Geophysical Research Abstracts
Vol. 21, EGU2019-5192, 2019
EGU General Assembly 2019

Combining variational data assimilation and particle filters: the variational mapping particle filter

Manuel Pulido (1,2) and Peter Jan van Leeuwen (1,3)

(1) University of Reading, Department of Meteorology, UK (m.a.pulido@reading.ac.uk), (2) Universidad Nacional del Nordeste, Physics, Corrientes, Argentina, (3) Colorado State University, US

Recent works in the machine learning community have started to combine two classical statistical concepts: Monte Carlo sampling and variational inference. In the traditional variational inference, including variational data assimilation, some parameters of a proposed posterior density are estimated through maximazing the marginal likelihood or via maximum a posteriori estimation. The idea for combining them is to use as optimization parameters in variational inference the Monte Carlo sample points, i.e. the particles. In this way, we seek for a set of particles that best represent the posterior density by, for instance, maximazing the marginal likelihood. Following this idea, we introduce a novel particle filter that is based on local optimal transport principles. The minimization of the Kullback-Leibler divergence between the intermediate density represented by the particles and the posterior density is conducted by a sequence of local maps. The transformations are embedded in a reproducing kernel Hilbert space which defines the steepest descent directions. The particles are required to follow these directions in order to minimize the Kullback-Leibler divergence or equivalently maximize the marginal likelihood. The optimization can be interpreted as a flow in which the particles are active tracers, they are moved along the steepest descent directions, i.e. streamlines of the flow, but at the same time they define the flow. Evaluation of the variational mapping particle filter will be shown with experiments using a 1000-variables Lorenz-96 system and a 1.5 layer quasi-geostrophic model with a resolution of 256x256 in which a preconditioning step in the variational optimization based on 3DVar schemes is introduced. We show that the deterministic mappings avoid the resampling step, i.e. the number of effective particles remains close to the total number of particles even for long recursive implementations. This variational particle filter framework inherits the well-known convergence properties and efficient implementations of optimization algorithms in high-dimensional state spaces. It sheds some light on the representation of sequential Monte Carlo methods in high-dimensional state-spaces. The challenges for further developments and implementation f the variational mapping particle filter will be discussed.