Balanced Truncation for Bayesian Inference

E. Qian[1], **J. M. Tabeart**[2], **C. Beattie**[3], **S. Gugercin**[3], **J. Jiang**[4], **P. Kramer**[5], **and A. Narayan**[6]

[1] *California Institute of Technology*
[2] *University of Edinburgh*
[3] *Virginia Tech*
[4] *ShanghaiTech University*
[5] *Rensselaer Polytechnic Institute*
[6] *University of Utah*

We consider the Bayesian inverse problem of inferring the initial condition of a linear dynamical system from noisy output measurements taken after the initial time. In practical applications, the large dimension of the dynamical system state poses a computational obstacle to computing the exact posterior distribution. Balanced truncation is a system-theoretic method for model reduction which obtains an efficient reduced-dimension dynamical system by projecting the system operators onto state directions which simultaneously maximize energies defined by reachability and observability Gramians. We show that in our inference setting, the prior covariance and Fisher information matrices can be naturally interpreted as reachability and observability Gramians, respectively. We use these connections to propose a balancing approach to model reduction for the inference setting. The resulting reduced model then inherits stability properties and error bounds from system theory, and yields an optimal posterior covariance approximation.