
Model Reduction of Linear Dynamical Systems via Balancing for Bayesian Inference

Elizabeth Qian · Jemima M. Tabeart · Christopher Beattie ·
 Serkan Gugercin · Jiahua Jiang · Peter R. Kramer · Akil
 Narayan ·

Abstract We consider the Bayesian approach to the linear Gaussian inference 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. *Model reduction* offers a variety of computational tools that seek to reduce this computational burden. In particular, *balanced truncation* is a system-theoretic approach to model reduction which obtains an efficient reduced-dimension dynamical system by projecting the system operators onto state directions which trade off the reachability and observability of state directions as expressed through the associated Gramians. We introduce Gramian definitions relevant to the inference setting and propose a balanced truncation approach based on these inference Gramians that yield a reduced dynamical system that can be used to cheaply approximate the posterior mean and covariance. Our definitions exploit natural connections between (i) the reachability Gramian and the prior covariance and (ii) the observability Gramian and the Fisher information. The resulting reduced model then inherits stability properties and error bounds from system theoretic considerations, and in some settings yields an optimal posterior covariance approximation. Numerical demonstrations on two benchmark problems in model

This material is based upon work supported by the National Science Foundation under Grant No. DMS-1439786 and by the Simons Foundation Grant No. 50736 while the authors were in residence at the Institute for Computational and Experimental Research in Mathematics in Providence, RI, during the “Model and dimension reduction in uncertain and dynamic systems” program. EQ was supported in part by the Fannie and John Hertz Foundation. JMT was partially supported by EPSRC grant EP/S027785/1. AN was partially supported by NSF DMS-1848508. CB and SG were partially supported by NSF DMS-1819110.

E. Qian
 California Institute of Technology

J.M. Tabeart
 The University of Edinburgh

C. Beattie
 Virginia Tech

S. Gugercin
 Virginia Tech

J. Jiang
 ShanghaiTech University

P.R. Kramer
 Rensselaer Polytechnic Institute

A. Narayan
 University of Utah

reduction show that our method can yield near-optimal posterior covariance approximations with order-of-magnitude state dimension reduction.

Keywords Bayesian inference · Balanced truncation · Model reduction

1 Introduction

The task of updating computational models to take into account empirical observations — a task that plays a central role in *data assimilation* — occurs in a diverse range of disciplines, including geophysics, oceanography, and atmospheric sciences [12, 27, 35, 43]. Specific applications include sea ice modeling [47, 58], geothermal reservoir model calibration [19, 20], and nitrogen concentration of water bodies [24]. In the Bayesian inferential approach to data assimilation, model parameters $\mathbf{p} \in \mathbb{R}^d$ are related to observed data $\mathbf{m} \in \mathbb{R}^{d_{\text{obs}}}$ through a measurement model that combines a *forward map* $\mathbf{G} : \mathbb{R}^d \rightarrow \mathbb{R}^{d_{\text{obs}}}$ with additive stochastic noise $\epsilon \in \mathbb{R}^{d_{\text{obs}}}$:

$$\mathbf{m} = \mathbf{G}(\mathbf{p}) + \epsilon. \quad (1.1)$$

Before observations are obtained, uncertainty about the model parameters is encoded in a *prior* probability distribution for \mathbf{p} . The probability distribution of the measurements conditioned on the parameter, $\mathbf{m}|\mathbf{p}$, is called the *likelihood*. After observations are obtained, the Bayesian approach computes an updated *posterior* distribution for $\mathbf{p}|\mathbf{m}$ which reflects the uncertainty in the parameters conditioned on the measured data and is proportional to the product of the prior and the likelihood of the data. Once the posterior distribution has been characterized, subsequent steps in application may require computing and utilizing posterior statistics. In the linear Gaussian case, this requires computing only a posterior mean and covariance, but for general non-Gaussian posteriors, these statistics are approximated using sampling procedures [52] (e.g., rejection sampling [25, 26, 42] or Markov chain Monte Carlo methods [4, 33, 39]). Modern overviews of inference techniques and algorithms can be found in excellent survey articles [22, 55] and textbooks [17, 37, 56].

In practical settings, several factors may pose computational challenges to the use of Bayesian inference. First, the parameter dimension d is often large; e.g., when \mathbf{p} represents the high-dimensional spatial discretization of a continuous field. Second, it is common to be interested in parameter values, \mathbf{p} , at an initial time, with observations, \mathbf{m} , taken after the initial time. In these settings, the forward model may involve simulation of a high-dimensional dynamical system, making the characterization of the posterior distribution computationally expensive.

To reduce this computational burden, one class of methods exploits the fact that in practice the data are only informative in a low-dimensional subspace of \mathbb{R}^d . We highlight the approach of [54] for the linear Gaussian setting, which defines the posterior covariance as a negative semi-definite update to the prior covariance. The work [54] shows that the optimal low-rank approximation to this update lies in a subspace spanned by directions identified from dominant eigenvalue-eigenvector pairs of a matrix pencil that encodes a trade-off between the prior uncertainty and the Fisher information of the measurement model. This low-rank approximation performs well when $d_{\text{obs}} \ll d$, as is common in data assimilation problems for weather and climate [12, 35]. More broadly, this approach is beneficial in cases where the eigenvalues of the matrix pencil exhibit rapid decay, which can occur even if $d_{\text{obs}} > d$: some examples are discussed in the introduction of [54]. Extensions of the approach in [54] have been used to develop low rank approximation strategies for inverse problems with nonlinear forward models [21], goal-oriented notions of optimality [53], and low-rank tensor-based approaches to reducing computational complexity in time as well as state [9].

A complementary class of methods employs *model reduction*, which seeks to reduce the expense of evaluating the forward model \mathbf{G} . In projection-based model reduction, a high-dimensional linear dynamical system is approximated by projecting the system operators onto a low-dimensional subspace, yielding an inexpensive low-dimensional dynamical system [2, 3, 8, 31]. One popular strategy known as *balanced truncation* [40, 41] defines the projection subspace as the span of the dominant eigenvectors of a matrix pencil that encodes a trade-off between states that the system can easily

reach and states that contribute more to the system output. States are retained if they are both highly observable and easily reachable; otherwise, they are discarded. The resulting reduced-order dynamics evolve in a way that approximate the input-output map of the original large-order dynamical system with high fidelity, frequently with far smaller system order [7, 31, 38, 44].

1.1 Contributions of this article

This work considers the Bayesian inference setting where the parameter p to be inferred is the initial state (at time $t = 0$) of a linear dynamical system, which is presumed to have a Gaussian prior distribution. A linear observer views the evolving state at selected times $t > 0$ in the presence of Gaussian measurement noise. The primary goal of this work is then to propose a new method of model reduction for linear dynamical systems in this setting that exploits natural connections between (i) the optimal posterior approximation strategy of [54] and (ii) the system-theoretic model reduction method of balanced truncation. Our contributions are:

1. We define a notion of prior *compatibility* with system dynamics that allows the prior covariance matrix to be interpreted as a reachability Gramian in the inference context.
2. We propose two strategies for choosing a compatible prior: (i) by subjecting the system to a stochastic driving force at negative times to “spin up” the system state to a stationary distribution, and (ii) by modifying an incompatible prior to make it compatible. The spin-up strategy is related to ideas from the work [46] which considers stochastic dynamical systems.
3. We define a *noisy observability Gramian* that modifies the usual system-theoretic observability Gramian definition to account for additive noise in measurements.
4. We argue that the noisy observability Gramian is a natural continuous-time analogue of the Fisher information matrix, and we introduce two observation processes under which the Fisher information matrix recovers (in expectation/in the continuum limit) the noisy observability Gramian up to a scaling constant.
5. We propose a model reduction approach for our inference setting that applies the method of balanced truncation using a compatible prior covariance and the noisy observability Gramian, which inherits system-theoretic stability and error guarantees.
6. We show that the proposed approach can recover the optimal posterior covariance approximation of [54] in certain cases.
7. We demonstrate numerically that the proposed approach yields near-optimal approximations when the noisy observability Gramian and Fisher information matrix define similar approximation subspaces, which occurs in our numerical experiments when measurements are closely spaced over a long observation interval. Even when measurements are limited and the approximation subspaces are not particularly close, the proposed approach yields approximations that achieve order-of-magnitude system state dimension reduction with low, albeit sub-optimal, errors.

We note that [54] shows that the optimal low-rank posterior update coincides with an oblique projection of the high-dimensional forward map \mathbf{G} . Our contributions extend this result in two ways: first, we show that the low-rank update is in fact immediately associated with a *low-dimensional* reduced model for the linear dynamical system. This enables efficient evolution of the reduced dynamical system for low rank inference in settings where the map \mathbf{G} is only available implicitly through a high-dimensional linear dynamical systems model. Second, we show that with certain observation models, this reduced model has bounded error and inherits stability. We also note an earlier work [38] developed a goal-oriented low-rank approach to classical and statistical inverse problems that was also inspired by balanced truncation but departs from the dynamical system setting: the parameter is inferred in a low-dimensional subspace defined by an ‘experiment observability Gramian’ and a goal-oriented ‘prediction observability Gramian’. In contrast, we show that balanced truncation for dynamical systems can be used to develop reduced dynamical systems with desirable properties.

Section 2 introduces our Bayesian inference problem for the initial condition of a linear dynamical system, summarizes the balanced truncation approach to model reduction, and presents the low rank

inference procedure via optimal posterior updates from [54]. Section 3 introduces our Gramian definitions for the inference setting, proposes the balanced truncation posterior approximation approach based on these Gramians, and proves results regarding the error, stability, and optimal posterior covariance approximation of the resulting model. These theoretical results are based on specific relationships between system theoretic Gramians and matrices appearing in inference that are discussed in detail in Section 4. Section 5 presents and discusses numerical demonstrations of the proposed method on two benchmark problems. Section 6 summarizes our discussion and considers directions for further investigation.

2 Setting and background

In this section, we introduce our notation (Section 2.1) and formulate the Bayesian inverse problem we consider for a high-dimensional linear dynamical system (Section 2.2). We summarize two approaches to dimension reduction relevant in our setting: the system-theoretic approach to model reduction of linear time-invariant systems (Section 2.3) and the low rank inference approach of [54] (Section 2.4).

2.1 Notation

Vectors will be denoted by boldface lowercase letters, e.g., \mathbf{v} , and matrices by boldface uppercase letters, e.g., \mathbf{V} . We will use $\|\mathbf{v}\|$ to denote the Euclidean (ℓ^2) norm, and $\|\mathbf{v}\|_M^2 = \mathbf{v}^\top \mathbf{M} \mathbf{v}$ for symmetric positive (semi)-definite \mathbf{M} to denote the squared weighted \mathbf{M} -(semi)norm. For square matrices \mathbf{A} and \mathbf{B} of the same size, $\mathbf{A} \succ \mathbf{B}$ and $\mathbf{A} \succeq \mathbf{B}$ mean that $\mathbf{A} - \mathbf{B}$ is positive definite and positive semidefinite, respectively, with analogous meanings for \prec and \preceq . The positive integers d , d_{in} , and d_{out} denote the dimensions of state, input, and output vectors of our dynamical system. Random vectors (i.e., vectors of random variables) will not be denoted differently from (deterministic) vectors, but their properties will be noted when they are introduced. The expectation of a random vector \mathbf{x} is denoted as $\mathbb{E}[\mathbf{x}]$ and the covariance of two random vectors \mathbf{x}, \mathbf{y} is denoted $\text{cov}[\mathbf{x}, \mathbf{y}]$.

2.2 Problem formulation

We consider the following linear dynamical system:

$$\frac{d\mathbf{x}}{dt} = \mathbf{A}\mathbf{x}, \quad (2.1)$$

where $\mathbf{x}(t) \in \mathbb{R}^d$ and $\mathbf{A} \in \mathbb{R}^{d \times d}$. The initial state $\mathbf{x}(0) = \mathbf{p}$ is unknown. Noisy measurements \mathbf{m}_i of the system output are taken at times $0 < t_1 < t_2 < \dots < t_n$ according to the following measurement model:

$$\mathbf{m}_i \equiv \mathbf{C}\mathbf{x}(t_i) + \boldsymbol{\epsilon}_i = \mathbf{C}e^{\mathbf{A}t_i}\mathbf{p} + \boldsymbol{\epsilon}_i, \quad (2.2)$$

where $\mathbf{C} \in \mathbb{R}^{d_{\text{out}} \times d}$ is the state-to-output operator and $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_\epsilon)$ represents independently and identically distributed additive Gaussian noise, with positive definite covariance $\boldsymbol{\Gamma}_\epsilon \in \mathbb{R}^{d_{\text{out}} \times d_{\text{out}}}$. In our analysis, we will consider both finite n and $n = \infty$ measurements, and we will allow the observations times $\mathbf{t} = (t_1, \dots, t_n)^\top$ to be either deterministic or randomly distributed. Random observation times are to be interpreted in the sense that they are uncertain at the beginning of the observation process, but become known as measurements are taken. The random observation model is introduced principally to facilitate discussion of the linkage between observability Gramians and Fisher information in Section 4.2.3.

Our task is to infer the unknown initial state $\mathbf{x}(0) = \mathbf{p}$ from the noisy measurements in eq. (2.2). We use a Bayesian statistical approach to do so; i.e., we take as given a Gaussian prior:

$$\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}_{\text{pr}}) \quad (2.3)$$

with prior covariance $\boldsymbol{\Gamma}_{\text{pr}} \in \mathbb{R}^{d \times d}$. The Gaussian prior, combined with our linear Gaussian measurement model (eq. (2.2)), yields the following Gaussian likelihood conditioned on the measurement times:

$$\mathbf{m} | (\mathbf{p}, \mathbf{t}) \sim \mathcal{N}(\mathbf{G}\mathbf{p}, \boldsymbol{\Gamma}_{\text{obs}}), \quad (2.4)$$

where $\mathbf{m} \in \mathbb{R}^{nd_{\text{out}}}$, $\mathbf{G} \in \mathbb{R}^{nd_{\text{out}} \times d}$, and $\boldsymbol{\Gamma}_{\text{obs}} \in \mathbb{R}^{nd_{\text{out}} \times nd_{\text{out}}}$ are given as follows:

$$\mathbf{m} \equiv \begin{bmatrix} \mathbf{m}_1 \\ \vdots \\ \mathbf{m}_n \end{bmatrix}, \quad \mathbf{G} \equiv \begin{bmatrix} \mathbf{C}e^{\mathbf{A}t_1} \\ \vdots \\ \mathbf{C}e^{\mathbf{A}t_n} \end{bmatrix}, \quad \boldsymbol{\Gamma}_{\text{obs}} \equiv \begin{bmatrix} \boldsymbol{\Gamma}_\epsilon & & & \\ & \boldsymbol{\Gamma}_\epsilon & & \\ & & \ddots & \\ & & & \boldsymbol{\Gamma}_\epsilon \end{bmatrix}. \quad (2.5)$$

The Bayesian posterior distribution (conditioned on the measurement times) for such a prior and likelihood is again Gaussian:

$$\mathbf{p} | (\mathbf{m}, \mathbf{t}) \sim \mathcal{N}(\boldsymbol{\mu}_{\text{pos}}, \boldsymbol{\Gamma}_{\text{pos}}), \quad (2.6)$$

where

$$\boldsymbol{\mu}_{\text{pos}} = \boldsymbol{\Gamma}_{\text{pos}} \mathbf{G}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \mathbf{m}, \quad \boldsymbol{\Gamma}_{\text{pos}} = (\mathbf{H} + \boldsymbol{\Gamma}_{\text{pr}}^{-1})^{-1}, \quad (2.7)$$

where $\mathbf{H} \in \mathbb{R}^{d \times d}$ denotes the Fisher information matrix of \mathbf{m} with respect to the parameter \mathbf{x}_0 :

$$\mathbf{H} \equiv \mathbf{G}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \mathbf{G} = \sum_{i=1}^n \mathbf{G}_i^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{G}_i = \sum_{i=1}^n e^{\mathbf{A}^\top t_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_i}. \quad (2.8)$$

In practical settings, the high state dimension d can pose challenges to the computation of the conditional posterior statistics (2.7), especially when \mathbf{G} is available only implicitly through a computationally expensive model that evolves the high-dimensional dynamical system (2.1). Our procedure addresses this challenge by providing a computationally efficient strategy for emulating the dynamical system through model reduction.

2.3 System-theoretic model reduction via balanced truncation

Balanced truncation [40, 41] is a system-theoretic method for projection-based model reduction that yields a cheaply evaluated low-dimensional reduced model whose input-output map approximates that of the original high-dimensional system. Section 2.3.1 describes the system theory setting and Section 2.3.2 presents the balanced truncation method.

2.3.1 System theory background

The usual setting in which balanced truncation is considered involves a linear time-invariant (LTI) system given by

$$\frac{dx}{dt} = Ax + Bu(t) \quad (2.9a)$$

$$y = Cx \quad (2.9b)$$

where $u(t)$ is a d_{in} -dimensional input which drives the system filtered through the input port, $B \in \mathbb{R}^{d \times d_{in}}$, and $C \in \mathbb{R}^{d_{out} \times d}$ describes the d_{out} -dimensional observer output, $y(t)$. This system-theoretic setting differs from our inference setting in eq. (2.1) through the presence of an input $u(t)$ and the lack of noise in the output $y(t)$ (compared to the measurements m_i defined in eq. (2.2)).

We assume that the system (2.9) is stable, i.e., the eigenvalues of A lie in the open left half-plane. The stable system (2.9) has infinite *reachability* Gramian $P_\infty \in \mathbb{R}^{d \times d}$ and *observability* Gramian $Q_y \in \mathbb{R}^{d \times d}$, given by

$$P_\infty \equiv \int_0^\infty e^{At} BB^\top e^{A^\top t} dt, \quad Q_y \equiv \int_0^\infty e^{A^\top t} C^\top Ce^{At} dt. \quad (2.10)$$

We use subscripts \cdot_∞ and \cdot_y to denote the specific Gramians defined in (2.10), contrasting with generic reachability and observability Gramians, P and Q , which could be defined either as above or in other ways. Note in particular that Q_y denotes an observability Gramian tied to the LTI output y (2.9), whereas we will later introduce a Gramian Q_m for noisy measurements m (2.2). The Gramians P_∞, Q_y are unique solutions to the dual Lyapunov equations

$$AP_\infty + P_\infty A^\top = -BB^\top, \quad A^\top Q_y + Q_y A = -C^\top C. \quad (2.11)$$

These infinite-horizon Gramians for the dynamical system (2.9) are associated with a reachability energy and an observability energy, defined as quadratic forms with respect to a state $\hat{x} \in \mathbb{R}^d$, respectively, as

$$\|\hat{x}\|_{P_\infty^{-1}}^2 \equiv \hat{x}^\top P_\infty^{-1} \hat{x}, \quad \|\hat{x}\|_{Q_y}^2 \equiv \hat{x}^\top Q_y \hat{x}. \quad (2.12)$$

The reachability energy $\|\hat{x}\|_{P_\infty^{-1}}^2$ is the minimum L^2 -norm $\int_0^\infty \|u(t)\|^2 dt$ of a control input, $u(t)$, required to steer the system from the origin ($x(0) = 0$) to \hat{x} , potentially over an infinite time horizon. Similarly, observability energy $\|\hat{x}\|_{Q_y}^2$ is the L^2 -norm $\int_0^\infty \|y(t)\|^2 dt$ of the (noise-free) output signal over an infinite time horizon generated by the free evolution of the system initialized at $x(0) = \hat{x}$ with null input, $u(t) = 0$. While the Gramian definitions in eq. (2.10) are perhaps the most commonly used, finite-horizon and discretized/approximate Gramian definitions are also sometimes employed, the latter especially for large-scale problems where solving eq. (2.11) is computationally intractable. For details, see, e.g., [2, 6, 10, 31].

A system is *reachable* if P_∞ is positive definite (denoted as $P_\infty \succ 0$) and a system is *observable* if $Q_y \succ 0$. The system is *minimal* if it is both reachable and observable. In all that follows, we assume that the system is minimal.

2.3.2 Balanced truncation

In projection-based model reduction, a reduced model is obtained by projecting the system operators onto low-dimensional subspaces. In balanced truncation, these subspaces are determined by the dominant eigenvectors of the matrix pencil (Q, P^{-1}) , i.e., the v_i associated with the dominant eigenvalues, δ_i^2 , satisfying

$$Q v_i = \delta_i^2 P^{-1} v_i, \quad (2.13)$$

where \mathbf{P} is a reachability Gramian and \mathbf{Q} is an observability Gramian: the standard balanced truncation approach takes $\mathbf{P} = \mathbf{P}_\infty$ and $\mathbf{Q} = \mathbf{Q}_y$, but we will in Section 3 introduce approaches that use different Gramians. The scalars $\{\delta_i\}_{i=1}^d$ with $\delta_1 > \delta_2 > \dots > \delta_d$ are called the *Hankel singular values* and are the nonzero singular values of the underlying Hankel operator. We assume that the Hankel singular values are distinct to simplify the presentation. The dominant eigenvectors, \mathbf{v}_i , maximize the following Rayleigh quotient:

$$\frac{\mathbf{v}^\top \mathbf{Q} \mathbf{v}}{\mathbf{v}^\top \mathbf{P}^{-1} \mathbf{v}} = \left(\frac{\|\mathbf{v}\|_{\mathbf{Q}}}{\|\mathbf{v}\|_{\mathbf{P}^{-1}}} \right)^2. \quad (2.14)$$

Thus, directions retained by balanced truncation will be simultaneously easy to reach (low reachability energy) and easy to observe (high observability energy), or will realize some trade-off between these two desirable properties.

The balanced truncation procedure has the following explicit algebraic formulation. Suppose $\mathbf{P} = \mathbf{R}\mathbf{R}^\top$ and $\mathbf{Q} = \mathbf{L}\mathbf{L}^\top$ (e.g., through a Cholesky factorization), and let $\mathbf{U}\Delta\mathbf{Z}^\top = \mathbf{L}^\top \mathbf{R}$ be the singular value decomposition of $\mathbf{L}^\top \mathbf{R}$. Then Δ contains the Hankel singular values, $\Delta = \text{diag}(\delta_1, \dots, \delta_d)$. Define $\mathbf{T} = \mathbf{R}\mathbf{Z}\Delta^{-\frac{1}{2}} \in \mathbb{R}^{d \times d}$ and note that \mathbf{T} is *invertible* with $\mathbf{T}^{-1} = \Delta^{-\frac{1}{2}} \mathbf{U}^\top \mathbf{L}^\top$. The ordered generalized eigenvectors, \mathbf{v}_i , of $(\mathbf{Q}, \mathbf{P}^{-1})$ are the columns of $\mathbf{V} = \mathbf{R}\mathbf{Z}$, and through a renormalization by $\Delta^{-\frac{1}{2}}$ become the columns of \mathbf{T} .

\mathbf{T} is a *balancing* transformation, meaning that in the transformed coordinates $\tilde{\mathbf{x}} = \mathbf{T}^{-1}\mathbf{x}$ the reachability and observability Gramians are equal diagonal matrices. The transformed LTI system is given by,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \tilde{\mathbf{A}}\tilde{\mathbf{x}} + \tilde{\mathbf{B}}\mathbf{u}(t) \quad (2.15a)$$

$$\tilde{\mathbf{y}} = \tilde{\mathbf{C}}\tilde{\mathbf{x}}, \quad (2.15b)$$

with

$$\mathbf{T}\tilde{\mathbf{x}} = \mathbf{x}, \quad \tilde{\mathbf{A}} = \mathbf{T}^{-1}\mathbf{A}\mathbf{T}, \quad \tilde{\mathbf{B}} = \mathbf{T}^{-1}\mathbf{B}, \quad \tilde{\mathbf{C}} = \mathbf{C}\mathbf{T},$$

and the system (2.15) is called (principal-axis) *balanced* because its reachability and observability Gramians $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{Q}}$ are equal and diagonal, satisfying

$$\mathbf{T}^{-1}\mathbf{P}\mathbf{T}^{-\top} = \tilde{\mathbf{P}} = \Sigma = \tilde{\mathbf{Q}} = \mathbf{T}^\top\mathbf{Q}\mathbf{T}. \quad (2.16)$$

Order- r balanced truncation obtains a reduced dynamical system by transforming to the balancing coordinates $\tilde{\mathbf{x}}$ and truncating to the leading r components of this balanced state. Denote by $\mathbf{U}_r, \mathbf{Z}_r \in \mathbb{R}^{d \times r}$ the leading r columns of \mathbf{U}, \mathbf{Z} , respectively, and let Δ_r denote the upper-left $r \times r$ block of Δ . Then, let \mathbf{T}_r denote the leading r columns of \mathbf{T} and let \mathbf{S}_r^\top denote the leading r rows of \mathbf{T}^{-1} , i.e.,

$$\mathbf{T}_r = \mathbf{R}\mathbf{Z}_r\Delta_r^{-\frac{1}{2}} \in \mathbb{R}^{d \times r}, \quad \mathbf{S}_r^\top = \Delta_r^{-\frac{1}{2}}\mathbf{U}_r^\top \mathbf{L}^\top \in \mathbb{R}^{r \times d}. \quad (2.17)$$

Note that $\mathbf{S}_r^\top \mathbf{T}_r = \mathbf{I}_r$. Then, the order- r balanced truncation approximation of (2.9) is a Petrov-Galerkin approximation given by restricting the state to lie in $\text{range}(\mathbf{T}_r)$, and by enforcing (2.9) on $\text{range}(\mathbf{S}_r)^\perp$. This results in the reduced-order approximation of eq. (2.9),

$$\frac{d\mathbf{x}_r}{dt} = \mathbf{A}_r \mathbf{x}_r + \mathbf{B}_r \mathbf{u}(t), \quad (2.18a)$$

$$\mathbf{y}_r = \mathbf{C}_r \mathbf{x}_r, \quad (2.18b)$$

where $\mathbf{A}_r = \mathbf{S}_r^\top \mathbf{A} \mathbf{T}_r \in \mathbb{R}^{r \times r}$, $\mathbf{B}_r = \mathbf{S}_r^\top \mathbf{B} \in \mathbb{R}^{r \times d_{\text{in}}}$, $\mathbf{C}_r = \mathbf{C} \mathbf{T}_r \in \mathbb{R}^{d_{\text{out}} \times r}$, and $\mathbf{T}_r \mathbf{x}_r \approx \mathbf{x}$.

Balanced truncation provides guaranteed stability and *a priori* L_2 -error bounds on the approximation error expressed in terms of the Hankel singular values in (2.13):

Lemma 2.1 (Thm. 7.9 of [2]) Consider the system (2.9) with output \mathbf{y} and reachability and observability Gramians \mathbf{P}_∞ and $\mathbf{Q}_\mathbf{y}$ defined in (2.10). If the original system (2.9) is stable, then the reduced model (2.18) is both stable and balanced: the infinite-horizon reachability and observability Gramians of the reduced state, \mathbf{x}_r , are diagonal and equal, and given by

$$\mathbf{P}_r \equiv \int_0^\infty e^{\mathbf{A}_r t} \mathbf{B}_r \mathbf{B}_r^\top e^{\mathbf{A}_r^\top t} dt = \int_0^\infty e^{\mathbf{A}_r^\top t} \mathbf{C}_r^\top \mathbf{C}_r e^{\mathbf{A}_r t} dt \equiv \mathbf{Q}_{\mathbf{y}_r} = \text{diag}(\delta_1, \dots, \delta_r) \in \mathbb{R}^{r \times r}. \quad (2.19)$$

Furthermore, if $\mathbf{x}(0) = 0$, then the output \mathbf{y}_r of eq. (2.18) commits an error bounded by twice the sum of the truncated Hankel singular values:

$$\|\mathbf{y} - \mathbf{y}_r\|_{L^2(\mathbb{R})} \leq \left(2 \sum_{j=r+1}^d \delta_j \right) \|\mathbf{u}\|_{L^2(\mathbb{R})}. \quad (2.20)$$

2.4 Low rank inference via optimal posterior covariance approximation

We now summarize the main results of [54] for dimension reduction of the linear Gaussian Bayesian inverse problem by optimal approximation of the posterior covariance and mean. This approach is applicable even when \mathbf{G} does not entail the evolution of a dynamical system, and thus complements the model reduction approach of Section 2.3.

Let (τ_i^2, \mathbf{w}_i) denote the ordered generalized eigenpairs of the pencil $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$, i.e.,

$$\mathbf{H}\mathbf{w}_i = \tau_i^2 \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{w}_i, \quad (2.21)$$

with $\mathbf{w}_i^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{w}_j = \delta_{i,j}$, where $\delta_{i,j}$ denotes the Kronecker delta, and $\tau_1 \geq \tau_2 \geq \dots \geq \tau_d$. Then, note that the posterior covariance, $\boldsymbol{\Gamma}_{\text{pos}}$, satisfies

$$\boldsymbol{\Gamma}_{\text{pos}} = (\boldsymbol{\Gamma}_{\text{pr}}^{-1} + \mathbf{H})^{-1} = \boldsymbol{\Gamma}_{\text{pr}} - \sum_{i=1}^d \frac{\tau_i^2}{1 + \tau_i^2} \mathbf{w}_i \mathbf{w}_i^\top. \quad (2.22)$$

When there are only a few non-zero eigenvalues, τ_i^2 (e.g., when $d_{\text{obs}} \ll d$), or when τ_i^2 decays rapidly with i , the sum in (2.22) can be well-approximated by a low-rank symmetric positive semidefinite matrix. The work in [54] thus considers the optimal approximation of $\boldsymbol{\Gamma}_{\text{pos}}$ within the following class of rank- r negative semidefinite updates to $\boldsymbol{\Gamma}_{\text{pr}}$:

$$\mathcal{M}_r := \left\{ \mathbf{M} \in \mathbb{R}^{d \times d} \mid \mathbf{M} = \boldsymbol{\Gamma}_{\text{pr}} - \mathbf{K} \mathbf{K}^\top \succ 0 \text{ with } \text{rank}(\mathbf{K}) \leq r \right\}. \quad (2.23)$$

The distance between the true posterior covariance and an approximation within \mathcal{M}_r is measured by the Förstner metric [28] between symmetric positive semidefinite matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{d \times d}$:

$$d_{\mathcal{F}}(\mathbf{A}, \mathbf{B}) = \sum_{i=1}^d \ln^2(\sigma_i), \quad (2.24)$$

where σ_i are the eigenvalues of the pencil (\mathbf{A}, \mathbf{B}) . The Förstner distance is invariant to similarity transformation, and $d_{\mathcal{F}}(\mathbf{A}, \mathbf{B}) = d_{\mathcal{F}}(\mathbf{A}^{-1}, \mathbf{B}^{-1})$. The Förstner distance is widely used in statistics, being the natural Riemannian metric on the manifold of positive-definite (e.g., full-rank covariance) matrices [14, Section 6.1] that can be generalized with appropriate modifications to rank-deficient manifolds of semi-definite matrices [15].

The Förstner-optimal approximation $\hat{\boldsymbol{\Gamma}}_{\text{pos}}$ to $\boldsymbol{\Gamma}_{\text{pos}}$ within the class of rank- r updates \mathcal{M}_r ,

$$\hat{\boldsymbol{\Gamma}}_{\text{pos}} = \arg \min_{\mathbf{M} \in \mathcal{M}_r} d_{\mathcal{F}}(\boldsymbol{\Gamma}_{\text{pos}}, \mathbf{M}) \quad (2.25)$$

is then given by [54]:

$$\widehat{\boldsymbol{\Gamma}}_{\text{pos}} = \boldsymbol{\Gamma}_{\text{pr}} - \boldsymbol{K}_* \boldsymbol{K}_*^\top \quad \text{where} \quad \boldsymbol{K}_* \boldsymbol{K}_*^\top = \sum_{i=1}^r \frac{\tau_i^2}{1 + \tau_i^2} \boldsymbol{w}_i \boldsymbol{w}_i^\top, \quad (2.26)$$

with optimal Förstner distance

$$d_{\mathcal{F}}(\boldsymbol{\Gamma}_{\text{pos}}, \widehat{\boldsymbol{\Gamma}}_{\text{pos}}) = \min_{\boldsymbol{M} \in \mathcal{M}_r} d_{\mathcal{F}}(\boldsymbol{\Gamma}_{\text{pos}}, \boldsymbol{M}) = \sum_{i=r+1}^d \ln^2 \left(\frac{1}{1 + \tau_i^2} \right). \quad (2.27)$$

An alternative analysis in [54] allows one to associate the optimal approximation (2.26) with the posterior covariance of a projected forward map. Note that the pairs $(\frac{1}{1 + \tau_i^2}, \tilde{\boldsymbol{w}}_i)$ with $\tilde{\boldsymbol{w}}_i = \boldsymbol{\Gamma}_{\text{pr}}^{-1} \boldsymbol{w}_i$ are the generalized eigenpairs of the pencil $(\boldsymbol{\Gamma}_{\text{pos}}, \boldsymbol{\Gamma}_{\text{pr}})$. Let $\boldsymbol{\Pi}_r = \sum_{i=1}^r \tilde{\boldsymbol{w}}_i \tilde{\boldsymbol{w}}_i^\top$ denote the oblique spectral projector for $\boldsymbol{\Gamma}_{\text{pr}}^{-1} \boldsymbol{\Gamma}_{\text{pos}}$ onto $\text{span}\{\tilde{\boldsymbol{w}}_1, \tilde{\boldsymbol{w}}_2, \dots, \tilde{\boldsymbol{w}}_r\}$. Then, let

$$\widehat{\boldsymbol{G}} = \boldsymbol{G} \boldsymbol{\Pi}_r^\top, \quad \widehat{\boldsymbol{H}} = \widehat{\boldsymbol{G}}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \widehat{\boldsymbol{G}}, \quad (2.28)$$

denote the projected forward map and projected Fisher information, respectively. Then, the optimal posterior covariance approximation (2.26) can be equivalently expressed as

$$\widehat{\boldsymbol{\Gamma}}_{\text{pos}} = (\boldsymbol{\Gamma}_{\text{pr}}^{-1} + \widehat{\boldsymbol{H}})^{-1}. \quad (2.29)$$

The work in [54] further considers optimal approximation to the posterior mean $\boldsymbol{\mu}_{\text{pos}}$ in two different classes of mean approximations: a class of low-rank (LR) mean approximations given by

$$\mathcal{V}_r^{\text{LR}} := \{\boldsymbol{v} = \boldsymbol{N} \boldsymbol{m} \mid \boldsymbol{N} \in \mathbb{R}^{d \times nd_{\text{out}}}, \text{rank}(\boldsymbol{N}) \leq r\} \quad (2.30)$$

and a class of mean approximations given by low-rank update (LRU) approximations to the posterior covariance:

$$\mathcal{V}_r^{\text{LRU}} := \{\boldsymbol{v} = (\boldsymbol{\Gamma}_{\text{pr}} - \boldsymbol{N}) \boldsymbol{G}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \boldsymbol{m} \mid \boldsymbol{N} \in \mathbb{R}^{d \times d}, \text{rank}(\boldsymbol{N}) \leq r\}. \quad (2.31)$$

Within $\mathcal{V}_r^{\text{LR}}$, the class of low-rank mean approximations (2.30), the following mean approximation is optimal in the sense that it minimizes the $\|\cdot\|_{\boldsymbol{\Gamma}_{\text{pos}}^{-1}}$ -norm error averaged over the joint distribution of the initial condition (with marginal given by the prior) and the data (e.g., Bayes risk):

$$\widehat{\boldsymbol{\mu}}_{\text{pos}}^{\text{LR}} = \sum_{i=1}^r \frac{\tau_i}{1 + \tau_i^2} \boldsymbol{w}_i \boldsymbol{w}_i^\top = \widehat{\boldsymbol{\Gamma}}_{\text{pos}} \widehat{\boldsymbol{G}}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \boldsymbol{m}. \quad (2.32)$$

Within $\mathcal{V}_r^{\text{LRU}}$, the class of low-rank update mean approximations (2.31), the optimal mean approximation in the $\|\cdot\|_{\boldsymbol{\Gamma}_{\text{pos}}^{-1}}$ -norm Bayes risk is given by [54]:

$$\widehat{\boldsymbol{\mu}}_{\text{pos}}^{\text{LRU}} = \widehat{\boldsymbol{\Gamma}}_{\text{pos}} \boldsymbol{G}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \boldsymbol{m}. \quad (2.33)$$

Note that in a dynamical systems setting, computation of the posterior approximations $\widehat{\boldsymbol{\Gamma}}_{\text{pos}}$, $\widehat{\boldsymbol{\mu}}_{\text{pos}}^{\text{LR}}$, and $\widehat{\boldsymbol{\mu}}_{\text{pos}}^{\text{LRU}}$ requires evolving the full high-dimensional dynamics of \boldsymbol{G} and then projecting the result to obtain $\widehat{\boldsymbol{G}}$. In the next section, we propose a model reduction approach that yields a low-dimensional dynamical systems model by combining elements from this posterior covariance approximation approach and from balanced truncation.

3 Balanced truncation for Bayesian inference

We now explore the application of the balanced truncation procedure of Section 2.3 to obtain a reduced model for the linear dynamical system (2.1) in an effort to reduce computational cost in our Bayesian inverse problem. To do this, we first negotiate two technical hurdles which frustrate a direct connection: the dynamical system (2.1) lacks an input and thus a notion of reachability, and the standard infinite-horizon observability Gramian \mathbf{Q}_y accounts for neither the additive measurement noise nor the specific observation times of the measurement model (2.2). To reconcile these discrepancies with the inference problem, we introduce suitably modified Gramian definitions in Section 3.1. In Section 3.2, we then present two variations of our proposed balanced truncation approach to Bayesian inference. In Section 3.3, we establish conditions under which one variation of our proposed method will approximately recover the optimal approximation qualities of the low rank inference procedure established in [54], with a faster computation of the forward maps via a stable and balanced reduced approximation of the dynamical system. The second variation we present may appear more natural, though it does not inherit the stability guarantees of a proper balanced truncation. Both variations will be explored in numerical examples in Section 5.

3.1 Inference Gramians

We begin by introducing a definition of *prior-compatibility* which allows the prior covariance, $\mathbf{\Gamma}_{\text{pr}}$, to be interpreted as a reachability Gramian of a forced system related to (2.1).

Definition 3.1 The prior covariance, $\mathbf{\Gamma}_{\text{pr}}$, will be said to be *prior-compatible* with state dynamics induced by \mathbf{A} if $\mathbf{A}\mathbf{\Gamma}_{\text{pr}} + \mathbf{\Gamma}_{\text{pr}}\mathbf{A}^\top$ is negative semidefinite.

If a prior covariance $\mathbf{\Gamma}_{\text{pr}}$ satisfies Definition 3.1, then there exists some $\mathbf{B} \in \mathbb{R}^{d \times d_{\text{in}}}$ with $d_{\text{in}} = \text{rank}(\mathbf{A}\mathbf{\Gamma}_{\text{pr}} + \mathbf{\Gamma}_{\text{pr}}\mathbf{A}^\top)$ such that the Lyapunov equation (2.11) $\mathbf{A}\mathbf{\Gamma}_{\text{pr}} + \mathbf{\Gamma}_{\text{pr}}\mathbf{A}^\top = -\mathbf{B}\mathbf{B}^\top$ is satisfied. Thus, the prior-compatibility with state dynamics means that the prior covariance matrix $\mathbf{\Gamma}_{\text{pr}}$ is the reachability Gramian of an LTI system with dynamics specified by \mathbf{A} and an input port matrix \mathbf{B} . It is not generally necessary to identify a particular port matrix, \mathbf{B} , for Definition 3.1 to hold, although in some applications \mathbf{B} may be naturally provided (for example if the end goal of the inference problem is to determine a control input). Discussions of how a compatible prior might be either naturally defined or obtained by modifying a non-compatible prior are provided in Section 4.1.

We make use of a modified observability Gramian that is motivated by the noisy measurement model for \mathbf{m} . Recall that the standard observability Gramian \mathbf{Q}_y defines an energy that is the squared L^2 -norm of the output signal $\mathbf{y}(t)$. In our inference context, where the measurement at time t_i , \mathbf{m}_i , is an output polluted by measurement noise $\epsilon_i \sim \mathcal{N}(0, \mathbf{\Gamma}_\epsilon)$, we quantify the discernability of the signal from the noise by the Mahalanobis distance from the equilibrium state $\mathbf{0}$ to the conditional distribution $\mathcal{N}(\mathbf{C}\mathbf{e}^{\mathbf{A}t_i}\mathbf{p}, \mathbf{\Gamma}_\epsilon)$ of $\mathbf{m}_i | (\mathbf{p}, t_i)$:

$$D(\mathbf{0}, \mathcal{N}(\mathbf{C}\mathbf{e}^{\mathbf{A}t_i}\mathbf{p}, \mathbf{\Gamma}_\epsilon)) := \|\mathbf{C}\mathbf{e}^{\mathbf{A}t_i}\mathbf{p}\|_{\mathbf{\Gamma}_\epsilon^{-1}}. \quad (3.1)$$

One way to understand why this is a natural way to measure the distance between a distribution (the signal with observation noise) and a point (the stable equilibrium) is to note that for the case of a Gaussian distribution, neighborhoods of the mean of the distribution defined by the Mahalanobis distance contain the most probability relative to any other set with the same Lebesgue measure (Appendix B). This motivates our definition of a modified observability Gramian for noisy measurements.

Definition 3.2 The noisy observability Gramian $\mathbf{Q}_m \in \mathbb{R}^{d \times d}$ is given by

$$\mathbf{Q}_m \equiv \int_0^\infty \mathbf{e}^{\mathbf{A}^\top t} \mathbf{C}^\top \mathbf{\Gamma}_\epsilon^{-1} \mathbf{C} \mathbf{e}^{\mathbf{A}t} dt. \quad (3.2)$$

Thus, \mathbf{Q}_m defines an energy based on integrating over all positive time the squared Mahalanobis distance between the equilibrium and the noise-polluted outputs. Note that the Fisher information matrix \mathbf{H} (2.8) can be viewed as defining an analogous energy based on summing over discrete measurement times. This analogy motivates an approximation of \mathbf{H} in terms of \mathbf{Q}_m discussed in subsequent sections.

3.2 Method

We propose a model reduction approach for the inference problem defined in Section 2.2 that applies the balanced truncation procedure of Section 2.3 by taking first a compatible prior covariance, $\boldsymbol{\Gamma}_{\text{pr}}$, and then using as inference Gramians $\mathbf{P} = \boldsymbol{\Gamma}_{\text{pr}}$ and $\mathbf{Q} = \mathbf{Q}_m$. This leads to the following reduced-order model for the linear dynamics (2.1):

$$\frac{d\mathbf{x}_r}{dt} = \mathbf{A}_r \mathbf{x}_r, \quad \mathbf{x}_r(0) = \mathbf{S}_r^\top \mathbf{p}, \quad (3.3)$$

where $\mathbf{A}_r = \mathbf{S}_r^\top \mathbf{A} \mathbf{T}_r$, and the following reduced-order approximation to the measurement model (2.2):

$$\mathbf{m}_i \approx \mathbf{C}_r \mathbf{x}_r(t_i) + \boldsymbol{\epsilon}_i, \quad (3.4)$$

where $\mathbf{C}_r = \mathbf{C} \mathbf{T}_r$, and $\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \boldsymbol{\Gamma}_\epsilon)$ as before. The reduced dynamics (3.3) and measurement model (3.4) define a reduced forward map from the reduced state to the noise-free output:

$$\mathbf{G}_{BT} \equiv \begin{bmatrix} \mathbf{C}_r e^{\mathbf{A}_r t_1} \mathbf{S}_r^\top \\ \vdots \\ \mathbf{C}_r e^{\mathbf{A}_r t_n} \mathbf{S}_r^\top \end{bmatrix} = \begin{bmatrix} \mathbf{C}_r e^{\mathbf{A}_r t_1} \\ \vdots \\ \mathbf{C}_r e^{\mathbf{A}_r t_n} \end{bmatrix} \mathbf{S}_r^\top \in \mathbb{R}^{nd_{\text{out}} \times d}. \quad (3.5)$$

We then define a corresponding reduced-order approximate Fisher information matrix,

$$\mathbf{H}_{BT} \equiv \mathbf{G}_{BT}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \mathbf{G}_{BT} = \mathbf{S}_r \left(\sum_{i=1}^n e^{\mathbf{A}_r^\top t_i} \mathbf{C}_r^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C}_r e^{\mathbf{A}_r t_i} \right) \mathbf{S}_r^\top, \quad (3.6)$$

and associated posterior mean and covariance approximations:

$$\boldsymbol{\mu}_{\text{pos}, BT} = \boldsymbol{\Gamma}_{\text{pos}, BT} \mathbf{G}_{BT}^\top \boldsymbol{\Gamma}_{\text{obs}}^{-1} \mathbf{m}, \quad \boldsymbol{\Gamma}_{\text{pos}, BT} = (\mathbf{H}_{BT} + \boldsymbol{\Gamma}_{\text{pr}}^{-1})^{-1}. \quad (3.7)$$

As previously discussed at the end of Section 3.1, the noisy observability Gramian \mathbf{Q}_m can be viewed as a continuous-time analogue of \mathbf{H} . This suggests a variation of the proposed method that uses the inference Gramians $\mathbf{P} = \boldsymbol{\Gamma}_{\text{pr}}$ and $\mathbf{Q} = \mathbf{H}$ to define the balanced truncation. We label the two variations of our proposed method as follows:

1. **BT-Q**, which uses the inference Gramians $(\mathbf{P}, \mathbf{Q}) = (\boldsymbol{\Gamma}_{\text{pr}}, \mathbf{Q}_m)$ in order to define the posterior approximations in (3.7), and
2. **BT-H**, which uses the inference Gramians $(\mathbf{P}, \mathbf{Q}) = (\boldsymbol{\Gamma}_{\text{pr}}, \mathbf{H})$ in order to define the posterior approximations in (3.7).

While the BT-H approach may seem to be a more natural choice for our inference setting, we will show in Section 3.3 that the BT-Q approach inherits desirable properties from standard system-theoretic balanced truncation that the BT-H approach generally does not enjoy.

3.3 Properties

We note that while the principal goal of the balanced truncation approach in the present work is to provide a reduced model approximation to posterior quantities, the proposed BT-Q approach inherits guarantees from the system-theoretic LTI setting, as described in our first result.

Lemma 3.1 *Suppose Γ_{pr} is prior-compatible with the state dynamics (2.1) as defined in Definition 3.1. Then, there exists a $\mathbf{B} \in \mathbb{R}^{d \times d_{\text{in}}}$ so that Γ_{pr} and \mathbf{Q}_m are the infinite reachability and observability Gramians of the following LTI system:*

$$\frac{d\mathbf{x}}{dt} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}(t), \quad \bar{\mathbf{y}}(t) = \bar{\mathbf{C}}\mathbf{x}(t), \quad (3.8)$$

where $\bar{\mathbf{C}} = \Gamma_{\text{obs}}^{-\frac{1}{2}} \mathbf{C}$.

Additionally, the balanced reduced LTI system (2.18) based on the inference Gramians Γ_{pr} and \mathbf{Q}_m is balanced, stable, and satisfies the following Hankel singular value tail bound: if $\mathbf{x}(0) = 0$, then

$$\|\bar{\mathbf{C}}\mathbf{x}(t) - \bar{\mathbf{C}}_r\mathbf{x}_r(t)\|_{L^2(\mathbb{R})} \leq \left(2 \sum_{j=r+1}^d \delta_j \right) \|\mathbf{u}\|_{L^2(\mathbb{R})}, \quad (3.9)$$

where δ_j are the Hankel singular values associated with the inference Gramians Γ_{pr} and \mathbf{Q}_m .

Proof The first part of Lemma 3.1 follows immediately from the compatibility condition in Definition 3.1 and from replacing \mathbf{C} with $\bar{\mathbf{C}}$ in eq. (2.10). The second part then follows from direct application of Lemma 2.1 to (3.8). \square

Corollary 3.1 *The proposed balanced truncation reduced model (3.3) for the inference dynamical system (2.1) is stable.*

Corollary 3.1 has practical significance for the proposed BT-Q approach of Section 3.2, because computing the approximate posterior quantities requires evolving the reduced dynamics given by (2.1): Corollary 3.1 guarantees that these reduced dynamics preserve the stability of the full dynamics. On the other hand, while the error and balancing guarantees of Lemma 3.1 for LTI systems are not directly applicable to the inference setting we consider here, they may serve as the basis for future extension to inference problems with control inputs.

We now turn our attention to the quality of the proposed posterior covariance approximation eq. (3.7): again, we focus on the BT-Q approach. The BT-Q posterior covariance approximation must tautologically be suboptimal relative to the optimal approximation scheme of [54] discussed in Section 2.4. A natural question is: when does covariance approximation approach optimality? In comparing the covariance approximations (3.7) and (2.29), it is clear that as $r \rightarrow d$, both $\mathbf{H}_{\text{BT}} \rightarrow \mathbf{H}$ and $\widehat{\mathbf{H}} \rightarrow \mathbf{H}$; i.e., both the balanced truncation Fisher information approximation and the projected Fisher information of [54] converge to the true Fisher information and thus the two posterior covariance approximations both converge to the full posterior covariance. However, since our goal is a low rank approximation with $r \ll d$, we would like some sense of when we can expect \mathbf{H}_{BT} to be close to $\widehat{\mathbf{H}}$ for general r .

For \mathbf{H}_{BT} resulting from the BT-Q approach, approximate agreement with $\widehat{\mathbf{H}}$ requires that the dominant r generalized eigenvectors of the pencil $(\mathbf{H}, \Gamma_{\text{pr}}^{-1})$ be close to the dominant r generalized eigenvectors of the pencil $(\mathbf{Q}_m, \Gamma_{\text{pr}}^{-1})$. This would follow when the Fisher information \mathbf{H} (2.8) is approximately proportional to the noisy observability Gramian \mathbf{Q}_m in Definition 3.2 that is $\mathbf{Q}_m \approx h\mathbf{H}$ for a scaling factor h . This would seem to apply to situations in which the observations are of sufficient frequency and duration for the discrete sum in \mathbf{H} to provide, up to normalization, a good approximation to the integral in \mathbf{Q}_m . In Section 4.2, we formalize this idea and identify the scaling factor h as the (average) inter-observation interval and establish two limits of observation protocols under which we can obtain \mathbf{Q}_m as a rescaled limit of \mathbf{H} .

Even when the spectral projections based on \mathbf{H} and \mathbf{Q}_m are close, there remains a difference between $\widehat{\mathbf{H}}$ and \mathbf{H}_{BT} because the former involves the evolution of the full dynamics represented by matrix \mathbf{A} , whereas the proposed balanced truncation approximations require only the evolution of a computationally inexpensive reduced system \mathbf{A}_r . Remarkably, we can formulate a relation between the balanced truncation approximation to the Fisher information and the optimal low rank inference approximation $\widehat{\mathbf{H}}$ without a further assumption about the size of r .

To express this connection precisely, we consider the following noisy observability Gramian for the balanced truncation approximation:

$$\mathbf{Q}_m^{\text{BT}} \equiv \mathbf{S}_r \left(\int_0^\infty e^{\mathbf{A}_r^\top t} \mathbf{C}_r^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C}_r e^{\mathbf{A}_r t} dt \right) \mathbf{S}_r^\top \in \mathbb{R}^{d \times d}. \quad (3.10)$$

The quantity \mathbf{Q}_m^{BT} can be viewed as a continuous-time analogue of \mathbf{H}_{BT} : it lifts the reduced Gramian for the reduced state $\mathbf{x}_r \in \mathbb{R}^r$ to the full state $\mathbf{x} \in \mathbb{R}^d$ via the map \mathbf{S}_r . Thus if the sum in \mathbf{H} well-approximates the integral in \mathbf{Q}_m , i.e., $h\mathbf{H} \approx \mathbf{Q}_m$, we would also expect $h\mathbf{H}_{\text{BT}} \approx \mathbf{Q}_m^{\text{BT}}$ because \mathbf{H}_{BT} and \mathbf{Q}_m^{BT} have the same sum-to-integral relation as \mathbf{H} and \mathbf{Q}_m , differing only in the use of reduced operators \mathbf{C}_r and \mathbf{A}_r vs. full operators \mathbf{C} and \mathbf{A} .

Proposition 3.1 *Let (τ_i^2, \mathbf{w}_i) and $(\delta_i^2, \mathbf{v}_i)$, respectively, denote the ordered generalized eigenpairs of the pencil $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$ and $(\mathbf{Q}_m, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$, respectively, with normalization $\mathbf{w}_i^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{w}_i = 1$ and $\mathbf{v}_i^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{v}_i = 1$. Then, if the r dominant eigenpairs match up to a scaling constant $h > 0$, i.e., if*

$$\delta_i^2 = h \tau_i^2, \quad \mathbf{w}_i = \mathbf{v}_i, \quad \forall i \leq r, \quad (3.11)$$

then $\mathbf{Q}_m^{\text{BT}} = h \widehat{\mathbf{H}}$.

Proof The integral in eq. (3.10) is the noisy observability Gramian of the balanced truncation reduced system and thus, by Lemma 3.1 it is balanced, i.e.,

$$\int_0^\infty e^{\mathbf{A}_r^\top t} \mathbf{C}_r^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C}_r e^{\mathbf{A}_r t} dt = \text{diag}\{\delta_1, \delta_2, \dots, \delta_r\} \equiv \boldsymbol{\Delta}_r. \quad (3.12)$$

Define $\mathbf{V}_r = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r]$ and note from its definition in (2.17) that $\mathbf{S}_r^\top = \boldsymbol{\Delta}_r^{\frac{1}{2}} \mathbf{V}_r^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1}$. Thus, eq. (3.10) is equivalent to

$$\mathbf{Q}_m^{\text{BT}} = \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{V}_r \boldsymbol{\Delta}_r^2 \mathbf{V}_r^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1}.$$

Similarly defining $\mathbf{W}_r = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_r]$, we have from (2.28),

$$\widehat{\mathbf{H}} = \boldsymbol{\Gamma}_{\text{pr}}^{-1} \mathbf{W}_r \boldsymbol{\Sigma}_r^2 \mathbf{W}_r^\top \boldsymbol{\Gamma}_{\text{pr}}^{-1}$$

with $\boldsymbol{\Sigma}_r = \text{diag}\{\tau_1, \tau_2, \dots, \tau_r\}$. Under the assumption (3.11), this gives us $\mathbf{Q}_m^{\text{BT}} = h \widehat{\mathbf{H}}$. \square

Proposition 3.1 provides us with a bridge between \mathbf{H}_{BT} and $\widehat{\mathbf{H}}$ for any r , including $r \ll d$. It suggests the following argument: as noted above, when $\mathbf{Q}_m \approx h\mathbf{H}$, generally $\mathbf{Q}_m^{\text{BT}} \approx h\mathbf{H}_{\text{BT}}$ will also be true, and Proposition 3.1 would imply $\mathbf{Q}_m^{\text{BT}} \approx h \widehat{\mathbf{H}}$, so that $\mathbf{H}_{\text{BT}} \approx \widehat{\mathbf{H}}$. That is, when \mathbf{Q}_m and \mathbf{H} are approximately proportional, the BT-Q posterior covariance approximation converges to the optimal one without further assumption on r . Rigorizing such an argument would require a formulation of Proposition 3.1 with approximate rather than precise equalities of the dominant eigenpairs. Rather than pursuing this cumbersome exercise, we instead support this argument by numerical experiments in Section 5. We confirm that in observation scenarios where the Fisher information matrix \mathbf{H} and noisy observability Gramian \mathbf{Q}_m are approximately proportional, the BT-Q approach leads to near-optimal posterior covariance approximations.

The preceding discussion provides intuition for what types of measurement models will lead to posterior covariance approximations that are close to the optimal low-rank approach of [54]. Our

analysis and discussion do not address the quality of the posterior mean approximation in (3.7): while the balanced truncation mean approximation does belong to the class of low rank mean approximations (2.30) considered in [54], the balanced truncation mean approximation differs in the use of the balanced truncation adjoint approximation vs. the projected full adjoint used in the optimal low-rank mean (2.32) of [54]. Numerical experiments in Section 5 compare empirically the balanced truncation means with the optimal means of [54]. We emphasize that in the dynamical system setting, both optimal approximations of [54] require the evolution of the full dynamics, whereas the proposed balanced truncation approximation approach requires only the evolution of a computationally inexpensive reduced system.

Until now we have focused on properties of the BT-Q approach, and the result of Proposition 3.1 relies on the assumption (3.11) that the dominant eigenspaces of \mathbf{Q}_m and \mathbf{H} match, and that the associated eigenvalues match up to a constant scale factor. Comparable guarantees for the BT-H approach, which involves replacing \mathbf{Q}_m with \mathbf{H} in the balanced truncation approximation, appear unlikely for general choices of Fisher information matrix, since they cannot generally be interpreted as infinite-horizon observability Gramians for any minimal LTI system, and so Lemma 2.1 does not directly apply. However, when the dominant generalized eigenvectors of the pencils $(\mathbf{H}, \Gamma_{\text{pr}}^{-1})$ and $(\mathbf{Q}_m, \Gamma_{\text{pr}}^{-1})$ coincide, then the BT-H and BT-Q reduced models will be equivalent, as stated in our next result (note that no conditions are imposed on the generalized eigenvalues here).

Proposition 3.2 *Let (τ_i^2, \mathbf{w}_i) and $(\delta_i^2, \mathbf{v}_i)$, respectively, denote the ordered generalized eigenpairs of the pencil $(\mathbf{H}, \Gamma_{\text{pr}}^{-1})$ and $(\mathbf{Q}_m, \Gamma_{\text{pr}}^{-1})$, respectively, with normalization $\mathbf{w}_i^\top \Gamma_{\text{pr}}^{-1} \mathbf{w}_i = 1$ and $\mathbf{v}_i^\top \Gamma_{\text{pr}}^{-1} \mathbf{v}_i = 1$. Then, if the r dominant eigenvectors match, i.e., if $\mathbf{w}_i = \mathbf{v}_i$, for all $i \leq r$, then the reduced model, $\mathbf{A}_r = \mathbf{S}_r^\top \mathbf{A} \mathbf{T}_r$, $\mathbf{B}_r = \mathbf{S}_r^\top \mathbf{B}$, $\mathbf{C}_r = \mathbf{C} \mathbf{T}_r$, produced using the state space projection $\mathbf{T}_r \mathbf{S}_r^\top$ defined by the pencil $(\mathbf{H}, \Gamma_{\text{pr}}^{-1})$, replicates (up to a diagonal scaling of basis) the reduced model produced by standard balanced truncation using the pencil $(\mathbf{Q}_m, \Gamma_{\text{pr}}^{-1})$.*

Proof Suppose the truncation order is r , and let the columns of $\mathbf{V}_r = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r]$ collect the dominant generalized eigenvectors of the pencil $(\mathbf{Q}_m, \Gamma_{\text{pr}}^{-1})$ which are presumed to be shared with the pencil $(\mathbf{H}, \Gamma_{\text{pr}}^{-1})$. Taking $\Delta_r = \text{diag}\{\delta_1, \delta_2, \dots, \delta_r\}$ and $\Sigma_r = \text{diag}\{\tau_1, \tau_2, \dots, \tau_r\}$, we may write $\mathbf{Q}_m \mathbf{V}_r = \Gamma_{\text{pr}}^{-1} \mathbf{V}_r \Delta_r^2$ and $\mathbf{H} \mathbf{V}_r = \Gamma_{\text{pr}}^{-1} \mathbf{V}_r \Sigma_r^2$. Defining $\Gamma_{\text{pr}} = \mathbf{R} \mathbf{R}^\top$, $\mathbf{Q}_m = \mathbf{L} \mathbf{L}^\top$, and $\mathbf{H} = \tilde{\mathbf{L}} \tilde{\mathbf{L}}^\top$, we determine singular value decompositions $\mathbf{U} \Delta \mathbf{Z}^\top = \mathbf{L}^\top \mathbf{R}$ and $\tilde{\mathbf{U}} \Sigma \tilde{\mathbf{Z}}^\top = \tilde{\mathbf{L}}^\top \mathbf{R}$, with associated truncated matrices, \mathbf{U}_r , Δ_r , \mathbf{Z}_r , $\tilde{\mathbf{U}}_r$, Σ_r , and $\tilde{\mathbf{Z}}_r$. The truncated balancing projection determined by the standard Gramians \mathbf{Q}_m and Γ_{pr} are given by $\mathbf{T}_r = \mathbf{R} \mathbf{Z}_r \Delta_r^{-\frac{1}{2}}$ and $\mathbf{S}_r^\top = \Delta_r^{-\frac{1}{2}} \mathbf{U}_r^\top \mathbf{L}^\top$, producing a standard balanced truncation reduced model defined as $\mathbf{A}_r = \mathbf{S}_r^\top \mathbf{A} \mathbf{T}_r$, $\mathbf{B}_r = \mathbf{S}_r^\top \mathbf{B}$, $\mathbf{C}_r = \mathbf{C} \mathbf{T}_r$. The corresponding projection determined by the inference Gramians \mathbf{H} and Γ_{pr} are given by $\tilde{\mathbf{T}}_r = \mathbf{R} \tilde{\mathbf{Z}}_r \Sigma_r^{-\frac{1}{2}}$ and $\tilde{\mathbf{S}}_r^\top = \Sigma_r^{-\frac{1}{2}} \tilde{\mathbf{U}}_r^\top \tilde{\mathbf{L}}^\top$, producing a reduced model defined as $\tilde{\mathbf{A}}_r = \tilde{\mathbf{S}}_r^\top \tilde{\mathbf{A}} \tilde{\mathbf{T}}_r$, $\tilde{\mathbf{B}}_r = \tilde{\mathbf{S}}_r^\top \mathbf{B}$, $\tilde{\mathbf{C}}_r = \mathbf{C} \tilde{\mathbf{T}}_r$. Notice $\mathbf{V}_r = \mathbf{R} \mathbf{Z}_r = \mathbf{R} \tilde{\mathbf{Z}}_r$, so that $\mathbf{Z}_r = \tilde{\mathbf{Z}}_r$. Hence with $\mathbf{D} = \Sigma_r^{-\frac{1}{2}} \Delta_r^{\frac{1}{2}}$, we have

$$\tilde{\mathbf{T}}_r = \mathbf{R} \tilde{\mathbf{Z}}_r \Sigma_r^{-\frac{1}{2}} = \mathbf{R} \mathbf{Z}_r \Delta_r^{-\frac{1}{2}} \mathbf{D} = \mathbf{T}_r \mathbf{D}$$

Likewise, $\mathbf{L}^\top \mathbf{V}_r = \mathbf{U}_r \Delta_r$, and $\tilde{\mathbf{L}}^\top \mathbf{V}_r = \tilde{\mathbf{U}}_r \Sigma_r$, which implies $\mathbf{L} \mathbf{U}_r \Delta_r = \mathbf{Q}_m \mathbf{V}_r = \Gamma_{\text{pr}}^{-1} \mathbf{V}_r \Delta_r^2$, and $\tilde{\mathbf{L}} \tilde{\mathbf{U}}_r \Sigma_r = \mathbf{H} \mathbf{V}_r = \Gamma_{\text{pr}}^{-1} \mathbf{V}_r \Sigma_r^2$ so that

$$\tilde{\mathbf{S}}_r^\top = \Sigma_r^{-\frac{1}{2}} \tilde{\mathbf{U}}_r^\top \tilde{\mathbf{L}}^\top = \Sigma_r^{\frac{1}{2}} \mathbf{V}_r^\top \Gamma_{\text{pr}}^{-1} = \mathbf{D}^{-1} \Delta_r^{\frac{1}{2}} \mathbf{V}_r^\top \Gamma_{\text{pr}}^{-1} = \mathbf{D}^{-1} \Delta_r^{-\frac{1}{2}} \mathbf{U}_r^\top \mathbf{L}^\top = \mathbf{D}^{-1} \mathbf{S}_r^\top$$

Thus, the two state space projections are equal: $\tilde{\mathbf{T}}_r \tilde{\mathbf{S}}_r^\top = \mathbf{T}_r \mathbf{S}_r^\top$ and the two reduced models are state-space equivalent: $\tilde{\mathbf{A}}_r = \mathbf{D}^{-1} \mathbf{A}_r \mathbf{D}$, $\tilde{\mathbf{B}}_r = \mathbf{D}^{-1} \mathbf{B}_r$, and $\tilde{\mathbf{C}}_r = \mathbf{C}_r \mathbf{D}$.

Proposition 3.2 tells us that when the BT-H and BT-Q generalized eigenproblems define the same dominant r directions, then guarantees applicable to the BT-Q approach will transfer to the BT-H approach as well. We note that this result relies crucially on the fact that both approaches use

\mathbf{I}_{pr} as the reachability Gramian – matching of the dominant subspaces is not sufficient to guarantee state-space equivalent reduced models if different reachability Gramians are used. We numerically compare the BT-Q and BT-H posterior approximations in Section 5; our results show that while the BT-H approach often leads to similar results as the BT-Q approach, the BT-Q approach is more robust.

4 Relationships between System-Theoretic and Bayesian Inference Gramians

The proposed approach in Section 3 hinges on two relationships between the canonical system-theoretic Gramians and the inference Gramians that we define in Section 3.1. We discuss now in more depth some theoretical and practical issues pertaining to these connections. In Section 4.1 we discuss Definition 3.1 for prior compatibility in more detail, and provide two strategies for constructing priors with the desired compatibility properties which are also suitable for the Bayesian inference task. In Section 4.2, we explore the relationship between the Fisher information matrix and system-theoretic observability Gramians, and describe two limiting settings in which the Fisher information matrix recovers the noisy observability Gramian up to a multiplicative constant.

4.1 Prior covariance compatibility

As briefly discussed in Section 3.1, if a prior covariance matrix \mathbf{I}_{pr} satisfies Definition 3.1, then there exists some $\mathbf{B} \in \mathbb{R}^{d \times d_{\text{in}}}$ for which \mathbf{I}_{pr} is exactly the reachability Gramian for the LTI system (2.9). We emphasize that not all covariance matrices will be prior-compatible with respect to dynamics determined by an independently assigned stable matrix, \mathbf{A} . One example of this takes $\mathbf{I}_{\text{pr}} = \mathbf{I}$ and \mathbf{A} to be any stable matrix whose numerical range extends into the right half plane, i.e., a matrix \mathbf{A} having eigenvalues in the complex open left half-plane but such that $\mathbf{A} + \mathbf{A}^T$ has at least one *positive* eigenvalue. With $d = 2$, for example, $\mathbf{A} = \begin{bmatrix} -1 & 3 \\ 0 & -2 \end{bmatrix}$ would have this property. We now present two strategies for formulating compatible prior covariances: Section 4.1.1 describes how a prior covariance may naturally be defined as the result of a spin-up process if an input port matrix is available. Section 4.1.2 describes how an incompatible prior covariance may be modified to ensure compatibility.

4.1.1 Spin-up

Given \mathbf{A} and a matrix $\mathbf{B} \in \mathbb{R}^{d \times d_{\text{in}}}$, a compatible prior may naturally be defined by stochastically exciting or “spinning up” the dynamical system (2.1) from $t = -\infty$ to $t = 0$ as follows:

$$d\mathbf{x} = \begin{cases} \mathbf{A}\mathbf{x} dt + \mathbf{B} d\mathbf{W}(t), & t < 0, \\ \mathbf{A}\mathbf{x} dt, & t \geq 0, \end{cases} \quad (4.1)$$

where the usual control-theoretic input [5, 23, 29] has been replaced by white noise. Here the noise is represented by differentials $d\mathbf{W}(t)$, where \mathbf{W} is a vector of dimension d_{in} containing independent copies of standard Brownian motion [30]. The matrix \mathbf{B} acts as the port matrix distributing the white noise driving onto the state components.

The stationary distribution of the system (4.1) during the stochastic driving period has statistics given in the following proposition.

Proposition 4.1 *At any finite time $t \leq 0$, the state $\mathbf{x}(t)$ governed by eq. (4.1) is normally distributed with constant mean and covariance given by*

$$\mathbb{E}[\mathbf{x}(t)] = 0, \quad \text{cov}(\mathbf{x}(t)) = \mathbb{E}[\mathbf{x}(t)\mathbf{x}^T(t)] = \int_0^\infty e^{\mathbf{A}\tau} \mathbf{B} \mathbf{B}^T e^{\mathbf{A}^T \tau} d\tau \quad (4.2)$$

Moreover, $\boldsymbol{\Gamma} \equiv \text{cov}(\boldsymbol{x}(0))$ is a solution to the Lyapunov equation, (2.11),

$$\mathbf{A}\boldsymbol{\Gamma} + \boldsymbol{\Gamma}\mathbf{A}^\top = -\mathbf{B}\mathbf{B}^\top. \quad (4.3)$$

We provide the proof in Appendix A. Equating this stationary distribution with the prior distribution (2.3) on $\boldsymbol{x}_0 = \boldsymbol{x}(0)$ gives the relations:

$$\boldsymbol{\mu}_{\text{pr}} = 0, \quad \boldsymbol{\Gamma}_{\text{pr}} = \int_0^\infty e^{\mathbf{A}t} \mathbf{B}\mathbf{B}^\top e^{\mathbf{A}^\top t} dt. \quad (4.4)$$

This is consistent with the practice of “spinning up” atmospheric models so that their starting state (at least before data from observations are incorporated) reflects a plausible “background covariance” [34,50]. Thus, eq. (4.4) provides a natural way of specifying a prior distribution satisfying Definition 3.1 if an input port matrix is specified. Indeed, [46] define the (infinite-time) reachability Gramian for a stochastic system such as eq. (4.1) as the stationary covariance of the state.

4.1.2 Compatibility modification

We now suppose that we are given a prior covariance $\boldsymbol{\Gamma}_{\text{pr}}$ that does *not* satisfy Definition 3.1 and provide a direct computational procedure that modifies the prior so that it is compatible. Suppose $\boldsymbol{\Gamma}_0$ is a given positive definite matrix that we take as an initial approximation for a prior covariance matrix that fails to be prior-compatible, i.e.,

$$\mathbf{M}_0 = \mathbf{A}\boldsymbol{\Gamma}_0 + \boldsymbol{\Gamma}_0\mathbf{A}^\top \not\preceq \mathbf{0},$$

so that, say, \mathbf{M}_0 is indefinite. If we write the spectral factorization of \mathbf{M}_0 as

$$\mathbf{M}_0 = [\mathbf{U}_+ \ \mathbf{U}_-] \begin{bmatrix} \Lambda_+ & \\ & \Lambda_- \end{bmatrix} \begin{bmatrix} \mathbf{U}_+^\top \\ \mathbf{U}_-^\top \end{bmatrix},$$

where $\Lambda_+ = \Omega_+^2$ contains the strictly positive eigenvalues of \mathbf{M}_0 while Λ_- contains the complementary nonpositive (i.e., negative and zero) eigenvalues of \mathbf{M}_0 . The *nearest* negative semidefinite matrix to \mathbf{M}_0 (with respect to either the spectral norm or the Frobenius norm) is $\mathbf{M}_- = \mathbf{U}_- \Lambda_- \mathbf{U}_-^\top$, and the modification to $\boldsymbol{\Gamma}_0$ that achieves this nearest negative semidefinite residual satisfies:

$$\mathbf{A}(\boldsymbol{\Gamma}_0 + \Delta) + (\boldsymbol{\Gamma}_0 + \Delta)\mathbf{A}^\top = \mathbf{U}_- \Lambda_- \mathbf{U}_-^\top \preceq \mathbf{0}$$

or equivalently,

$$\mathbf{A}\Delta + \Delta\mathbf{A}^\top + \mathbf{U}_+ \Lambda_+ \mathbf{U}_+^\top = 0, \quad (4.5)$$

which will have a unique positive definite solution, $\Delta = \mathbf{E}^\top \mathbf{E}$. Since $\mathbf{U}_+ \Lambda_+ \mathbf{U}_+^\top$ has an immediately accessible square root, namely $\mathbf{U}_+ \Omega_+$, the Cholesky factor, \mathbf{E} , of Δ can be computed directly using an approach introduced by Hammarling in [32]. This in particular allows one to compute a unique compatible prior $\boldsymbol{\Gamma}_{\text{pr}} := \boldsymbol{\Gamma}_0 + \Delta$. The modification, Δ , to $\boldsymbol{\Gamma}_0$ is *not* generally the smallest modification to $\boldsymbol{\Gamma}_0$ making it prior-compatible. Rather, it is the uniquely determined modification to $\boldsymbol{\Gamma}_0$ producing the nearest negative semidefinite residual. See Appendix D for a routine in MATLAB that implements this compatibility computation.

4.2 Observability Gramians and the Fisher information matrix

Our proposed approach is motivated by connections between the Fisher information matrix and system-theoretic observability Gramians. In contrast to the connection between the prior covariance and the reachability Gramian discussed in Section 4.1, we describe in Section 4.2.1 a fundamental obstacle to establishing a precise relation between the Fisher information matrix and an observability Gramian under a general given observation protocol. However, the Fisher information matrix \mathbf{H} can approach the noisy observability Gramian \mathbf{Q}_m , up to a scaling factor of sampling frequency, in certain limits. Note that this scaling factor of inverse time arises because \mathbf{Q}_m is a time integral and thus has an extra dimension of time when compared to \mathbf{H} , which is a sum. Section 4.2.2 considers the limit of continuous measurements and Section 4.2.3 considers infinite discrete-time measurements taken at times distributed according to a Poisson process. These approaches are related to asymptotics of covariance approximations by Fisher information matrices [1], and are also related to models of infill and domain asymptotics in statistics [18, 59, 60].

While infinite-horizon definitions (2.10) of system-theoretic Gramians are most commonly used, and indeed are required for the results of Lemmas 2.1 and 3.1 to hold, other Gramian definitions exist [2, 6, 31]. For example, one can define the Gramians over a finite time interval $[t_1, t_2]$, leading to time-limited Gramians. Furthermore, in practical balanced truncation settings, the infinite time-horizon Gramians are sometimes approximated via numerical quadrature [13, 16, 40, 48, 57] or via approximate solution methods to the Lyapunov equations that yield low-rank Gramian approximations [11, 36, 51]. In this context, the scaled Fisher information matrix can be viewed as an approximation to the infinite noisy Gramian of Definition 3.2, see, e.g., [45, 49]. When \mathbf{H} is a sufficiently close approximation to \mathbf{Q}_m so that their dominant generalized eigenpairs with $\mathbf{\Gamma}_{\text{pr}}^{-1}$ coincide, then Proposition 3.1 applies.

4.2.1 Incompatibility of the Fisher information matrix

From formal duality, one could formulate a *likelihood-compatibility* condition, in analogy to Definition 3.1 for the prior covariance, by saying a Fisher information matrix \mathbf{H} is *likelihood-compatible* with state dynamics induced by \mathbf{A} , if $\mathbf{H}\mathbf{A} + \mathbf{A}^\top\mathbf{H}$ is negative semidefinite. That would imply the existence of some $\mathbf{C} \in \mathbb{R}^{d_{\text{out}} \times d}$ such that \mathbf{H} solves a dual Lyapunov equation (2.11) $\mathbf{A}^\top\mathbf{H} + \mathbf{H}\mathbf{A} = -\mathbf{C}^\top\mathbf{C}$, and is thus the infinite-horizon observability Gramian of an LTI dynamical system (2.9) with output port matrix \mathbf{C} . Replacing the noisy observability Gramian \mathbf{Q}_m with a likelihood-compatible \mathbf{H} in the BT-Q approach of Section 3.2 (which requires a compatible prior) would then yield a reduced model that inherits the theoretical guarantees of Lemma 3.1 and Proposition 3.1.

Unfortunately, Fisher information matrices for observation models that occur in plausible inference settings will be extremely unlikely to be likelihood-compatible. In particular, the Fisher information as defined in (2.8), $\mathbf{H} = \sum_{i=1}^n e^{\mathbf{A}^\top t_i} \mathbf{C}^T \mathbf{\Gamma}_{\text{obs}}^{-1} \mathbf{C} e^{\mathbf{A} t_i}$, is often rank-deficient, and in that case in order for \mathbf{H} also to be likelihood compatible with respect to the dynamics determined by \mathbf{A} , it is necessary for $\ker(\mathbf{H})$ to be an invariant subspace for \mathbf{A} , which is extraordinarily unlikely to occur with measurement models specified in practice. One may include adjustments to the Fisher information intended to enforce likelihood-compatibility (analogously to the prior compatibility modification of Section 4.1.2) however such adjustments appear difficult to justify since they would represent significant *ad hoc* changes to the observation protocol that at best would be unrealistic.

4.2.2 Continuum Limit of Regularly Spaced Observations

We proceed next to show that despite the frustrated equivalence between the Fisher information matrix and a continuous-time observability Gramian under general observation protocols, we can establish a formal proportionality between these quantities in an idealized continuum limit of finely

spaced perpetual observations. Suppose the measurement times $\{t_i\}_{i=1}^n$, with $n \geq 2$, in the definition of the Fisher information matrix (2.8) are equispaced in the interval $[t_{\text{start}}, t_{\text{end}}]$, i.e.,

$$t_i := t_{\text{start}} + (i-1)h, \quad h := \frac{t_{\text{end}} - t_{\text{start}}}{n-1}, \quad (4.6)$$

for $i = 1, \dots, n$. If we scale the Fisher information matrix by the sampling time interval h , then the large- n energy limits to an integral (we make the dependence of \mathbf{H} on n explicit here):

$$\hat{\mathbf{H}}_\infty := \lim_{n \rightarrow \infty} h \mathbf{H}_n = \lim_{n \rightarrow \infty} h \sum_{i=1}^n e^{\mathbf{A}^\top t_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_i} = \int_{t_{\text{start}}}^{t_{\text{end}}} e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} dt. \quad (4.7)$$

In this finite observation interval setting, $\hat{\mathbf{H}}_\infty$ satisfies the modified Lyapunov equation,

$$\mathbf{A}^\top \hat{\mathbf{H}}_\infty + \hat{\mathbf{H}}_\infty \mathbf{A} = e^{\mathbf{A}^\top t_{\text{end}}} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_{\text{end}}} - e^{\mathbf{A}^\top t_{\text{start}}} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_{\text{start}}}. \quad (4.8)$$

We recover the noisy infinite observability Gramian of Definition 3.2 if we take the observation start time $t_{\text{start}} \downarrow 0$ and observation end time $t_{\text{end}} \uparrow \infty$ together with the finer spacing h of the observations, provided the normalization of the energy scales appropriately with the sampling interval.

Proposition 4.2 *Let a stable matrix \mathbf{A} and an output matrix \mathbf{C} be given. Define a collection of observation times $\{t_i\}_{i=1}^n$ by $t_i = (i-1)h(n)$ where $h(n) \equiv \frac{T(n)}{n-1}$ with the observation time interval $T(n)$ increasing sublinearly with the number of observations so that $T(n) \rightarrow \infty$ while $h(n) = \frac{T(n)}{n-1} \rightarrow 0$ as $n \rightarrow \infty$. Then,*

$$\lim_{n \rightarrow \infty} h(n) \sum_{i=1}^n e^{\mathbf{A}^\top t_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_i} = \int_0^\infty e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} dt. \quad (4.9)$$

Proof The integrand on the right hand side and the summand on the left hand side are continuous matrix functions whose norm is bounded and decays exponentially. Truncating both the sum and the integral to times $0 \leq t, t_i \leq \tau$ leads to a Riemann sum convergence on this truncated interval. The truncated component $t_i, t > \tau$ can be bounded in norm by a function of the form $k_1 e^{-k_2 \tau}$ for positive constants k_1 and k_2 depending on the norm of \mathbf{C} , the norm bound on $\boldsymbol{\Gamma}_\epsilon^{-1}$, and the stability properties of the matrix \mathbf{A} . By taking $\tau \uparrow \infty$, we achieve the stated result. \square

4.2.3 Observation Models at Random Times

Observability Gramians with an integral form can also be associated to observation models that take place at discrete, but randomly distributed, moments of time. For example, the observation times in a field experiment might not be perfectly predictable due to considerations of weather, availability of human resources, and/or availability of equipment. We do not propose randomly distributed observation times as an important practical innovation, but rather to provide another means of interpreting the relation between the observability Gramian and the Fisher information.

As in the previous section, we specify a time t_{start} when observations begin and a time t_{end} when observations end. If we have n observations to be taken at independent and uniformly distributed times over this interval $t_i \sim U(t_{\text{start}}, t_{\text{end}})$, we would have a random Fisher information due to the randomness in the t_i . However, if we average over the realizations of the uniformly distributed observation times, we would obtain the mean Fisher information:

$$\begin{aligned} \bar{\mathbf{H}} &\equiv (t_{\text{end}} - t_{\text{start}})^{-n} \int_{t_{\text{start}}}^{t_{\text{end}}} \int_{t_{\text{start}}}^{t_{\text{end}}} \dots \int_{t_{\text{start}}}^{t_{\text{end}}} \sum_{i=1}^n e^{\mathbf{A}^\top t'_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t'_i} dt'_1 dt'_2 \dots dt'_n \\ &= \frac{n}{t_{\text{end}} - t_{\text{start}}} \int_{t_{\text{start}}}^{t_{\text{end}}} e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} dt. \end{aligned}$$

The same mean Fisher information would result if we divided the observation period into n equal subintervals of length $h = (t_{\text{end}} - t_{\text{start}})/n$, and took one observation at a time t_i that is uniformly distributed within each size- h subinterval:

$$\begin{aligned}\bar{\mathbf{H}} &= \sum_{i=1}^n \frac{1}{h} \int_{t_{\text{start}}+(i-1)h/n}^{t_{\text{start}}+ih/n} e^{\mathbf{A}^\top t'_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t'_i} dt'_i \\ &= \frac{1}{h} \int_{t_{\text{start}}}^{t_{\text{end}}} e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} dt.\end{aligned}\quad (4.10)$$

Notice we have n subintervals in this formulation, as compared to the $n - 1$ subintervals in Section 4.2.2 because we have here n observations distributed in the interior, rather than the endpoints, of the subintervals.

To obtain the standard infinite horizon observability Gramian (2.10) under a similar paradigm, we consider a countably infinite sequence of observations at randomly distributed times ($n \rightarrow \infty$). Under the appropriate limiting models of random measurements we can recover, in expectation, an infinite-horizon Gramian multiplied by the sampling frequency.

Proposition 4.3 *Let a stable matrix \mathbf{A} and an output matrix \mathbf{C} be given. Let $\{t_i\}_{i \geq 1}$ be a random sequence of times defined by either of the following models:*

- (a) *For some $h > 0$, $t_i \sim U((i-1)h, ih)$.*
- (b) *For some $h > 0$, the $\{t_i\}_{i \geq 1}$ are the realization of a Poisson point process on $[0, \infty)$ with intensity h^{-1} .*

Then, as $n \uparrow \infty$, the expected Fisher information (4.10) limits to an infinite-time observability Gramian scaled by the mean sampling frequency:

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\sum_{i=1}^n e^{\mathbf{A}^\top t_i} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t_i} \right] = h^{-1} \int_0^\infty e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} dt. \quad (4.11)$$

Proof Define

$$f(t) := \boldsymbol{\xi}^\top \cdot e^{\mathbf{A}^\top t} \mathbf{C}^\top \boldsymbol{\Gamma}_\epsilon^{-1} \mathbf{C} e^{\mathbf{A} t} \cdot \boldsymbol{\xi}.$$

where $\boldsymbol{\xi} \in \mathbb{R}^d$. Under model (a), a similar computation as in eq. (4.10) yields,

$$\mathbb{E} \left[\sum_{i=1}^n f(t_i) \right] = h^{-1} \int_0^{hn} f(t) dt. \quad (4.12)$$

Since f is non-negative, continuous for all $t \geq 0$, and decays exponentially for large t , we also have,

$$\lim_{n \rightarrow \infty} \int_0^{hn} f(t) dt = \int_0^\infty f(t) dt. \quad (4.13)$$

As $\boldsymbol{\xi} \in \mathbb{R}^d$ was arbitrary, we obtain the result (4.11). Under model (b), note that f is integrable for $t \geq 0$. Then Campbell's theorem for Poisson processes asserts that

$$\mathbb{E} \left[\sum_{i=1}^\infty f(t_i) \right] = h^{-1} \int_0^\infty f(t) dt.$$

This, along with,

$$\mathbb{E} \left[\lim_{n \rightarrow \infty} \sum_{i=1}^n f(t_i) \right] = \lim_{n \rightarrow \infty} \mathbb{E} \left[\sum_{i=1}^n f(t_i) \right], \quad (4.14)$$

which follows from the monotone convergence theorem, and again the arbitrary choice of $\boldsymbol{\xi} \in \mathbb{R}^d$, proves the desired result. \square

Remark. The standard observability Gramian (2.10) would correspond to an infinite sequence of observations at unit frequency, either according to a standard Poisson point process or a randomized sampling within regularly spaced intervals. The above result also implies that the large- n expectation of the Fisher information matrix (2.8) equals the (scaled) infinite-time observability Gramian, assuming the observations t_i are distributed according to one of the models in Proposition 4.3.

5 Numerical experiments

We now demonstrate the balanced truncation approaches on two test problems based on LTI systems that are benchmarks within the system-theoretic model reduction community¹: a heat equation example and a structural model for the ISS1R module. For each LTI system, the inference problem is to infer the initial system state from output measurements made after $t = 0$ corrupted by Gaussian noise. For each system, we will use the system dynamics to determine a compatible prior covariance matrix as defined in Definition 3.1. Then, we will compare the posterior mean and covariance approximation performance of three approaches:

1. from [54], the optimal low-rank update (**OLRU**) approximations of the posterior covariance (2.26) and mean (2.33), and the optimal low-rank (**OLR**) posterior mean approximation,
2. the **BT-Q** posterior mean and covariance approximation of Section 3.2 based on the pencil $(\mathbf{Q}_m, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$, and
3. the **BT-H** posterior mean and covariance approximation of Section 3.2 based on the pencil $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$.

For both LTI systems, we consider both a ‘good’ observation model where the final observation time, T , and number of measurements, n , are both large, as well as a ‘bad’ observation model where T and n are both small. Sections 5.1 and 5.2 describe the heat equation and ISS1R problems, respectively, and present numerical results for the aforementioned posterior approximations. Section 5.3 discusses the numerical results from both examples.

5.1 Heat equation

We consider the heat equation in a one-dimensional rod with homogeneous Dirichlet boundary conditions. The state and output dimensions are $d = 200$ and $d_{\text{out}} = 1$, respectively, and the single output is the temperature measured at $2/3$ of the rod length. We define the prior covariance $\boldsymbol{\Gamma}_{\text{pr}}$ to be the infinite reachability Gramian (2.10) that would result from taking $\mathbf{B} = \mathbf{I}$. The true initial condition is drawn from $\mathcal{N}(0, \boldsymbol{\Gamma}_{\text{pr}})$ and used to generate noisy output measurements at $t_i = ih$, for $i = 1, 2, \dots, n$. The noise values at each measurement are independently and identically distributed: $\boldsymbol{\Gamma}_{\epsilon_i} = \boldsymbol{\Gamma}_{\epsilon} = \sigma_{\text{obs}}^2 \mathbf{I}$, where the noise standard deviation is chosen to be roughly 10% of maximum magnitude of the noiseless output, $\sigma_{\text{obs}} = 0.008$.

We consider in Figure 5.1 an observation model with $n = 5 \times 10^5$ measurements spaced $h = 10^{-4}$ apart, for a final observation time $T = 50$. The spectral abscissa of \mathbf{A} for this system is ≈ -0.1 , so this T value is about five times the slowest characteristic timescale. We therefore expect that the finely spaced observations in this model can capture most of the system dynamics and approximate the infinite integral. Indeed, the relative Frobenius norm difference between $h\mathbf{H}$ and \mathbf{Q} is 0.1% in this case. In contrast, Figure 5.2 considers a limited observation model with $n = 100$ measurements spaced $h = 0.1$ apart, for a final observation time of $T = 10$. In this case the relative Frobenius norm difference between $h\mathbf{H}$ and \mathbf{Q} is 15%.

¹ LTI system matrices and documentation for both examples can be found online at <http://slicot.org/20-site/126-benchmark-examples-for-model-reduction>

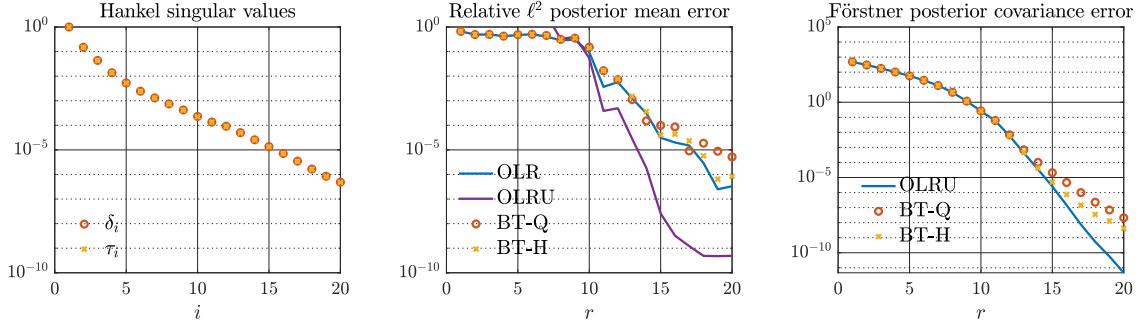


Fig. 5.1: Results for the heat equation model with measurements spaced $h = 10^{-4}$ apart until $T = 50$, for which Q_m and hH have a relative error of 0.1%. In the left panel we plot τ_i and δ_i , the square roots of the generalized eigenvalues of the pencils $(H, \Gamma_{\text{pr}}^{-1})$ and $(Q_m, \Gamma_{\text{pr}}^{-1})$, respectively. These are normalized relative to δ_1, τ_1 , respectively, and coincide with the Hankel singular values. In the middle and right panels, we compare respectively the relative errors of the posterior means of the various low-rank approximation approaches in the ℓ^2 norm, and the errors of the posterior covariances in the Förstner distance.

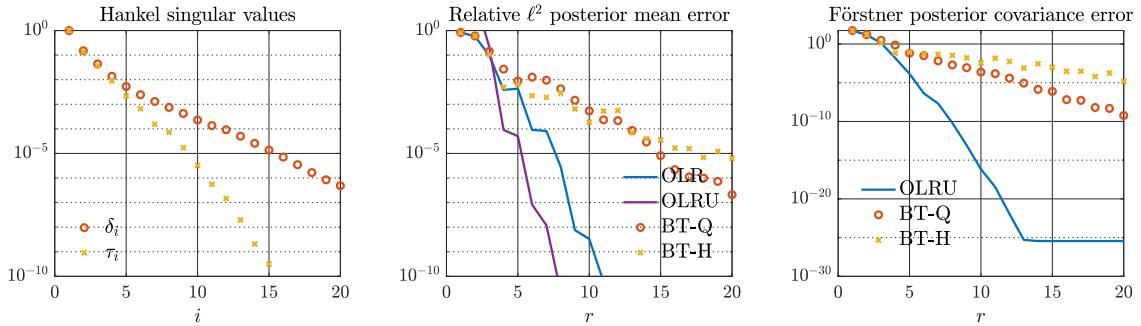


Fig. 5.2: Results for the heat equation model with measurements spaced $h = 10^{-1}$ apart until $T = 10$, for which Q_m and hH have a relative error of 15%. The quantities plotted are as defined in Figure 5.1

5.2 ISS1R module

This benchmark LTI system is a structural model for the flex modes of the Zvezna service module of the International Space Station [2]. The state and output dimensions are $d = 270$ and $d_{\text{out}} = 3$, respectively; the outputs correspond to roll, pitch, and yaw gyroscope measurements. An input matrix B with $d_{\text{in}} = 3$ is provided as part of the model that corresponds to roll, pitch, and yaw jet controls. This B is used to generate a compatible prior Γ_{pr} . The true initial condition is drawn from $\mathcal{N}(0, \Gamma_{\text{pr}})$ and used to generate noisy output measurements. The noise covariance $\Gamma_{\epsilon} = \text{diag}\{0.0025^2, 0.0005^2, 0.0005^2\}$ is set so that the noise standard deviation of each output is roughly 10% of the maximum output signal as before.

Figure 5.3 considers an observation model with $n = 3000$ measurements spaced $h = 0.1$ apart for a final observation time $T = 300$. In this case, the relative difference between hH and Q_m is 1%. Figure 5.4 considers $n = 10$ measurements spaced $h = 1$ apart for a final observation time $T = 10$. In this case, the relative difference between hH and Q_m is 53%.

5.3 Discussion

In both test problems, when the observation model has a large number of observations, n , with high temporal resolution hH is a close approximation of Q_m in the Frobenius norm (Figures 5.1 and 5.3). In this case we observe good agreement in τ_i and δ_i , the square roots of the generalized eigenvalues of

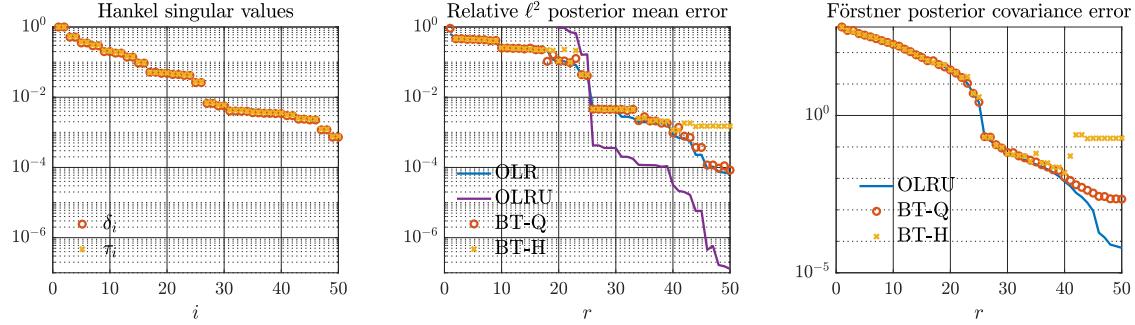


Fig. 5.3: Results for the ISS1R model with measurements spaced $h = 0.1$ apart until $T = 300$, for which \mathbf{Q}_m and $h\mathbf{H}$ have a relative error of 1%. The quantities plotted are as defined in Figure 5.1

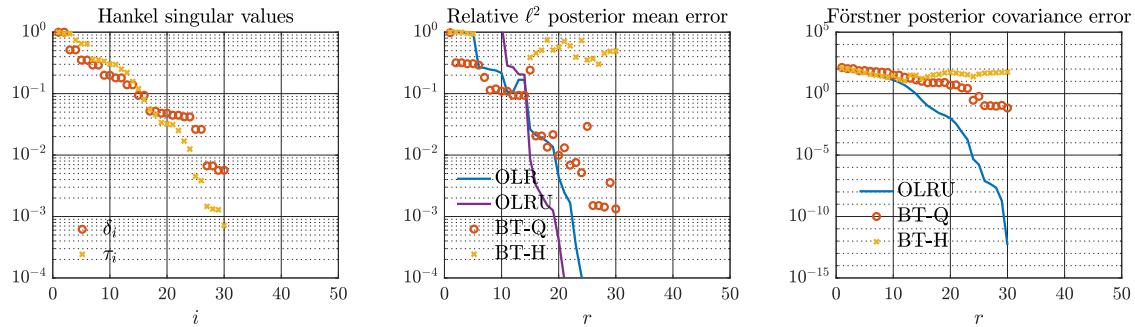


Fig. 5.4: Results for the ISS1R model with measurements spaced $h = 1$ apart until $T = 10$, for which \mathbf{Q}_m and $h\mathbf{H}$ have a relative error of 53%. Because $r = 30 = nd_{\text{out}}$, the OLR/U approximations correspond to the full posterior in this case. The quantities plotted are as defined in Figure 5.1

the pencils $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$ and $(\mathbf{Q}_m, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$, respectively. The posterior covariance approximations using both BT approaches are optimal for low-to-middling values of r , indicating agreement not only in the generalized eigenvalues but also in the leading generalized eigendirections. At higher r , the balanced truncation posterior covariance approximations diverge from the OLR/U approximation, indicating that the trailing eigendirections do not match as well, but the balanced truncation approximations exhibit respectable errors nevertheless. In contrast, when the observation model has smaller n , with lower temporal resolution of observations (Figures 5.2 and 5.4) and shorter observation intervals, we observe poor agreement in δ_i and τ_i and significantly sub-optimal posterior covariance approximations from balanced truncation approaches for all but the smallest r -values. This indicates generally poor agreement in the generalized eigenpairs of $(\mathbf{Q}_m, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$ and $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$, in all but the leading few directions.

Across all test cases, the BT-Q mean error exhibits similar decay to that of the square root of the generalized pencil eigenvalues δ_i (these are the Hankel singular values of the system (3.8)). We note that while it may be possible to bound the posterior mean error using system-theoretic results, the error bound of Lemma 3.1 does not automatically apply to the posterior mean. Showing that this or a similar bound holds is a direction for future work. In contrast, while the BT-H mean error performs similarly to the BT-Q mean error when $h\mathbf{H}$ and \mathbf{Q}_m are close; in cases where $h\mathbf{H}$ is far from \mathbf{Q}_m the BT-H error is higher, significantly so in the ISS1R case of Figure 5.4. We note that while both the OLR/U mean approximations can achieve lower errors than the balanced truncation approach, the OLR/U mean approximations require evolving the full adjoint dynamics in the dynamical system setting, while the balanced truncation approaches provide a low-dimensional adjoint to evolve. Finally, we observe that in some cases (particularly in Figure 5.4) the BT-Q mean error is smaller than that of the OLR approach. This does not contradict the optimality results of [54],

as the OLR mean is optimal in the Bayes risk, not the ℓ^2 -norm in which we compare posterior mean approximations.

Our results show that our proposed approach of Section 3.2 provides an efficient low-dimensional reduced model that can be used to compute posterior mean and covariance approximations that are near-optimal when the observation model allows the $(\mathbf{H}, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$ eigenpairs to closely approximate the $(\mathbf{Q}_m, \boldsymbol{\Gamma}_{\text{pr}}^{-1})$ eigenpairs. However, our results show that even when these eigenpairs are far apart, the BT-Q approach of Section 3.2 provides approximations with reasonably low (albeit sub-optimal) error.

6 Conclusions

We consider the linear Gaussian Bayesian inverse problem for inferring the initial condition of a high-dimensional linear dynamical system from noisy measurements taken at $t > 0$. In this setting, we propose a new model reduction method that exploits natural connections between the system-theoretic method of balanced truncation and the optimal posterior approximation work of [54]. To adapt balanced truncation to the inference setting, we define new Gramians for the inference context. We introduce a notion of *prior compatibility* that allows the prior covariance to be interpreted as an infinite reachability Gramian, and we propose two ways of constructing a compatible prior: (i) as the stationary covariance of a system that has been spun-up for $t < 0$ by white noise and (ii) via a modification to an incompatible prior. We also define a *noisy observability Gramian* which measures observability energy in a way that accounts for the additive noise in the inference measurements and show that this Gramian coincides with the (scaled) Fisher information matrix in two limits: (i) the limit of continuous observations and (ii) the expectation of countably infinite Poisson-distributed measurement times.

Our proposed approach uses the prior covariance and noisy observability Gramian as inference Gramians in the method of balanced truncation to obtain a reduced model that inherits system-theoretic stability and error guarantees. Additionally, when the dominant generalized eigenpairs of the pencil defined by these inference Gramians match those of the pencil defined by the Fisher information matrix and the prior covariance, our reduced model can additionally recover the optimal posterior covariance approximation of [54]. Our model yields a reduced dynamical system which can be cheaply evolved to approximate the full system: in this way, our approach extends the work of [54], which does not automatically yield a reduced dynamical system. However, the posterior mean approximations considered in [54] and in our approach are different in general.

Numerical experiments on two benchmark problems from the model reduction community demonstrate that our proposed approach yields near-optimal posterior mean *and* covariance approximations when the noisy observability Gramian is close to the scaled Fisher information matrix. We demonstrate theoretically and by numerical examples that these matrices are close when the observations are of suitable frequency and duration. When these two matrices differ significantly, the proposed approach yields sub-optimal posterior mean and covariance approximations that nevertheless exhibit low errors that may be acceptable in some applications.

The connections between linear Gaussian Bayesian inference problems and model reduction via balanced truncation which we have established here are intended as a potential means for transferring perspectives between the two research communities. Some open questions about the proposed method are directions for future work, including: how to effectively leverage the proposed reduced model for computational gains in data assimilation applications; analysis of the posterior mean approximation; and extension to nonlinear settings. Another natural extension of the proposed approach is to Bayesian inference of the initial state of a system driven stochastically for all time, including through the observation period, which is observed at possibly countably infinite and possibly random times. See Appendix C for a discussion of this problem.

A Proof of Proposition 4.1

Proof The equation 4.1 can be solved by integrating factors in the standard way because the stochastic product rule applied to $d(e^{\mathbf{A}t}x)$ produces no corrections:

$$\mathbf{x}(t) = \int_{-\infty}^t e^{\mathbf{A}(t-s)} \mathbf{B} d\mathbf{W}(s)$$

With the Ito isometry, which is equivalent to the formal rule $\mathbb{E}[d\mathbf{W}(s)d\mathbf{W}(s')] = I\delta(s-s')dsds'$, we compute:

$$\begin{aligned} \mathbb{E}[\mathbf{x}(t)\mathbf{x}^\top(t)] &= \int_{-\infty}^t e^{\mathbf{A}(t-s)} \mathbf{B} (e^{\mathbf{A}(t-s)} \mathbf{B})^\top ds = \int_{-\infty}^t e^{\mathbf{A}(t-s)} \mathbf{B} \mathbf{B}^\top e^{\mathbf{A}^\top(t-s)} ds \\ &= \int_0^\infty e^{\mathbf{A}s} \mathbf{B} \mathbf{B}^\top e^{\mathbf{A}^\top s} ds, \end{aligned} \quad (\text{A.1})$$

which also shows that $d\mathbb{E}[\mathbf{x}\mathbf{x}^\top] = 0$. The method of moments gives:

$$\begin{aligned} d(\mathbf{x}\mathbf{x}^\top) &= (\mathbf{A}\mathbf{x} dt + \mathbf{B} d\mathbf{W}(t))\mathbf{x}^\top + \mathbf{x}(\mathbf{x}^\top \mathbf{A}^\top + d\mathbf{W}^\top(t)\mathbf{B}^\top) + \mathbf{B} d\mathbf{W}(t)\mathbf{d}\mathbf{W}(t)^\top \mathbf{B}^\top \\ &= \mathbf{A}\mathbf{x}\mathbf{x}^\top dt + \mathbf{x}\mathbf{x}^\top \mathbf{A}^\top dt + \mathbf{B} d\mathbf{W}(t)\mathbf{x}^\top + \mathbf{x} d\mathbf{W}(t)^\top \mathbf{B}^\top + \mathbf{B}\mathbf{B}^\top dt \end{aligned} \quad (\text{A.2})$$

Taking expectations and using $\mathbb{E}[d\mathbf{W}\mathbf{x}] = 0$, we recover the desired Lyapunov equation 4.3:

$$d\mathbb{E}[\mathbf{x}\mathbf{x}^\top] = 0 = (\mathbf{A}\mathbb{E}[\mathbf{x}\mathbf{x}^\top] + \mathbb{E}[\mathbf{x}\mathbf{x}^\top]\mathbf{A}^\top + \mathbf{B}\mathbf{B}^\top)dt.$$

□

B Optimality Property of Mahalanobis Distance for Gaussian Distributions

Proposition B.1 *The neighborhood $M \equiv \{z \in \mathbb{R}^d : D(z, \mathcal{N}(\mu, \Sigma)) < \delta\}$ defined by the Mahalanobis distance 3.1 satisfies the following optimality condition:*

$$M = \operatorname{argmax}_{S \in \mathcal{O}: \lambda(S) = \lambda(M)} \mathbb{P}(S).$$

where \mathbb{P} denotes the probability measure on \mathbb{R}^d corresponding to a multivariate Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$, \mathcal{O} is the collection of open sets in \mathbb{R}^d and $\lambda(S)$ denotes the Lebesgue measure of set S .

Proof M can be verified by definition to correspond to an ellipsoidal level set of the multivariate Gaussian density, with the probability density strictly greater on the interior of M than on the complement. Consequently, for any $S \in \mathcal{O}$, the probability density on $S \setminus M$ is strictly greater than the probability density on $M \setminus S$. If M and S have the same Lebesgue measure, then so do $S \setminus M$ and $M \setminus S$. Thus $P(S \setminus M) \leq P(M \setminus S)$, with equality holding only if $P(S \setminus M) = 0$, which can only happen if $S = M$ since both sets are open. □

C Stochastic Forcing Continuing Through Positive Times

We consider here how the connection between balanced truncation and the optimal low-rank posterior update procedure described in 54 is affected if we were to consider the more natural case of a stochastic linear dynamical system driving consistently for both positive and negative times:

$$d\mathbf{x} = \mathbf{A}\mathbf{x} dt + \mathbf{B} d\mathbf{W}(t), \quad -\infty < t < \infty. \quad (\text{C.1})$$

Now the observations involve noise not only from the measurement process, but from the dynamical noise at previous times. We can write in place of Eq. 2.2:

$$\mathbf{m}_i = \mathbf{C}e^{\mathbf{A}t_i} \mathbf{x}_0 + \boldsymbol{\Xi}_i,$$

where

$$\boldsymbol{\Xi}_i = \int_0^{t_i} \mathbf{C}e^{\mathbf{A}(t_i-s)} \mathbf{B} d\mathbf{W}(s) + \boldsymbol{\epsilon}_i.$$

In the expression of the abstract linear Bayesian inference problem:

$$\mathbf{m} = \mathbf{G}\mathbf{x}_0 + \boldsymbol{\Xi},$$

the noisy component $\boldsymbol{\Xi} \in \mathbb{R}^{d_{\text{obs}}}$ of the measurements is now correlated across observation times. Thus, $\boldsymbol{\Xi} \sim \mathcal{N}(0, \boldsymbol{\Gamma}_{\text{tot}})$ where $\boldsymbol{\Gamma}_{\text{tot}} \in \mathbb{R}^{d_{\text{obs}} \times d_{\text{obs}}}$ has block structure:

$$\boldsymbol{\Gamma}_{\text{obs}} = \begin{bmatrix} \boldsymbol{\Gamma}_{\epsilon_1, \epsilon_1} & \boldsymbol{\Gamma}_{\epsilon_1, \epsilon_2} & \cdots & \boldsymbol{\Gamma}_{\epsilon_1, \epsilon_n} \\ \boldsymbol{\Gamma}_{\epsilon_2, \epsilon_1} & \boldsymbol{\Gamma}_{\epsilon_2, \epsilon_2} & \cdots & \boldsymbol{\Gamma}_{\epsilon_2, \epsilon_n} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{\Gamma}_{\epsilon_n, \epsilon_1} & \boldsymbol{\Gamma}_{\epsilon_n, \epsilon_2} & \cdots & \boldsymbol{\Gamma}_{\epsilon_n, \epsilon_n} \end{bmatrix}. \quad (\text{C.2})$$

with covariances of measurements at different times given explicitly by;

$$\boldsymbol{\Gamma}_{\epsilon_i, \epsilon_j} = \boldsymbol{\Gamma}_{\epsilon_i} \delta_{ij} + \int_{|t_i - t_j|}^{\max(t_i, t_j)} \mathbf{C} e^{\mathbf{A}s} \mathbf{B} \mathbf{B}^\top e^{\mathbf{A}^\top s} \mathbf{C}^\top ds$$

with Kronecker delta function δ_{ij} . The reason for the correlations is the influence of the common past driving noise. These correlations are indeed mitigated if the observation times are spaced sufficiently far apart so the minimal eigenvalue of $-\mathbf{A}|t_i - t_j|$ becomes large, but this is not a terribly interesting case as it would imply also weak dependence of the observations on the state \mathbf{x}_0 . The correlation in the effective measurement noise would lead to a very different expression for the Fisher information matrix in Eq. (2.8), as $\boldsymbol{\Gamma}_{\text{obs}}^{-1}$ would no longer be block diagonal, and we would thereby lose interpretability of the Fisher information matrix as a discretized version of an observability Gramian.

D Modification of covariance to induce prior-compatibility

This appendix complements the discussion in Section 4.1.2. We provide below a Matlab routine that, given a system matrix \mathbf{A} and a prior covariance $\mathbf{Gamma0}$, returns the upper-triangular Cholesky factor \mathbf{R}_Gam of a modified prior that is compatible with the state dynamics of \mathbf{A} . Our MATLAB function listed below utilizes the `lyapchol` routine from the Control System Toolbox to solve (4.5) for the Cholesky factor \mathbf{E} of the prior modification Δ .

```
function R_Gam=minmodPriorCompat(Gamma0,A)
% Find a minimal residual modification producing
% a modified Gamma0 that is prior-compatible with A
% Input: Gamma0 - initial covariance matrix (n x n positive def matrix)
%         A - dynamic system matrix (n x n strictly stable matrix)
% Returns: R_Gam - the Cholesky factor of a modification of Gamma0
% making it compatible with A: Gamma1 = R_Gam'*R_Gam = Gamma0 + E'*E
% such that A*Gamma1+Gamma1*A' is the *nearest* negative semidefinite
% matrix to A*Gamma0+Gamma0*A' - Note that E'*E is *not* the minimal
% perturbation to Gamma0 required for prior-compatibility.
%
% (Requires lyapchol from control system toolbox).
%
R0=chol(Gamma0);
M0=A*Gamma0+Gamma0*A';
[V,D]=eig(M0); eigval=real(diag(D));
flagPos=(eigval>0);
% If Gamma0 is prior-compatible then flagPos is empty
if isempty(flagPos)
    R_Gam=R0;
else
    d_pos=sqrt(eigval(flagPos));
    V_pos=V(:,flagPos);
    E = lyapchol(A,V_pos*diag(d_pos));
    [~,R_Gam]=qr([R0;E],0);
end
```

References

1. Abt, M., Welch, W.J.: Fisher information and maximum-likelihood estimation of covariance parameters in Gaussian stochastic processes. *Canadian Journal of Statistics* **26**(1), 127–137 (1998). DOI 10.2307/3315678. URL <https://onlinelibrary.wiley.com/doi/abs/10.2307/3315678>
2. Antoulas, A.C.: Approximation of large-scale dynamical systems. SIAM (2005)

3. Antoulas, A.C., Beattie, C.A., Gugercin, S.: Interpolatory Methods for Model Reduction. Computational Science & Engineering. Society for Industrial and Applied Mathematics (2020). DOI 10.1137/1.9781611976083. URL <https://pubs.siam.org/doi/book/10.1137/1.9781611976083>
4. Apte, A., Hairer, M., Stuart, A.M., Voss, J.: Sampling the posterior: An approach to non-Gaussian data assimilation. *Physica D: Nonlinear Phenomena* **230**(1-2), 50–64 (2007)
5. Baur, U., Benner, P., Feng, L.: Model order reduction for linear and nonlinear systems: a system-theoretic perspective. *Archives of Computational Methods in Engineering* **21**(4), 331–358 (2014)
6. Benner, P., Breiten, T.: Chapter 6: Model Order Reduction Based on System Balancing, pp. 261–295
7. Benner, P., Goyal, P.: Balanced truncation model order reduction for quadratic-bilinear control systems. arXiv preprint arXiv:1705.00160 (2017)
8. Benner, P., Gugercin, S., Willcox, K.: A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. *SIAM Review* **57**(4), 483–531 (2015). DOI 10.1137/130932715. URL <http://pubs.siam.org/doi/abs/10.1137/130932715>
9. Benner, P., Qiu, Y., Stoll, M.: Low-Rank Eigenvector Compression of Posterior Covariance Matrices for Linear Gaussian Inverse Problems. *SIAM/ASA Journal on Uncertainty Quantification* **6**(2), 965–989 (2018). DOI 10.1137/17M1121342. URL <http://pubs.siam.org/doi/abs/10.1137/17M1121342>. Publisher: Society for Industrial and Applied Mathematics
10. Benner, P., Saak, J.: Numerical solution of large and sparse continuous time algebraic matrix Riccati and Lyapunov equations: a state of the art survey. *GAMM-Mitteilungen* **36**(1), 32–52 (2013)
11. Benner, P., Saak, J.: Numerical solution of large and sparse continuous time algebraic matrix Riccati and Lyapunov equations: A state of the art survey. *GAMM Mitteilungen* **36**(1), 32–52 (2013)
12. Bennett, A.F.: Inverse Modeling of the Ocean and Atmosphere. Cambridge University Press (2005)
13. Bertram, C., Faßbender, H.: Lyapunov and Sylvester equations: A quadrature framework. Tech. rep., arXiv preprint available at <https://arxiv.org/abs/1903.05383> (2020)
14. Bhatia, R.: Positive Definite Matrices. Princeton University Press (2007)
15. Bonnabel, S., Sepulchre, R.: Riemannian Metric and Geometric Mean for Positive Semidefinite Matrices of Fixed Rank. *SIAM Journal on Matrix Analysis and Applications* **31**(3), 1055–1070 (2010). DOI 10.1137/080731347. URL <https://pubs.siam.org/doi/abs/10.1137/080731347>
16. Breiten, T.: Structure-preserving model reduction for integro-differential equations. *SIAM Journal on Control and Optimization* **54**(6), 2992–3015 (2016)
17. Calvetti, D., Somersalo, E.: An Introduction to Bayesian Scientific Computing: Ten Lectures on Subjective Computing. Surveys and Tutorials in the Applied Mathematical Sciences. Springer-Verlag, New York (2007). DOI 10.1007/978-0-387-73394-4
18. Chen, H.S., Simpson, D.G., Ying, Z.: INFILL ASYMPTOTICS FOR A STOCHASTIC PROCESS MODEL WITH MEASUREMENT ERROR. *Statistica Sinica* **10**(1), 141–156 (2000). URL <https://www.jstor.org/stable/24306709>. Publisher: Institute of Statistical Science, Academia Sinica
19. Cui, T., Fox, C., O’Sullivan, M.J.: Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm. *Water Resources Research* **47**(10) (2011). DOI <https://doi.org/10.1029/2010WR010352>
20. Cui, T., Fox, C., O’Sullivan, M.J.: A posteriori stochastic correction of reduced models in delayed-acceptance mcmc, with application to multiphase subsurface inverse problems. *International Journal for Numerical Methods in Engineering* **118**(10), 578–605 (2019)
21. Cui, T., Martin, J., Marzouk, Y.M., Solonen, A., Spantini, A.: Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems* **30**(11), 114015 (2014)
22. Dashti, M., Stuart, A.M.: The Bayesian Approach to Inverse Problems. In: R. Ghanem, D. Higdon, H. Owhadi (eds.) *Handbook of Uncertainty Quantification*, pp. 311–428. Springer International Publishing, Cham (2017). DOI 10.1007/978-3-319-12385-1_7
23. Datta, B.: Numerical methods for linear control systems, vol. 1. Academic Press (2004)
24. De Oliveira, V., Ecker, M.D.: Bayesian hot spot detection in the presence of a spatial trend: application to total nitrogen concentration in chesapeake bay. *Environmetrics: The official journal of the International Environmetrics Society* **13**(1), 85–101 (2002)
25. Deligiannidis, G., Doucet, A., Rubenthaler, S.: Ensemble rejection sampling. arXiv preprint arXiv:2001.09188 (2020)
26. Devroye, L.: General principles in random variate generation. In: Non-uniform random variate generation, pp. 27–82. Springer (1986)
27. Evensen, G.: Data Assimilation: The Ensemble Kalman Filter, 2nd ed. edn. Springer (2009)
28. Förstner, W., Moonen, B.: A Metric for Covariance Matrices. In: E.W. Grafarend, F.W. Krumm, V.S. Schwarze (eds.) *Geodesy-The Challenge of the 3rd Millennium*, pp. 299–309. Springer, Berlin, Heidelberg (2003). DOI 10.1007/978-3-662-05296-9_31
29. Fortmann, T.E., Hitz, K.L.: An introduction to linear control systems. Crc Press (1977)
30. Gardiner, C.: Stochastic Methods: A Handbook for the Natural and Social Sciences, 4th ed. 2009 edition edn. Springer, Berlin (2009)
31. Gugercin, S., Antoulas, A.C.: A survey of model reduction by balanced truncation and some new results. *International Journal of Control* **77**(8), 748–766 (2004)
32. Hammarling, S.J.: Numerical solution of the stable, non-negative definite Lyapunov equation. *IMA Journal of Numerical Analysis* **2**(3), 303–323 (1982)

33. Hensman, J., Matthews, A.G., Filippone, M., Ghahramani, Z.: MCMC for variationally sparse Gaussian processes. In: *Advances in Neural Information Processing Systems*, pp. 1648–1656 (2015)
34. Johns, T.C., Carnell, R.E., Crossley, J.F., Gregory, J.M., Mitchell, J.F.B., Senior, C.A., Tett, S.F.B., Wood, R.A.: The second Hadley Centre coupled ocean-atmosphere GCM: model description, spinup and validation. *Climate Dynamics* **13**(2), 103–134 (1997)
35. Kalnay, E.: *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, Cambridge (2002). DOI 10.1017/CBO9780511802270
36. Kürschner, P.: Efficient low-rank solution of large-scale matrix equations. Ph.D. thesis, Otto von Guericke Universität Magdeburg (2016)
37. Law, K., Stuart, A., Zygalakis, K.: Data Assimilation, *Texts in Applied Mathematics*, vol. 62. Springer International Publishing, Cham (2015)
38. Lieberman, C., Willcox, K.: Goal-oriented inference: Approach, linear theory, and application to advection diffusion. *SIAM Review* **55**(3), 493–519 (2013)
39. Miroshnikov, A., Wei, Z., Conlon, E.M.: Parallel Markov chain Monte Carlo for non-Gaussian posterior distributions. *Stat* **4**(1), 304–319 (2015)
40. Moore, B.C.: Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Transactions on Automatic Control* **26**(1), 17–32 (1981)
41. Mullis, C.T., Roberts, R.A.: Synthesis of minimum roundoff noise fixed point digital filters. *Circuits and Systems, IEEE Transactions on* **23**(9), 551–562 (1976)
42. von Neumann, J.: Various techniques used in connection with random digits. In: A.S. Householder, G.E. Forsythe, H.H. Germond (eds.) *Monte Carlo Method, National Bureau of Standards Applied Mathematics Series*, vol. 12, chap. 13, pp. 36–38. US Government Printing Office, Washington, DC (1951)
43. Palacios, M.B., Steel, M.F.J.: Non-Gaussian Bayesian geostatistical modeling. *Journal of the American Statistical Association* **101**(474), 604–618 (2006)
44. Petreczky, M., Wisniewski, R., Leth, J.: Theoretical analysis of balanced truncation for linear switched systems. *IFAC Proceedings Volumes* **45**(9), 240–247 (2012)
45. Powel, N.D., Morgansen, K.A.: Empirical observability Gramian rank condition for weak observability of nonlinear systems with control. In: 2015 54th IEEE Conference on Decision and Control (CDC), pp. 6342–6348 (2015). DOI 10.1109/CDC.2015.7403218
46. Redmann, M., Freitag, M.A.: Balanced model order reduction for linear random dynamical systems driven by Lévy noise. *Journal of Computational Dynamics* **5**(1&2), 33 (2018). DOI 10.3934/jcd.2018002. URL <https://www.aims.org/article/doi/10.3934/jcd.2018002>
47. Rounce, D.R., Khurana, T., Short, M.B., Hock, R., Shean, D.E., Brinkerhoff, D.J.: Quantifying parameter uncertainty in a large-scale glacier evolution model using Bayesian inference: application to High Mountain Asia. *Journal of Glaciology* **66**(256), 175–187 (2020). DOI 10.1017/jog.2019.91
48. Rowley, C.W.: Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurcat. Chaos* **15**(3), 997–1013 (2005)
49. Roy, P., Cela, A., Hamam, Y.: On the relation of FIM and Controllability Gramian. In: 2009 IEEE International Symposium on Industrial Embedded Systems, pp. 37–41 (2009). DOI 10.1109/SIES.2009.5196189. ISSN: 2150-3117
50. Séférian, R., Gehlen, M., Bopp, L., Resplandy, L., Orr, J.C., Marti, O., Dunne, J.P., Christian, J.R., Doney, S.C., Ilyina, T., et al.: Inconsistent strategies to spin up models in CMIP5: implications for ocean biogeochemical model performance assessment. *Geoscientific Model Development* **9**(5), 1827–1851 (2016)
51. Simoncini, V.: Computational methods for linear matrix equations. *SIAM Review* **58**(3), 377–441 (2016)
52. Soize, C., Ghanem, R.G., Desceliers, C.: Sampling of Bayesian posteriors with a non-Gaussian probabilistic learning on manifolds from a small dataset. *Statistics and Computing* pp. 1–25 (2020)
53. Spantini, A., Cui, T., Willcox, K., Tenorio, L., Marzouk, Y.: Goal-oriented optimal approximations of Bayesian linear inverse problems. *SIAM Journal on Scientific Computing* **39**(5), S167–S196 (2017)
54. Spantini, A., Solonen, A., Cui, T., Martin, J., Tenorio, L., Marzouk, Y.: Optimal Low-rank Approximations of Bayesian Linear Inverse Problems. *SIAM Journal on Scientific Computing* **37**(6), A2451–A2487 (2015). DOI 10.1137/140977308
55. Stuart, A.M.: Inverse Problems: A Bayesian Perspective. *Acta Numerica* **19**, 451–559 (2010). DOI 10.1017/S0962492910000061
56. Tarantola, A.: *Inverse Problem Theory and Methods for Model Parameter Estimation*. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics (2005). DOI 10.1137/1.9780898717921
57. Willcox, K., Peraire, J.: Balanced model reduction via the proper orthogonal decomposition. *AIAA JOURNAL* **40**(11), 2323–2330 (2002)
58. Yang, F., Zheng, Y., Duan, J., Fu, L., Wiggins, S.: The tipping times in an arctic sea ice system under influence of extreme events. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **30**(6), 063125 (2020)
59. Zhang, H., Zimmerman, D.L.: Towards Reconciling Two Asymptotic Frameworks in Spatial Statistics. *Biometrika* **92**(4), 921–936 (2005). URL <https://www.jstor.org/stable/20441245>. Publisher: [Oxford University Press, Biometrika Trust]
60. Zhu, Z., Zhang, H.: Spatial sampling design under the infill asymptotic framework. *Environmetrics* **17**(4), 323–337 (2006). DOI 10.1002/env.772. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.772>