
Alessio Spantini
**Optimal low-rank approximations of Bayesian linear
inverse problems**

888 MASSACHUSETTS Avenue Apt 318

Cambridge

02139 (MA)

USA

`spantini@mit.edu`

Solonen Antti

Cui Tiangang

James Martin

Luis Tenorio

Youssef Marzouk

In this talk we present statistically optimal and computationally efficient dimensionality reduction techniques for large-scale linear Gaussian inverse problems. These approximations of the Gaussian posterior distribution are at the heart of many state-of-the-art algorithms for nonlinear Bayesian inference. In particular, we study structure-exploiting approximations of the posterior covariance matrix as a low-rank update of the prior covariance matrix and prove optimality of the low-rank update for various metrics. These approximations are particularly useful when the data are informative relative to the prior only about a low-dimensional subspace of the parameter space. We also propose fast optimal approximations of the posterior mean that are particularly useful when repeated posterior mean evaluations are required for multiple sets of data (e.g., online inference). We extend these optimal approximations to the important case of goal-oriented inference where the quantity of interest (QoI) is a linear function of the inference parameters. We show that the posterior distribution of the QoI can be computed avoiding the explicit characterization of the full posterior distribution of the parameters. In particular, we focus on directions in the parameter space that are informed by the data, relative to the prior, and that are relevant to the QoI.