Gradient Flows for Sampling: Mean-Field Models, Gaussian Approximations and Affine Invariance

Andrew Stuart, Caltech

Sampling a probability distribution with an unknown normalization constant is a fundamental problem in computational science and engineering. This task may be cast as an optimization problem over all probability measures, by choice of a suitable energy function. Then an initial distribution can be evolved to the desired minimizer (the target distribution) via a gradient flow with respect to a chosen metric. The choice of the energy and the metric lead to different approaches and it is of interest to understand their role. We provide theoretical insights into these choices.

Having chosen an energy and a metric, development of an actionable algorithm requires approximation of the gradient flow. Mean-field models, whose law is governed by the gradient flow in the space of probability measures, may be identified; particle approximations of these mean-field models form the basis of algorithms. The gradient flow approach is also the basis of algorithms for variational inference, in which the optimization is performed over a parameterized family of probability distributions such as Gaussians or Gaussian mixtures; the underlying gradient flow is restricted to the parameterized family. Numerical results are presented to illustrate the resulting methodologies.

Joint work with Y. Chen (NYU), D.Z. Huang (PKU), J. Huang (U Penn) and S. Reich (Potsdam).