generative models, and her work has been applied in industrial systems and implemented in deep learning frameworks (e.g. Tensorflow Probability and Pyro). She regularly gives tutorials and lectures on probabilistic ML and generative models at machine learning research summer schools, as well as invited tutorials on Advances in Approximate Inference at NeurIPS 2020 and UAI 2025. She was a co-organiser of the Advances in Approximate Bayesian Inference (AABI) symposium in 2020-2023, as well as many NeurIPS/ICML/ICLR workshops on topics related to probabilistic learning. She is a Program Chair for AISTATS 2024 and a General Chair for AISTATS 2025 and 2026. Her work on Bayesian ML has also been recognised in AAAI 2023 New Faculty Highlights.

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

Minh-Ngoc Tran

Associate Professor

The University of Sydney

Title: Intrinsic natural gradient: definition, computation and applications

Abstract: The Euclidean natural gradient method is a widely used tool in statistical optimization. We extend the concept of natural gradient to a Riemannian manifold and develop a computationally efficient method for computing it. The method is fully intrinsic, avoiding reliance on an ambient Euclidean space. We apply the method to the Variational Bayes and Maximum Likelihood Estimation problems, where the parameters lie on a manifold. We establish theoretical convergence guarantees comparable to those of standard stochastic gradient descent, and validate the robustness and efficiency of our approach across a range of learning tasks. Joint work with Dario Draca.

Bio: Minh Ngoc Tran is an Associate Professor of Business Analytics at the University of Sydney Business School. He earned his BSc and MSc in Mathematics from Vietnam National University, Hanoi, before completing a PhD in Statistics at the National University of Singapore in 2012.

His research spans statistical methodology and applied statistics. On the methodological side, he develops Bayesian computation and machine learning techniques, with a particular emphasis on variational Bayes and the integration of emerging quantum computation methods into data analysis. In applied work, he leverages modern statistical approaches to advance studies in cognitive science, consumer behaviour and financial econometrics.

Minh Ngoc's contributions have been recognised nationally and internationally, with publications in leading statistical journals and conferences. His research has attracted over A$5 million in funding, including three competitive ARC grants, and he is a sought after speaker at both national and international meetings. As an enthusiastic educator, Minh Ngoc adopts a research led, student focused teaching style that consistently earns outstanding feedback.

Matias Quiroz

Senior Lecturer
University of Technology Sydney

Title: When Bayes gets hard: Inference in doubly intractable models

Abstract: Bayes gets hard when the likelihood includes an intractable normalising factor, giving rise to doubly intractable models. We propose a signed pseudo-marginal Metropolis–Hastings algorithm with an unbiased block-Poisson likelihood estimator, which handles negativity via importance sampling and comes with finite-sample guarantees. The estimator supports vectorisation and parallelisation, is compatible with correlated pseudo-marginal strategies, and admits heuristic guidelines for tuning its hyperparameters. We demonstrate its superior performance on the Ising model for spatial data and the Kent model for directional data.

Bio: Matias Quiroz is a Senior Lecturer in Statistics at the University of Technology Sydney, a position he has held since 2023. He previously held appointments at Stockholm University and the Central Bank of Sweden. He serves as Associate Editor for the journals Computational Statistics and Data Analysis and Econometrics and Statistics. His research focuses on Bayesian computation and statistical machine learning, with publications in leading journals including Journal of the American Statistical Association, Journal of Computational and Graphical Statistics, and Journal of Machine Learning Research, as well as top conference proceedings such as ICML and AISTATS.

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

Jiajie Zhu

Associate Professor
KTH Royal Institute of Technology

Title: Gradient Flows in the Joint and Marginal Entropy-Regularized Transport Geometry for Sampling and Inference

Abstract: Many problems in machine learning and statistics can be cast as abstract flows of probability measures, such as Otto's Wasserstein gradient flows. Such problems appear in sampling, variational inference, generative modeling, etc. In this talk, I will focus on the computational aspects of gradient flows, building upon the mathematical foundation of optimal transport and PDE analysis, specifically on entropy-regularized transport for gradient flows. Historically, there are two ways of regularizing transport using entropy: joint and marginal. The former gave rise to the Sinkhorn iteration while the latter generates the Wasserstein-Fisher-Rao (a.k.a. Hellinger-Kantorovich) distance. In this talk, I will present some concrete computational algorithms motivated by the analysis of entropy-regularized transport gradient flows, with applications to sampling, inference, and optimization.

Bio: Jia-Jie Zhu (https://jj-zhu.github.io/) is an applied mathematician, machine learner. He heads a research group at the Weierstrass Institute, Berlin, Germany, and is an incoming tenured associate professor in mathematics at the KTH Royal Institute of Technology in Stockholm, Sweden. Previously, he worked as a postdoctoral researcher in machine learning at the Max-Planck-Institute for Intelligent Systems, Tübingen. He received his Ph.D. training in optimization at the University of Florida, USA, and a B.Sc. degree in mathematics from Fudan University, China. He is interested in the intersection of machine learning, analysis, and optimization, on topics such as (PDE) gradient flows of probability measures, optimal transport, and robustness of learning and optimization algorithms.

Chi Zhang

Professor

Institute of Vertebrate Paleontology and Paleoanthropology, Chinese Academy of Sciences

Title: Bayesian tip dating approach

Abstract: Dating the tree of life is fundamental for understanding the evolutionary process and co-evolution of life and environment. Traditional stepwise approach does not take full use of the data and fails to account for uncertainties from various sources. The recently developed Bayesian total-evidence tip-dating approach has the advantage of combining morphological characters and geological ages from fossils and morphological and molecular data from extant taxa, and properly incorporates uncertainties from data and parameters in the statistical inference. The method utilizes the fossilized birth-death process for the timetree, the relaxed clock model for the evolutionary rates, and the Markov process for the character changes. This tip-dating approach has been productively applied to various biological groups to study their evolution, including birds, pterosaurs, and mammals.

Bio: Chi Zhang is a research professor in the Institute of Vertebrate Paleontology and Paleoanthropology, Chinese Academy of Sciences. His research interest is applying Bayesian statistics in paleontology and molecular evolution to infer evolutionary histories (e.g., divergence times, evolutionary rates and diversification patterns), and evaluating the properties and performance of the methods using biological data and simulations. He develops phylogenetic methods in the software packages MrBayes and BEAST2 and studies the microevolution of several vertebrate groups.

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

Kamélia Daudel

Assistant Professor

ESSEC Business School

Title: Learning with Importance Weighted Variational Inference

Abstract: Several popular variational bounds involving importance weighting ideas have been proposed to generalize and improve on the Evidence Lower BOund (ELBO) in the context of maximum likelihood optimization, such as the Importance Weighted Auto-Encoder (IWAE) and the Variational Rényi (VR) bounds. The methodology to learn the parameters of interest using these bounds typically amounts to running gradient-based variational inference algorithms that incorporate the reparameterization trick. However, the way the choice of the variational bound impacts the outcome of variational inference algorithms can be unclear. In this talk, we will present and motivate the VR-IWAE bound, a novel variational bound that unifies the ELBO, IWAE and VR bounds methodologies. In particular, we will provide asymptotic analyses for the VR-IWAE bound and its reparametrized gradient estimator, which reveal the advantages and limitations of the VR-IWAE bound methodology while enabling us to compare of the ELBO, IWAE and VR bounds methodologies. Our work advances the understanding of importance weighted variational inference methods and we will illustrate our theoretical findings empirically.

Bio: Dr. Daudel is an Assistant Professor of Statistics at ESSEC Business School. Prior to that, she was a postdoctoral researcher in the Department of Statistics at the University of Oxford working with Arnaud Doucet. Her research lies in the field of Approximate Inference. In particular, she is interested in Variational Inference methods which go beyond the commonly-used parametric variational distribution framework and which involve flexible variational bounds. Dr. Daudel received the first prize of IP Paris Best Thesis Award 2022 for her thesis.

Alexandre Bouchard Côté

Professor

University of British Columbia

Title: Computing Normalization Constants on GPUs

Abstract: Over the years, several algorithms have been developed to tackle normalization constant estimation. A handful of those have passed the test of time thanks to their capacity to beat the curse of dimensionality in many realistic scenarios: I will focus here on Annealed Importance Sampling (AIS) and Sequential Monte Carlo (SMC) methods. I will describe new theoretical analysis of these algorithms, and methodological developments supported by this analysis. Of particular interest is a novel approach to GPU deployment.

Bio: Alexandre Bouchard is a professor of statistics at the University of British Columbia. He received his PhD in computer science from the University of California, Berkeley. His research focuses on computational Bayesian methods and applications in cancer genomics and phylogenetics. https://www.stat.ubc.ca/~bouchard/index.html

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

Xin Tong

Associate Professor

National University of Singapore

Title: Diffusion models for high dimensional distributions

Abstract: Diffusion model is a popular tool to generate new data samples. However, rigorous understanding of diffusion model is still lacking. One issue is how to train these models for high dimensional problems as score function estimation is subject to the curse of dimension. Another issue is how to avoid the memorization effect, where the diffusion model is bound to generate an exact copy from the training data. We will provide solutions to the first issue by focusing on high dimensional distributions with sparse dependence. We will leverage the sparse dependence to provide a local estimation of the score functions. As for the second issue, we will modify the diffusion model in the final stage and generate new samples close to the same manifold where the training data is originated.

Bio: Xin Tong is an associate professor in the Department of Mathematics at the National University of Singapore. He earned his Ph.D. from Princeton University in 2013. Prior to his current role, he was a postdoctoral researcher at the Courant Institute of New York University. As an applied mathematician, his research spans uncertainty quantification, machine learning, and operations research. His primary expertise lies in the development and analysis of stochastic algorithms.

Yuling Jiao

Professor
Wuhan University

Title: Provable Diffusion Posterior Sampling for Bayesian Inversion

Abstract: Diffusion model has been proved to be a powerfull tool for Bayeian inverse problem. Instead of approximate the time-dependent likelihood adhockly in the literature, we directly calculate the posterior score utilizing a restricted Gaussian oracle via Langevin Monte Carlo dynamics. The resulted posterior score can then be used to generate samples by solving the inverse diffusion process. We bound the error of the estimated posteroir distribution and underlying truth in TV distance. We numerically demonstrate the effectiveness of our method in solving various image inverse problem, including linear and nonlinear deblur problems. The results reveal a better performance in PSNR and SSIM of our method, compared with both traditional variational method and the recent diffusion posterior sampling method.

Bio: Yuling Jiao is a Professor and PhD Supervisor serving as the Associate Dean at the School of Artificial Intelligence, Wuhan University. He has been selected for the National Young Talents Program. His primary research focuses on machine learning and scientific computing, with recent emphasis on the mathematical foundations of deep learning. He has published over thirty papers in flagship journals and conferences across computational mathematics, applied mathematics, statistics, electronic engineering, and artificial intelligence, including:SIAM journals (5 papers)、Appl.Comput. Harmon. Anal.(2 papers)、Inverse Probl. (2 papers)、Ann. Stat.(3 papers)、J.Amer. Statist. Assoc. (2 papers)、IEEE Trans. Inf. Theory (5 papers)、IEEE Trans. Signal Process. (3 papers)、J. Mach. Learn. Res. (7 papers)、ICML (3 papers)、NeurIPS(3 papers, including one Oral and one Spotlight)、Nat. Commun. He has led several research projects, including a sub-task of the National Key R&D Program of China, an NSFC General Program project, and a series of industry collaboration projects with Huawei.

贝叶斯分析与人工智能前沿研讨会
WORKSHOP ON BAYESIAN ANALYSIS AND ARTIFICIAL INTELLIGENCE

Harrison Bo Hua Zhu

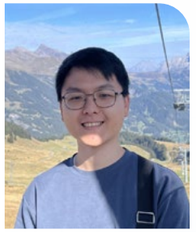

Assistant Professor
University of Copenhagen

Title: Recurrent Memory for Online Interdomain Gaussian Processes

Abstract: We propose a novel online Gaussian process (GP) model that is capable of capturing long-term memory in sequential data in an online learning setting. Our model, Online HiPPO Sparse Variational Gaussian Process (OHSVGP), leverages the HiPPO (High-order Polynomial Projection Operators) framework, which is popularized in the RNN domain due to its long-range memory modeling capabilities. We interpret the HiPPO time-varying orthogonal projections as inducing variables with time-dependent orthogonal polynomial basis functions, which allows the SVGP inducing variables to memorize the process history. We show that the HiPPO framework fits naturally into the interdomain GP framework and demonstrate that the kernel matrices can also be updated online in a recurrence form based on the ODE evolution of HiPPO. We evaluate OHSVGP with online prediction for 1D time series, continual learning in discriminative GP model for data with multidimensional inputs, and deep generative modeling with sparse Gaussian process variational autoencoder, showing that it outperforms existing online GP methods in terms of predictive performance, long-term memory preservation, and computational efficiency.

Bio: Harrison Bo Hua Zhu is an assistant professor in the Health Data Science and AI Section at the University of Copenhagen. Previously, he did his PhD at Imperial College London under the supervision of Professor Seth Flaxman and Professor Yingzhen Li. Harrison's research focuses on Bayesian machine learning methodology with applications to public health, particularly infectious diseases.

PEKING UNIVERSITY MAP

Red Star= Zhihua building (where the conference will be held).
Note: please enter through **the west gate of Zhihua building** (all other gates require face scanning).

Blue star= The Lakeview Hotel Beijing

Red line=How you go from your hotel to Zhihua Building.

Green line=How you go from Zhihua Building to Chinese Restaurant at Shao Yuan.