

GEOMETRIC SUBSPACE UPDATES WITH APPLICATIONS TO ONLINE ADAPTIVE NONLINEAR MODEL REDUCTION*

RALF ZIMMERMANN[†], BENJAMIN PEHERSTORFER[‡], AND KAREN WILLCOX[§]

Abstract. In many scientific applications, including model reduction and image processing, subspaces are used as ansatz spaces for the low-dimensional approximation and reconstruction of the state vectors of interest. We introduce a procedure for adapting an existing subspace based on information from the least-squares problem that underlies the approximation problem of interest such that the associated least-squares residual vanishes exactly. The method builds on a Riemannian optimization procedure on the Grassmann manifold of low-dimensional subspaces, namely the Grassmannian Rank-One Update Subspace Estimation (GROUSE). We establish for GROUSE a closed-form expression for the residual function along the geodesic descent direction. Specific applications of subspace adaptation are discussed in the context of image processing and model reduction of nonlinear partial differential equation systems.

Key words. online adaptive model reduction, dimension reduction, Grassmann manifold, Grassmannian Rank-One Update Subspace Estimation (GROUSE), discrete empirical interpolation method (DEM), gappy proper orthogonal decomposition (POD), masked projection, rank-one updates, optimization on manifolds, subspace fitting, least-squares, image processing

AMS subject classifications. 65F30, 65F20, 49M15, 65K10, 37M99

DOI. 10.1137/17M1123286

1. Introduction. Dimension reduction techniques play an important role in the application of computational methods—identifying inherent low-dimensional structure in the problem at hand can often lead to significant reductions in computational complexity. Consider a set of state vectors embedded in the n -dimensional Euclidean space \mathbb{R}^n , $n \in \mathbb{N}$. The goal of dimension reduction is to restrict the space of state vector candidates to a subspace of \mathbb{R}^n of low dimension $p \ll n$. In doing so, the n -degree-of-freedom problem of computing full-scale state vectors is replaced by the task of determining the p coefficients of a basis expansion in the reduced subspace. If, for example, the state vectors are solutions of a computational model, then this dimension reduction underlies the derivation of a projection-based reduced model. As another example, the state vectors might represent experimental data or other system samples such as representations of an image. In those cases, the dimension reduction seeks an efficient compression of the data and a low-dimensional subspace in which to reconstruct unknown states. When n is large, dimension reduction often leads to a tremendous reduction in computational complexity; however, acceptable accuracy is only retained if the full state vectors can be approximated well in the p -dimensional

*Received by the editors March 29, 2017; accepted for publication (in revised form) by P.-A. Absil November 20, 2017; published electronically February 15, 2018.

<http://www.siam.org/journals/simax/39-1/M112328.html>

Funding: The work of the first author was funded via a research fellowship awarded by the German Research Foundation (DFG), grant Zi 1250/2-1. The work of the second and third authors was supported by the U.S. Department of Energy, awards DE-FG02-08ER2585 and DE-SC0009297, as part of the DiaMonD Multifaceted Mathematics Integrated Capability Center, and the AFOSR MURI grant FA9550-15-1-0038.

[†]Department of Mathematics and Computer Science, SDU Odense, Odense DK-5230, Denmark, and Department of Aeronautics & Astronautics, MIT, Cambridge, MA 02139 (zimmermann@imada.sdu.dk).

[‡]Department of Mechanical Engineering, University of Wisconsin-Madison, Madison, WI 53706 (peherstorfer@wisc.edu).

[§]Department of Aeronautics & Astronautics, MIT, Cambridge, MA 02139 (kwillcox@mit.edu).

subspace. Thus, the identification and numerical representation of subspaces plays a critical role.

In classical projection-based model reduction, the reduced subspace is determined once in a so-called offline phase. Subsequently, it stays fixed while the reduced model is evaluated during the so-called online phase. Online adaptive model reduction breaks this division, and modifies the subspace during the evaluation process to better meet the current conditions for the reduced state vector prediction.

Online subspace adaptation can be approached from a geometric perspective: The set of all subspaces $\mathcal{U} \subset \mathbb{R}^n$ of a certain fixed dimension p forms the Grassmann manifold [2]. Subspaces are spatial locations on this manifold and are represented in numerical schemes by column-orthogonal matrices in $\mathbb{R}^{n \times p}$. One-parameter subspace modifications correspond to curves on the Grassmannian.

In the special case, where the subspace adaptation is based on a linear least-squares residual function, the *Grassmannian Rank-One Update Subspace Estimation* (GROUSE, [8]) applies: When approximating an unsampled state vector in the subspace \mathcal{U} based on partial information, the associated least-squares residual is then related to a velocity vector of a geodesic curve on the Grassmannian. GROUSE shows that this geodesic curve corresponds to a matrix curve of rank-one modifications on the underlying column-orthogonal matrices that act as subspace representatives.

Main contributions. We show that the GROUSE geodesic of rank-one updates crosses a subspace \mathcal{U}^* that allows for an *exact* representation of the given partial information. Mathematically, this is a nonlinear root-finding problem on the Grassmann manifold. We derive a closed-form expression for the residual with respect to the partial information along the GROUSE geodesic. In particular, this allows us to read off the root, but it may be of potential use in general when analyzing GROUSE with other step size schemes. As an auxiliary, we establish a general formula for the rank-one update of orthogonal projectors. Moreover, we generalize the method to subspace adaptation based on general least-squares systems and to the adaptation of a subspace of the subspace in question.

In the results section, we demonstrate that the proposed method applies in combination with the following well-established dimension reduction techniques: gappy proper orthogonal decomposition (gappy POD, [27, 18]) and discrete empirical interpolation method (DEIM, [22]). More precisely, we consider an application to gappy POD image processing, and we combine the subspace adaptation with the DEIM to construct an adaptive reduced model for the time-dependent nonlinear FitzHugh–Nagumo partial differential equation system, which models the electrical activity in a neuron. In contrast to the standard use case in the GROUSE literature [8, 49], our focus is not on estimating a subspace from scratch based on potentially noisy data but to adapt a given subspace of valid approximations based on incomplete but noise-free observations. In the DEIM setting, it is not the final subspace that is of main interest but rather the enhanced approximation capabilities after each adaptation.

Context and related work. The Grassmann manifold can be represented as a matrix manifold. For comprehensive background information on optimization on matrix manifolds, we refer the reader to [2] and its extensive bibliography. Matrix manifolds appear frequently in image processing and computer vision [35], where they often take the form of subspace identification problems. A related field of application is low-rank matrix factorizations, which arise in data analysis problems of various kinds, among them matrix completion [8, 15]. The GROUSE method was introduced in [8] as a tool for both subspace identification from incomplete and/or noisy data and the matrix completion problem and was further developed and analyzed in [10, 31, 48, 49].

A recent survey of model reduction methods for parametric systems is [14]. Most online adaptive model reduction techniques rely on precomputed quantities that restrict the way the reduced space can be changed online. One example is parametric model reduction based on the interpolation of reduced models, where reduced operators are interpolated on matrix manifolds [3, 23, 38, 4, 36, 50]. There are also dictionary approaches [30, 34] that construct a reduced space online from a subset of a large number of precomputed basis vectors, and localized reduced modeling techniques [5, 40, 26, 24] that select online one of several precomputed reduced models.

In contrast, we are interested here in online adaptive model reduction methods that derive updates to the reduced model with information that is obtained from the full model in the online phase; thus, the adaptation uses information that is unavailable in the offline phase. There are several approaches that generate new data from the full model in the online phase, or derive new reduced basis vectors with an h -refinement [21] based on an adjoint model of the full model, and then rebuild the reduced model [37, 39, 44, 45]; however, this is often computationally expensive. An efficient online adaptation that uses new data online was presented in [46, 6] for localized reduced models. A reference state is subtracted from the snapshots of localized reduced models. It is shown that this corresponds to a rank-one update of the reduced space corresponding to the localized reduced models; however, this is only a limited form of adapting a reduced model because each snapshot receives the same change. In [42, 41], dynamic reduced models are introduced that adapt to changes in the full model without requiring access to the high-fidelity operators; however, the approach is limited to linear problems and to problems where high-resolution sensor information is available that provides approximations of the full state vectors. For nonlinear problems, an adaptive DEIM was presented in [43], which derives low-rank updates to the DEIM basis from sparse data of nonlinear terms. In this paper we draw on the theory of Grassman manifolds and subspace updates to introduce a more flexible method for adaptive model reduction that applies to nonlinear problems and reproduces the inputted sparse data exactly.

Notation and preliminaries. The $(p \times p)$ -identity matrix is denoted by $I_p \in \mathbb{R}^{p \times p}$. If the dimension is clear, we will simply write I . The $(p \times p)$ -orthogonal group, i.e., the set of all square orthogonal matrices, is denoted by

$$O_p = \{R \in \mathbb{R}^{p \times p} \mid R^T R = R R^T = I_p\}.$$

For a matrix $U \in \mathbb{R}^{n \times p}$, the subspace spanned by the columns of U is denoted by $\mathcal{U} := \text{colspan}(U) := \{U\alpha \in \mathbb{R}^n \mid \alpha \in \mathbb{R}^p\} \subset \mathbb{R}^n$. The set of all p -dimensional subspaces $\mathcal{U} \subset \mathbb{R}^n$ forms the *Grassmann manifold*

$$Gr(n, p) := \{\mathcal{U} \subset \mathbb{R}^n \mid \mathcal{U} \text{ subspace, } \dim(\mathcal{U}) = p\}.$$

The *Stiefel manifold* is the compact matrix manifold of all column-orthogonal rectangular matrices

$$St(n, p) := \{U \in \mathbb{R}^{n \times p} \mid U^T U = I_p\}.$$

The Grassmann manifold can be realized as a quotient manifold of the Stiefel manifold

$$(1) \quad Gr(n, p) = St(n, p)/O_p = \{[U] \mid U \in St(n, p)\},$$

where $[U] = \{UR \mid R \in O_p\}$ is the *orbit*, or *equivalence class* of U under actions of the orthogonal group. Hence, by definition, two matrices $U, \tilde{U} \in St(n, p)$ are in the same

O_p -orbit if they differ by a $(p \times p)$ -orthogonal matrix:

$$[U] = [\tilde{U}] :\Leftrightarrow \exists R \in O_p : U = \tilde{U}R.$$

A matrix $U \in St(n, p)$ is called a *matrix representative* of a subspace $\mathcal{U} \in Gr(n, p)$ if $\mathcal{U} = \text{colspan}(U)$. We will also consider the orbit $[U]$ and the subspace $\mathcal{U} = \text{colspan}(U)$ as the same object. As in [25], we will make use throughout of the quotient representation (1) of the Grassmann manifold with matrices in $St(n, p)$ acting as representatives in numerical computations. From the manifold perspective, each p -dimensional subspace of \mathbb{R}^n is a *single point* on $Gr(n, p)$.

For a rectangular, full column rank matrix $X \in \mathbb{R}^{n \times p}$, the *orthogonal projection* onto the column span of X is

$$(2) \quad \Pi_X : \mathbb{R}^n \rightarrow \text{colspan } X, \quad y \mapsto X(X^T X)^{-1} X^T y.$$

We will consider special orthogonal projectors associated with the Cartesian coordinate directions. Let $e_j \in \mathbb{R}^n$ denote the j th canonical unit vector, $j = 1, \dots, n$. Given a subset of $m \in \mathbb{N}$ indices $J = \{j_1, \dots, j_m\} \subset \{1, \dots, n\}$, the (column-orthogonal) matrix $P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{n \times m}$ is called *the mask matrix corresponding to the index set J* . Left-multiplication of a vector with the transpose of P realizes the projection onto the selected components in the same order as listed in J , i.e., $P^T y = (y_{j_1}, \dots, y_{j_m})^T \in \mathbb{R}^m$ for all $y \in \mathbb{R}^n$. The matrix PP^T is the canonical orthogonal projection onto the coordinate axes j_1, \dots, j_m .

Throughout, whenever a mask matrix $P \in \mathbb{R}^{n \times m}$ is applied to a subspace representative $U \in St(n, p)$, we assume that $m > p$ and that the matrix of selected rows $P^T U \in \mathbb{R}^{m \times p}$ has full column rank p .

Organization. Section 2 recaps the GROUSE approach and transfers the idea of the geometric subspace adaptation to the context of model reduction. It also reviews the essentials on the numerical treatment of Grassmann manifolds. Section 3 presents the core methodological contributions of this paper, where we derive a closed-form of the Grassmann rank-one update that solves the underlying least-squares residual equation exactly. Example applications in the context of adaptive model reduction and image processing are presented in section 4, and section 5 concludes the paper.

2. Problem statement. In this section, we first summarize GROUSE following [8]. We then develop the connection between the theory of GROUSE and the task of adapting a low-dimensional subspace for model reduction. Lastly, we discuss relevant concepts in the numerical treatment of Grassmann manifolds.

2.1. GROUSE. Let $P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{n \times m}$ be a mask matrix, let $\mathcal{U}_0 \subset \mathbb{R}^n$ be a p -dimensional subspace with matrix representation $\mathcal{U}_0 = [U_0]$, $U_0 \in St(n, p)$, and let $b \in \mathbb{R}^m$ be a given data vector, $p < m < n$. GROUSE considers the masked least-squares problem

$$(3) \quad y(\mathcal{U}_0) := \arg \min_{\tilde{y} \in \mathcal{U}_0} \|P^T \tilde{y} - b\|_2^2,$$

which features the (subspace dependent) unique solution

$$(4) \quad y(\mathcal{U}_0) = U_0 \alpha(\mathcal{U}_0) \in \mathbb{R}^n, \quad \alpha(\mathcal{U}_0) = (U_0^T P P^T U_0)^{-1} U_0^T P b \in \mathbb{R}^p.$$

The corresponding residual vector $r(\mathcal{U}_0) := b - P^T y(\mathcal{U}_0)$ is, in general, nonzero. For a fixed mask matrix P and a fixed right-hand side b , the residual vector is associated

with a differentiable function on $Gr(n, p)$, the residual norm function

$$(5) \quad F_{P,b} : Gr(n, p) \rightarrow \mathbb{R}, \quad \mathcal{U} \mapsto \|r(\mathcal{U})\|_2^2 = b^T b - b^T P^T U (U^T P P^T U)^{-1} U^T P b;$$

see [8, eqs. (2) and (3)]. (The matrix U in the definition of $F_{P,b}$ can be any representative $U \in St(n, p)$ of the subspace \mathcal{U} ; see (1). The subscripts P, b will be dropped when clear from the context.) Given a sequence of incomplete observations in the form of data vectors $b_s \in \mathbb{R}^m, s = 1, 2, \dots$, with corresponding mask matrices P_s , GROUSE adapts the initial subspace such that the objective

$$(6) \quad \mathcal{U} \mapsto \sum_{s=1}^{\infty} F_{P_s, b_s}(\mathcal{U}) = \sum_{s=1}^{\infty} \|P_s^T y(\mathcal{U}) - b_s\|^2$$

is minimized; see [8, eq. (5)].¹

The GROUSE algorithm works sequentially by addressing one data vector b_s at a time. It performs a step along the geodesic line on $Gr(n, p)$ [25], [2, section 4] in the direction of steepest descent, which is given by the negative of the gradient of (5) with respect to the subspace $\mathcal{U}_0 = [U_0]$. The gradient is represented by the rank-one matrix $G = -2P(b_s - P^T U_0 \alpha_s) \alpha_s^T$ with $\alpha_s = (U_0^T P P^T U_0)^{-1} U_0^T P b_s$; see [8, eq. (9)], [25, eq. (2.70)]. The direction of steepest descent is $H = -G$. Because H is rank-one, its thin SVD $H = \Phi \Sigma V^T$ reduces to $H = \frac{Pr}{\|r\|} (\sigma_1) v^T$, where r is the residual vector, $v = \frac{\alpha_s}{\|\alpha_s\|}$ and $\sigma_1 = 2\|r\|\|\alpha\|$ is the single nonzero singular value of H . Evaluating the Grassmann geodesic [25, section 2.5.1] along this descent direction leads to

$$(7) \quad t \mapsto U_0(t) = U_0 + \left((\cos(t\sigma_1) - 1)U_0 v + \sin(t\sigma_1) \frac{Pr}{\|r\|} \right) v^T =: U_0 + \hat{x}(t)v^T;$$

see [8, eqs. (11) and (12)]. At each iteration $s = 1, 2, \dots$, the GROUSE algorithm [8, Alg. 1] chooses a step size $t = \eta_s$ and replaces the previous subspace representative U_{s-1} by $U_s = U_{s-1}(\eta_s)$ according to (7). Local and global convergence results are given in [9, 48, 49].

2.2. Subspace adaptation and model reduction. We consider here projection-based model reduction methods. These methods make use of a subspace $\mathcal{U}_0 \subset \mathbb{R}^n$ of comparatively low dimension $\dim(\mathcal{U}_0) = p \ll n$ that is assumed to contain the essential information about a set $\mathcal{X} \subset \mathbb{R}^n$ of state vectors over a range of operating conditions. More precisely, the fundamental assumption underlying the dimension reduction is that the n -dimensional state vectors $y \in \mathcal{X}$ may be approximated up to sufficient accuracy with only p degrees of freedom via

$$(8) \quad y \approx \tilde{y}(\alpha) = U_0 \alpha, \quad \alpha \in \mathbb{R}^p,$$

where $U_0 \in St(n, p)$ is a matrix representative of \mathcal{U}_0 . The standard case in model reduction is that the set of state vectors \mathcal{X} is the solution manifold of a parametric partial differential equation (PDE).

In the following, we consider the special case that only incomplete information on a state vector $y \in \mathcal{X}$ is available. This case is encountered in the model reduction techniques gappy POD [27] and DEIM [22]. The incomplete data imposes equality

¹For complete data vectors $b_s \in \mathbb{R}^n$, (6) is the same as [49, eq. (2)].

constraints on the $m < n$ components y_{j_1}, \dots, y_{j_m} of a state vector $y \in \mathcal{X}$ via the equation

$$(9) \quad P^T y = \begin{pmatrix} y_{j_1} \\ \vdots \\ y_{j_m} \end{pmatrix} =: b, \quad P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{n \times m}.$$

Under the requirement that y be contained in \mathcal{U}_0 , the underdetermined equation (9) translates into the overdetermined masked least-squares problem (3) with corresponding solution (4). This establishes a direct link to the GROUSE approach.

The objective of our work is to find a subspace $\mathcal{U}^* \in Gr(n, p)$ close to \mathcal{U}_0 such that the best subspace-restricted least-squares solution $y(\mathcal{U}^*)$ features an *exact zero residual*, $\|r(\mathcal{U}^*)\|_2 = 0$. In solving this equation for the unknown \mathcal{U}^* , we adapt the original reduced subspace \mathcal{U}_0 according to the least-squares problem arising from the new (partial) information about y . The requirement of \mathcal{U}^* being close to \mathcal{U}_0 is important in the context of model reduction because we want the approximation (8) to remain valid for \mathcal{U}^* .

We formalize the objective. Define the *feasibility set*

$$(10) \quad \mathcal{Z} := \left\{ \mathcal{U} \in Gr(n, p) \mid \min_{\tilde{y} \in \mathcal{U}} \|P^T \tilde{y} - b\|_2 = 0 \right\}.$$

The set \mathcal{Z} is nonempty.² From GROUSE, it is known that the geodesic curve $t \mapsto \mathcal{U}(t)$ that starts in $\mathcal{U}(0) = \mathcal{U}_0$ with velocity given by the direction of steepest descent of the residual norm function (5) is a matrix curve of rank-one updates on the initial subspace \mathcal{U}_0 ; see (7). We will show that this curve crosses the feasibility set \mathcal{Z} and determine the first intersection point. By writing the residual vector as $r(\mathcal{U}_0) = b - \Pi_{P^T \mathcal{U}_0} b$, where $\Pi_{P^T \mathcal{U}_0}$ is the orthogonal projection (2) onto $\text{colspan}(P^T \mathcal{U}_0)$, this objective becomes a *nonlinear equation* on the Grassmann manifold:

$$(11) \quad \text{solve } b - \Pi_{P^T \mathcal{U}(t^*)} b = 0 \text{ for } t^* \in \mathbb{R}.$$

The condition $b - \Pi_{P^T \mathcal{U}(t^*)} b = 0$ is equivalent to $[\mathcal{U}(t^*)] \in \mathcal{Z}$.

A contribution of this paper is an explicit formula for the time-dependent residual $r(\mathcal{U}(t)) = b - \Pi_{P^T \mathcal{U}(t)} b$ derived in section 3, from which the solution to (11) can be read off in closed form. In contrast to GROUSE, whose overall aim is the iterative global minimization of (6), we focus on the single adaptation steps and the nonlinear residual equation on $Gr(n, p)$. We arrive in this way at the same formula for t^* that was obtained in [49, Alg. 1, section 3.1, App. C] as the optimal greedy step size in an iterative subspace updating scheme based on complete right-hand side vectors.

In summary, our approach is a method for determining a subspace \mathcal{U}^* contained in the set \mathcal{Z} from (10) that can be reached via a geodesic path along the descent direction starting in \mathcal{U}_0 . Figure 1 and section SM1 from the supplement illustrate this principle. In subsection 3.3, we show that this is not restricted to the special case of masked least-squares problems $\|P^T \tilde{y} - b\|_2$ but can be generalized to arbitrary underdetermined systems $\|A \tilde{y} - b\|_2$, $A \in \mathbb{R}^{m \times n}$.

2.3. Numerical aspects of the Grassmann manifold. Our approach to solve (11) is presented in section 3 and builds on geometric concepts on the Grassmann manifold $Gr(n, p)$. This subsection reviews a few essential aspects of the numerical treatment of Grassmann manifolds. We refer the reader to [1, 2, 25] for details.

²Any subspace \mathcal{U} that contains a vector $y = Pb + v$, where $v \in \mathbb{R}^n$ is in the $(n - m)$ -dimensional kernel of P^T is in \mathcal{Z} .




FIG. 1. Graphical illustration of the geometric subspace adaptation: The sphere visualizes the Grassmann manifold $Gr(n, p)$. The solid line marks the set \mathcal{Z} of all subspaces in $Gr(n, p)$ that contain zero-residual solutions to the least-squares problem (3). The black triangle shows the initial subspace \mathcal{U}_0 . The dashed line is the geodesic starting in \mathcal{U}_0 with velocity given by minus the gradient of the least-squares residual function. Our goal is to compute the subspace \mathcal{U}^* , where the geodesic meets the set \mathcal{Z} .

Tangent spaces and normal coordinates. The tangent space $T_{\mathcal{U}}Gr(n, p)$ at a point $\mathcal{U} \in Gr(n, p)$ can be thought of as the space of velocity vectors of differentiable curves on $Gr(n, p)$ passing through \mathcal{U} . For any matrix representative $U \in St(n, p)$ of $\mathcal{U} \in Gr(n, p)$ the tangent space of $Gr(n, p)$ at \mathcal{U} is represented by

$$T_{\mathcal{U}}Gr(n, p) = \{\Delta \in \mathbb{R}^{n \times p} \mid U^T \Delta = 0\} \subset \mathbb{R}^{n \times p},$$

its canonical metric being $\langle \Delta, \tilde{\Delta} \rangle_{Gr} = \text{tr}(\Delta^T \tilde{\Delta})$ [25, section 2.5]. Endowing each tangent space with this metric turns $Gr(n, p)$ into a *Riemannian manifold*. A geodesic $t \mapsto \mathcal{U}(t)$ on $Gr(n, p)$ is a locally length-minimizing curve. A geodesic is uniquely determined by its starting point $\mathcal{U}(0)$ and its starting velocity $\dot{\mathcal{U}}(0) = \Delta \in T_{\mathcal{U}_0}Gr(n, p)$, [2, p. 102].

The corresponding Riemannian *exponential mapping* is

$$\text{Exp}_{\mathcal{U}_0} : T_{\mathcal{U}_0}Gr(n, p) \rightarrow Gr(n, p), \quad \Delta \mapsto \text{Exp}_{\mathcal{U}_0}(\Delta) := \mathcal{U}(1).$$

The Riemannian exponential maps a tangent vector $\Delta \in T_{\mathcal{U}_0}Gr(n, p)$ to the endpoint $\mathcal{U}(1)$ of a geodesic path $\mathcal{U} : [0, 1] \rightarrow Gr(n, p)$ starting at $\mathcal{U}(0) = \mathcal{U}_0 \in Gr(n, p)$ with velocity $\Delta \in T_{\mathcal{U}_0}Gr(n, p)$.

An efficient algorithm for evaluating the Grassmann exponential is derived in [25, section 2.5.1]. The explicit form of the associated geodesic is

$$(12) \quad \mathcal{U}(t) = \text{Exp}_{\mathcal{U}_0}(t\Delta) = [U_0 V \cos(t\Sigma)V^T + \Phi \sin(t\Sigma)V^T], \quad \Delta \stackrel{\text{SVD}}{=} \Phi \Sigma V^T.$$

The exponential mapping gives a local parametrization from the (flat, Euclidean) tangent space to the manifold. This is also referred to as representing the manifold in *normal coordinates* [32, section III.8], [33, Lem. 5.10].

Distance between subspaces. Given two subspaces $[U], [\tilde{U}] \in Gr(n, p)$, the *ith canonical or principal angle* between $[U]$ and $[\tilde{U}]$ is $\theta_i := \arccos(\sigma_i) \in [0, \frac{\pi}{2}]$, where σ_i is the *ith-largest singular value* of $U^T \tilde{U} \in \mathbb{R}^{p \times p}$ [29, section 12.4.3].

The Riemannian distance between $[U], [\tilde{U}] \in Gr(n, p)$ is

$$(13) \quad \text{dist}([U], [\tilde{U}]) := \|\Theta\|_2, \quad \Theta = (\theta_1, \dots, \theta_p) \in \mathbb{R}^p.$$

Normal coordinates are radially isometric with respect to the Riemannian distance on $Gr(n, p)$ and the canonical metric on $T_{\mathcal{U}_0}Gr(n, p)$ in the following sense: the length of a tangent vector Δ as measured by the metric in $T_{\mathcal{U}_0}Gr(n, p)$ is the same as the Riemannian distance $\text{dist}(\mathcal{U}_0, \text{Exp}_{\mathcal{U}_0}(\Delta))$ on $Gr(n, p)$, provided that Δ is in a neighborhood of $0 \in T_{\mathcal{U}_0}Gr(n, p)$, where the exponential is invertible [33, Lem. 5.10 and Cor. 6.11].

The Grassmann manifold is a compact homogeneous space [32]. In particular, by [47, Thm. 8(b)], any two points on $Gr(n, p)$ can be connected by a geodesic of length $\leq \frac{\sqrt{p}}{2}\pi$. This is related to the so-called *injectivity radius* of the Grassmann manifold [47], which is the maximal radius ρ such that the exponential map at *any point* $[U] \in Gr(n, p)$ is a diffeomorphism onto the open ball $B(0, \rho) \subset T_{[U]}Gr(n, p)$ around the origin in the corresponding tangent space. The injectivity radius of the Grassmann manifold is $\rho = \frac{\pi}{2}$ [47]. This concept is relevant to the step of conducting the line search within Grassmann optimization schemes. We make the following observation: Using the explicit formulas for the exponential mapping and its (local) inverse, called the logarithmic mapping $\text{Log}_{[U]}$ (see [12, section 3]), one can show that $\text{Log}_{[U]} \circ \text{Exp}_{[U]}(\Delta) = \Delta$ for all tangent vectors Δ of *spectral* norm $\|\Delta\|_2 = \sigma_1(\Delta) < \pi/2$, where $\sigma_1(\Delta)$ is the largest singular value of Δ . As a consequence, we have the following observation.

Observation 1. For all $[U] \in Gr(n, p)$, let

$$B_{[U], \text{spec}}(0, \pi/2) := \left\{ \Delta \in T_{[U]}Gr(n, p) \mid \sigma_1(\Delta) < \frac{\pi}{2} \right\}.$$

Then the exponential mapping $\text{Exp}_{[U]}$ is a radial isometry on $B_{[U], \text{spec}}(0, \pi/2)$.

This observation is important for numerical computations because

$$B_{[U], \text{spec}}(0, \pi/2) \supset \left\{ \Delta \in T_{[U]}Gr(n, p) \mid \sqrt{\langle \Delta, \Delta \rangle_{Gr}} = \|(\sigma_1, \dots, \sigma_p)^T\|_2 < \frac{\pi}{2} \right\},$$

i.e., the spectral $\pi/2$ -ball in the tangent space encloses the canonical $\pi/2$ -ball in the tangent space. The above observation leads to the next proposition which has implications on the uniqueness of solutions to (11).

PROPOSITION 1. *Let $[U] \in Gr(n, p)$, $\Delta \in T_{[U]}Gr(n, p)$, and $\tilde{U} = \text{Exp}_{[U]}(\Delta)$. If $\|\Delta\|_2 < \frac{\pi}{2}$, then $\text{dist}([U], [\tilde{U}]) = \|\Delta\|_{Gr}$. In particular, the length of the geodesic path starting in $[U]$ and ending in $[\tilde{U}]$ is less than $\frac{\sqrt{p}}{2}\pi$.*

Proof. Let $\Delta \stackrel{\text{SVD}}{=} \Phi \Sigma V^T$ with $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_p)$ and let $\sigma_1 = \|\Delta\|_2 < \frac{\pi}{2}$. The exponential projection of Δ onto $Gr(n, p)$ is $[\tilde{U}] = \text{Exp}_{[U]}(\Delta) = [UV \cos(\Sigma)V^T + \Phi \sin(\Sigma)V^T]$.

The SVD of $U^T \tilde{U}$ is $V \cos(\Sigma)V^T$, so that $0 \leq \theta_k := \arccos(\cos(\sigma_k)) = \sigma_k < \frac{\pi}{2}$. Hence, $(\sigma_1, \dots, \sigma_p)^T = (\theta_1, \dots, \theta_p)^T := \Theta \in \mathbb{R}^p$ is precisely the vector of canonical angles between $[U]$ and $[\tilde{U}]$ (when listing the canonical angles in descending order); see (13). As a consequence,

$$\text{dist}([U], [\tilde{U}]) = \|\Theta\|_2 = \sqrt{\text{tr}(\Sigma^2)} = \sqrt{\text{tr}(\Delta^T \Delta)} = \|\Delta\|_{Gr}.$$

Since $\sigma_1 < \frac{\pi}{2}$, we have $\|\Delta\|_{Gr} = (\sum_{i=1}^p \sigma_p^2)^{1/2} < \frac{\sqrt{p}}{2}\pi$. □

A subtlety of Proposition 1 is that the length condition on Δ is with respect to the spectral norm rather than the canonical norm.

3. Solving the Grassmann residual equation. We now return to our goal formulated in subsection 2.2: the solution of (11). In subsection 3.1, we derive a general update formula for orthogonal projectors under rank-one modifications. Subsection 3.2 derives an explicit time-dependent expression for the Grassmann residual along the GROUSE geodesic. In particular, this allows us to read off the closed-form solution to (11). A generalization to least-squares systems featuring arbitrary matrices rather than mask matrices as operators is given in subsection 3.3. Subsequently, subsection 3.4 introduces an extension for performing the Grassmann subspace adaptation over selected directions of the subspace only.

3.1. A closed-form rank-one update for orthogonal projectors. In this subsection, we derive a formula for orthogonal projectors under rank-one updates that turns out to be an essential building block in solving (11). As this result is also of independent interest, we state it in a more general setting.

Let $X \in \mathbb{R}^{m \times p}$. Recall from (2) that the orthogonal projection onto $\text{colspan } X$ is $\Pi_X = X(X^T X)^{-1} X^T$. Let $X \stackrel{\text{SVD}}{=} Q \Sigma R^T$ be the thin SVD of X with $Q \in St(m, p)$, $\Sigma \in \mathbb{R}^{p \times p}$ diagonal, $R \in O_p$ orthogonal. Then Π_X is expressed alternatively as $\Pi_X = Q Q^T$.

Let $x \in \mathbb{R}^m$, $v \in \mathbb{R}^p$, and consider the rank-one update

$$X_{new} = X + x v^T \in \mathbb{R}^{m \times p}.$$

We are interested in an expression $\Pi_{X_{new}} = Q_{new} Q_{new}^T$, where $Q_{new} \in St(m, p)$. One standard way to approach this is via rank-one SVD updates [19, 17]. However, this requires an auxiliary SVD of a $(p \times p)$ -matrix. Here, we can avoid this, since we are not interested in the fully updated $X_{new} \stackrel{\text{SVD}}{=} Q_{new} \Sigma_{new} R_{new}^T$ or even in Q_{new} alone but only in $Q_{new} Q_{new}^T$.

LEMMA 2. *As in the above setting, let $X \stackrel{\text{SVD}}{=} Q \Sigma R^T$, $X_{new} = X + x v^T$, and define*

$$(14a) \quad \tilde{q} = x - Q Q^T x, \quad q = \frac{\tilde{q}}{\|\tilde{q}\|_2} \in \mathbb{R}^m,$$

$$(14b) \quad g = \begin{pmatrix} g_p \\ g_{p+1} \end{pmatrix} = \begin{pmatrix} -\Sigma^{-1} R^T v \\ \frac{1}{\|\tilde{q}\|_2} (1 + x^T Q \Sigma^{-1} R^T v) \end{pmatrix} \in \mathbb{R}^{p+1}.$$

Then the orthogonal projection onto $\text{colspan}(X_{new})$ is

$$(15) \quad \Pi_{X_{new}} = (Q, q) \begin{pmatrix} Q^T \\ q^T \end{pmatrix} - \frac{1}{\|g\|_2^2} (Q, q) g g^T \begin{pmatrix} Q^T \\ q^T \end{pmatrix}.$$

Proof. We start with a decomposition inspired by [17, eq. (3)]. Note that $(Q, q) \in St(m, p+1)$ by construction. It holds that

$$X + x v^T = (Q, q) \begin{pmatrix} \Sigma R^T + Q^T x v^T \\ \|\tilde{q}\|_2 v^T \end{pmatrix} =: (Q, q) M,$$

where $M \in \mathbb{R}^{(p+1) \times p}$. Let $M \stackrel{\text{SVD}}{=} \tilde{Q} \tilde{\Sigma} \tilde{R}^T$ be the thin SVD of M , i.e., $\tilde{Q} \in St(p+1, p)$, $\tilde{\Sigma}, \tilde{R}^T \in \mathbb{R}^{p \times p}$. Formally, the updated SVD is

$$X + x v^T = \left((Q, q) \tilde{Q} \right) \tilde{\Sigma} \tilde{R}^T =: Q_{new} \Sigma_{new} R_{new}^T.$$

Let $g \in \mathbb{R}^{p+1}$ be such that $(\tilde{Q}, \frac{g}{\|g\|}) \in O_{p+1}$ is an orthogonal completion of \tilde{Q} . Because of $I_{p+1} = (\tilde{Q}, \frac{g}{\|g\|})(\tilde{Q}, \frac{g}{\|g\|})^T$, we have

$$\tilde{Q}\tilde{Q}^T = I_{p+1} - \frac{1}{\|g\|^2}gg^T$$

and, as a consequence,

$$(16) \quad Q_{new}Q_{new}^T = (Q, q)\tilde{Q}\tilde{Q}^T \begin{pmatrix} Q^T \\ q^T \end{pmatrix} = (Q, q) \left(I_{p+1} - \frac{1}{\|g\|^2}gg^T \right) \begin{pmatrix} Q^T \\ q^T \end{pmatrix}.$$

Hence, it is sufficient to determine g , which is characterized up to a scalar factor by $\tilde{Q}^Tg = 0$. Since $\text{colspan}(M) = \text{colspan}(\tilde{Q})$, this condition is equivalent to $M^Tg = 0$. Let $g_p \in \mathbb{R}^p$ denote the first p components of g and let $g_{p+1} \in \mathbb{R}$ be the last entry such that $g^T = (g_p^T, g_{p+1})$. When writing the equation $g^TM = 0$ as

$$(g_p^T, g_{p+1}) \begin{pmatrix} \Sigma & Q^T x \\ 0 & \|\tilde{q}\|_2 \end{pmatrix} \begin{pmatrix} R^T \\ v^T \end{pmatrix} = 0,$$

it is straightforward to show that

$$g = \begin{pmatrix} -\Sigma^{-1}R^T v \\ \frac{1}{\|\tilde{q}\|_2}(1 + x^T Q \Sigma^{-1} R^T v) \end{pmatrix} \in \mathbb{R}^{p+1}$$

and any scalar multiple of this vector is a valid solution. Using this vector in (16) proves the lemma. \square

3.2. An explicit expression for the Grassmann residual function along the GROUSE geodesic. We now state our main theorem on the solution of the nonlinear equation (11).

THEOREM 3. *Let $U_0 = [U_0] \in Gr(n, p)$ be represented by $U_0 \in St(n, p)$. Let $P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{(n \times m)}$ be a mask matrix. Moreover, let $b \in \mathbb{R}^m$ and suppose that $U_0^T P b \neq 0$.*

Let $\alpha = (U_0^T P P^T U_0)^{-1} U_0^T P b$ be the optimal coefficient vector corresponding to the masked least-squares problem

$$\min_{\tilde{\alpha} \in \mathbb{R}^p} \|P^T U_0 \tilde{\alpha} - b\|^2$$

and let $r = b - P^T U_0 \alpha$ the associated residual vector. Set $v = \frac{\alpha}{\|\alpha\|_2}$ and $s_1 = 2\|r\|_2 \|\alpha\|_2$. Moreover, write $P^T U_0 \stackrel{SVD}{=} Q \Sigma R^T \in \mathbb{R}^{m \times p}$ and $g_p = -\Sigma^{-1} R^T v$. The t -dependent residual vector along the geodesic descent direction is

$$r([U(t)]) = b - \Pi_{P^T U(t)} b = \frac{\|r\|_2 - \|\alpha\|_2 \tan(ts_1)}{1 + \tan^2(ts_1) \|g_p\|^2} \left(\frac{r}{\|r\|_2} - \frac{\tan(ts_1)}{\|\alpha\|_2} Q \Sigma^{-2} Q^T b \right).$$

Proof. Reconsider (7) and let

$$x(t) = P^T \hat{x}(t) = (\cos(ts_1) - 1) P^T U_0 v + \sin(ts_1) \frac{r}{\|r\|_2},$$

$$v = \frac{\alpha}{\|\alpha\|_2}, \quad \alpha = (U_0^T P P^T U_0)^{-1} U_0^T P b,$$

so that

$$P^T U(t) = P^T U_0 + x(t)v^T.$$

Since $P^T U(t)$ is a rank-one update of $P^T U_0$, Lemma 2 applies. Introducing $P^T U_0 \stackrel{\text{SVD}}{=} Q\Sigma R^T \in \mathbb{R}^{m \times p}$, we obtain $r = b - QQ^T b$ and $\alpha = R\Sigma^{-1}Q^T b$. The t -dependent orthogonal projection onto $\text{colspan}(P^T U(t))$ is

$$(18) \quad \Pi_{P^T U(t)} = (Q, q(t)) \begin{pmatrix} Q^T \\ q^T(t) \end{pmatrix} - \frac{1}{\|g(t)\|_2^2} (Q, q(t)) g(t) g^T(t) \begin{pmatrix} Q^T \\ q^T(t) \end{pmatrix},$$

where

$$\begin{aligned} \tilde{q}(t) &= x(t) - QQ^T x(t), \quad q(t) = \frac{\tilde{q}(t)}{\|\tilde{q}(t)\|_2} \in \mathbb{R}^m, \\ g(t) &= \begin{pmatrix} g_p \\ g_{p+1}(t) \end{pmatrix} = \begin{pmatrix} -\Sigma^{-1} R^T v \\ \frac{1}{\|\tilde{q}(t)\|_2} (1 + x^T(t) Q \Sigma^{-1} R^T v) \end{pmatrix} \in \mathbb{R}^{p+1}. \end{aligned}$$

We have $Q^T r = 0$ and thus $Q^T x(t) = \frac{\cos(ts_1) - 1}{\|\alpha\|_2} Q^T b$. This leads to $\tilde{q}(t) = \frac{\sin(t)}{\|r\|_2} r$ and $\|\tilde{q}(t)\|_2 = |\sin(t)|$ as well as $q(t) = \text{sign}(\sin(t)) \frac{r}{\|r\|_2} = \pm q$, where we standardize $q = \frac{r}{\|r\|_2}$. Moreover,

$$x^T(t) Q \Sigma^{-1} R^T v = \frac{1}{\|\alpha\|_2^2} (\cos(ts_1) - 1) \underbrace{b^T Q \Sigma^{-2} Q^T b}_{\|\alpha\|_2^2} = (\cos(ts_1) - 1),$$

so that $g(t)$ is

$$g(t) = \begin{pmatrix} -\frac{1}{\|\alpha\|_2} \Sigma^{-2} Q^T b \\ \frac{\cos(ts_1)}{|\sin(ts_1)|} \end{pmatrix} \in \mathbb{R}^{p+1}.$$

It holds that $\frac{\cos(ts_1)}{|\sin(ts_1)|} q(t) = \frac{\cos(ts_1)}{\sin(ts_1)} q$. Hence, according to (18), we may consistently work with $+q$ and $\frac{\cos(ts_1)}{\sin(ts_1)} = \cot(ts_1)$. In order to evaluate the updated projection (18), we compute

$$\begin{aligned} (Q, q) \begin{pmatrix} g_p \\ g_{p+1}(t) \end{pmatrix} &= -\frac{1}{\|\alpha\|_2} Q \Sigma^{-2} Q^T b + \cot(ts_1) q, \\ g_p^T Q^T b &= -\frac{1}{\|\alpha\|_2} b^T Q \Sigma^{-2} Q^T b = -\|\alpha\|_2, \text{ and} \\ q^T b &= \frac{1}{\|r\|_2} r^T b = \frac{1}{\|r\|_2} \underbrace{(b^T b - b^T Q Q^T b)}_{\|r\|_2^2} = \|r\|_2. \end{aligned}$$

Substituting these identities in (18), we arrive at

$$(20) \quad \begin{aligned} r([U(t)]) &= b - \Pi_{P^T U(t)} b = b - QQ^T b - q q^T b \\ &\quad + \frac{1}{\|g(t)\|_2} (Q g_p + \cot(ts_1) q) (g_p^T Q^T b + \cot(ts_1) q^T b) \\ &= \frac{\cot(ts_1) \|r\|_2 - \|\alpha\|_2}{\|g(t)\|_2^2} \left(\cot(ts_1) \frac{r}{\|r\|_2} - \frac{1}{\|\alpha\|_2} Q \Sigma^{-2} Q^T b \right), \end{aligned}$$

as was claimed. \square

Note that the only special property of P that is exploited in the proof is that $P^T P r = r$. Hence, the result holds when P is replaced with an arbitrary column-orthogonal matrix.

There is a number of conclusions that can be drawn from Theorem 3.

COROLLARY 4.

1. The t -dependent residual norm along the steepest descent direction is

$$(21) \quad \|r([U(t)])\|_2 = \|b - \Pi_{P^T U(t)} b\|_2 = \frac{|\|r\|_2 - \|\alpha\|_2 \tan(ts_1)|}{\sqrt{1 + \|g_p\|_2^2 \tan^2(ts_1)}}.$$

2. The residual norm function is continuous and $\frac{\pi}{s_1}$ -periodic along the steepest descent direction with

$$\|r([U(0)])\|_2 = \|r\|_2 \quad \text{and} \quad \left\| \left[U \left(\frac{\pi}{2s_1} \right) \right] \right\|_2 = \frac{\|\alpha\|_2}{\|g_p\|_2} = \frac{\|\alpha\|_2^2}{\|Q\Sigma^{-2}Q^T b\|_2}.$$

3. The first root along the geodesic descent direction is at

$$(22) \quad t^* = \frac{1}{s_1} \arctan \left(\frac{\|r\|_2}{\|\alpha\|_2} \right) \in \left(0, \frac{\pi}{2s_1} \right).$$

The associated matrix $U^* := U_0 + ((\cos(t^* s_1) - 1)U_0 v + \sin(t^* s_1) \frac{Pr}{\|r\|})v^T$ is such that the subspace $\mathcal{U}^* := [U^*]$ is contained in the set \mathcal{Z} from (10), i.e.,

$$(23) \quad F(\mathcal{U}^*) = \min_{\tilde{\alpha} \in \mathbb{R}^p} \|P^T U^* \tilde{\alpha} - b\|^2 = 0.$$

Stated differently, it holds that b is contained in $\text{colspan}(P^T U^*)$, that is, $b = \Pi_{P^T U^*} b$.

4. The coefficient vector associated with $\mathcal{U}^* = [U^*] = (23)$ is $\alpha^* = \sqrt{1 + \frac{\|r\|_2^2}{\|\alpha\|_2^2}} \alpha$. The associated $y^* \in \mathbb{R}^n$ is $y^* = U^* \alpha^* = U_0 \alpha + Pr = U_0 \alpha + P(b - P^T U_0 \alpha)$. Hence, y^* can be readily obtained without computing any of t^*, α^*, U^* .
5. The first maximum along the geodesic descent direction is at

$$t_{max} = \frac{1}{s_1} \left(\pi - \arctan \left(\frac{\|\alpha\|_2}{\|r\|_2 \|g_p\|_2^2} \right) \right) \in \left(\frac{\pi}{2s_1}, \frac{\pi}{s_1} \right)$$

with corresponding value $\|r([U(t_{max})])\|_2 = \sqrt{\|r\|_2^2 + \frac{\|\alpha\|_2^2}{\|g_p\|_2^2}}$.

Proof. By taking into account that r is orthogonal to $\text{colspan}(Q)$, Pythagoras' theorem gives $\|(\cot(ts_1) \frac{r}{\|r\|_2} - \frac{1}{\|\alpha\|_2} Q\Sigma^{-2}Q^T b)\|_2 = \sqrt{\cot^2(ts_1) + \|g_p\|_2^2} = \|g(t)\|_2$. The formula (21) is now an immediate consequence of (20). From (21), statements 2, 3, and 5 of the corollary are straightforward.

On statement 4: From statement 3, we know that there exists $\alpha^* \in \mathbb{R}^p$ such that $P^T U^* \alpha^* - b = 0$. After plugging in the explicit expression for U^* , we obtain the equation

$$P^T U_0 \left(\alpha^* - \frac{\alpha^T \alpha^*}{\|\alpha\|_2^2} \alpha \right) + \left(\frac{\alpha^T \alpha^*}{\|\alpha\|_2 \sqrt{\|\alpha\|_2^2 + \|r\|_2^2}} - 1 \right) b = 0.$$

If the unmodified least-squares problem (3) features a nonzero residual, then b is not contained in $\text{colspan } P^T U_0$. Hence, both quantities in the round brackets must be zero, which leads to $\alpha^* = \frac{\alpha^T \alpha^*}{\|\alpha\|_2^2} \alpha = \frac{\sqrt{\|\alpha\|_2^2 + \|r\|_2^2}}{\|\alpha\|_2} \alpha$. The calculation of y^* is, therefore, straightforward. \square

Appendix A features a shortcut to statements 3 and 4 of Corollary 4. An example of a plot of the residual norm function (21) from a practical application is displayed in Figure 5.

Remark 5. The GROUSE convergence analysis in [10] is based on local considerations and a step length of $\tilde{t} = \frac{1}{s_1} \arcsin\left(\frac{\|r\|_2}{\|\alpha\|_2}\right)$, which matches the t^* in (22) up to terms of third order, when the residual and, therefore, the ratio $\|r\|_2/\|\alpha\|_2$ are small. In the fully sampled case, that is, when complete right-hand side data is available, [49] shows that the same t^* of (22) is also the greedy-optimal step with respect to the determinant-similarity and the Frobenius norm discrepancy of two subspaces in an iterative subspace updating scheme; see [49, section 3.1 and App. C]. In contrast, we arrived at t^* from the independent approach of solving the nonlinear equation (11) and with a different proof that relies on Lemma 2. Combining these facts shows that the subspace discrepancy is maximal if and only if the subspace update is such that the residual vanishes exactly.

The proof of Proposition 1 shows that the distance between the subspaces $[U_0]$ and $[U^*]$ is $t^*s_1 = \arctan\left(\frac{\|r\|_2}{\|\alpha\|_2}\right) < \frac{\pi}{2}$. Hence, when performing the t^* -optimal rank-one update on $[U_0]$ according to Corollary 4, we stay within the injectivity radius. As a consequence from general differential geometry, the geodesic $t \mapsto [U(t)]$ is length-minimizing, that is, there is no shorter curve on $Gr(n, p)$ that connects $[U_0]$ and $[U^*]$.³

We emphasize that the update formula of Lemma 2 for orthogonal projectors under rank-one modifications was used as an intermediate theoretical fact in proving Theorem 3 but that it is not required to actually compute the rank-one update and the associated quantities Q, q, g in order to obtain the optimal t^* and the subspace $[U^*] = [U(t^*)]$. MATLAB code that considers this fact is in the supplement section SM4.

We draw a corollary that corresponds to the special case where the mask matrix P is the identity I_n , i.e., the case where complete data is available. Recall that the best least-squares approximation to a given vector b that is contained in a subspace \mathcal{U}_0 is the orthogonal projection $U_0U_0^T b$ of b onto \mathcal{U}_0 , with an associated residual of $r = b - U_0U_0^T b$. The SVD of $P^T U_0$ is now trivially $P^T U_0 = Q\Sigma R^T = U_0 I_p I_p^T$ so that the expressions involving Q, Σ, R simplify.

COROLLARY 6. *Let $\mathcal{U}_0 = [U_0] \in Gr(n, p)$ be represented by $U_0 \in St(n, p)$. Let $b \in \mathbb{R}^n$ and suppose that $\alpha := U_0^T b \neq 0$. Set $v = \frac{\alpha}{\|\alpha\|_2}$ and $s_1 = 2\|r\|_2\|\alpha\|_2$. Then the t -dependent residual norm is*

$$\|r([U(t)])\|_2 = \|b - \Pi_{U(t)} b\|_2 = \frac{|\|r\|_2 - \|\alpha\|_2 \tan(ts_1)|}{\sqrt{1 + \tan^2(ts_1)}}.$$

Define

$$t^* = \frac{1}{s_1} \arctan\left(\frac{\|r\|_2}{\|\alpha\|_2}\right).$$

Then $U^ := U(t^*) := U_0 + ((\cos(t^*s_1) - 1)U_0 v + \sin(t^*s_1)\frac{r}{\|r\|})v^T$ is such that b is contained in the subspace $\mathcal{U}^* := [U^*]$, i.e., $b = \Pi_{U^*} b$.*

Remark 7. Corollary 6 has a connection with rank-one SVD updates as considered in [19, 16, 17]. One application in [17, Table 1] is to *revise* an existing SVD

³This does not necessarily mean that there is no other point $[\tilde{U}^*] \in \mathcal{Z}$ that is closer to $[U_0]$.

$U_0 \Sigma_0 V_0^T = (X, c)$ such that the column c is replaced with a column b in the modified SVD $U' \Sigma' V'^T = (X, b)$. In terms of the associated orthogonal projectors, we have $U' U'^T b = b$. With Corollary 6, we obtain a subspace $[U^*]$ that also contains b . Yet, this is not achieved by explicitly exchanging a column c of the original data matrix for the new column b . Rather, via the update $U_0 + ((\cos(t^* s_1) - 1)U_0 v + \sin(t^* s_1) \frac{r}{\|r\|}) v^T =: U_0 + x^* v^T$, the missing residual part is distributed over all columns of the original representative U_0 . In order to emulate this with the “revise”-approach of [17, Table 1], one first has to rotate the subspace representative with $\Phi = (v^\perp, v) \in O_p$, so that $(U_0 + x^* v^T) \Phi = U_0 \Phi + (0, \dots, 0, x^*)$, i.e., the rank-one update acts on a single direction of the new representative $U_0 \Phi$. Allowing for rotations of the representative U_0 in the update scheme enables more general updates than when working with a fixed representative U_0 . Hence, we expect that $\text{dist}([U_0], [U^*]) \leq \text{dist}([U_0], [U'])$. This is confirmed in the example featured in subsection 4.2.

Another relation between GROUSE and the incremental SVD of [16] was exposed in [9]. The approach considered in [9] corresponds to first attaching new column data to a given subspace representative. Then, the SVD update is performed on the augmented matrix representative and consequently retruncated to its original dimensions. It is shown that this procedure can be emulated via GROUSE when a specific step size is chosen for the rank-one increment. However, the modified U' obtained in this way does not feature the property $U' U'^T b = b$, i.e., it does not correspond to a subspace that reproduces b exactly. More details can be found in supplemental section SM2.

3.3. The general case. When the operator in the underlying least-squares problem (3) is not a mask matrix but an arbitrary real matrix, then the Grassmann gradient associated with the residual function is still rank-one so that GROUSE continues to apply. Convergence results for GROUSE with arbitrary sampling matrices are given in [48].

Mind that Corollary 4 remains valid with the same proof, when the mask matrix P is replaced with an arbitrary column-orthogonal matrix. For general subspace-restricted least-squares problems

$$\min_{\tilde{\alpha} \in \mathbb{R}^p} \|AU_0 \tilde{\alpha} - b\|^2,$$

where the operator $A \in \mathbb{R}^{m \times n}$, $m \leq n$, is arbitrary but such that AU_0 has full column rank, we can proceed as follows. Let $QR = A^T$ be the thin qr-decomposition of A^T with $Q \in St(n, m)$, $R \in \mathbb{R}^{p \times p}$. Then

$$\|AU_0 \tilde{\alpha} - b\|^2 = \|R^T (Q^T U_0 \tilde{\alpha} - (R^T)^{-1} b)\|^2.$$

Since Q is column-orthogonal, we may apply Theorem 3 and Corollary 4 to the least-squares problem

$$\min_{\tilde{\alpha}} \|Q^T U_0 \tilde{\alpha} - (R^T)^{-1} b\|^2$$

to produce a modified U^* such that $\hat{\alpha} := \arg \min_{\tilde{\alpha} \in \mathbb{R}^p} \|Q^T U^* \tilde{\alpha} - (R^T)^{-1} b\|^2$ fulfills $0 = \|Q^T U^* \hat{\alpha} - (R^T)^{-1} b\|^2$. As a consequence, $\|AU^* \hat{\alpha} - b\|^2 = 0$. In summary, we present the following theorem.

THEOREM 8. *Let $p < m \leq n$. Consider the general subspace restricted least-squares problem*

$$\min_{\tilde{\alpha} \in \mathbb{R}^p} \|AU_0 \tilde{\alpha} - b\|^2, \quad A \in \mathbb{R}^{m \times n}, \quad b \in \mathbb{R}^m, \quad [U_0] \in Gr(n, p), \quad \text{rank}(AU_0) = p.$$

Let $QR = A^T$ and suppose that R is regular. Then there exists a subspace $[U^*] \in Gr(n, p)$ such that

$$\min_{\tilde{\alpha} \in \mathbb{R}^p} \|AU^*\tilde{\alpha} - b\|^2 = 0 \quad \text{and} \quad \text{dist}([U_0], [U^*]) = \arctan\left(\frac{\|r\|_2}{\|\alpha\|_2}\right) = \tau^*,$$

where $\alpha = \arg \min_{\tilde{\alpha} \in \mathbb{R}^p} \|Q^T U_0 \tilde{\alpha} - (R^T)^{-1} b\|^2$, $r = (R^T)^{-1} b - Q^T U_0 \alpha$. The subspace U^* is given by

$$U^* = U_0 + \left((\cos(\tau^*) - 1) U_0 \frac{\alpha}{\|\alpha\|_2} + \sin(\tau^*) \frac{Qr}{\|r\|_2} \right) \frac{\alpha^T}{\|\alpha\|_2}.$$

3.4. Adapting a subspace of a subspace. There are many applications where it might be desirable to keep some directions of a given subspace fixed while adapting the remaining ones. In the context of adaptive model reduction, such situations are likely to occur if the columns spanning the subspace in question stem from a principal component analysis or proper orthogonal decomposition (POD), and are thus ordered by information content. In these cases, the user might want to keep the most dominant subspace directions fixed, while adapting the portion of the subspace spanned by the less important basis vectors. This subsection describes the modifications to the methodology for doing so; a sample application is presented in subsection 4.2.

Let $f : Gr(n, p) \rightarrow \mathbb{R}$, $[U] \mapsto f([U])$ be a differentiable function. Let us divide the column set of a subspace representative $U \in St(n, p)$ into a constant portion $U_c \in St(n, p-l)$ and a portion $U_l \in St(n, l)$ that is subject to change, so that $U = (U_c, U_l) \in St(n, p-l) \times St(n, l)$. By fixing U_c , we obtain a function $f_l : Gr(n, l) \rightarrow \mathbb{R}$, $f_l([U_l]) = f([U_c, U_l])$ with gradient $G_l := \nabla f_l([U_l]) \in \mathbb{R}^{n \times l}$. The gradient induces the search direction $H_l = -G_l$. The geodesic associated with the search direction $H_l \stackrel{\text{SVD}}{=} \Phi_l S_l V_l^T \in \mathbb{R}^{n \times l}$ is represented by

$$(24) \quad U_l(t) = \text{Exp}_{[U_l]}(tH_l) = U_l V_l \cos(tS_l) V_l^T + \Phi_l \sin(tS_l) V_l^T.$$

Note that S_l and V_l are $(l \times l)$ -matrices. For each t , the matrix $U_l(t) \in St(n, l)$ is a feasible *orthogonal* subspace representative. Yet, we have to consider the possibility that the compound matrix $(U_c, U_l(t))$ ceases to be a valid subspace representative in $St(n, p)$.⁴ It is even conceivable that $[U_l(t)]$ moves towards the subspace $[U_c]$ spanned by the fixed basis vectors so that the compound matrix $(U_c, U_l(t))$ not only loses the orthogonal-columns property but even becomes rank deficient. One way to avoid this is to reorthogonalize $U_l(t)$ against U_c , say, by conducting an extra Gram–Schmidt procedure. However, Proposition 9 below implies that the orthogonality between the columns of the matrices $U_l(t)$ and the constant columns of the matrix block U_c is preserved along the geodesic path in *direction of the least-squares gradient*, so that in this case, the corresponding compound matrix $(U_c, U_l(t))$ is also an orthogonal matrix representative in $St(n, p)$ and a Gram–Schmidt reorthogonalization is unnecessary.

PROPOSITION 9. Let $f : Gr(n, p) \rightarrow \mathbb{R}$ be differentiable. Suppose that

$$(25) \quad T_{[U]} Gr(n, p) \ni \nabla_{[U]} f = (\nabla_{[U_c]} f_c, \nabla_{[U_l]} f_l) \in (T_{[U_c]} Gr(n, p-l)) \times (T_{[U_l]} Gr(n, l)),$$

where it is understood that $\nabla_{[U_c]} f_c$ and $\nabla_{[U_l]} f_l$ denote the gradients of the restrictions $f_c : [U_c] \mapsto f([U_c, U_l])$ and $f_l : [U_l] \mapsto f([U_c, U_l])$, respectively.

Let $[U_0] = [U_c, U_{l,0}] \in Gr(n, p)$. Let $t \mapsto [U_l(t)] \subset Gr(n, l)$ be the geodesic path along the descent direction $-\nabla_{[U_{l,0}]} f_l$. Then $U_c^T U_l(t) = 0$ for all t .

⁴Appendix B shows that this actually may happen even along search directions of rank one.

Therefore, the corresponding curve of concatenated matrices $(U_c, U_l(t)) \subset \mathbb{R}^{n \times p}$ is a curve of orthogonal matrices in $St(n, p)$. Hence, for each t , $[(U_c, U_l(t))] \in Gr(n, p)$, consistent with the quotient space viewpoint (1).

Proof. Let $U_0 = (U_c, U_{l,0}) \in St(n, p)$, where $U_c \in St(n, p-l)$ and $U_{l,0} \in St(n, l)$. The gradient with respect to f is a tangent vector in $T_{[U_0]}Gr(n, p)$, hence $U_0^T \nabla_{[U_0]} f = 0$. By (25),

$$(26) \quad 0 = U_0^T \nabla_{[U_0]} f = \begin{pmatrix} U_c^T \\ U_{l,0}^T \end{pmatrix} (\nabla_{[U_c]} f_c, \nabla_{[U_{l,0}]} f_l).$$

In particular, $U_c^T \nabla_{[U_{l,0}]} f_l = 0$. Writing $\nabla_{[U_{l,0}]} f_l \stackrel{\text{SVD}}{=} -\Phi_l S_l V_l^T \in \mathbb{R}^{n \times l}$, we have $U_c^T \Phi_l = 0$, since the columns of Φ_l span the same space as the columns of $\nabla_{[U_{l,0}]} f_l$. Hence, the geodesic at t , $U_l(t) = U_{l,0} V_l \cos(t S_l) V_l^T + \Phi_l \sin(t S_l) V_l^T$ is also orthogonal to U_c , i.e., $U_c^T U_l(t) = 0$. \square

As can be seen from the proof, the proposition is not specific to the GROUSE context nor does it depend on the rank of the gradient. It holds, in general, whenever the gradient splitting of (25) holds. This, however, is not always the case; see Appendix B. The objective function F of (5) features this property: When allowing only the last l directions of (U_c, U_l) to vary, we obtain a differentiable $F_l : Gr(n, l) \rightarrow \mathbb{R}$ with

$$F_l([U_l]) = b^T b - b^T P^T (U_c, U_l) \left(\begin{pmatrix} U_c^T \\ U_l^T \end{pmatrix} P P^T (U_c, U_l) \right)^{-1} \begin{pmatrix} U_c^T \\ U_l^T \end{pmatrix} P b.$$

The associated gradient, now a rank-one $(n \times l)$ -matrix, reads

$$G_l := \nabla_{[U_l]} F_l = -2P (b - P^T U \alpha) \alpha^T \begin{pmatrix} 0_{(p-l) \times l} \\ I_l \end{pmatrix}, \quad \alpha = (U^T P P^T U)^{-1} U^T P b,$$

where $U = (U_c, U_l)$. The next corollary transfers the result of Corollary 4 to the setting of adapting only the last l columns of a given subspace representative.

COROLLARY 10. *Let $\mathcal{U}_0 = [U_0] \in Gr(n, p)$ be represented by $U_0 \in St(n, p)$. Let $P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{(n \times m)}$ be a mask matrix and let $b \in \mathbb{R}^m$.*

Let $\alpha = (U_0^T P P^T U_0)^{-1} U_0^T P b$ be the optimal coefficient vector corresponding to the masked least-squares problem

$$\min_{\tilde{\alpha} \in \mathbb{R}^p} \|P^T U_0 \tilde{\alpha} - b\|^2$$

and let $r = b - P^T U_0 \alpha$ be the associated residual vector. Let $l \in \mathbb{N}$, $l \leq p$ and write columnwise $U_0 = (U_c, U_{l,0})$, $U_c = (u_0^1, \dots, u_0^{p-l})$, $U_{l,0} = (u_0^{p-l+1}, \dots, u_0^p)$. Moreover, let $\alpha_l = (0_{(p-l) \times l}, I_l) \alpha$ and $v_l = \frac{\alpha_l}{\|\alpha_l\|_2} \in \mathbb{R}^l$.

Set $s_1 = 2\|r\|_2 \|\alpha_l\|_2$ and define

$$t^* = \frac{1}{s_1} \arctan \left(\frac{\|r\|_2}{\|\alpha_l\|_2} \right)$$

and $U_l(t^) = U_{l,0} + ((\cos(t^* s_1) - 1) U_l v_l + \sin(t^* s_1) \frac{P r}{\|r\|}) v_l^T$.*

Then $U^ := U(t^*) := (U_c, U_l(t^*))$ is such that the subspace $\mathcal{U}^* := [U^*]$ is contained in the set \mathcal{Z} from (10), i.e.,*

$$F(U^*) = \min_{\tilde{\alpha} \in \mathbb{R}^p} \|P^T U^* \tilde{\alpha} - b\|^2 = 0,$$

which means that t^ solves (11).*

Proof. According to Proposition 9, the concatenated matrix $(U_c, U_l(t))$ is a valid subspace representative in $St(n, p)$ for all t . Applying the mask operator P to $(U_c, U_l(t))$ leads to the matrix curve

$$P^T U(t) = P^T(U_c, U_l(t)) = P^T(U_c, U_l V_l \cos(tS_l) V_l^T + \Phi_l \sin(tS_l) V_l^T).$$

Because Φ_l, S_l, V_l stem from an SVD of the rank-one gradient $-G_l$, we have that $S_l = \text{diag}(s_1, 0, \dots, 0)$, $s_1 = 2\|r\|_2\|\alpha_l\|_2$. It follows that

$$\begin{aligned} P^T U(t) &= P^T(U_c, U_{l,0}) + \left(0_{n \times (p-l)}, (\cos(ts_1) - 1)P^T U_l v_l + \sin(ts_1) \frac{r}{\|r\|_2}\right) \\ &= P^T(U_c, U_{l,0}) + x(t) (0_{1 \times (p-l)}, v_l^T), \end{aligned}$$

where $v_l = \frac{\alpha_l}{\|\alpha_l\|_2}$ is the first column of V_l . This is again a rank-one update on $P^T U(t)$ and the rest of the proof is analogous to the proof of Theorem 3. \square

Remark 11. When we are adapting only the last column u_0^p of the initial matrix $U_0 = (u_0^1, \dots, u_0^p) \in St(n, p)$, then the resulting U^* is given by $(u_0^1, \dots, u_0^{p-1}, u_0^p(t^*))$, where the last column evaluates to $u_0^p(t^*) = \frac{1}{\sqrt{\|r\|_2^2 + |\alpha_p|^2}} (u_0^p \alpha_p + Pr)$. This is precisely the same result that is obtained by replacing the last column of U_0 with the vector $U_0 \alpha + Pr$ and reorthogonalization the new last column against the columns of $U_0^{p-1} := (u_0^1, \dots, u_0^{p-1})$ via a single Gram–Schmidt step $(I - U_0^{p-1} (U_0^{p-1})^T)(U_0 \alpha + Pr) = (u_0^p \alpha_p + Pr)$. In this case, the t^* -GROUSE update applied to the last column of the subspace representative U_0 and the ($[U]$ -part of the) “revise” SVD update of [17, Table 1, p. 23] coincide; cf. Remark 7. For more details, see supplemental section SM2.

4. Application to adaptive model reduction. This section applies the geometric rank-one subspace update in the specific contexts of online adaptive model reduction and image reconstruction. For each application, we describe how the subspace adaptation is employed and we demonstrate the method with numerical examples.

4.1. Adaptation for POD-DEIM reduced models. We present an online adaptive DEIM that is based on our geometric rank-one subspace update. In contrast to the standard use case in the GROUSE literature [8, 49], the focus here is not on estimating a subspace from scratch based on a global objective function (6) but adapting a subspace that is already a good approximant for the underlying simulation process during the online phase.

We first formulate our online adaptive DEIM for nonlinear dynamical systems and then present numerical results for the FitzHugh–Nagumo system. To ease exposition and to focus on benchmarking our online adaptive DEIM reduced models, we consider dynamical systems without parameters and inputs. Thus, the aim of the following reduced models is to reproduce well the solution of the full-order dynamical system, instead of predicting solutions for new parameters and inputs. We note, however, that the following POD-DEIM and our online adaptive POD-DEIM reduced models are applicable to parametrized models and models with inputs; see [43, 14].

4.1.1. POD-DEIM-Galerkin reduced models. Consider a nonlinear dynamical system in the time interval $[0, T] \subset \mathbb{R}$, with end time $T > 0$. Let $t_0, t_1, \dots, t_K \in [0, T] \subset \mathbb{R}$ be $K + 1 \in \mathbb{N}$ time steps with $t_0 = 0$ and $t_K = T$. Discretizing with, e.g., the forward Euler method leads to the system of equations

$$(27) \quad E y_i = A y_{i-1} + f(y_{i-1}), \quad i = 1, \dots, K,$$

corresponding to the time steps t_1, \dots, t_K , respectively. Let $n \in \mathbb{N}$ denote the dimension of the discrete state space. We have the system matrices $A \in \mathbb{R}^{n \times n}$ and $E \in \mathbb{R}^{n \times n}$. The nonlinear function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ corresponds to the nonlinear terms of the dynamical system. The state vector at time step t_i is denoted as $y_i \in \mathbb{R}^n$. The initial condition is $y_0 \in \mathbb{R}^n$. We consider here the case where the nonlinear function f is evaluated componentwise at the state vector y_i ; see, e.g., [22]. We further assume the well-posedness of (27).

To derive a reduced model of the full model (27), we select a set of $n_s \in \mathbb{N}$ snapshots $\{y_{j_1}, \dots, y_{j_{n_s}}\} \subset \{y_1, \dots, y_K\}$ at the time steps $t_{j_1}, \dots, t_{j_{n_s}}$ with indices $j_1, \dots, j_{n_s} \in \{1, \dots, K\}$. POD constructs orthonormal basis vectors $v_1, \dots, v_{n_r} \in \mathbb{R}^n$ of the n_r -dimensional POD space that is the solution to the minimization problem

$$\min_{v_1, \dots, v_{n_r} \in \mathbb{R}^n} \sum_{i=1}^{n_s} \left\| y_{j_i} - \sum_{l=1}^{n_r} (v_l^T y_{j_i}) v_l \right\|_2^2.$$

The POD basis $V = (v_1, \dots, v_{n_r})$ is formed of the left-singular vectors of the snapshot matrix $Y = (y_{j_1}, \dots, y_{j_{n_s}}) \in \mathbb{R}^{n \times n_s}$ corresponding to the n_r largest singular values. The POD-Galerkin reduced model of (27) is

$$(28) \quad \tilde{E} \tilde{y}_i = \tilde{A} \tilde{y}_{i-1} + V^T f(V \tilde{y}_{i-1}),$$

where $\tilde{y}_i \in \mathbb{R}^{n_r}$ is the reduced state vector at time step t_i for $i = 1, \dots, K$, and $\tilde{E} = V^T E V$, $\tilde{A} = V^T A V$ are the reduced operators.

Solving (28) requires evaluating the nonlinear function $f(V \tilde{y}_{i-1})$ at the n -dimensional vector $V \tilde{y}_{i-1} \in \mathbb{R}^n$, which can be computationally expensive. DEIM derives an approximation of $f(V \tilde{y}_{i-1})$ to avoid evaluating f at all n components of $V \tilde{y}_{i-1}$. To this end, DEIM constructs $p \in \mathbb{N}$ DEIM basis vectors $u^1, \dots, u^p \in \mathbb{R}^n$ using POD on the nonlinear snapshots $f(y_{j_1}), \dots, f(y_{j_{n_s}}) \in \mathbb{R}^n$. The DEIM basis vectors are the columns of the DEIM basis matrix $U = (u_1, \dots, u_p) \in \mathbb{R}^{n \times p}$. Additionally, DEIM selects $p \in \mathbb{N}$ DEIM interpolation points $q_1, \dots, q_p \in \{1, \dots, n\}$ using a greedy strategy; see [22]. The DEIM mask matrix is $P = (e_{q_1}, \dots, e_{q_p}) \in \{0, 1\}^{n \times p}$. The DEIM interpolant is the pair (U, P) . The DEIM approximation of the nonlinear function f evaluated at the vector $V \tilde{y}_i$ is given as

$$(29) \quad f(V \tilde{y}_i) \approx U(P^T U)^{-1} P^T f(V \tilde{y}_i).$$

The POD-DEIM-Galerkin reduced model of (27) at a time step $t_i, i = 1, \dots, K$, is

$$(30) \quad \tilde{E} \tilde{y}_i = \tilde{A} \tilde{y}_{i-1} + V^T U(P^T U)^{-1} P^T f(V \tilde{y}_{i-1}).$$

The reduced model (30) is often orders of magnitude faster to solve than the full model (27) and the reduced state vectors $\tilde{y}_1, \dots, \tilde{y}_K \in \mathbb{R}^{n_r}$ lead to accurate approximations $V \tilde{y}_1, \dots, V \tilde{y}_K \in \mathbb{R}^n$ of the full state vectors $y_1, \dots, y_K \in \mathbb{R}^n$, respectively.

4.1.2. Online adaptive model reduction. We adapt the DEIM interpolant of the nonlinear function f in the online phase, i.e., we adapt the DEIM basis U and the DEIM mask matrix P during the time stepping. We proceed as follows. Let U_0 denote the DEIM basis matrix, which is derived using POD as discussed in subsection 4.1.1. Further, let $q_1^0, \dots, q_p^0 \in \{1, \dots, n\}$ be the DEIM interpolation points and $P_0 = (e_{q_1^0}, \dots, e_{q_p^0})$ the mask matrix that are derived with the DEIM procedure in the offline phase; see subsection 4.1.1. Consider now the online phase at time step

t_1 . To compute the reduced state vector \tilde{y}_1 , we first adapt the DEIM basis matrix U_0 and the mask matrix P_0 to U_1 and P_1 , respectively, and then use the adapted DEIM interpolant (U_1, P_1) in the reduced model (30) to compute the reduced state vector \tilde{y}_1 . The DEIM basis matrix U_0 is adapted to U_1 using the GROUSE rank-one update, as we will discuss in detail in subsection 4.1.3. This process is continued iteratively, i.e., at time step t_i , we adapt U_{i-1} and P_{i-1} to obtain U_i and P_i , respectively, and then use the adapted interpolant (U_i, P_i) for computing the reduced state vector \tilde{y}_i at time step t_i . Note that the POD basis matrix V and the reduced linear operators \tilde{E} and \tilde{A} are kept unchanged online (although in principle they too could be adapted).

4.1.3. Subspace adaptation in online adaptive model reduction. We use the GROUSE rank-one update with the residual-annihilating step size (22) to adapt the DEIM basis matrix. Consider time step t_i for $i = 1, \dots, K$. To adapt the DEIM basis matrix U_{i-1} to U_i at time step t_i , we follow [43] and *oversample* the DEIM approximation. Let $\{q_{p+1}^i, \dots, q_{p+s}^i\} \subset \{1, \dots, n\} \setminus \{q_1^{i-1}, \dots, q_p^{i-1}\}$ be a set of $s \in \mathbb{N}$ additional indices that are drawn uniformly from the set $\{1, \dots, n\} \setminus \{q_1^{i-1}, \dots, q_p^{i-1}\}$, where $q_1^{i-1}, \dots, q_p^{i-1}$ are the DEIM interpolation points of the previous time step t_{i-1} . The extended mask matrix $S_i \in \{0, 1\}^{n \times m}$, $m = p + s$, is assembled from the points in the set $\{q_1^{i-1}, \dots, q_p^{i-1}, q_{p+1}^i, \dots, q_{p+s}^i\}$ as $S_i = (e_{q_1^{i-1}}, \dots, e_{q_p^{i-1}}, e_{q_{p+1}^i}, \dots, e_{q_{p+s}^i})$. The matrix S_i corresponds to $m = p + s > p$ point indices, and therefore, the interpolation problem (29) of the classical DEIM approximation with the interpolant (U_{i-1}, P_{i-1}) becomes an overdetermined least-squares problem using the extended mask matrix S_i

$$(31) \quad \alpha = \arg \min_{\alpha \in \mathbb{R}^p} \|S_i^T U_{i-1} \tilde{\alpha} - S_i^T f(V \tilde{y}_{i-1})\|_2^2$$

with

$$f(V \tilde{y}_{i-1}) \approx U_{i-1} \alpha.$$

The solution α of (31) is

$$\alpha = (U_{i-1}^T S_i S_i^T U_{i-1})^{-1} U_{i-1}^T S_i S_i^T f(V \tilde{y}_{i-1}).$$

The regression problem (31) fits into the framework of the GROUSE subspace adaptation approach of subsection 2.2, so that we can find the adapted DEIM basis matrix U_i with the low-rank update derived in Corollary 4. In addition to updating the DEIM basis matrix, the DEIM interpolation points $q_1^{i-1}, \dots, q_p^{i-1}$ are updated to q_1^i, \dots, q_p^i . For this task we use the algorithm introduced in [43, section 4]. The entire DEIM online adaptivity procedure is summarized in Algorithm 1.

4.1.4. Example of DEIM subspace adaptation. We apply the online subspace adaptation to the DEIM interpolant of a reduced model of the FitzHugh–Nagumo system. The FitzHugh–Nagumo system is used in the original DEIM paper [22] as a benchmark example. The number of time steps is $K = 10^6$ and the dimension of the discretized state space is $n = 2048$. The state vectors $y_0, y_{1000}, y_{2000}, \dots, y_K \in \mathbb{R}^n$ at every 1000th time step are used as snapshots to construct $n_r = 10$ POD basis vectors and the corresponding POD basis matrix $V \in \mathbb{R}^{n \times n_r}$. The nonlinear function is evaluated at the snapshot time instances to obtain the nonlinear snapshots $f(y(t_0)), f(y(t_{1000})), f(y(t_{2000})), \dots, f(y(t_K))$.

We compare the error of a static reduced model without online subspace adaptation to the error of an adaptive reduced model as in Algorithm 1. We report the average of the relative L_2 error of the approximation $V \tilde{y}_i$ to the reference y_i at the

Algorithm 1. Time stepping a reduced model with online adaptive DEIM.

Input: System matrices E, A , nonlinear function f , initial condition y_0 , POD basis matrix V , DEIM basis matrix U_0 , DEIM interpolation points matrix P_0 , number of sampling points s , adaptation interval l

```


1: Set  $\tilde{y}_0 = V^T y_0$ 
2: for  $i = 1, \dots, K$  do
3:   if  $\text{mod}(i, l) == 0$  then
4:     {Adapt DEIM interpolant every  $l$ th time step}
5:     Set  $q_1^{i-1}, \dots, q_p^{i-1}$  to the interpolation points of  $P_{i-1}$ 
6:     Draw  $q_{p+1}^i, \dots, q_{p+s}^i$  uniformly from  $\{1, \dots, n\} \setminus \{q_1^{i-1}, \dots, q_p^{i-1}\}$ 
7:     Construct mask matrix  $S_i$  from points  $q_1^{i-1}, \dots, q_p^{i-1}, q_{p+1}^i, \dots, q_{p+s}^i$ 
8:     Evaluate nonlinear function at sampling points  $b = S_i^T f(V\tilde{y}_{i-1})$ 
9:     {Employ Corollary 4 to adapt  $U_{i-1}$ }
10:    Set  $\alpha = (U_{i-1}^T S_i S_i^T U_{i-1})^{-1} U_{i-1}^T S_i b$  and  $r = b - S_i^T U_{i-1} \alpha$ 
11:    Set  $v = \alpha / \|\alpha\|_2$ ,  $s_1 = 2\|r\|_2 / \|\alpha\|_2$ , and  $t^* = s_1^{-1} \arctan(\|r\|_2 / \|\alpha\|_2)$ 
12:    Adapt basis matrix
         $U_i = U_{i-1} + ((\cos(t^* s_1) - 1) U_{i-1} v + \sin(t^* s_1) (S_i r) / \|r\|_2) v^T$ 
13:    Adapt interpolation points matrix  $P_{i-1}$  to  $P_i$  with [43, Algorithm 2]
14:  else
15:    Set  $U_i = U_{i-1}$  and  $P_i = P_{i-1}$  {No adaptation}
16:  end if
17:   $\tilde{f}_i = V^T U_i (P_i^T U_i)^{-1} P_i^T f(V\tilde{y}_{i-1})$  {Approximate nonlinear function}
18:  Solve reduced model  $\tilde{E}\tilde{y}_i = \tilde{A}\tilde{y}_{i-1} + \tilde{f}_i$  for  $\tilde{y}_i$ 
19: end for
Output: Reduced states  $\tilde{y}_0, \dots, \tilde{y}_K$ 

```

time steps $t_{500}, t_{1500}, \dots, t_{K-500}$. Thus, the error is measured at time steps other than where the snapshots were taken.

Figure 2(a) compares the L_2 error of the states of the reduced model (30) with a static DEIM interpolant to the error of the reduced model with an adaptive DEIM interpolation. The dimension of the DEIM subspace is varied over the range $p \in \{2, 4, 6, 8, 10\}$. The DEIM subspace and the DEIM interpolation points are adapted every 50th time step, which means that we set $l = 50$ in Algorithm 1. At each adaptation step, the geometric rank-one update of Corollary 4 is performed to adapt the DEIM basis matrix based on $s \in \{200, 400, 600\}$ sampling points. Note that the computational costs of the rank-one update are bounded by $\mathcal{O}(np)$. The error of the static and the online adaptive reduced model decreases with the DEIM dimension, which shows that the POD space, which is static and derived from snapshots taken over the whole time interval, approximates well the full-order state vectors; see subsection 4.1.1. The online adaptive DEIM interpolant can further reduce the error by about an order of magnitude. Figure 2(b) reports results for the online adaptive reduced model, where the DEIM interpolant is adapted every 50th, 100th, and 200th time step with a fixed number of $s = 200$ samples. This means that Algorithm 1 is run with $l = 50, 100, 200$, respectively. The results confirm that increasing the number of adaptivity steps increases the accuracy of the results.

Figure 3 shows results for the online adaptive DEIM interpolant where different step sizes are used. We compare four different step size selections in Figure 3. The curve with the label “adapt, optimal” refers to the residual annihilator t^* , which is



(a) Adapt every 50th time step with varying sample size.

(b) Adapt with $s = 200$ samples with varying adaptation interval.

FIG. 2. The average relative L_2 error of a static reduced model is compared to the error of a reduced model with an online adaptive DEIM interpolant. The online adaptation based on the low-rank updates achieves up to an order of magnitude improvement in the L_2 error compared to the static DEIM interpolant.

derived in Corollary 4 and implemented in Algorithm 1. The curve with label “adapt, asym. optimal” corresponds to the step size $\tilde{t} = \frac{1}{s_1} \arcsin\left(\frac{\|r\|_2}{\|\alpha\|_2}\right)$ that is discussed in the GROUSE convergence analysis of [10]; see also Remark 5. We additionally compare to the constant step size 0.05 in “adapt, constant” and a decaying step size $0.05/i$ in “adapt, decaying step size,” as in, e.g., the GROUSE numerical experiments in [8], where i is the counter variable in the for-loop in Algorithm 1. The number of samples is set to $s = 400$ and the DEIM subspace and the DEIM interpolation points are adapted every 50th time step. The optimal and the asymptotically optimal step size lead to similar results (the curves are on top of each other), which was to be expected, since the functions \arctan and \arcsin match up to terms of third order. The less sophisticated choices “adapt, constant” and “adapt, decaying step size” lead to poor results which are even worse than those produced by the static subspace for DEIM basis dimensions of 8 and 10. This shows that for the application at hand, it is crucial to select a residual-related step size based on the ratio $\frac{\|r\|_2}{\|\alpha\|_2}$, e.g., the minimizer t^* from Corollary 4.

4.2. Subspace adaptation for gappy POD image reconstruction. In this section, the geometric subspace update is applied in combination with the method of gappy POD [27, 18] on an image processing problem, where we use the method to implant a new feature into a given subspace.

We briefly summarize gappy POD. Given a set of snapshots $\{y_k \mid k = 1, \dots, n_s\} \subset \mathbb{R}^n$, let $\mathcal{U} = \text{colspan}(U)$ be the associated POD subspace represented by $U \in St(n, p)$ with $p \leq n_s$. Further, let $y^g \in \mathbb{R}^n$ be an *incomplete snapshot* associated with an index set $J = \{j_1, \dots, j_m\} \subset \{1, \dots, n\}$ of cardinality $m \in \mathbb{N}$; y^g is incomplete in the sense that only components with indices in J are considered as accurate information. Gappy POD computes a vector contained in \mathcal{U} that best fits the incomplete snapshot y^g in a least-squares sense. Employing the mask matrix $P = (e_{j_1}, \dots, e_{j_m}) \in \{0, 1\}^{n \times m}$,

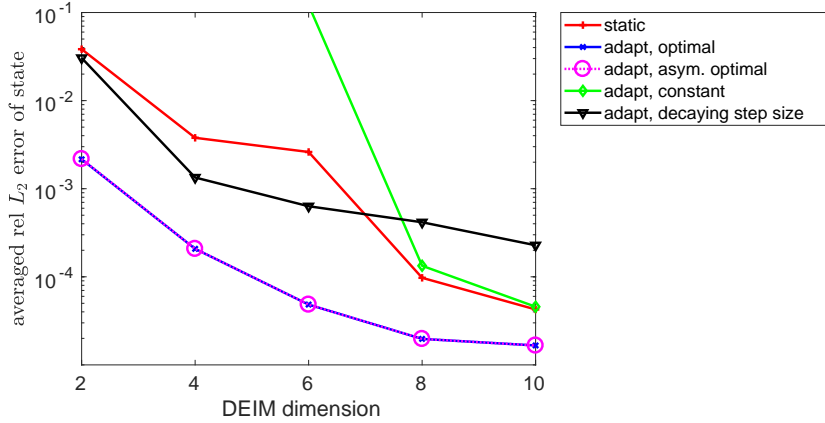


FIG. 3. The plot reports the error of the online adaptive POD-DEIM reduced model for different step sizes. The label “adapt, optimal” refers to the residual annihilator derived in Corollary 4, “adapt, asym. optimal” refers to the step size \tilde{t} mentioned in Remark 5, “adapt, constant” refers to the constant step size 0.05, and “adapt, decaying step size” refers to the step size $0.05/i$, where i is the counter variable in Algorithm 1. Note that the curves of “adapt, optimal” and “adapt, asym. optimal” are on top of each other.



FIG. 4. Face database used for gappy POD example.

the gappy POD approximation $y^{gpod} \in \mathbb{R}^n$ is determined by the masked least-squares minimization problem

$$(32) \quad y^{gpod} = U\alpha_{gpod}, \quad \alpha_{gpod} = \arg \min_{\alpha \in \mathbb{R}^p} \|P^T U\alpha - P^T y^g\|^2.$$

(Notice the similarities to the DEIM approach from subsection 4.1.4. Reference [28] exposes further details on the relation between gappy POD and the Empirical Interpolation Method (EIM, [11]), which predates DEIM.) In our concrete example of image processing, the snapshot set is taken from the so-called *Yale Database* [13]; see also [20, section 5.2].⁵ Representing each image as a snapshot vector $y_k \in \mathbb{R}^n$, $n = 4096$, yields a snapshot matrix of dimension $Y \in \mathbb{R}^{4096 \times 10}$. The snapshots are displayed in Figure 4. The single image with glasses has been deliberately omitted from the snapshot set, so that no picture in the snapshot ensemble features the property “glasses-on.” The “glasses”-detail from this picture, displayed in the lower left corner of Figure 6, acts as a vector of gappy data $y^g \in \mathbb{R}^{4096}$ with $m = 1336$ nonzero entries and corresponding mask matrix P . The gappy POD objective is to find the linear combination of snapshots that comes closest to representing the “glasses”-feature in a least-squares sense. The resulting image is displayed in the second column of Figure 6 with the top picture showing the gappy POD solution and the bottom picture showing the reference image projected onto the subspace spanned by the POD modes. The

⁵More precisely, we have used row 11 of the set of 165 Yale images in (64×64) -MATLAB format provided by Deng Cai at <http://www.cad.zju.edu.cn/home/dengcai/Data/FaceData.html>.

gappy POD reconstruction is a poor approximation of the reference picture because the POD space does not contain any information required to represent glasses.




FIG. 5. Plot of two periods of residual norm function (21) corresponding with the gappy POD subspace adaptation. The circle locates the first root and the star indicates the global maximum.




FIG. 6. Gappy POD approximation of a picture excerpt. To be read columnwise: Reference picture and training excerpt. Gappy POD reconstruction based on the excerpt and projection of complete reference onto the POD subspace. Gappy POD reconstruction using an adapted POD subspace and projection of complete reference thereon. Gappy POD reconstruction after adapting only the last column of POD subspace and projection of complete reference thereon.

Now, we use the GROUSE rank-one update combined with Corollary 4 to annihilate the gappy POD residual, which corresponds to solving the nonlinear equation (11) on $Gr(n, p) = Gr(4096, 10)$. The input data are the mask matrix $P \in \mathbb{R}^{n \times m}$ associated with the picture excerpt, the corresponding right-hand side $b = P^T y^g \in \mathbb{R}^m$, and the subspace representative $U_0 \in St(n, p)$ stemming from a POD of the input snapshots. A plot of the residual norm function along the rank-one update is displayed in Figure 5.

The update leads to a subspace representative U^* that allows for a perfect reproduction of the picture excerpt but also makes use of the information that was




FIG. 7. Initial face data set (bottom row) and its projection onto the corresponding POD subspace with only the last column adapted to the training excerpt (middle row) and its projection onto the fully adapted POD subspace (top row).

previously sampled. We repeat the exercise while modifying only the last column of the initial POD subspace representative U_0 according to Corollary 10.

The gappy POD approximations using the adapted subspaces are shown in the last two columns of Figure 6, again in comparison with the projection of the reference image onto the respective subspace, as is clear from Corollary 4.⁶

The important thing is how the adapted subspaces have changed. This can be visualized by projecting the initial snapshot ensemble onto the adapted subspaces; see Figure 7. Apart from the fact that the bright white spots in the original data set are reproduced in a graying way when projected onto the last-column adapted subspace, these two data sets look almost the same (Figure 7, bottom rows). In contrast, the original data set projected onto the fully adapted subspace features the property “glasses-on” throughout (Figure 7, top row). Nevertheless, the subspace distance between $[U_0]$ and the fully adapted $[U^*]$ is 0.1273, while the distance between $[U_0]$ and the subspace $[U^*]$ with only the last column adjusted is 1.2828, more than ten times as large. Recall from Remark 7 that the latter $[U^*]$ corresponds to an SVD update with respect to a column-replacement in the original subspace representative U_0 .

Additional experiments are featured in section SM3 from the supplement. The supplement also includes MATLAB code for the adapted gappy POD examples discussed here.

5. Summary and conclusion. Subspace update problems arise in model reduction, machine learning, pattern recognition, and computer vision. This paper focuses on the particular use case of subspace adaptation in combination with the model reduction methods of gappy POD and DEIM. These methods have in common that a mask matrix is utilized to extract the features deemed most important to the underlying problem. In both cases, the objective of the downstream subspace adaptation is to produce subspaces that contain elements that match the selected components. We have formalized this objective as a nonlinear equation on the Grass-

⁶This transfers in an analogous form to the sub-subspace setting of subsection 3.4 in which both reconstructed images coincide since they both correspond to copying the training set to the respective entries of the unmodified gappy POD solution.

mann manifold and have provided a closed-form solution that builds on the GROUSE approach [8, 49].

In the DEIM test case, discussed in subsection 4.1.4, the mask matrix operates on vectors contained in the subspace that represents the nonlinear terms of the underlying discretized PDE. In the gappy POD test cases, discussed in subsection 4.2, the mask matrix selects the important components from vectors contained in the subspace of state vector solution candidates. In the test case of DEIM-based model reduction, the Grassmann subspace update is used as an online adaptation method to improve the fit of the components sampled from the nonlinear term. The reduced model with online subspace updating achieves an average error of about one order of magnitude lower than a classical reduced model without the adaptation. In the gappy POD image processing example, the Grassmann subspace update is applied to implement a new feature in the subspace of solution candidates that is not contained in the sample data set. We expect the method to show similar advantages when used in combination with the missing point estimation [7, 51], because of the similarities to DEIM and gappy POD.

Appendix A. A direct solution of the Grassmann residual equation (11). This appendix features a short solution of (11). Obviously, (11) is solved if we can find $t^* \in \mathbb{R}$ and $\alpha^* \in \mathbb{R}^p$ such that $P^T U(t^*) \alpha^* = b$, where $U(t^*) := U_0 + ((\cos(t^* s_1) - 1)U_0 v + \sin(t^* s_1) \frac{Pr}{\|r\|})v^T$. (All occurring quantities to be understood were introduced in Theorem 3.) Mind that $v = \alpha / \|\alpha\|_2$. Using an additional real parameter λ and the ansatz $\alpha^* = \lambda \alpha = \lambda v \|\alpha\|_2$ leads to the equation

$$(33) \quad \lambda \cos(t^* s_1) \left(\left(1 - \tan(t^* s_1) \frac{\|\alpha\|_2}{\|r\|_2} \right) P^T U_0 \alpha + \tan(t^* s_1) \frac{\|\alpha\|_2}{\|r\|_2} b \right) = b.$$

By setting $t^* = \frac{1}{s_1} \arctan \left(\frac{\|r\|_2}{\|\alpha\|_2} \right)$, the terms involving $P^T U_0$ cancel, which leaves an equation for λ :

$$\lambda \left(\arctan \left(\frac{\|r\|_2}{\|\alpha\|_2} \right) \right) b = b.$$

The solution is $\lambda = \frac{1}{\cos(\arctan(\frac{\|r\|_2}{\|\alpha\|_2}))} = \sqrt{1 + \frac{\|r\|_2^2}{\|\alpha\|_2^2}}$.

In addition to its concision, this approach has the advantage that it simultaneously gives both t^* and the associated vector of coefficients $\alpha^* = \sqrt{\left(\frac{\|r\|_2}{\|\alpha\|_2} + 1\right)} \alpha \in \mathbb{R}^p$. On the other hand it does not allow one to keep track of the residual depending on t , because for $t \neq t^*$, a defining equation is missing and $\alpha(t)$ and α are not collinear.

Nevertheless, we remark that the above shortcut approach may be adapted to apply also in the setting of Corollary 10 from subsection 3.4. In this case, one can work from the ansatz $\alpha^* = (\alpha_1, \dots, \alpha_{p-l}, \lambda(\alpha_{p-l+1}, \dots, \alpha_p))^T$.

One may also start by first applying the orthogonal coordinate transformation $\Phi = (v, Z) \in O_p$ to the subspace representative U_0 , where $Z \in \mathbb{R}^{p \times (p-1)}$ contains an arbitrary orthonormal basis of v^\perp , and then work with $U_0 \Phi$, $U(t^*) \Phi$. This course of action essentially leads to (33) appearing in the first column of $U(t^*) \Phi$ and the rest of the argument is analogous. See [49, App. C, Proof of Lemma 4] for related considerations.

Appendix B. Addendum to subsection 3.4. A simple example of a differentiable Grassmann objective function for which Proposition 9 does not hold is

$$f : Gr(n, p) \rightarrow \mathbb{R}, \quad [U] \mapsto x^T U U^T y,$$

where $x, y \in \mathbb{R}^n$ are not orthogonal to $[U]$.

By using the basic fact that $D_X(v^T X w) = \left(\frac{\partial}{\partial x_{ij}} v^T X w\right)_{ij} = v w^T$ and the product rule, we see that the Grassmann gradient is

$$\nabla_{[U]} f = (I - UU^T) D_U f = (I - UU^T) (xy^T + yx^T) U,$$

where $D_U f = \left(\frac{\partial f}{\partial u_{i,j}}\right)_{i,j} \in \mathbb{R}^{n \times p}$; see [25, eq. (2.70)]. (Note that $\nabla_{[U]} f$ is of rank two in general, but of rank one if $x = y$.) Introducing $U = (U_1, U_2)$ with $U_1 \in St(n, p-l)$, $U_2 \in St(n, l)$, we may write $UU^T = U_1 U_1^T + U_2 U_2^T$. By fixing U_1 , f becomes a function $f_2 : Gr(n, l) \rightarrow \mathbb{R}$, $[U_2] \mapsto x^T U_1 U_1^T y + x^T U_2 U_2^T y$. The gradient is

$$\nabla_{[U_2]} f_2 = (I - U_2 U_2^T) (xy^T + yx^T) U_2 \in \mathbb{R}^{n \times l}.$$

Likewise, for $f_1 : Gr(n, p-l) \rightarrow \mathbb{R}$, $[U_1] \mapsto x^T U_1 U_1^T y + x^T U_2 U_2^T y$, we obtain

$$\nabla_{[U_1]} f_1 = (I - U_1 U_1^T) (xy^T + yx^T) U_1 \in \mathbb{R}^{n \times l}.$$

Splitting up the original gradient into an $(n \times (p-l))$ and an $(n \times l)$ matrix gives

$$\begin{aligned} \nabla_{[U]} f &= ((I - UU^T)(xy^T + yx^T)U_1, (I - UU^T)(xy^T + yx^T)U_2) \\ &\neq ((I - U_1 U_1^T)(xy^T + yx^T)U_1, (I - U_2 U_2^T)(xy^T + yx^T)U_2) \\ &= (\nabla_{[U_1]} f_1, \nabla_{[U_2]} f_2). \end{aligned}$$

In particular, $U_1^T \nabla_{[U_2]} f_2 = U_1^T xy^T U_2 + U_1^T yx^T U_2 \neq 0$ and the geodesic $U_2(t)$ in $Gr(n, l)$ along the gradient direction $\nabla_{[U_2]} f_2$ is not orthogonal to U_1 ,

$$U_1^T U_2(t) \neq 0.$$

A sufficient condition for (25) and Proposition 9 to hold is $(I - UU^T) D_U f = D_U f$ or, in short, $U^T D_U f = 0$.

Acknowledgments. The authors would like to thank Laura Balzano (University of Michigan) for helpful conversations concerning GROUSE. Moreover, the authors thank the reviewers and Associate Editor for inspiring comments that improved this manuscript.

REFERENCES

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Riemannian geometry of Grassmann manifolds with a view on algorithmic computation*, Acta Appl. Math., 80 (2004), pp. 199–220.
- [2] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, NJ, 2008.
- [3] D. AMSALLEM AND C. FARHAT, *Interpolation method for adapting reduced-order models and application to aeroelasticity*, AIAA J., 46 (2008), pp. 1803–1813.
- [4] D. AMSALLEM AND C. FARHAT, *An online method for interpolating linear parametric reduced-order models*, SIAM J. Sci. Comput., 33 (2011), pp. 2169–2198, <https://doi.org/10.1137/100813051>.
- [5] D. AMSALLEM, M. ZAHR, AND C. FARHAT, *Nonlinear model order reduction based on local reduced-order bases*, Internat. J. Numer. Methods Engrg., 92 (2012), pp. 891–916.
- [6] 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), Adv. Comput. Math., 41 (2015), pp. 1187–1230.
- [7] P. ASTRID, S. WEILAND, K. WILLCOX, AND T. BACKX, *Missing points estimation in models described by proper orthogonal decomposition*, IEEE Trans. Automat. Control, 53 (2008), pp. 2237–2251.

- [8] L. BALZANO, R. NOWAK, AND B. RECHT, *Online identification and tracking of subspaces from highly incomplete information*, in 2010 48th Annual Allerton Conference on Communication, Control, and Computing, Allerton, IEEE, 2010, pp. 704–711.
- [9] L. BALZANO AND S. J. WRIGHT, *On GROUSE and incremental SVD*, in 2013 5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013, pp. 1–4.
- [10] L. BALZANO AND S. J. WRIGHT, *Local convergence of an algorithm for subspace identification from partial data*, *Found. Comput. Math.*, 15 (2015), pp. 1279–1314.
- [11] M. BARRAULT, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *An empirical interpolation method: Application to efficient reduced-basis discretization of partial differential equations*, *C. R. Math. Acad. Sci. Paris*, 339 (2004), pp. 667–672.
- [12] E. BEGELFOR AND M. WERMAN, *Affine invariance revisited*, 2012 IEEE Conference on Computer Vision and Pattern Recognition, 2 (2006), pp. 2087–2094.
- [13] P. N. BELHUMEUR, J. P. HESPANHA, AND D. J. KRIEGMAN, *Eigenfaces vs. fisherfaces: Recognition using class specific linear projection*, *IEEE Trans. on PAMI*, 19 (1997), pp. 711–720.
- [14] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A survey of projection-based model reduction methods for parametric dynamical systems*, *SIAM Rev.*, 57 (2015), pp. 483–531, <https://doi.org/10.1137/130932715>.
- [15] N. BOUMAL AND P.-A. ABSIL, *Low-rank matrix completion via preconditioned optimization on the Grassmann manifold*, *Linear Algebra Appl.*, 475 (2015), pp. 200–239.
- [16] M. BRAND, *Incremental singular value decomposition of uncertain data with missing values*, in *Computer Vision ECCV 2002*. ECCV 2002, A. Heyden, G. Sparr, M. Nielsen, P. Johansen, eds, *Lecture Notes in Comput. Sci.* 2350 Springer, Berlin, Heidelberg, 2002, pp. 707–720.
- [17] M. BRAND, *Fast low-rank modifications of the thin singular value decomposition*, *Linear Algebra Appl.*, 415 (2006), pp. 20–30.
- [18] T. BUI-THANH, M. DAMODARAN, AND K. WILLCOX, *Proper orthogonal decomposition extensions for parametric applications in transonic aerodynamics*, *AIAA J.*, 42 (2004), pp. 1505–1516.
- [19] J. R. BUNCH AND C. P. NIELSEN, *Updating the singular value decomposition*, *Numer. Math.*, 31 (1978), pp. 111–129.
- [20] D. CAI, X. HE, J. HAN, AND H.-J. ZHANG, *Orthogonal Laplacianfaces for face recognition*, *IEEE Trans. Image Process.*, 15 (2006), pp. 3608–3614.
- [21] K. CARLBERG, *Adaptive h-refinement for reduced-order models*, *Internat. J. Numer. Methods Engrg.*, 102 (2015), pp. 1192–1210.
- [22] S. CHATURANTABUT AND D. C. SORENSEN, *Nonlinear model reduction via discrete empirical interpolation*, *SIAM J. Sci. Comput.*, 32 (2010), pp. 2737–2764, <https://doi.org/10.1137/090766498>.
- [23] J. DEGROOTE, J. VIERENDEELS, AND K. WILLCOX, *Interpolation among reduced-order matrices to obtain parameterized models for design, optimization and probabilistic analysis*, *Internat. J. Numer. Methods Fluids*, 63 (2010), pp. 207–230.
- [24] M. DIHLMANN, M. DROHMANN, AND B. HAASDONK, *Model reduction of parametrized evolution problems using the reduced basis method with adaptive time-partitioning*, in *Proceedings of the International Conference on Adaptive Modeling and Simulation*, D. Aubry, P. Dez, B. Tie, and N. Pars, eds., 2011, pp. 156–167.
- [25] A. EDELMAN, T. A. ARIAS, AND S. T. SMITH, *The geometry of algorithms with orthogonality constraints*, *SIAM J. Matrix Anal. Appl.*, 20 (1999), pp. 303–353, <https://doi.org/10.1137/S0895479895290954>.
- [26] J. EFTANG AND B. STAMM, *Parameter multi-domain hp empirical interpolation*, *Internat. J. Numer. Methods Engrg.*, 90 (2012), pp. 412–428.
- [27] R. EVERSON AND L. SIROVICH, *Karhunen-Loève procedure for gappy data*, *J. Opt. Soc. Am.*, 12 (1995), pp. 1657–1664.
- [28] D. GALBALLY, K. FIDKOWSKI, K. WILLCOX, AND O. GHATTAS, *Non-linear model reduction for uncertainty quantification in large-scale inverse problems*, *Internat. J. Numer. Methods Engrg.*, 81 (2010), pp. 1581–1608.
- [29] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, 1996.
- [30] S. KAULMANN AND B. HAASDONK, *Online greedy reduced basis construction using dictionaries*, in *Proceedings of the 7th Vienna International Conference on Mathematical Modelling*, I. Troch and F. Breiteneker, eds., 2012, pp. 112–117.
- [31] R. KENNEDY, L. BALZANO, S. J. WRIGHT, AND C. J. TAYLOR, *Online algorithms for factorization-based structure from motion*, *Comput. Vis. Image Underst.*, 150 (2016), pp. 139–152.

- [32] S. KOBAYASHI AND K. NOMIZU, *Foundations of Differential Geometry*, Vol. I, Interscience Tracts in Pure and Applied Mathematics 15, John Wiley & Sons, New York, London, Sidney, 1963.
- [33] J. LEE, *Riemannian Manifolds: An Introduction to Curvature*, Springer Verlag, New York, Berlin, Heidelberg, 1997.
- [34] Y. MADAY AND B. STAMM, *Locally adaptive greedy approximations for anisotropic parameter reduced basis spaces*, SIAM J. Sci. Comput., 35 (2013), pp. A2417–A2441, <https://doi.org/10.1137/120873868>.
- [35] Y. MAN LUI, *Advances in matrix manifolds for computer vision*, Image Vis. Comput., 30 (2012), pp. 380–388.
- [36] T. S. NGUYEN, *A real time procedure for affinely dependent parametric model order reduction using interpolation on Grassmann manifolds*, Internat. J. Numer. Methods Engrg., 93 (2013), pp. 818–833.
- [37] M. OHLBERGER AND F. SCHINDLER, *Error control for the localized reduced basis multi-scale method with adaptive on-line enrichment*, SIAM J. Sci. Comput., 37 (2015), pp. A2865–A2895, <https://doi.org/10.1137/151003660>.
- [38] H. PANZER, J. MOHRING, R. EID, AND B. LOHMANN, *Parametric model order reduction by matrix interpolation*, Automatisierungstechnik, 58 (2010), pp. 475–484.
- [39] A. PAUL-DUBOIS-TAINE AND D. AMSALLEM, *An adaptive and efficient greedy procedure for the optimal training of parametric reduced-order models*, Internat. J. Numer. Methods Engrg., 102 (2015), pp. 1262–1292.
- [40] B. PEHERSTORFER, D. BUTNARU, K. WILLCOX, AND H.-J. BUNGARTZ, *Localized discrete empirical interpolation method*, SIAM J. Sci. Comput., 36 (2014), pp. A168–A192, <https://doi.org/10.1137/130924408>.
- [41] B. PEHERSTORFER AND K. WILLCOX, *Detecting and adapting to parameter changes for reduced models of dynamic data-driven application systems*, in International Conference on Computational Science, Procedia Computer Science 51, Elsevier, 2015, pp. 2553–2562.
- [42] B. PEHERSTORFER AND K. WILLCOX, *Dynamic data-driven reduced-order models*, Comput. Methods Appl. Mech. Engrg., 291 (2015), pp. 21–41.
- [43] B. PEHERSTORFER AND K. WILLCOX, *Online adaptive model reduction for nonlinear systems via low-rank updates*, SIAM J. Sci. Comput., 37 (2015), pp. A2123–A2150, <https://doi.org/10.1137/140989169>.
- [44] L. PENG AND K. MOHSENI, *An online manifold learning approach for model reduction of dynamical systems*, SIAM J. Numer. Anal., 52 (2014), pp. 1928–1952, <https://doi.org/10.1137/130927723>.
- [45] M.-L. RAPÙN AND J. VEGA, *Reduced order models based on local POD plus Galerkin projection*, J. Comput. Phys., 229 (2010), pp. 3046–3063.
- [46] K. WASHBAUGH, D. AMSALLEM, M. ZAHR, AND C. FARHAT, *Nonlinear model reduction for CFD problems using local reduced-order bases*, in 42nd AIAA Fluid Dynamics Conference and Exhibit, Fluid Dynamics and Co-located Conferences, AIAA Paper 2012-2686, 2012, pp. 1–16.
- [47] Y.-C. WONG, *Differential geometry of Grassmann manifolds*, Proc. Nat. Acad. Sci. U.S.A., 57 (1967), pp. 589–594.
- [48] D. ZHANG AND L. BALZANO, *Convergence of a Grassmannian Gradient Descent Algorithm for Subspace Estimation from Undersampled Data*, preprint, <https://arxiv.org/abs/1610.00199>, 2016.
- [49] D. ZHANG AND L. BALZANO, *Global convergence of a Grassmannian gradient descent algorithm for subspace estimation*, in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS, Cadiz, Spain, 2016, pp. 1460–1468.
- [50] R. ZIMMERMANN, *A locally parametrized reduced order model for the linear frequency domain approach to time-accurate computational fluid dynamics*, SIAM J. Sci. Comput., 36 (2014), pp. B508–B537, <https://doi.org/10.1137/130942462>.
- [51] R. ZIMMERMANN AND K. WILLCOX, *An accelerated greedy missing point estimation procedure*, SIAM J. Sci. Comput., 38 (2016), pp. A2827–A2850, <https://doi.org/10.1137/15M1042899>.