DATA-INSPIRE Workshop:

Monte Carlo, Dynamic Systems and Robotics

Sponsored by the TRIPODS DATA-INSPIRE Institute, a joint collaboration of DIMACS and the Rutgers Departments of Computer Science, Mathematics, and Statistics
http://robotics.cs.rutgers.edu/data-inspire/

Date: *Friday, April 23, 2021,*

10:00 pm – 12:00 pm (Eastern Time)

Zoom Link:
https://rutgers.zoom.us/j/95568605240?pwd=TGIwMjNrbUc3eExXazB0dDY3emlXZz09

Program

10:00-10:50: Speaker: ***Arnaud Doucet*** (Oxford University)

Title: Differentiable Particle Filtering via Entropy-Regularized Optimal Transport with Applications in Robot Localization

10:50-11:00: Q&A

11:00-11:50: Speaker: ***Jun S. Liu*** (Harvard University)

Title: Optimal Resampling for Sequential Monte Carlo

11:50-12:00: Q&A

Abstracts

Arnaud Doucet (Oxford University)

Title: Differentiable Particle Filtering via Entropy-Regularized Optimal Transport with Applications in Robot Localization

Abstract: Particle Filtering (PF) methods are an established class of procedures for performing inference in non-linear state-space models. Resampling is a key ingredient of PF, necessary to obtain low variance likelihood and states estimates. However, traditional resampling methods result in PF-based loss functions being non-differentiable with respect to model and PF parameters. In a variational inference context, resampling also yields high variance gradient estimates of the PF-based evidence lower bound. By leveraging optimal transport ideas, we introduce a principled differentiable particle filter and provide convergence results. We demonstrate this novel method on a variety of applications including robot localization.

Joint work with Adrien Corenflos, James Thornton and George Deligiannidis

Bio: Arnaud Doucet obtained his PhD degree in 1997 from University Paris-Sud Orsay. Ever since he has held faculty positions at Melbourne University, Cambridge University, the University of British Columbia and the Institute of Statistical Mathematics. He joined the department of Statistics of Oxford University in 2011 where he is currently Professor and also holds a part-time appointment as Senior Research Scientist at Google DeepMind since 2019. He was an IMS Medallion Lecturer in 2016, was elected an IMS Fellow in 2017 and was awarded the Guy Medal in Silver from the Royal Statistical Society in 2020.

Jun S. Liu (Harvard University)

Title: Optimal Resampling for Sequential Monte Carlo

Abstract: Sequential Monte Carlo algorithms have been widely accepted as a powerful computational tool for making inference with dynamical systems. A key step in sequential Monte Carlo is resampling, which plays a role of steering the algorithm towards the future dynamics. Several strategies have been used in practice, including multinomial resampling, residual resampling, optimal resampling, stratified resampling, and optimal transport resampling. We will review some of these approaches and show that in one-dimensional cases the optimal transport resampling is equivalent to stratified resampling on the sorted particles, and they both minimize the resampling variance as well as the expected squared energy distance between the original and resampled empirical distributions. In general d-dimensional cases, if the particles are first sorted using the Hilbert curve, we show that the variance of stratified resampling is $O\left(m^{-\left(1+2d^{-1}\right)}\right)$ with m being the number of resampled particles and that this is optimal for ordered stratified resampling , as conjectured in Gerber et al. (2019). In light of these results, we show that, for dimension $d > 1$, the mean square error of sequential quasi-Monte Carlo with n particles can be $O\left(n^{-1-\frac{4}{d(d+4)}}\right)$ if Hilbert curve resampling is used and a specific low-discrepancy set is chosen. To our knowledge, this is the first known convergence rate lower than $o(n^{-1})$. The presentation is based on the joint work with Wenshuo Wang, Yichao Li, and Ke Deng.

Jun Liu is Professor of Statistics at Harvard University. Dr. Liu received his BS degree in mathematics in 1985 from Peking University and Ph.D. in statistics in 1991 from the University of Chicago. Dr. Liu received the NSF CAREER Award in 1995 and the Mitchell Award in 2000. He was selected as a Medallion Lecturer in 2002, a Bernoulli Lecturer in 2004, a Kuwait Lecturer of Cambridge University in 2008; and elected to Fellow of the IMS and ASA. In 2002, he won the prestigious COPSS Presidents' Award. In 2010, he was awarded the Morningside Gold Medal in Applied Mathematics. He was honored with the Outstanding Achievement Award in 2012, and the Pao-Lu Hsu Award in 2016 by the International Chinese

Dr. Liu and his collaborators introduced the statistical missing data formulation and Gibbs sampling strategies for biological sequence motif analysis in early 1990s. The resulting algorithms for protein sequence alignments, gene regulation analyses, and genetic studies have been adopted by many researchers as standard computational biology tools. Dr. Liu has made fundamental contributions to statistical computing and Bayesian modeling. He pioneered sequential Monte Carlo (SMC) methods and invented novel Markov chain Monte Carlo (MCMC) techniques. His theoretical and methodological studies on SMC and MCMC algorithms have had a broad impact in many areas. Dr. Liu has also pioneered novel Bayesian modeling techniques for discovering nonlinear and interactive effects in high-dimensional data, and led the developments of theory and methods for sufficient dimension reduction in high-dimensions. Dr. Liu has served on numerous government's grant review panels and editorial boards of leading statistical journals, including the co-editorship of JASA from 2011-2014. Dr. Liu has published more than 250 research articles in leading scientific journals, mentioned 30 postdoctoral fellows, and graduated more than 30 PhD students.