

## The horseshoe prior for edge-preserving Bayesian inversion

**Speaker:** Yiqiu Dong, Technical University of Denmark

**Abstract:** In many large-scale inverse problems characterization of sharp edges in the solution is desired. In the Bayesian approach to inverse problems, edge-preservation is often achieved using Markov random field priors based on heavy-tailed distributions. Another strategy, popular in sparse statistics, is the application of hierarchical shrinkage priors. An advantage of this formulation lies in expressing the prior as a conditionally Gaussian distribution depending on heavy-tailed distributed hyperparameters. In this presentation, we revisit the shrinkage horseshoe prior and introduce its formulation for edge-preserving settings. We discuss a Gibbs sampling framework to solve the Bayesian inverse problem. Applications from imaging science show that our computational procedure is able to compute sharp edge-preserving posterior point estimates with reduced uncertainty.