RESEARCH

Dynamic data-driven model reduction: Adapting reduced models from incomplete data

Benjamin Peherstorfer* and Karen Willcox

*Correspondence: pehersto@mit.edu Department of Aeronautics & Astronautics, MIT, 77 Massachusetts Avenue, 02139 Cambridge, USA Full list of author information is available at the end of the article

Abstract

This work presents a data-driven online adaptive model reduction approach for systems that undergo dynamic changes. Classical model reduction constructs a reduced model of a large-scale system in an offline phase and then keeps the reduced model unchanged during the evaluations in an online phase; however, if the system changes online, the reduced model may fail to predict the behavior of the changed system. Rebuilding the reduced model from scratch is often too expensive in time-critical and real-time environments. We introduce a dynamic data-driven adaptation approach that adapts the reduced model from incomplete sensor data obtained from the system during the online computations. The updates to the reduced models are derived directly from the incomplete data, without recourse to the full model. Our adaptivity approach approximates the missing values in the incomplete sensor data with gappy proper orthogonal decomposition. These approximate data are then used to derive low-rank updates to the reduced basis and the reduced operators. In our numerical examples, incomplete data with 30-40% known values are sufficient to recover the reduced model that would be obtained via rebuilding from scratch.

Keywords: model reduction; online adaptivity; dynamic data-driven reduced models; incomplete sensor data; gappy proper orthogonal decomposition; dynamic data-driven application systems

1 Introduction

Dynamic online (near real-time) capability estimation is a pivotal component of future autonomous systems to dynamically observe, orient, decide, and act in complex and changing environments. We consider the situation where the dynamics of the system are modeled by a parametrized partial differential equation (PDE) and sensor data are generated that provide information on the current state of the system. The system dynamics are approximated by a large-scale parametrized computer model, the so-called *full model*, resulting from the discretization of the underlying PDE. We rely on (projection-based) model reduction [7, 29, 45] to derive a low-cost reduced model of the full model to meet the real-time demands of online capability estimation. Reduced models are typically built with one-time high-computational costs in an offline phase and then stay unchanged while they are repeatedly evaluated in an online phase. However, in changing environments, the properties and the behavior of the system might change even during the online phase. Rebuilding the reduced model from scratch to take into account the changes in the system is often too time consuming. We therefore rely on dynamic data-driven reduced models, as introduced in [42]. Dynamic data-driven reduced models adapt directly from sensor



Figure 1: The system depends on observable parameters, which are inputs to the system, and latent parameters, which model the external influence on the system and cannot be controlled. Dynamic data-driven reduced models adapt directly from sensor data to changes in the latent parameters (i.e., external influence), without recourse to the full model. (Figure from [42].)

data to changes in the underlying system, without recourse to the full model; however, the dynamic data-driven approach as presented in [42] requires sensor samples that measure the full large-scale state of the system.

Here, we present an extension to the dynamic data-driven approach that handles incomplete sensor samples. We consider the situation where we might have the ability to sense the full large-scale state of the system, but where we can afford to process only a subset of the sensor data. For example, new sensor technologies (e.g., "sensor skins") provide high-resolution sensor data of an entire component (e.g., an aircraft wing) but processing these tremendous amounts of data online is computationally challenging. Note that this is in contrast to settings where we have sparse sensors that are in fixed locations. Our methodology processes a selection of the sensor data—an incomplete sensor sample—that contains the essential information for updating the reduced model. Furthermore, we can dynamically change this selection of the sensor data during the online phase, so that at each step we process the subset of sensor data that are most informative to the event at hand.

To model changes in the system, the parameters of the system are split into observable and latent parameters, see Figure 1. The observable parameters are inputs to the system and therefore the values of these parameters are known. Latent parameters describe external influences on the system (e.g., damage, fatigue, erosion). The values of the latent parameters are unknown, except for the nominal latent parameters that describe the nominal state of the system (e.g., no-damage condition). Since the values of the latent parameters are unknown, a reduced model can be built in the offline phase for the nominal latent parameter only. If the latent parameters change online (e.g., the system gets damaged), the reduced model fails to predict the behavior of the system. Rebuilding the reduced model from scratch requires inferring the value of the changed latent parameters from the sensor data with a model of the changed system, then assembling the full model operators



corresponding to the inferred latent parameters, and deriving the reduced model. Rebuilding from scratch therefore is often too expensive in the context of online capability estimation, see, e.g., [1, 33, 42, 37] for a discussion. The dynamic datadriven approach introduced in [42] exploits the sensor data of the system to adapt the reduced model to changes in the latent parameters online, without the computationally expensive inference step and without assembling the full model operators for the inferred latent parameters, see Figure 2.

There are several online adaptation approaches for reduced models. We distinguish between approaches that solely rely on pre-computed quantities for the adaptation and approaches that adapt the reduced model from new data that are generated during the online phase. Interpolation between reduced operators and reduced models [39, 18, 2, 51], localization approaches [20, 36, 21, 3, 19, 11, 40, 9, 46], and dictionary approaches [30, 35] rely on pre-computed quantities but do not incorporate information from new data into the reduced model online. In [4], local reduced models are adapted from partial data online to smooth the transition between the local models. In [12], an h-adaptive refinement is presented that splits basis vectors based on an unsupervised learning algorithm and residuals that become available online. The online adaptive approach [43] adapts the approximation of nonlinear terms from sparse data of the full model. There is also a body of work that rebuilds reduced models from scratch, e.g., in optimization [27, 50, 32], inverse problem [17, 25], and multiscale methods [38]. We also mention that reduced models have been used in the context of dynamical data-driven application systems (DDDAS), which dynamically incorporate data into an executing application, and, in reverse, dynamically steer the measurement process. In [26], proper generalized decomposition [16] is used in a DDDAS to recover from device malfunctions by reconfiguring the simulation process. In [28], online parameter identification from measurements is considered for DDDAS with proper generalized decomposition. The work [1, 33, 37, 34] considers model reduction for structural health monitoring in DDDAS.

Our extension to handle incomplete sensor samples in the dynamic data-driven reduced model adaptation builds on gappy proper orthogonal decomposition (POD), which is a method to approximate unknown or missing values in vector-valued data [22]. Gappy POD reconstructs the unknown values by representing the data vector as a linear combination of POD basis vectors. Applications of gappy POD in model reduction include flow field reconstruction [49, 10], acceleration of efficient approximations of nonlinear terms [5, 13, 24], and forecasting for time-dependent problems [14]. In our adaptation approach, we first construct a gappy POD basis from incomplete sensor samples using an incremental POD basis generation algorithm. The missing values of the incomplete sensor samples are then approximate sensor samples are used in the obtained gappy POD basis. These approximate sensor samples are used in the dynamic data-driven adaptation to derive updates to the reduced model.

This paper is organized as follows. Section 2 introduces the full model and the dynamic data-driven adaptation. Section 3 defines incomplete sensor samples and describes the problem setup in detail. Section 4 introduces the extension to the dynamic data-driven adaptation approach that handles incomplete sensor samples. The numerical results in Section 5 demonstrate that in our examples 30-40% of the values of the sensor samples are sufficient to recover reduced models that accurately capture the changes in the latent parameters. Section 6 gives concluding remarks.

2 Preliminaries and adaptation from complete data

This section briefly discusses model reduction for systems with observable and latent parameters and summarizes the dynamic data-driven adaptation approach presented in [42].

2.1 Systems with latent parameters

Consider a parametrized system of equations stemming from the discretization of a parametrized PDE

$$A_{\eta}(\mu)y_{\eta}(\mu) = f(\mu).$$
(1)

The full model (1) depends on the observable parameter $\boldsymbol{\mu} \in \mathcal{D}$, where $\mathcal{D} \subset \mathbb{R}^d$ with $d \in \mathbb{N}$, and on the latent parameter $\boldsymbol{\eta} \in \mathcal{E}$, where $\mathcal{E} \subset \mathbb{R}^{d'}$ with $d' \in \mathbb{N}$. In general, the value of the latent parameter is unknown, only the value of a nominal latent parameter $\boldsymbol{\eta}_0 \in \mathcal{E}$ is known, see Section 1. The linear operator $\boldsymbol{A}_{\boldsymbol{\eta}}(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$ is an $N \times N$ matrix, where $N \in \mathbb{N}$ is the number of degrees of freedom of the full model (1). The linear operator $\boldsymbol{A}_{\boldsymbol{\eta}}(\boldsymbol{\mu})$ depends on the observable and on the latent parameter. The operator $\boldsymbol{A}_{\boldsymbol{\eta}}(\boldsymbol{\mu})$ has an affine parameter dependence with respect to the observable parameter

$$oldsymbol{A}_{oldsymbol{\eta}}(oldsymbol{\mu}) = \sum_{i=1}^{l_A} \Theta^{(i)}_A(oldsymbol{\mu}) oldsymbol{A}^{(i)}_{oldsymbol{\eta}}$$

where $l_A \in \mathbb{N}$ and $\Theta_A^{(1)}, \ldots, \Theta_A^{(l_A)} : \mathcal{D} \to \mathbb{R}$. The linear operators $A_{\eta}^{(1)}, \ldots, A_{\eta}^{(l_A)} \in \mathbb{R}^{N \times N}$ are independent of the observable parameter. Note that an affine parameter dependence with respect to μ can be approximated with sparse sampling methods, e.g., [22, 6, 5, 15, 13]. Note further that no affine parameter dependence with respect to the latent parameter is required. The state $y_{\eta}(\mu) \in \mathbb{R}^N$ is an N-dimensional

vector. The right-hand side $f(\mu) \in \mathbb{R}^N$ depends on the observable parameter but is independent of the latent parameter. The right-hand side has an affine parameter dependence with respect to μ

$$oldsymbol{f}(oldsymbol{\mu}) = \sum_{i=1}^{l_f} \Theta_f^{(i)}(oldsymbol{\mu}) oldsymbol{f}^{(i)}\,,$$

with $l_f \in \mathbb{N}, \Theta_f^{(1)}, \dots, \Theta_f^{(l_f)} : \mathcal{D} \to \mathbb{R}$, and the μ -independent vectors $f^{(1)}, \dots, f^{(l_f)} \in \mathbb{R}^N$.

2.2 Classical model reduction for systems with latent parameters

Let $\mathbf{Y}_{\eta_0} \in \mathbb{R}^{N \times M}$ be the snapshot matrix that contains as columns $M \in \mathbb{N}$ state vectors $\mathbf{y}_{\eta_0}(\boldsymbol{\mu}_1), \ldots, \mathbf{y}_{\eta_0}(\boldsymbol{\mu}_M) \in \mathbb{R}^N$ of the full model (1) corresponding to the observable parameters $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_M \in \mathcal{D}$ and the nominal latent parameter $\eta_0 \in \mathcal{E}$. The POD basis matrix $\mathbf{V}_{\eta_0} \in \mathbb{R}^{N \times n}$ contains as columns the first $n \in \mathbb{N}$ left-singular vectors of the snapshot matrix \mathbf{Y}_{η_0} that correspond to the largest singular values. The POD basis vectors, i.e., the columns of the POD basis matrix \mathbf{V}_{η_0} , span the *n*-dimensional POD space \mathcal{V}_{η_0} .

The reduced linear operator $\tilde{A}_{\eta_0}(\mu) \in \mathbb{R}^{n \times n}$ is obtained via Galerkin projection of the equations of the full model onto the POD space \mathcal{V}_{η_0} . Consider therefore the projected μ -independent operators

$$ilde{oldsymbol{A}}_{oldsymbol{\eta}_0}^{(i)} = oldsymbol{V}_{oldsymbol{\eta}_0}^T oldsymbol{A}_{oldsymbol{\eta}_0}^{(i)} oldsymbol{V}_{oldsymbol{\eta}_0} \,, \qquad i=1,\ldots,l_A \,.$$

By exploiting the affine parameter dependence of the linear operator $A_{\eta_0}(\mu)$ on the observable parameter $\mu \in \mathcal{D}$, the reduced linear operator $\tilde{A}_{\eta_0}(\mu)$ is

$$ilde{oldsymbol{A}}_{oldsymbol{\eta}_0}(oldsymbol{\mu}) = \sum_{i=1}^{l_A} \Theta^{(i)}_A(oldsymbol{\mu}) ilde{oldsymbol{A}}^{(i)}_{oldsymbol{\eta}_0}$$

Similarly, the reduced right-hand side is

$$ilde{oldsymbol{f}}_{oldsymbol{\eta}_0}(oldsymbol{\mu}) = \sum_{i=1}^{l_f} \Theta_f^{(i)}(oldsymbol{\mu}) ilde{oldsymbol{f}}^{(i)}$$
 ,

where $\tilde{f}^{(i)} = V_{\eta_0}^T f^{(i)} \in \mathbb{R}^n$ for $i = 1, ..., l_f$. The reduced model for the latent parameter η_0 is

$$\tilde{\boldsymbol{A}}_{\boldsymbol{\eta}_0}(\boldsymbol{\mu})\tilde{\boldsymbol{y}}_{\boldsymbol{\eta}_0}(\boldsymbol{\mu}) = \tilde{\boldsymbol{f}}_{\boldsymbol{\eta}_0}(\boldsymbol{\mu}), \qquad (2)$$

where $\tilde{\boldsymbol{y}}_{\boldsymbol{\eta}_0}(\boldsymbol{\mu}) \in \mathbb{R}^n$ is the reduced state. The reduced right-hand side $\tilde{\boldsymbol{f}}_{\boldsymbol{\eta}_0}(\boldsymbol{\mu}) \in \mathbb{R}^n$ in the reduced model (2) depends on the latent parameter $\boldsymbol{\eta}_0$ because of the projection onto the POD space $\mathcal{V}_{\boldsymbol{\eta}_0}$, in contrast to the right-hand side vector $\boldsymbol{f}(\boldsymbol{\mu}) \in \mathbb{R}^N$ in the full model (1).

2.3 Dynamic data-driven adaptation for reduced models

The reduced model (2) is derived from snapshots with the latent parameter $\eta = \eta_0$ set to the nominal latent parameter η_0 . This means that if the latent parameter changes online, the reduced model (2) cannot predict the behavior of the system. In [42, 41], a dynamic data-driven adaptation approach is presented that successively adapts a reduced model in $M' \in \mathbb{N}$ adaptivity steps to changes in the latent parameter. Consider therefore the $h = 1, \ldots, M'$ adaptivity steps, in which the reduced model is adapted from the nominal latent parameter η_0 to the changed latent parameter, say, $\eta_1 \in \mathcal{E}^{[1]}$. In each adaptivity step h, a sensor sample $\hat{y}_{\eta_1}(\mu_{M+h}) \in \mathbb{R}^N$ is received. The sensor sample $\hat{y}_{\eta_1}(\mu_{M+h})$ is an approximation of the state $y_{\eta_1}(\mu_{M+h})$ for the changed latent parameter η_1 and an observable parameter $\mu_{M+h} \in \mathcal{D}$. The difference $\|\hat{y}_{\eta_1}(\mu_{M+h}) - y_{\eta_1}(\mu_{M+h})\|$ between the sensor sample and the state in a norm $\|\cdot\|$ is noise, measurement error, and the discrepancy of the full model and reality (model discrepancy [31]). At step h, the sensor samples matrix $S_h \in \mathbb{R}^{N \times h}$ contains the received sensor samples $\hat{y}_{\eta_1}(\mu_{M+1}), \ldots, \hat{y}_{\eta_1}(\mu_{M+h}) \in \mathbb{R}^N$ as columns

$$oldsymbol{S}_h = [\hat{oldsymbol{y}}_{oldsymbol{\eta}_1}(oldsymbol{\mu}_{M+1}), \dots, \hat{oldsymbol{y}}_{oldsymbol{\eta}_1}(oldsymbol{\mu}_{M+h})] \in \mathbb{R}^{N imes h}$$
 .

At each adaptivity step h = 1, ..., M', the dynamic data-driven adaptation first adapts the POD basis and then the reduced operators. Consider the POD basis adaptation first. At step h = 1, the first snapshot, i.e., the first column, in the snapshot matrix Y_{η_0} is replaced with the sensor sample $\hat{y}_{\eta_1}(\mu_{M+1}) \in \mathbb{R}^N$ and the snapshot matrix at step h = 1 is obtained

$$oldsymbol{Y}_1 = [\hat{oldsymbol{y}}_{oldsymbol{\eta}_1}(oldsymbol{\mu}_{M+1}), oldsymbol{y}_{oldsymbol{\eta}_0}(oldsymbol{\mu}_2), \dots, oldsymbol{y}_{oldsymbol{\eta}_0}(oldsymbol{\mu}_M)] \in \mathbb{R}^{N imes M}$$
 .

Note that there is no particular ordering of the snapshots in the snapshots matrix. We replace the first column of Y_{η_0} because we are at step h = 1. By reordering the columns of Y_{η_0} , any other snapshot can be replaced at step h = 1. The matrix Y_1 is the result of an additive rank-one update to the snapshot matrix Y_{η_0} . Let $e_i \in \{0,1\}^N$ be the canonical unit vector with 1 at component *i* and 0 at all other components for $i = 1, \ldots, N$. Then, the snapshot matrix Y_1 is

$$\boldsymbol{Y}_1 = \boldsymbol{Y}_{\boldsymbol{\eta}_0} + \boldsymbol{a} \boldsymbol{e}_1^T$$
,

where $\boldsymbol{a} = \hat{\boldsymbol{y}}_{\eta_1}(\boldsymbol{\mu}_{M+1}) - \boldsymbol{y}_{\eta_0}(\boldsymbol{\mu}_1) \in \mathbb{R}^N$. Therefore, the POD basis matrix $\boldsymbol{V}_1 \in \mathbb{R}^{N \times n}$ corresponding to the snapshot matrix \boldsymbol{Y}_1 can be approximately derived from \boldsymbol{V}_{η_0} via the adaptation algorithm [8]. The algorithm extracts the components $\boldsymbol{\alpha} = \boldsymbol{a} - \boldsymbol{V}_{\eta_0} \boldsymbol{V}_{\eta_0}^T \boldsymbol{a}$ and $\boldsymbol{\beta} = \boldsymbol{e}_1 - \boldsymbol{V}_{\eta_0} \boldsymbol{V}_{\eta_0}^T \boldsymbol{e}_1$ of \boldsymbol{a} and \boldsymbol{e}_1 , respectively, that are orthogonal to \boldsymbol{V}_{η_0} . The vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are used to derive a rotation matrix $\boldsymbol{V}' \in \mathbb{R}^{n \times n}$ of size $n \times n$ and an additive rank-one update $\boldsymbol{\gamma} \boldsymbol{\delta}^T$ with $\boldsymbol{\gamma} \in \mathbb{R}^N$ and $\boldsymbol{\delta} \in \mathbb{R}^n$. Computing the rotation matrix and the rank-one update requires computing the singular value decomposition (SVD) of an $(n+1) \times (n+1)$ matrix. The adapted POD basis matrix \boldsymbol{V}_1 is then given by

$$oldsymbol{V}_1 = oldsymbol{V}_{oldsymbol{\eta}_0}oldsymbol{V}' + oldsymbol{\gamma}oldsymbol{\delta}^T$$
 .

^[1]Note that the adaptation can be repeated to adapt from η_1 to $\eta_2 \in \mathcal{E}$ and so on.

Note that an SVD of a typically small $(n + 1) \times (n + 1)$ matrix is required by the adaptation algorithm, instead of the SVD of an $N \times M$ matrix if the POD basis matrix were computed directly from \mathbf{Y}_1 without reusing \mathbf{V}_{η_0} . We refer to [8] for details on the adaptation of the POD basis matrix. The adaptation algorithm is summarized in [42, Algorithm 1] for the case of the dynamic data-driven adaptation.

Consider now the adaptation of the operators at step h = 1. The goal is to approximate the reduced operators

$$\tilde{\boldsymbol{A}}_{\boldsymbol{\eta}_1}^{(i)} = \boldsymbol{V}_1^T \boldsymbol{A}_{\boldsymbol{\eta}_1}^{(i)} \boldsymbol{V}_1, \qquad i = 1, \dots, l_A,$$

without assembling the full operators $A_{\eta_1}^{(1)}, \ldots, A_{\eta_1}^{(l_A)} \in \mathbb{R}^{N \times N}$ corresponding to the changed latent parameter η_1 . Therefore, at adaptivity step h = 1, the operators

$$\bar{A}_{1}^{(i)} = V_{1}^{T} A_{\eta_{0}}^{(i)} V_{1}, \qquad i = 1, \dots, l_{A}, \qquad (3)$$

are constructed. The operator $\bar{A}_1^{(i)}$ is the full operator $A_{\eta_0}^{(i)}$ for latent parameter η_0 projected onto the adapted POD space \mathcal{V}_1 with the adapted POD basis matrix V_1 , for $i = 1, \ldots, l_A$. Note that (3) projects the full operators corresponding to the nominal latent parameter η_0 , and not the operators corresponding to the changed latent parameter η_1 . Then, additive updates $\delta \tilde{A}_1^{(1)}, \ldots, \delta \tilde{A}_1^{(l_A)} \in \mathbb{R}^{n \times n}$ are derived from the sensor sample matrix S_1 with the optimization problem

$$\min_{\delta \tilde{\boldsymbol{A}}_{h}^{(1)},...,\delta \tilde{\boldsymbol{A}}_{h}^{(l_{A})} \in \mathbb{R}^{n \times n}} \sum_{j=1}^{h} \left\| \sum_{i=1}^{l_{A}} \Theta_{A}^{(i)}(\boldsymbol{\mu}_{M+j}) \left(\bar{\boldsymbol{A}}_{h}^{(i)} + \delta \tilde{\boldsymbol{A}}_{h}^{(i)} \right) \boldsymbol{V}_{h}^{T} \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_{1}}(\boldsymbol{\mu}_{M+j}) - \tilde{\boldsymbol{f}}_{h}(\boldsymbol{\mu}_{M+j}) \right\|_{2}^{2}$$

$$\tag{4}$$

where $\tilde{f}_h(\mu_{M+j}) \in \mathbb{R}^n$ is the reduced right-hand side with respect to the POD basis V_h . Note that the optimization problem (4) is formulated for general $h \geq 1$, and not only for h = 1. The solution of the optimization problem (4) are the updates $\delta \tilde{A}_h^{(1)}, \ldots, \delta \tilde{A}_h^{(l_A)}$ that best-fit the sensor samples in the sensor sample matrix S_h . The optimization problem (4) is a least-squares problem that can be solved with, e.g., the QR decomposition. For $h < l_A n$, the least-squares problem is underdetermined, and only low-rank approximations of $\delta \tilde{A}_h^{(1)}, \ldots, \delta \tilde{A}_h^{(l_A)}$ are computed [42]. At step h = 1, the adapted operators are

$$ilde{A}_{1}^{(i)} = ar{A}_{\eta_{0}}^{(i)} + \delta ilde{A}_{1}^{(i)}, \qquad i = 1, \dots, l_{A},$$

and the adapted reduced operator $\tilde{A}_1(\mu) \in \mathbb{R}^{n \times n}$ can be assembled using the affine parameter dependence as

$$ilde{m{A}}_1(m{\mu}) = \sum_{i=1}^{l_A} \Theta_{l_A}^{(i)}(m{\mu}) ilde{m{A}}_1^{(i)} \, .$$

In each adaptivity step h = 1, ..., M', this POD basis and operator adaptation is repeated. This means, at step h, the POD basis matrix is adapted from V_{h-1} to V_h by exploiting that the snapshot matrix Y_h at step h is the result of a rank-one update to the snapshot matrix Y_{h-1} from the previous step. The adapted reduced operator $\tilde{A}_h(\mu)$ is derived via the additive rank-one updates $\delta \tilde{A}_h^{(1)}, \ldots, \delta \tilde{A}_h^{(l_A)} \in \mathbb{R}^{n \times n}$, which are obtained via optimization from the sensor samples matrix $S_h = [\hat{y}_{\eta_1}(\mu_{M+1}), \ldots, \hat{y}_{\eta_1}(\mu_{M+h})] \in \mathbb{R}^{N \times h}$. For sufficiently many sensor samples, and if the sensor samples are noise-free, the reduced operator $\tilde{A}_{\eta_1}(\mu)$ with respect to the POD basis matrix V_h equals the adapted reduced operator $\tilde{A}_h(\mu)$, see [42].

3 Incomplete sensor samples

The dynamic data-driven adaptation derives updates to a reduced model from sensor samples. We consider here the situation where we receive incomplete sensor samples, i.e., partial measurements of the state. This section mathematically defines incomplete sensor samples, and the next section develops the extension to the dynamic data-driven adaptation to handle incomplete sensor samples.

Let $\hat{y}_{\eta_1}(\boldsymbol{\mu}_{M+h}) \in \mathbb{R}^N$ be the (complete) sensor sample that is received at adaptivity step h. Let $k \in \mathbb{N}$ with k < N and let $p_1^h, \ldots, p_k^h \in \{1, \ldots, N\}$ be pairwise distinct indices of the sensor sample $\hat{y}_{\eta_1}(\boldsymbol{\mu}_{M+h}) \in \mathbb{R}^N$. The indices p_1^h, \ldots, p_k^h give rise to a point selection matrix

$$oldsymbol{P}_h = [oldsymbol{e}_{p_1^h}, \dots, oldsymbol{e}_{p_k^h}] \in \mathbb{R}^{N imes k}$$

The point selection matrix P_h selects the components with indices p_1^h, \ldots, p_k^h . For example, consider the vector $\boldsymbol{x} = [x_1, \ldots, x_N]^T \in \mathbb{R}^N$, then we have

$$\begin{bmatrix} x_{p_1^h} \\ \vdots \\ x_{p_k^h} \end{bmatrix} = \boldsymbol{P}_h^T \boldsymbol{x}$$

From the point selection matrix \boldsymbol{P}_h , we derive the matrix $\boldsymbol{Q}_h \in \mathbb{R}^{N \times (N-k)}$ that selects the components of the (complete) sensor sample $\hat{\boldsymbol{y}}_{\eta_1}(\boldsymbol{\mu}_{M+h})$ that are missing in the incomplete sensor sample $\hat{\boldsymbol{y}}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h})$. The matrices \boldsymbol{P}_h and \boldsymbol{Q}_h lead to the decomposition

$$\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}(\boldsymbol{\mu}_{M+h}) = \boldsymbol{P}_h \boldsymbol{P}_h^T \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}(\boldsymbol{\mu}_{M+h}) + \boldsymbol{Q}_h \boldsymbol{Q}_h^T \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}(\boldsymbol{\mu}_{M+h}).$$

The matrix $\boldsymbol{P}_h \boldsymbol{P}_h^T$ selects all components that correspond to the indices p_1^h, \ldots, p_k^h and sets the components at all other indices $\{1, \ldots, N\} \setminus \{p_1^h, \ldots, p_k^h\}$ to zero. The matrix $\boldsymbol{Q}_h \boldsymbol{Q}_h^T$ has the opposite effect and selects all components with indices in $\{1, \ldots, N\} \setminus \{p_1^h, \ldots, p_k^h\}$ and sets the components with indices $\{p_1^h, \ldots, p_k^h\}$ to zero.

We define the incomplete sensor sample $\hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+h})$ of the (complete) sensor sample $\hat{y}_{\eta_1}(\mu_{M+h})$ corresponding to the point selection matrix P_h as

$$\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h}) = \boldsymbol{P}_h \boldsymbol{P}_h^T \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}(\boldsymbol{\mu}_{M+h}) \in \mathbb{R}^N.$$
(5)

The values at the components of the incomplete sensor sample $\hat{y}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h})$ with indices p_1^h, \ldots, p_k^h are set to the corresponding components of the (complete) sensor sample $\hat{y}_{\eta_1}(\boldsymbol{\mu}_{M+h})$. All other components are missing in the incomplete sensor sample and their values in $\hat{y}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h})$ are zero through the definition (5).

4 Dynamic data-driven adaptation from incomplete sensor samples

We propose an extension to the dynamic data-driven adaptation approach that handles incomplete sensor samples. Consider the adaptation from the nominal latent parameter η_0 to the latent parameter η_1 in the M' adaptivity steps $h = 1, \ldots, M'$. At each adaptivity step $h = 1, \ldots, M'$, we receive incomplete sensor samples $\hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+h}) \in \mathbb{R}^N$ and the corresponding point selection matrices P_h . The point selection matrix depends on h and might change at each adaptivity step, see the discussion on future sensor technologies in Section 1. The number of known components k is independent of h and stays constant for all $h = 1, \ldots, M'$.

We split the adaptivity steps $M' = M^{\text{basis}} + M^{\text{update}}$ into $M^{\text{basis}} \in \mathbb{N}$ and $M^{\text{update}} \in \mathbb{N}$ steps. In the first $h = 1, \ldots, M^{\text{basis}}$ steps, a gappy POD basis is derived from the incomplete sensor samples $\hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+1}), \ldots, \hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+M^{\text{basis}}}) \in \mathbb{R}^N$. At the subsequent M^{update} steps $h = M + M^{\text{basis}} + 1, \ldots, M'$, the missing values of the incomplete sensor samples $\hat{y}_{\eta_1}(\mu_{M+M^{\text{basis}}+h}) \in \mathbb{R}^N$ are approximated using gappy POD with the obtained gappy POD basis. The approximations of the missing values and the components in the incomplete sensor sample are combined to approximate the complete sensor sample. The dynamic data-driven adaptation is then applied to these approximate sensor samples to update the reduced model. Section 4.1 discusses the construction of the gappy POD basis and Section 4.2 presents the adaptation of the reduced model from the approximate sensor samples. Section 4.3 summarizes the procedure and presents Algorithm 1.

4.1 Deriving the gappy POD basis

In the first $h = 1, \ldots, M^{\text{basis}}$ adaptivity steps, we derive a gappy POD basis from the incomplete sensor samples. Let $r \in \mathbb{N}$ be the dimension of the gappy POD basis with gappy POD basis matrix $U_h \in \mathbb{R}^{N \times r}$. The initial gappy POD basis matrix $U_0 \in \mathbb{R}^{N \times r}$ contains as columns the *r*-dimensional POD basis vectors corresponding to the snapshot matrix Y_{η_0} .

At step h = 1, we receive the incomplete sensor sample $\hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+1})$ and the corresponding point selection matrix $P_1 \in \mathbb{R}^{N \times k}$ with $Q_1 \in \mathbb{R}^{N \times (N-k)}$. We use the initial gappy POD basis matrix U_0 to derive the approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+1}) \in \mathbb{R}^N$ using gappy POD [22, 49, 10]

$$\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_{1}}^{\text{apprx}}(\boldsymbol{\mu}_{M+1}) = \boldsymbol{Q}_{1}\boldsymbol{Q}_{1}^{T}\boldsymbol{U}_{0}(\boldsymbol{P}_{1}^{T}\boldsymbol{U}_{0})^{+}\boldsymbol{P}_{1}^{T}\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_{1}}^{\text{incp}}(\boldsymbol{\mu}_{M+1}) + \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_{1}}^{\text{incp}}(\boldsymbol{\mu}_{M+1}) \,. \tag{6}$$

The matrix $(\boldsymbol{P}_1^T \boldsymbol{U}_0)^+ \in \mathbb{R}^{r \times k}$ is the Moore-Penrose pseudoinverse of the matrix $\boldsymbol{P}_1^T \boldsymbol{U}_0 \in \mathbb{R}^{k \times r}$. Since $\boldsymbol{P}_1^T \hat{\boldsymbol{y}}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+1}) = \boldsymbol{P}_1^T \hat{\boldsymbol{y}}_{\eta_1}(\boldsymbol{\mu}_{M+1})$, we have that $(\boldsymbol{P}_1^T \boldsymbol{U}_0)^+ \boldsymbol{P}_1^T \hat{\boldsymbol{y}}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+1})$ is the solution of the regression problem

$$\underset{\boldsymbol{c}\in\mathbb{R}^{r}}{\arg\min}\|\boldsymbol{P}_{1}^{T}(\boldsymbol{U}_{0}\boldsymbol{c}-\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_{1}}(\boldsymbol{\mu}_{M+1}))\|_{2}^{2}.$$
(7)

Note that the regression problem is overdetermined and has a unique solution if the matrix $\boldsymbol{P}_1^T \boldsymbol{U}_0$ has full column rank, which we typically ensure by selecting k > r. Therefore, the vector $\boldsymbol{U}_0(\boldsymbol{P}_1^T \boldsymbol{U}_0)^+ \boldsymbol{P}_1^T \hat{\boldsymbol{y}}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+1}) \in \mathbb{R}^N$ is the best approximation with respect to (7) of the complete sensor sample $\hat{\boldsymbol{y}}_{\eta_1}(\boldsymbol{\mu}_{M+1})$ in the space spanned by the columns of the POD basis matrix U_0 . The approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+1})$ combines this best approximation and the known values in the incomplete sensor sample. The values at the components corresponding to the missing components of the incomplete sensor sample are set to the best approximation, and the values at all other components are set to the values obtained from the incomplete sensor sample.

We then use the approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+1})$ to adapt the gappy POD basis from U_0 to U_1 . Consider therefore the snapshot matrix Y_0 and note that U_0 is the k-dimensional POD basis derived from Y_0 . We adapt the snapshot matrix Y_0 to $Y_1 \in \mathbb{R}^{N \times M}$ via a rank-one update that replaces column 1 of Y_0 with the approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+1}) \in \mathbb{R}^N$. Since Y_1 is the result of a rank-one update to Y_0 , the k-dimensional POD basis corresponding to Y_1 can be approximated in a computationally efficient manner using the incremental POD algorithm [8]. Note that this is the same approach as used in the dynamic datadriven adaptation, see Section 2.3. Thus, the adapted gappy POD basis matrix U_1 can be derived cheaply from the basis matrix U_0 .

At step h = 2, the approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+2})$ is constructed with the gappy POD basis matrix U_1 , which is then used to adapt from U_1 to U_2 . This process is continued until step $h = M^{\text{basis}}$, where the gappy POD basis matrix $U_{M^{\text{basis}}}$ is derived. Note that the number of columns in the snapshot matrix is fixed and that columns are replaced following the first-in-first-out principle if h > M.

4.2 Dynamic data-driven adaptation from approximate sensor samples

In the M^{update} steps $h = M + M^{\text{basis}} + 1, \ldots, M'$, we adapt the reduced model from approximate sensor samples using the dynamic data-driven adaptation. Consider therefore an adaptivity step $h > M^{\text{basis}}$, at which the incomplete sensor sample $\hat{y}_{\eta_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h}) \in \mathbb{R}^N$ and the corresponding point selection matrix $\boldsymbol{P}_h \in \mathbb{R}^{k \times N}$ are received. We use the gappy POD basis $\boldsymbol{U}_{M^{\text{basis}}}$ to derive the approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\boldsymbol{\mu}_{M+h})$ of the complete sensor sample with the gappy POD basis $\boldsymbol{U}_{M^{\text{basis}}}$. The approximate sensor sample $\hat{y}_{\eta_1}^{\text{apprx}}(\boldsymbol{\mu}_{M+h})$ is then used to adapt the reduced model with the dynamic data-driven adaptation as described in Section 2.3.

4.3 Computational procedure

Algorithm 1 summarizes the dynamic data-driven adaptation that can handle incomplete sensor samples. Inputs of Algorithm 1 are the POD basis matrix V_{h-1} , the operators $\tilde{A}_{h-1}^{(1)}, \ldots, \tilde{A}_{h-1}^{(l_A)}$, and the right-hand sides $\tilde{f}_{h-1}^{(1)}, \ldots, \tilde{f}_{h-1}^{(l_f)}$ derived at the previous adaptivity step h-1. If $h \leq M^{\text{basis}}$, the algorithm adapts the gappy POD basis from U_{h-1} to U_h using the approach presented in Section 4.1. First, the approximate sensor sample is constructed with gappy POD. Then, the adapted basis matrix U_h is computed with the incremental POD algorithm [8]. Only the gappy POD basis is adapted and the reduced model is returned unchanged. If $h > M^{\text{basis}}$, the approximate sensor sample is derived with gappy POD and $U_{M^{\text{basis}}}$. The approximate sensor sample is then used with the dynamic data-driven adaptation to derive the adapted POD basis V_h , the adapted operators $\tilde{A}_h^{(1)}, \ldots, \tilde{A}_h^{(l_A)}$, and the adapted right-hand sides $\tilde{f}_h^{(1)}, \ldots, \tilde{f}_h^{(l_f)}$.

Algorithm 1 Dynamic data-driven adaptation from incomplete sensor samples

1: procedure ADAPTINCOMPLETE($M^{\text{basis}}, U_{h-1}, V_{h-1}, \tilde{A}_{h-1}^{(1)}, \dots, \tilde{A}_{h-1}^{(l_A)}, \tilde{f}_{h-1}^{(1)}, \dots, \tilde{f}_{h-1}^{(l_f)}$) 2: Receive incomplete sensor sample $\hat{y}_{\eta_1}^{\text{incp}}(\mu_{M+h}) \in \mathbb{R}^N$ and point selection matrix $P_h \in \mathbb{R}^{k \times N}$ Construct matrix $oldsymbol{Q}_h \in \mathbb{R}^{(N-k) imes N}$ 3: if $h \leq M^{\text{basis}}$ then 4: $\stackrel{-}{\mathsf{Compute}}$ approximate sensor sample using gappy POD basis matrix U_{h-1} 5: $\hat{y}_{\eta_{1}}^{\text{apprx}}(\mu_{M+h}) = Q_{h}Q_{h}^{T}U_{h-1}(P_{h}^{T}U_{h-1})^{+}P_{h}^{T}\hat{y}_{\eta_{1}}^{\text{incp}}(\mu_{M+h}) + \hat{y}_{\eta_{1}}^{\text{incp}}(\mu_{M+h}) \,.$ Update snapshot matrix Y_{h-1} with $\hat{y}_{\eta_1}^{\text{apprx}}(\mu_{M+h})$ to obtain Y_h Derive adapted gappy POD basis matrix U_h from U_{h-1} using [8] 6: 7: Set $ilde{m{A}}_h^{(i)} = ilde{m{A}}_{h-1}^{(i)}$ for $i=1,\ldots,l_A$ 8: Set $\tilde{f}_{h}^{(i)} = \tilde{f}_{h-1}^{(i)}$ for $i = 1, \dots, l_f$ Set $V_h = V_{h-1}$ 9: 10: 11: else Compute approximate sensor sample using basis $oldsymbol{U}_{M^{ extsf{basis}}}$ 12: $\hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}^{\text{apprx}}(\boldsymbol{\mu}_{M+h}) = \boldsymbol{Q}_h \boldsymbol{Q}_h^T \boldsymbol{U}_{M^{\text{basis}}}(\boldsymbol{P}_h^T \boldsymbol{U}_{M^{\text{basis}}})^+ \boldsymbol{P}_h^T \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h}) + \hat{\boldsymbol{y}}_{\boldsymbol{\eta}_1}^{\text{incp}}(\boldsymbol{\mu}_{M+h}) \,.$ Get $V_h, ilde{A}_h^{(1)}, \dots, ilde{A}_h^{(l_A)}, ilde{f}_h^{(1)}, \dots, ilde{f}_h^{(l_f)}$ from dynamic adaptation with $\hat{y}_{\eta_1}^{\mathsf{apprx}}(\mu_{M+h})$ 13: 14: end if return $V_h, ilde{m{A}}_h^{(1)}, \dots, ilde{m{A}}_h^{(l_A)}, ilde{m{f}}_h^{(1)}, \dots, ilde{m{f}}_h^{(l_f)}$ 15: 16: end procedure

5 Numerical results

This section demonstrates the dynamic data-driven adaptation from incomplete sensor samples on a model of a bending plate. The latent parameter describes damage of the plate. The damage is a local decrease of the thickness of the plate. The model is based on the Mindlin plate theory [23, 47] that takes into account transverse shear deformations but neglects important nonlinear effects such as postbuckling behavior. Therefore, the model that we use in this section is a simple description of a plate in bending. We use the plate model only to provide a proof of concept of our adaptation approach. More advanced plate models are used in real-world engineering applications. We refer to Section 6 for a discussion on further applications of our adaptation approach.

We first build a reduced model for the nominal problem, i.e., the latent parameter is set to the nominal latent parameter $\eta_0 \in \mathcal{D}$ that corresponds to the no-damage condition. We then decrease the thickness of the plate stepwise and adapt the reduced model. After each change in the latent parameter, synthetic incomplete sensor samples are computed with the full model, which are used to adapt the reduced model. The following sections give details on the problem setup and report the numerical results.

5.1 Plate problem

We consider the static analysis of a plate in bending. The plate is clamped into a frame and a load is applied. Our problem is an extension of the plate problems introduced in [23, 42, 44]. The geometry of our plate problem is shown in Figure 3a. The spatial domain $\Omega \in [0,1]^2 \subset \mathbb{R}^2$ is split into two subdomains $\Omega = \Omega_1 \cup \Omega_2$. The problem has three observable parameters $\boldsymbol{\mu} = [\mu_1, \mu_2, \mu_3]^T \in \mathcal{D}$ with $\mathcal{D} = [0.05, 0.1]^2 \times [1, 100]$. The observable parameters μ_1 and μ_2 control the nominal thickness of the plate in the subdomain Ω_1 and Ω_2 , respectively. The third observable parameter μ_3 defines the load on the plate.



Figure 3: The spatial domain of the plate problem is split into two subdomains as shown in (a). The plot (b) shows the decay of the singular values of the snapshot matrix for nominal latent parameter η_0 (no-damage condition).

The latent parameter $\boldsymbol{\eta} = [\eta_1, \eta_2]^T \in \mathcal{E}$ controls the damage of the plate, i.e., the latent parameter defines the local decrease of the thickness that corresponds to the damage. The domain of the latent parameter is $\mathcal{E} = [0, 0.2] \times (0, 0.05]$. The thickness of the plate at position $\boldsymbol{x} \in \Omega$ is given by the function $t : \Omega \times \mathcal{D} \times \mathcal{E} \to \mathbb{R}$ with

$$t(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\eta}) = t_0(\boldsymbol{x};\boldsymbol{\mu}) \left(1 - \eta_1 \exp\left(-\frac{1}{2\eta_2^2} \|\boldsymbol{x} - \boldsymbol{z}\|_2^2\right) \right) \,,$$

and

$$t_0(\boldsymbol{x}; \boldsymbol{\mu}) = egin{cases} \mu_1 & ext{if } x_1 > 0.5 \ \mu_2 & ext{if } x_1 \leq 0.5 \end{cases}$$

with position $\boldsymbol{z} = [0.7, 0.4]^T \in \Omega$. The function t is nonlinear in $\boldsymbol{x}, \boldsymbol{\mu}$ and $\boldsymbol{\eta}$. We set the nominal latent parameter $\boldsymbol{\eta}_0$ to $\boldsymbol{\eta}_0 = [0, 0.01]^T \in \mathcal{E}$ that corresponds to no local decrease of the thickness and therefore to the no-damage condition.

The full model of the plate problem is a finite element model, see [23]. The corresponding system of equations is of the form (1), where $l_A = 4$, $l_f = 1$, $\Theta_f^{(1)}(\boldsymbol{\mu}) = \mu_3$,

$$\Theta_A^{(1)}(\mu) = \mu_1^3, \qquad \Theta_A^{(2)}(\mu) = \mu_2^3,$$

and

$$\Theta_A^{(3)}(\mu) = \mu_1, \qquad \Theta_A^{(4)}(\mu) = \mu_2$$

The system of equations has N = 4719 degrees of freedom. The thickness of the plate with $\boldsymbol{\mu} = [0.08, 0.07, 50]^T \in \mathcal{D}$ and with $\boldsymbol{\eta} = \boldsymbol{\eta}_0$ is visualized in Figure 4a and the deflection in Figure 4c. The thickness and the deflection of the plate with a



damage up to 20%, i.e., a local decrease of the thickness of the plate at z by 20%, is shown in Figure 4b and 4d, respectively.

We draw M = 1000 observable parameters $\mu_1, \ldots, \mu_M \in \mathcal{D}$ uniformly in \mathcal{D} and compute the corresponding state vectors with the full model to assemble the snap-shot matrix

$$oldsymbol{Y}_{oldsymbol{\eta}_0} = [\hat{oldsymbol{y}}_{oldsymbol{\eta}_0}(oldsymbol{\mu}_1), \dots, \hat{oldsymbol{y}}_{oldsymbol{\eta}_0}(oldsymbol{\mu}_M)] \in \mathbb{R}^{N imes M}$$
 .

Note that the latent parameter $\eta = \eta_0$ is set to the nominal latent parameter η_0 . Figure 3b plots the decay of the singular values of the snapshot matrix Y_{η_0} . We construct a reduced model via Galerkin projection onto the space spanned by the first n = 8 POD basis vectors of Y_{η_0} .

5.2 Setup of numerical experiments

We now describe the details of our numerical experiments. We have ten latent parameters $\eta_0, \eta_1, \ldots, \eta_9 \in \mathcal{E}$, where η_0 is the nominal latent parameter corresponding to the no-damage condition and

$$\boldsymbol{\eta}_i = \left[\frac{2i}{90}, \frac{2i}{360}\right]^T \in \mathcal{E}, \qquad i = 1, \dots, 9.$$

This means that from latent parameter η_{i-1} to η_i the thickness at position z is decreased by a factor of two, for i = 1, ..., 9. After each change of the latent

parameter, the sensor window is flushed and $M' \in \mathbb{N}$ incomplete sensor samples are received to adapt the reduced model.

Number of sensor samples We receive incomplete sensor samples, and therefore we use the extension to the dynamic data-driven adaptation described in Section 4. This means that the adaptivity steps $h = 1, \ldots, M'$ required for adapting from latent parameter η_{i-1} to η_i are split into $M^{\text{basis}} \in \mathbb{N}$ steps to derive the gappy POD basis and $M^{\text{update}} \in \mathbb{N}$ steps to update the reduced model. We chose M^{basis} and M^{update} conservatively in the following, because we are primarily interested in studying the effect of the number of missing components in the incomplete sensor samples onto the adaptation, rather than the number of sensor samples; see [42] for studies on the effect of the number of samples. We set $M^{\text{basis}} = 5000$ and therefore derive the gappy POD basis from $M^{\text{basis}} = 5000$ incomplete sensor samples. We buffer 50 incomplete sensor samples and use them in the incremental basis generation procedure described in Section 4.1.

The theory of the dynamic data-driven adaptation with complete sensor samples gives guidance on the selection of M^{update} . In case of complete sensor samples, setting $M^{\text{update}} = l_A \times n$ is sufficient to recover the reduced model that would be obtained via rebuilding from scratch [42]. Note that $l_A = 4$ is the number of μ -independent operators and n = 8 the dimension of the POD basis space. We set $M^{\text{update}} = 5 \times l_A \times n = 160$ since we adapt from incomplete sensor samples and therefore expect that the approximation of the missing values introduces additional error into the adaptation. In total, we receive $M' = M^{\text{basis}} + M^{\text{update}} = 5160$ incomplete sensor samples to adapt from η_{i-1} to η_i for $i = 1, \ldots, 9$.

Sensor sample generation The number of missing components N - k in the incomplete sensor samples is controlled by the number of known components k. To discuss the effect of k on the adaptation, we introduce separate numbers of known components $k^{\text{basis}} \in \mathbb{N}$ and $k^{\text{update}} \in \mathbb{N}$ for the gappy POD basis construction and the update, respectively. Furthermore, we introduce the sensor rates

$$\rho^{\text{basis}} = \frac{k^{\text{basis}}}{N} \times 100, \qquad \rho^{\text{update}} = \frac{k^{\text{update}}}{N} \times 100,$$

which are the percent of the number of known components of the total number of components N in the incomplete sensor samples. Thus, for example, $\rho^{\text{basis}} = 100\%$ means that all components are known and therefore that we have a complete sensor sample.

We synthetically generate incomplete sensor samples with the full model at each step h = 1, ..., M'. We therefore first draw uniformly an observable parameter $\boldsymbol{\mu}_{M+h}$ in \mathcal{D} and compute the state vector $\boldsymbol{y}_{\boldsymbol{\eta}}(\boldsymbol{\mu}_{M+h})$ with the full model for the current latent parameter $\boldsymbol{\mu}$. We then draw $k \in \mathbb{N}$ unique indices uniformly in $\{1, ..., N\}$ and construct the point selection matrix $\boldsymbol{P}_h \in \mathbb{R}^{N \times k}$. The incomplete sensor sample is $\hat{\boldsymbol{y}}_{\boldsymbol{\eta}}^{\text{incp}}(\boldsymbol{\mu}_{M+h}) = \boldsymbol{P}_h \boldsymbol{P}_h^T \boldsymbol{y}_{\boldsymbol{\eta}}(\boldsymbol{\mu}_{M+h})$. *Error computation* We compare three reduced models:

- A static reduced model that is built as described in Section 2.2. The static reduced model is not adapted to changes in the latent parameter.
- A rebuilt reduced model that is derived as in Section 2.2 but from M^{update} complete sensor samples corresponding to the current changed latent parameter. This requires repeating the computation of the POD basis and the operator projections, which is prohibitively expensive to conduct online.
- An online adaptive reduced model that is adapted to changes in the latent parameter from incomplete sensor samples with the dynamic data-driven adaptation described in Algorithm 1.

To assess the quality of the reduced models quantitatively, we draw ten observable parameters $\mu'_1, \ldots, \mu'_{10} \in \mathcal{D}$ uniformly in \mathcal{D} and compute the relative L_2 error with respect to the full model

$$e_r = \frac{1}{10} \sum_{i=1}^{10} \frac{\|\bar{\boldsymbol{y}}_{\eta}(\boldsymbol{\mu}'_i) - \boldsymbol{y}_{\eta}(\boldsymbol{\mu}'_i)\|_2}{\|\boldsymbol{y}_{\eta}(\boldsymbol{\mu}'_i)\|_2},$$
(8)

where η is the current latent parameter and $\bar{y}_{\eta}(\mu'_1), \ldots, \bar{y}_{\eta}(\mu'_{10}) \in \mathbb{R}^n$ are the state vectors obtained with either the static, the rebuilt, or the adapted reduced model.

5.3 Gappy POD basis from complete sensor samples

We first consider the situation where $\rho^{\text{basis}} = 100\%$ is fixed and the sensor rate ρ^{update} varies. This means that we have available complete sensor samples (without missing components) for deriving the gappy POD basis in the first M^{basis} steps but incomplete sensor samples for updating the reduced model in the final M^{update} adaptivity steps.

Figures 5 and 6 demonstrate the effect of the sensor rate ρ^{update} on the dynamic data-driven adaptation. First consider the static reduced model. As the latent parameter changes from η_0 (no damage) to η_9 (20% decrease of thickness), the error of the static reduced model increases by three orders of magnitude. The steps in the error curve reflect the changes in the latent parameter. The error of the rebuilt reduced model stays near 10^{-4} . Consider now the adaptive reduced model. The dimension of the gappy POD basis is set to r = 30. Figure 5 shows that a sensor rate $\rho^{\text{update}} = 0.6\%$ leads to an adapted reduced model with large errors. A sensor rate of $\rho^{\text{update}} = 0.6\%$ means that $k^{\text{update}} = 29$ components of the incomplete sensor sample are known, and therefore $k^{\text{update}} < r$. This violates the condition of gappy POD that requires a full-column rank $P_h^T U_{M^{\text{basis}}}$, see Section 4.1. For a slightly larger sensor rate $\rho^{\text{update}} = 0.8\%$, and $k^{\text{update}} > r$, our dynamic data-driven adaptation from incomplete sensor samples recovers the rebuilt reduced model. Figure 6 indicates that increasing the sensor rate ρ^{update} reduces the error of the adapted reduced model in the first few adaptivity steps after a change in the latent parameter, cf. Figure 5.

Note that the adapted reduced model achieves a slightly lower error than the rebuilt reduced model in Figure 5 and Figure 6. The dynamic data-driven adaptation constructs the adapted operators with an optimization problem from the sensor samples projected onto the POD space. This projection and the optimization cause



Figure 5: The steps in the error curve corresponding to the static reduced model reflect the changes in the latent parameter. Changing the latent parameter from η_0 (no damage) to η_9 (20% decrease of thickness) increases the error of the static reduced model by three orders of magnitude. The adaptive reduced model uses the incomplete sensor samples to adapt to changes in the latent parameters. The gappy POD basis is derived from complete sensor samples (i.e., $\rho^{\text{basis}} = 100\%$), the sensor rate ρ^{update} for deriving the updates to the reduced model is set to $\rho^{\text{update}} = 0.6\%$ (a) and to $\rho^{\text{update}} = 0.8\%$ (b). The dimension of the reconstruction basis is set to r = 30. For $\rho^{\text{update}} = 0.6\%$, the number of known components in the incomplete sensor samples is $k^{\text{update}} = 29 < r$ and therefore the regression problem underlying gappy POD becomes underdetermined, see Section 4.1. This leads to large errors. Increasing the sensor rate to $\rho^{\text{update}} = 0.8\%$ leads to an overdetermined regression problem and therefore to lower errors of the adapted reduced model.

the difference in the error of the adapted and the rebuilt reduced model, if the dimension of the reduced model is low. The difference decreases if the dimension of the reduced model is increased, see [42, Theorem 1].

Figure 7 reports the error behavior of an adapted reduced model that uses a gappy POD basis of dimension r = 40. For $\rho^{\text{update}} = 0.6\%$ and $\rho^{\text{update}} = 0.8\%$, we again obtain the situation $k^{\text{update}} < r$ and therefore obtain an underdetermined least-squares problem that introduces large errors in the adaptation. However, if the sensor rate ρ^{update} is increased, the approximation quality of the adapted reduced model increases too. The results in Figure 6 for r = 30 are similar to the result obtained in Figure 7 for r = 40. This shows that a gappy POD basis with r = 30 dimensions is sufficient in this example.

5.4 Gappy POD basis from incomplete sensor samples

We now consider the situation where $\rho^{\text{basis}} < 100\%$ and $\rho^{\text{update}} < 100\%$, i.e., the gappy POD basis is derived from incomplete sensor samples and the updates to the reduced models are obtained from incomplete sensor samples as well. Figure 8 shows the effect of the sensor rate ρ^{basis} on the adaptation. Figures 8a and 8b demonstrate that a sensor rate $\rho^{\text{basis}} = 10\%$ is too low to recover the rebuilt reduced model with the adapted reduced model in this example. Even setting the sensor rate for the update to $\rho^{\text{update}} = 90\%$ (i.e., generating the gappy POD basis from incomplete samples with $\rho^{\text{basis}} = 10\%$ and updating the reduced model from approximate



Figure 6: The gappy POD basis is derived from complete sensor samples (i.e., $\rho^{\text{basis}} = 100\%$) but the sensor rate for the incomplete sensor samples received during the M^{update} update steps is set to $\rho^{\text{update}} = 1\%$ (a) and $\rho^{\text{update}} = 5\%$ (b). The dimension of the gappy POD basis is set to r = 30.

sensor samples with $\rho^{\text{update}} = 90\%$) cannot compensate the inadequate sensor rate $\rho^{\text{basis}} = 10\%$. Increasing the sensor rate for the gappy POD basis construction to $\rho^{\text{basis}} = 30\%$ leads to an adapted reduced model that recovers the rebuilt reduced model. However, with $\rho^{\text{basis}} = 30\%$ there still are outliers that lead to a reduced model with a large error. Figure 9 shows that increasing the sensor rate to $\rho^{\text{basis}} = 70\%$ reduces those outliers significantly. Again, increasing the dimension of the gappy POD basis from r = 30 to r = 40 only slightly reduces the error of the adapted reduced model, compare Figures 9a,c,e with Figures 9b,d,f.

Figure 10 reports the runtime of the dynamic data-driven adaptation for $\rho^{\text{basis}} =$ $30\%, \rho^{\text{update}} = 50\%$ and $\rho^{\text{basis}} = 30\%, \rho^{\text{update}} = 90\%$. The latent parameter changes from η_0 to η_9 in nine steps. For each of the nine latent parameters η_1, \ldots, η_9 , the gappy POD basis is derived and the reduced model is adapted in M^{update} steps to the incomplete sensor samples. Thus, in total, nine gappy POD bases are derived and $9 \times M^{\text{update}}$ adaptivity steps are performed for adapting from η_0 to η_9 . Figure 10 reports the total runtime split into the runtime of the gappy POD basis construction and the adaptation. The runtime of the dynamic data-driven adaptation is compared to the runtime of rebuilding the reduced model from scratch in each of the $9 \times M^{\text{update}}$ adaptivity steps. The runtime of rebuilding the reduced model is split into the runtime of inferring the latent parameter from the sensor samples and the runtime of the offline phase where the reduced operators are constructed, see Section 1. The dynamic data-driven approach achieves a speedup of about two orders of magnitude compared to rebuilding the reduced model from scratch. Increasing ρ^{update} from $\rho^{\text{update}} = 50\%$ to $\rho^{\text{update}} = 90\%$ only slightly changes the runtime of the dynamic data-driven adaptation. The runtime measurements were performed on an i5-3570 CPU.



Figure 7: Increasing the dimension of the reconstruction basis to r = 40 only slightly decreases the error of the adapted reduced model, compared to a reconstruction basis with r = 30 (see Figures 5 and 6). Note that $\rho^{\text{update}} > 0.8\%$ is required to obtain an overdetermined regression problem in gappy POD in this example.

6 Summary and future work

We proposed an extension to the dynamic data-driven adaptation that handles incomplete sensor samples, i.e., partial measurements of the large-scale state. In our approach, a gappy POD basis is derived from incomplete sensor samples. The missing values of the incomplete sensor samples are approximated with gappy POD in the space spanned by the gappy POD basis. The reduced model is then adapted using the gappy POD approximations of the complete sensor samples with the dynamic data-driven adaptation. The numerical results confirm that about 30-40% of the total number of components of the sensor samples are sufficient to recover the reduced model that would be obtained via rebuilding from scratch.

Future sensing technologies (e.g., "sensor skins") of next-generation engineering systems will provide high-resolution measurements. Processing these large data sets will be computationally challenging. In big data analytics, sublinear algorithms are currently developed that look at only a subset of the given data set to meet runtime requirements [48]. Our approach follows a similar paradigm. We selectively process sensor data that are most informative for deriving the update to the reduced model and ignore large parts of the received data that are irrelevant in the current situation. Our approach is applicable even if the selection of the high-resolution sensor data is dynamically changing online, e.g., due to new damage events.

We considered here real-time structural assessment and decision-making but sensor data are available in many other applications. For example, in control, the goal is to design a controller that stabilizes a dynamical system. However, if the dynamical system passes through multiple regimes with different system characteristics, a single controller might be insufficient to stabilize the system. If sensor data, e.g., sparse measurements of the state of the dynamical system, are available, the controller can be adapted to the sensor data to take into account the changes in the underlying dynamical system. We also mention system identification as a potential application of our adaptation approach. Instead of starting with reduced operators derived in an offline phase, one could start with initial operators that have all components set to zero, and then adapt these operators to the available data. Such a system identification approach would derive a reduced model directly from data. In general, our approach is applicable to DDDAS for which massive amounts of sensor data are available.

Acknowledgments

This work was supported in part by the AFOSR MURI on multi-information sources of multi-physics systems under Award Number FA9550-15-1-0038, program manager Jean-Luc Cambier, and by the United States Department of Energy Applied Mathematics Program, Awards DE-FG02-08ER2585 and DE-SC0009297, as part of the DiaMonD Multifaceted Mathematics Integrated Capability Center. Some of the numerical examples were computed on the computer cluster of the Munich Centre of Advanced Computing.

References

- D. Allaire, J. Chambers, R. Cowlagi, D. Kordonowy, M. Lecerf, L. Mainini, F. Ulker, and K. Willcox. An Offline/Online DDDAS capability for self-aware aerospace vehicles. *Procedia Computer Science*, 18(0):1959–1968, 2013.
- D. Amsallem and C. Farhat. An online method for interpolating linear parametric reduced-order models. SIAM Journal on Scientific Computing, 33(5):2169–2198, 2011.
- D. Amsallem, M. Zahr, and C. Farhat. Nonlinear model order reduction based on local reduced-order bases. International Journal for Numerical Methods in Engineering, 92(10):891–916, 2012.
- D. Amsallem, M. Zahr, and K. Washabaugh. Fast local reduced basis updates for the efficient reduction of nonlinear systems with hyper-reduction. *Special issue on Model Reduction of Parameterized Systems* (MoRePaS), Advances in Computational Mathematics, 2014. (accepted).
- P. Astrid, S. Weiland, K. Willcox, and T. Backx. Missing point estimation in models described by proper orthogonal decomposition. *Automatic Control, IEEE Transactions on*, 53(10):2237–2251, 2008.
- M. Barrault, Y. Maday, N. Nguyen, and A. Patera. An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathematique*, 339(9):667–672, 2004.
- P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. SIAM Review, 57(4):483–531, 2015.
- 8. M. Brand. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and its Applications*, 415(1):20–30, 2006.
- S. L. Brunton, J. H. Tu, I. Bright, and J. N. Kutz. Compressive sensing and low-rank libraries for classification of bifurcation regimes in nonlinear dynamical systems. *SIAM Journal on Applied Dynamical Systems*, 13(4):1716–1732, 2014.
- T. Bui-Thanh, M. Damodaran, and K. Willcox. Aerodynamic data reconstruction and inverse design using proper orthogonal decomposition. AIAA Journal, 42(8), 2004.
- J. Burkardt, M. Gunzburger, and H.-C. Lee. POD and CVT-based reduced-order modeling of Navier–Stokes flows. Computer Methods in Applied Mechanics and Engineering, 196(1–3):337–355, 2006.

- K. Carlberg. Adaptive h-refinement for reduced-order models. International Journal for Numerical Methods in Engineering, 102(5):1192–1210, 2015.
- K. Carlberg, C. Bou-Mosleh, and C. Farhat. Efficient non-linear model reduction via a least-squares petrov-galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods* in Engineering, 86(2):155–181, 2011.
- K. Carlberg, J. Ray, and B. van Bloemen Waanders. Decreasing the temporal complexity for nonlinear, implicit reduced-order models by forecasting. *Computer Methods in Applied Mechanics and Engineering*, 289:79 – 103, 2015.
- S. Chaturantabut and D. Sorensen. Nonlinear model reduction via discrete empirical interpolation. SIAM Journal on Scientific Computing, 32(5):2737–2764, 2010.
- F. Chinesta, P. Ladeveze, and E. Cueto. A short review on model order reduction based on proper generalized decomposition. Archives of Computational Methods in Engineering, 18(4):395–404, 2011.
- T. Cui, Y. M. Marzouk, and K. E. Willcox. Data-driven model reduction for the bayesian solution of inverse problems. *International Journal for Numerical Methods in Engineering*, 102(5):966–990, 2015.
- J. Degroote, J. Vierendeels, and K. Willcox. Interpolation among reduced-order matrices to obtain parameterized models for design, optimization and probabilistic analysis. *International Journal for Numerical Methods in Fluids*, 63(2):207–230, 2010.
- M. Dihlmann, M. Drohmann, and B. Haasdonk. Model reduction of parametrized evolution problems using the reduced basis method with adaptive time-partitioning. In D. Aubry, P. Díez, B. Tie, and N. Parés, editors, *Proceedings of the International Conference on Adaptive Modeling and Simulation*, pages 156–167, 2011.
- 20. J. Eftang and A. Patera. Port reduction in parametrized component static condensation: approximation and a posteriori error estimation. International Journal for Numerical Methods in Engineering, 96(5):269–302, 2013.
- J. Eftang and B. Stamm. Parameter multi-domain hp empirical interpolation. International Journal for Numerical Methods in Engineering, 90(4):412–428, 2012.
- R. Everson and L. Sirovich. Karhunen-Loève procedure for gappy data. Journal of the Optical Society of America A, 12(8):1657–1664, 1995.
- 23. A. Ferreira. MATLAB Codes for Finite Element Analysis. Springer, 2008.
- D. Galbally, K. Fidkowski, K. Willcox, and O. Ghattas. Non-linear model reduction for uncertainty quantification in large-scale inverse problems. *International Journal for Numerical Methods in Engineering*, 81(12):1581–1608, 2010.
- D. Garmatter, B. Haasdonk, and B. Harrach. A reduced landweber method for nonlinear inverse problems. Technical report, University of Stuttgart, 2015. Available at Arxiv.
- C. Ghnatios, F. Masson, A. Huerta, A. Leygue, E. Cueto, and F. Chinesta. Proper generalized decomposition based dynamic data-driven control of thermal processes. *Computer Methods in Applied Mechanics and Engineering*, 213–216:29 – 41, 2012.
- 27. C. Gogu. Improving the efficiency of large scale topology optimization through on-the-fly reduced order model construction. *International Journal for Numerical Methods in Engineering*, 101(4):281–304, 2015.
- D. González, F. Masson, F. Poulhaon, A. Leygue, E. Cueto, and F. Chinesta. Proper generalized decomposition based dynamic data driven inverse identification. *Mathematics and Computers in Simulation*, 82(9):1677 – 1695, 2012.
- S. Gugercin and A. Antoulas. A survey of model reduction by balanced truncation and some new results. International Journal of Control, 77(8):748–766, 2004.
- S. Kaulmann and B. Haasdonk. Online greedy reduced basis construction using dictionaries. In I. Troch and F. Breitenecker, editors, *Proceedings of 7th Vienna International Conference on Mathematical Modelling*, pages 112–117, 2012.
- M. Kennedy and A. O'Hagan. Bayesian calibration of computer models. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 63(3):425–464, 2001.
- O. Lass. Reduced order modeling and parameter identification for coupled nonlinear PDE systems. PhD thesis, University of Konstanz, 2014.
- Laura Mainini and Karen E. Willcox. A surrogate modeling approach to support real-time structural assessment and decision-making. In 10th AIAA Multidisciplinary Design Optimization Conference, AIAA SciTech. American Institute of Aeronautics and Astronautics, Jan. 2014.
- M. Lecerf, D. Allaire, and K. Willcox. Methodology for dynamic data-driven online flight capability estimation. AIAA Journal, 53(10):3073–3087, June 2015.
- Y. Maday and B. Stamm. Locally adaptive greedy approximations for anisotropic parameter reduced basis spaces. SIAM Journal on Scientific Computing, 35(6):A2417–A2441, 2013.
- L. Mainini and K. Willcox. Sensitivity analysis of surrogate-based methodology for real time structural assessment. In AIAA Modeling and Simulation Technologies Conference, AIAA SciTech 2015, AIAA Paper 2015-1362. AIAA, 2015.
- L. Mainini and K. Willcox. Surrogate Modeling Approach to Support Real-Time Structural Assessment and Decision Making. AIAA Journal, 53(6):1612–1626, Apr. 2015.
- M. Ohlberger and F. Schindler. Error control for the localized reduced basis multi-scale method with adaptive on-line enrichment. SIAM J. Sci. Comput., 2015. accepted.
- H. Panzer, J. Mohring, R. Eid, and B. Lohmann. Parametric model order reduction by matrix interpolation. at - Automatisierungstechnik, 58(8):475–484, 2010.
- B. Peherstorfer, D. Butnaru, K. Willcox, and H.-J. Bungartz. Localized discrete empirical interpolation method. SIAM Journal on Scientific Computing, 36(1):A168–A192, 2014.
- B. Peherstorfer and K. Willcox. Detecting and adapting to parameter changes for reduced models of dynamic data-driven application systems. *Procedia Computer Science*, 51:2553 – 2562, 2015.
- B. Peherstorfer and K. Willcox. Dynamic data-driven reduced-order models. Computer Methods in Applied Mechanics and Engineering, 291:21 – 41, 2015.
- 43. B. Peherstorfer and K. Willcox. Online adaptive model reduction for nonlinear systems via low-rank updates.

SIAM Journal on Scientific Computing, 37(4):A2123–A2150, 2015.

- 44. B. Peherstorfer, K. Willcox, and M. Gunzburger. Optimal model management for multifidelity Monte Carlo estimation. Technical Report 15-2, Aerospace Computational Design Laboratory, MIT, 2015.
- G. Rozza, D. Huynh, and A. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering*, 15(3):1–47, 2007.
- S. Sargsyan, S. L. Brunton, and J. N. Kutz. Nonlinear model reduction for dynamical systems using sparse sensor locations from learned libraries. *Phys. Rev. E*, 92:033304, 2015.
- 47. E. Ventsel and T. Krauthammer. Thin Plates and Shells. CRC Press, 2001.
- 48. D. Wang and Z. Han. Sublinear Algorithms for Big Data Applications. Springer, 2015.
- K. Willcox. Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. Computers & Fluids, 35(2):208 – 226, 2006.
- M. Zahr and C. Farhat. Progressive construction of a parametric reduced-order model for pde-constrained optimization. International Journal for Numerical Methods in Engineering, 102(5):1111–1135, 2015.
- 51. R. Zimmermann. A locally parametrized reduced-order model for the linear frequency domain approach to time-accurate computational fluid dynamics. *SIAM Journal on Scientific Computing*, 36(3):B508–B537, 2014.



Figure 8: The plots in (a)-(d) show that a sensor rate $\rho^{\text{basis}} < 30\%$ for the construction of the gappy POD basis is insufficient and that the rebuilt reduced model cannot be recovered with the adaptation from incomplete sensor samples in this example. Even increasing the sensor rate ρ^{update} to $\rho^{\text{update}} = 90\%$ cannot compensate the poor approximation quality of the obtained gappy POD basis. Increasing the sensor rate to $\rho^{\text{basis}} = 30\%$ leads to a gappy POD basis that approximates the complete sensor samples sufficiently well so that the dynamic data-driven adaptation recovers the rebuilt reduced model, see (e) and (f). The dimension of the gappy POD basis is set to r = 30.



Figure 9: This figure shows that a sensor rate of $\rho^{\text{basis}} = 70\%$ is sufficient to derive a gappy POD basis that avoids large errors in the first few adaptivity steps after a change in the latent parameter, cf. Figure 8 for $\rho^{\text{basis}} = 30\%$. The figure also confirms that increasing the dimension of the gappy POD basis from r = 30(a, c, e) to r = 40(b, d, f) has an insignificant effect on the adaptation.



Figure 10: This figure compares the runtime of the dynamic data-driven adaptation to rebuilding the reduced model from scratch. Adapting the reduced model using our approach is about two orders of magnitude faster than rebuilding the reduced model from scratch.