Journal of Computational Physics 352 (2018) 388–409

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Sparsity enabled cluster reduced-order models for control

Eurika Kaiser [a,*], Marek Morzyński [b], Guillaume Daviller [c], J. Nathan Kutz [d], Bingni W. Brunton [e], Steven L. Brunton [a,d]

[a] Department of Mechanical Engineering, University of Washington, Seattle, WA 98195, United States
[b] Chair of Virtual Engineering, Poznań University of Technology, 60-965 Poznań, Poland
[c] CERFACS, F-31057 Toulouse CEDEX 01, France
[d] Department of Applied Mathematics, University of Washington, Seattle, WA 98195, United States
[e] Department of Biology and Institute of Neuroengineering, University of Washington, Seattle, WA 98195, United States

ARTICLE INFO

Article history:
Received 24 December 2016
Received in revised form 18 August 2017
Accepted 29 September 2017
Available online 3 October 2017

Keywords:
Reduced-order modeling
Sensor placement
Compressed sensing
Clustering
Flow control

ABSTRACT

Characterizing and controlling nonlinear, multi-scale phenomena are central goals in science and engineering. Cluster-based reduced-order modeling (CROM) was introduced to exploit the underlying low-dimensional dynamics of complex systems. CROM builds a data-driven discretization of the Perron–Frobenius operator, resulting in a probabilistic model for ensembles of trajectories. A key advantage of CROM is that it embeds nonlinear dynamics in a linear framework, which enables the application of standard linear techniques to the nonlinear system. CROM is typically computed on high-dimensional data; however, access to and computations on this full-state data limit the *online* implementation of CROM for prediction and control. Here, we address this key challenge by identifying a small subset of critical measurements to learn an efficient CROM, referred to as sparsity-enabled CROM. In particular, we leverage compressive measurements to faithfully embed the cluster geometry and preserve the probabilistic dynamics. Further, we show how to identify fewer optimized sensor locations tailored to a specific problem that outperform random measurements. Both of these sparsity-enabled sensing strategies significantly reduce the burden of data acquisition and processing for low-latency in-time estimation and control. We illustrate this unsupervised learning approach on three different high-dimensional nonlinear dynamical systems from fluids with increasing complexity, with one application in flow control. Sparsity-enabled CROM is a critical facilitator for real-time implementation on high-dimensional systems where full-state information may be inaccessible.

1. Introduction

Nonlinear, multi-scale phenomena are ubiquitous in many fields in science and engineering; examples include the spread of infectious diseases, Earth's climate system, neural brain activity, autonomous behavior of robotic systems, sustainable energy production, and greener transport systems. High-dimensionality challenges our ability to understand and realistically model these systems. Moreover, low-latency real-time prediction and control is still a difficult endeavor, despite continually increasing computing power. The long history of model reduction exhibits numerous examples of compact representations of such high-dimensional systems, such as POD-Galerkin models [1,2], that successfully capture the principal mechanisms.

* Corresponding author.
E-mail address: eurika@uw.edu (E. Kaiser).

https://doi.org/10.1016/j.jcp.2017.09.057
0021-9991/

An alternative representation of nonlinear systems is based on infinite-dimensional linear operators acting on functions of the state space, such as the Koopman and Perron–Frobenius operators. A critical motivation for these operator-based approaches is the ability to apply powerful linear estimation and control techniques to nonlinear systems [3]. Cluster-based reduced-order modeling (CROM) [4] was recently introduced to approximate the Perron–Frobenius operator in an unsupervised manner from high-dimensional data yielding a low-dimensional, linear model in probability space. The present work combines CROM with sparsity-promoting techniques, particularly the sparse sensor placement optimization for classification (SSPOC) architecture [5], to enable real-time prediction and control. The *sparsity enabled CROM* identifies the probabilistic dynamics from few optimized measurements or compressed data, facilitating faster computations for the online estimation, prediction, and control of high-dimensional systems.

Reduced-order models (ROMs) aim to simplify a high-dimensional system by reducing the degrees of freedom, keeping only those that are important to model the phenomena of interest. The intrinsic coordinates, in which the system exhibits such a low-rank structure, are often computed by proper orthogonal decomposition (POD) [1], and low-dimensional dynamics are obtained via Galerkin projection. ROMs for parameterized systems are enabled by efficient evaluation of the nonlinear terms using sparse sampling techniques such as gappy POD [6]. The state of the art algorithm for principled sparse sampling of ROMs involves the empirical interpolation method [7,8], with variants including the addition of a genetic algorithm [9] and the use of pivot locations from the QR factorization [10]. More generally, sparsity-promoting techniques play an increasingly important role in dynamical systems for model identification [11,12], mode selection [13], and sensor placement [5,9,14–16] as well as for classification [17–20] and reconstruction [21,22].

Nonlinearities arising in standard ROMs remain challenging. For estimation and control purposes, a linear representation is highly advantageous, spurring considerable work on operator-theoretic embeddings of nonlinear dynamics; these embeddings are not to be confused with local linearization. Techniques for linear representation of dynamics include the operator methods of Koopman [23–27], Perron–Frobenius [28,29] and Fokker–Planck [30]. These infinite-dimensional operators act on functions of the state space, providing a globally linear description of the system. The practical computation of finite-dimensional approximations of the Koopman operator include dynamic mode decomposition (DMD) [31,32] and its variants [33–37]. The Perron–Frobenius operator is the adjoint of the Koopman operator, and it provides a probabilistic description of the dynamics. The continuous-time Liouville equation [38], the infinitesimal generator for the one-parameter family of Perron–Frobenius operators, governs the evolution of the probability density function (p.d.f.) in the state space (i.e., how an ensemble of trajectories evolves). A prominent example is Hopf's derivation of a Liouville equation for the Navier–Stokes equation [39].

The Perron–Frobenius operator has been extensively studied to analyze the global dynamical behavior of complex systems, in particular transport and mixing processes [40–42]. Examples range from fluid dynamics [4,43–47], to meteorological and atmospheric sciences [48,49], molecular dynamics [50,51], and engineering applications [52–56]. Data-driven approximations of the Perron–Frobenius operator are generally obtained via the Ulam–Galerkin method [29,57], which reduces this operator to a stochastic matrix. The operator's spectral decomposition contains valuable information on coherence of the flow and the long-time behavior. Lightly damped or invariant eigenfunctions of the Perron–Frobenius operator are particularly relevant for the characterization and control of dynamical systems, as these are associated with persistent dynamics, such as coherent or almost-invariant sets [41,42,58,59], and invariant or ergodic [60] distributions. In practice, Ulam's method involves a high-dimensional discretization of the state space using a box partition, which suffers from the *curse of dimensionality*. If time-series data is available, the transition probabilities between the boxes can be determined directly, but the computational burden is significant.

Cluster-based reduced-order modeling is a particular realization of Ulam's method where a low-dimensional discretization is obtained in an unsupervised manner from data using a clustering algorithm. This data-driven discretization enables an efficient partitioning, avoiding superfluous covering of regions where data is not available. In contrast to standard Perron–Frobenius approximations, CROM targets the development of very low-dimensional models suitable for control. Thus, exactness of the model may be traded for fast computations and real-time applicability. The simplest CROM uses k-means clustering [61] to learn an intrinsic partition structure directly from data by grouping similar observations [62]. Clustering algorithms have found a wide range of applications, e.g. for dimensionality reduction and library learning [63] and trust-region reduced-order modeling [64,65], to name a few. Analogous to CROM, coarse-grained dynamical models based on k-means have also been developed for the prediction of observables in oceanography [66]. CROM and related partition-based approaches generally rely on full-state measurements, which may be inaccessible in practice, and limits their use for real-time estimation and control. Thus, optimized sampling strategies may streamline these algorithms for online use in a wide range of modeling and control applications. Alternative approaches that enable efficient computations of the Perron–Frobenius operator rely on sparse grids, such as the sparse Ulam method [67], or employ spectral collocation [68,69] and radial basis function methods [70], and more recently, tensor decompositions [71]. However, these approaches do not target the development of data-driven ROMs, which is addressed in this study.

In this work, we leverage sparsity-promoting techniques to construct an efficient CROM from few measurements, referred to as sparsity-enabled CROM, which is a critical enabler for its *online* application. We first show that a sufficient, but small number of random measurements embed the cluster geometry and preserve the probabilistic dynamics. Further, we demonstrate the ability to learn a minimal set of optimized sensors, using the sparse sensor placement optimization for classification (SSPOC) architecture [5], that are tailored to the specific CROM and provide performance on par with the high-dimensional CROM. Sparsity-enabled CROM allows one to identify low-dimensional probabilistic dynamics of high-