Stochastic simulation and machine learning

Henrik Hult, Pierre Nyquist

This project aims at analysing and developing methods at the intersection of *stochastic simulation* and *machine learning*. Stochastic simulation has long been a fundamental tool in essentially all areas where stochastic models are used; examples include e.g. physics, computational chemistry and biology, finance and insurance, computer science, Bayesian statistics, and machine learning and artificial intelligence. Similarly, machine learning methods are becoming increasingly important tools in many areas of science, including mathematics. In this project we want to explore different aspects of the interplay between the two areas. Topics of interest include, but are not limited to:

I. Efficient simulation methods for statistical physics models.
II. Stochastic simulation methods in machine learning.
III. Machine learning methods for designing efficient Monte Carlo methods.

The focus of the project will be tailored to the interests of the postdoc. The project will be part of a collaboration with Prof. S. Juneja at TIFR Mumbai, funded by the Swedish Research Council, and can involve research visits to TIFR.

I. Efficient simulation for statistical physics models. For models from statistical physics, one is typically interested in computing quantities related to macroscopic properties from the microscopic description provided by the model. This often amounts to computing expectations with respect to the underlying distribution or differences in the associated free energy. The standard approach is to use MCMC: construct ergodic Markov processes with the correct stationary distribution and use such processes for sampling purposes. Standard methods often perform poorly, for example due to metastability, and we are interested in exploring *extended ensemble methods*—methods in which the constructed Markov process lives in an extended state space. The design of such methods is delicate and our aim is to give rigorous performance guarantees to both new and existing methods.

II. Stochastic simulation methods in machine learning. A fundamental aspect of most machine learning methods is the need for efficient stochastic simulation methods, often in the form of MCMC methods. The relevant machine learning methods can often be viewed as statistical physics models and the development of new simulation methods as described in (I), is therefore an important part of making more advanced methods feasible from a computational standpoint.

Similarly, the optimisation step in the training of machine learning methods often relies on methods for stochastic simulation, most notably stochastic gradient descent and its offsprings. Despite the enormous empirical success, a lot of fundamental questions remain open regarding the properties of different variations—e.g. batch-norm, Langevin dynamics, etc.—of standard stochastic approximation techniques. We are interested in addressing such questions using tools from stochastic control, large deviations, partial differential equations and optimal control.

III. Machine learning methods for designing efficient Monte Carlo methods. From work on rare-event simulation, it is well-known that efficient Monte Carlo methods are intrinsically tied to solutions of partial differential equations of Hamilton-Jacobi-Bellman (HJB) form. The design of efficient sampling methods can therefore be cast as the problem of finding solutions to such PDEs. However in most instances this is a difficult task in itself, especially in high-dimensional settings. Recent developments in using deep neural networks for solving HJB equations opens up a new avenue for designing rare-event simulation methods. In addition to showcasing the feasibility of this approach, topics of interest include the ability of neural networks to approximate *viscosity subsolutions* of HJB equations, error analysis, and how different approaches for obtaining such approximations affect performance and convergence properties of the corresponding Monte Carlo methods.