NORMALISING FLOW-BASED DIFFERENTIABLE PARTICLE FILTERS

YUNPENG LI

There has been a surge of interest in incorporating neural networks into particle filters, e.g. differentiable particle filters, to perform joint sequential state estimation and model learning for non-linear non-Gaussian state-space models in complex environments. Existing approaches primarily use vanilla neural networks which do not allow density estimation. As a result, they are often restricted to a bootstrap particle filtering framework or employ predefined distribution families (e.g. Gaussian distributions), limiting their performance in more complex real-world scenarios.
In this talk I will introduce a differentiable particle filtering framework that uses conditional normalising flows to build its dynamic model, proposal distribution, and measurement model. This approach not only enables valid probability densities but also allows the proposed method to adaptively learn these modules in a flexible way, without being restricted to predefined distribution families. I will discuss the theoretical properties of the proposed filters and present their performance through a series of numerical experiments.