On continuous-time Bayesian inference and the geometry of the probability manifold

Aimee Maurais
Massachusetts Institute of Technology, United States of America
maurais@mit.edu

Inference in the Bayesian setting can be viewed as the problem of transforming a prior probability measure into a posterior measure. This transformation is frequently performed in "one shot" by applying a single update to an empirical or parametric approximation of the prior (e.g., Kalman and ensemble Kalman transforms, or more general techniques based on measure transport). Yet the prior-to-posterior update can also be viewed as a continuous transformation, governed by some dynamics on state space indexed by an artificial time. There are infinitely many choices of such dynamics (both deterministic or stochastic), with either finite or infinite time horizons, and any choice is associated with a transport equation encoding the particular path of probability measures taken between prior and posterior. In computational schemes used to realize these continuous-time transformations, a representation of the prior is initialized at time zero and the dynamics are simulated until a stopping time is reached, at which point the resulting probability distribution should approximate the posterior, if not realize it exactly. Computational simulation raises further questions linked to the choice of dynamics: how to compute a "step" given available information, how to choose step sizes, and how to determine stopping times for dynamics with infinite time horizons. Yet it is not well understood how the underlying choice of dynamics influences our ability to realize complex prior-to-posterior updates efficiently. In this work we elucidate connections among various frameworks which have been proposed for continuous-time Bayesian inference, and how design choices therein interact with the geometry of the probability manifold to influence performance.

Joint work with Youssef Marzouk (Massachusetts Institute of Technology).