SAND2016-0828J

KRYLOV-SUBSPACE RECYCLING VIA THE POD-AUGMENTED CONJUGATE-GRADIENT METHOD

KEVIN CARLBERG*, VIRGINIA FORSTALL†, AND RAY TUMINARO*

Abstract. This work presents a new Krylov-subspace-recycling method for efficiently solving sequences of linear systems of equations characterized by varying right-hand sides and symmetric-positive-definite matrices. As opposed to typical truncation strategies used in recycling such as deflation, we propose a truncation method inspired by goal-oriented proper orthogonal decomposition (POD) from model reduction. This idea is based on the observation that model reduction aims to compute a low-dimensional subspace that contains an *accurate* solution; as such, we expect the proposed method to generate a low-dimensional subspace that is well suited for computing solutions that can satisfy *inexact* tolerances. In particular, we propose specific goal-oriented POD 'ingredients' that align the optimality properties of POD with the objective of Krylov-subspace recycling. To compute solutions in the resulting 'augmented' POD subspace, we propose a hybrid direct/iterative three-stage method that leverages 1) the optimal ordering of POD basis vectors, and 2) well-conditioned reduced matrices. Numerical experiments performed on solid-mechanics problems highlight the benefits of the proposed method over existing approaches for Krylov-subspace recycling.

Key words. Krylov-subspace recycling, proper orthogonal decomposition, augmented Krylov methods, model reduction, conjugate-gradient method

AMS subject classifications. 15A12, 15A18, 65F10, 65F15, 65F50

1. Introduction.

This work considers solving a sequence of linear systems of equations characterized by varying right-hand sides and symmetric-positive-definite matrices. Such problems arise in a variety of engineering and science applications, including structural optimization, nonlinear structural dynamics, unconstrained numerical optimization, and nonlinear electromagnetics. In particular, we consider solving these linear systems to *inexact tolerances* using the (preconditioned) conjugate-gradient method; further, we allow for a solution-dependent output to serve as a quantity of interest.

While each linear-system solve can be executed independently of previous solves, reusing data generated during these solves can lead to improved convergence; this observation has led to the emergence of *Krylov-subspace recycling* methods. These recycling methods can also be considered 'augmented Krylov subspace methods' [36] because they 'augment' the typical Krylov subspace with a subspace computed from previous data, and subsequently compute solutions via projection with this augmented subspace. Initially, researchers developed methods that employ the space spanned by *all* Krylov vectors generated during the solution of previous linear systems as the augmenting subspace. In this framework, researchers developed methods to treat the case of an invariant matrix and multiple right-hand sides available either simultaneously (i.e., block Krylov methods) [28] or sequentially [35, 16, 14]; the sequential case occurs, for example, when restarting Krylov-subspace methods. These ideas were also extended to solve multiple systems with varying matrices; approximate orthogonalization techniques [31, 34], projection methods [15, 17], and an efficient full orthogonalization method [33] were developed for this purpose.

However, retaining the accumulation of all previous Krylov vectors can be computationally expensive and memory intensive, particularly when convergence is slow, the number of previous linear systems is large, or the length of the Krylov vectors is not large (e.g., preconditioners based on domain decomposition). This has led to the development of *truncation methods* that retain a subspace of the span of previous Krylov vectors as the augmenting subspace. First, deflation

*Sandia National Laboratories (ktcarlb@sandia.gov, rstumin@sandia.gov).
†University of Maryland (vhfors@gmail.com).

techniques for sequences of systems with invariant [7, 38] and varying [32, 29] matrices were developed. These methods employ approximated eigenvectors associated with the smallest eigenvalues of the governing matrices as the augmenting subspace. As such, they are effective primarily in cases where convergence is hampered by a few small eigenvalues. An alternative approach computes the augmenting subspace as the subspace that most accurately represents the Krylov subspace in the orthogonalization step of the generalized conjugate residual (GCR) method; this was also developed for both the invariant-matrix [10] and varying-matrices [29] cases.

These truncated Krylov-subspace-recycling techniques do not target the efficient solution of *inexact solutions*, which is the focus of this work. To this end, we propose a new proper orthogonal decomposition (POD)-augmented conjugate-gradient method for Krylov-subspace recycling. This approach—which employs goal-oriented model reduction to truncate previous Krylov vectors—is inspired by the observation that model-reduction techniques aim to generate low-dimensional approximations that preserve high levels of accuracy (i.e., satisfy inexact tolerances); it is based on preliminary work presented in Ref. [3]. The paper consists of the following new contributions.

- We perform analyses that expose the close relationship between goal-oriented POD and Krylov-subspace recycling (Theorems 3.1–3.3 and Corollary 3.4).
- We propose goal-oriented POD ingredients for truncating previous Krylov vectors, including
  - *Snapshots* comprising all previous Krylov vectors,
  - *Metrics* induced by 1) the system matrix and 2) the output quantity of interest, and
  - *Snapshot weights* arising from 1) the linear system before truncation and 2) a radial-basis-function approximation of the solution.

  Further, we bound the distance between these goal-oriented POD subspaces and the corresponding 'ideal' subspaces (Theorem 3.5).
- We develop a novel 'three-stage' algorithm, which accelerates the solution over the augmenting subspace using a hybrid direct/iterative approach. The algorithm leverages 1) the optimal ordering of POD basis vectors and 2) well-conditioned reduced matrices. The algorithm comprises
  - *Stage 1*: Direct solution over the first few (high-energy) POD basis vectors,
  - *Stage 2*: Iterative solution over the full augmenting space using the augmented CG algorithm, and
  - *Stage 3*: Iterative solution over the full space using the augmented preconditioned CG algorithm. This stage is equipped with new strategies for efficiently orthogonalizing against the entire augmenting subspace.

Section 2 provides the problem formulation, Section 3 describes the proposed POD-augmented conjugate-gradient algorithm, Section 4 describes the three-stage algorithm, Section 5 provides numerical experiments, and Section 6 concludes the paper. Note that Appendix B provides proofs for theoretical results. In the remainder of this manuscript, we denote matrices by capitalized bold letters, vectors by lowercase bold letters, and scalars by lowercase letters. We denote the columns of a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ by $\mathbf{a}_i \in \mathbb{R}^m$, $i \in \mathbb{N}(n)$ with $\mathbb{N}(a) := \{1, \dots, a\}$ such that $\mathbf{A} := [\mathbf{a}_1 \ \cdots \ \mathbf{a}_n]$. We denote the scalar-valued matrix elements by $a_{ij} \in \mathbb{R}$ such that $\mathbf{a}_j \equiv [a_{1j} \ \cdots \ a_{mj}]^T$, $j \in \mathbb{N}(n)$. In addition, we denote the range of a matrix by its calligraphic counterpart, i.e., $\text{range}(\mathbf{A}) \equiv \mathcal{A}$; we sometimes refer to $\mathbf{A}$ as the 'basis' for $\mathcal{A}$, although—more precisely—it is the 'basis in matrix form' for $\mathcal{A}$.

2. Problem formulation. This work considers solving a sequence of linear systems with a varying matrix

$$\mathbf{A}_j \mathbf{x}_j^\star = \mathbf{b}_j, \quad j = 1, \dots, p. \tag{2.1}$$

Here, $\mathbf{A}_j \in \mathrm{SPD}(n)$ and $\mathbf{b}_j \in \mathbb{R}^n$ denote the jth sparse system matrix and right-hand side, respectively, with $\mathrm{SPD}(n)$ denoting the set of symmetric-positive-definite (SPD) $n \times n$ matrices. The quantity $\mathbf{x}_j^\star \in \mathbb{R}^n$ is implicitly defined as the (exact) solutions to Eqs. (2.1). Further, we assume that the primary objective is to compute an output quantity of interest $\mathbf{q}(\mathbf{x}_j)$, $j = 1, \ldots, p$ with

$$\begin{aligned} \mathbf{q} :&\mathbb{R}^n \to \mathbb{R}^q \\ &\mathbf{x} \mapsto \mathbf{C}\mathbf{x}, \end{aligned}$$

where $\mathbf{C} \in \mathbb{R}^{q \times n}$. We consider computing a sequence of *inexact* solutions $\mathbf{x}_j$, $j = 1, \ldots, p$ to Eqs. (2.1) that satisfy

$$\|\mathbf{b}_j - \mathbf{A}_j \mathbf{x}_j\|_2 \leq \epsilon_j, \quad j = 1, \ldots, p. \tag{2.2}$$

where $\epsilon_j \geq 0$, $j = 1, \ldots, p$ denotes the forcing sequence [12].

2.1. Conjugate-gradient method. To compute each inexact solution $\mathbf{x}_j$, we consider applying the preconditioned conjugate gradient (PCG) algorithm, which computes a sequence of solutions that minimize the energy norm of the error over the (current) Krylov subspace. For the jth linear system, these solutions satisfy

$$\mathbf{x}_j^{(k)} = \arg \min_{\mathbf{x} \in \mathbf{x}_j^{(0)} + K_j^{(k)}(\mathbf{x}_j^{(0)})} \left\| \mathbf{x}_j^\star - \mathbf{x} \right\|_{\mathbf{A}_j}, \quad k = 1, \ldots, k_j, \tag{2.3}$$

where $K_j^{(k)} : \mathbf{x} \mapsto K^{(k)}\left(\mathbf{M}_j^{-1}\mathbf{A}_j, \mathbf{M}_j^{-1}(\mathbf{b}_j - \mathbf{A}_j\mathbf{x})\right) \subseteq \mathbb{R}^n$ is the affine search subspace, $K^{(k)} : (\mathbf{A}, \mathbf{b}) \mapsto \mathrm{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \ldots, \mathbf{A}^{k-1}\mathbf{b}\} \subseteq \mathbb{R}^n$ is the Krylov subspace at iteration k, $\mathbf{M}_j \in \mathrm{SPD}(n)$ is a preconditioner, $\mathbf{x}_j^{(0)} \in \mathbb{R}^n$ is the initial approximation, $(\mathbf{x}, \mathbf{y})_{\mathbf{A}_j} := \mathbf{x}^T \mathbf{A}_j \mathbf{y}$ and $\|\mathbf{x}\|_{\mathbf{A}_j} := \sqrt{(\mathbf{x}, \mathbf{x})_{\mathbf{A}_j}}$ denote the $\mathbf{A}_j$-weighted inner product and norm, and k_j denotes the number of iterations required such that $\mathbf{x}_j := \mathbf{x}_j^{(k_j)}$ satisfies inequality (2.2).

Optimality property (2.3) can be interpreted in a number of ways. In particular, given a basis $\mathbf{V} \in \mathbb{R}_\star^{n \times y}$ and centering point $\bar{\mathbf{x}} \in \mathbb{R}^n$ for an affine subspace $\bar{\mathbf{x}} + \mathcal{V}$ (recall $\mathcal{V} \equiv \mathrm{range}(\mathbf{V})$), then the solution $\tilde{\mathbf{x}}$ that minimizes the $\mathbf{A}$-norm of the error with respect to the exact solution $\mathbf{x}^\star$—which is implicitly defined by $\mathbf{A}\mathbf{x}^\star = \mathbf{b}$ with $\mathbf{A} \in \mathrm{SPD}(n)$ and $\mathbf{b} \in \mathbb{R}^n$—over this subspace can be written as

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x} \in \bar{\mathbf{x}} + \mathcal{V}} \|\mathbf{x}^\star - \mathbf{x}\|_{\mathbf{A}}. \tag{2.4}$$

Here, $\mathbb{R}_\star^{m \times n}$ denotes the noncompact Stiefel manifold: the set of full-column-rank $m \times n$ matrices. By the definition of an orthogonal projection [37], Eq. (2.4) can be equivalently expressed as

$$\tilde{\mathbf{x}} = \mathbf{P}_{\mathbf{A}}^{\bar{\mathbf{x}} + \mathcal{V}}(\mathbf{x}^\star), \tag{2.5}$$

where $\mathbf{P}_{\mathbf{A}}^K(\mathbf{x})$ denotes the $\mathbf{A}$-orthogonal projection of the vector $\mathbf{x} \in \mathbb{R}^n$ onto the (possibly affine) subspace $K \subseteq \mathbb{R}^n$, i.e., $\mathbf{P}_{\mathbf{A}}^K(\mathbf{x}) \in K$, $\forall \mathbf{x} \in \mathbb{R}^n$; $\left(\mathbf{x} - \mathbf{P}_{\mathbf{A}}^K(\mathbf{x}), \mathbf{z}\right)_{\mathbf{A}} = 0$, $\forall \mathbf{x} \in \mathbb{R}^n$, $\forall \mathbf{z} \in K$; and $\mathbf{P}_{\mathbf{A}}^K\left(\mathbf{P}_{\mathbf{A}}^K(\mathbf{x})\right) = \mathbf{P}_{\mathbf{A}}^K(\mathbf{x})$ (idempotency). This orthogonal projector can be defined algebraically as

$$\mathbf{P}_{\mathbf{A}}^{\bar{\mathbf{x}} + \mathcal{V}} : \mathbf{x} \mapsto \bar{\mathbf{x}} + \mathbf{V}\left(\mathbf{V}^T\mathbf{A}\mathbf{V}\right)^{-1}\mathbf{V}^T\mathbf{A}(\mathbf{x} - \bar{\mathbf{x}}). \tag{2.6}$$

If a symmetric factorization $\mathbf{A} = (\mathbf{A}^{1/2})^T\mathbf{A}^{1/2}$ is available (where $\mathbf{A}^{1/2}$ need not be upper triangular), then an equivalent definition is

$$\mathbf{P}_{\mathbf{A}}^{\bar{\mathbf{x}} + \mathcal{V}} : \mathbf{x} \mapsto \bar{\mathbf{x}} + \mathbf{V}\left(\mathbf{A}^{1/2}\mathbf{V}\right)^{+}\mathbf{A}^{1/2}(\mathbf{x} - \bar{\mathbf{x}}). \tag{2.7}$$

where a superscript + denotes the Moore–Penrose pseudoinverse.

Substituting definition (2.6) in Eq. (2.5) reveals that the solution $\tilde{\mathbf{x}}$ can be computed via Galerkin projection

$$\mathbf{V}^T\mathbf{A}\mathbf{V}\hat{\mathbf{x}} = \mathbf{V}^T\mathbf{r}(\bar{\mathbf{x}}), \quad \tilde{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{V}\hat{\mathbf{x}}, \tag{2.8}$$

where $\mathbf{r} : \mathbf{x} \mapsto \mathbf{b} - \mathbf{A}\mathbf{x}$ denotes the residual operator and we have used $\mathbf{A}\mathbf{x}^\star = \mathbf{b}$. If $\mathbf{V}^T\mathbf{A}\mathbf{V}$ is invertible, then Eqs. (2.5) and (2.8) imply that a vector $\tilde{\mathbf{x}}$ is the $\mathbf{A}$-orthogonal projection of the exact solution $\mathbf{x}^\star$ to $\mathbf{A}\mathbf{x}^\star = \mathbf{b}$ onto an affine subspace $\bar{\mathbf{x}} + \mathcal{V}$ if and only if its associated residual is orthogonal to $\mathcal{V}$ in the Euclidean inner product, i.e.,

$$\mathbf{x} = \mathbf{P}_{\mathbf{A}}^{\bar{\mathbf{x}}+\mathcal{V}}(\mathbf{x}^\star) \Leftrightarrow \mathbf{V}^T\mathbf{r}(\mathbf{x}) = 0. \tag{2.9}$$

Thus, equivalent interpretations of approximate solutions $\mathbf{x}_j^{(k)}$ are 1) as the solution to minimization problem (2.3); 2) as the orthogonal projection (from Eq. (2.4))

$$\mathbf{x}_j^{(k)} = \mathbf{P}_{\mathbf{A}_j}^{\mathbf{x}_j^{(0)}+K_j^{(k)}(\mathbf{x}_j^{(0)})}\left(\mathbf{x}_j^\star\right), \quad k = 1, \ldots, k_j;$$

3) as the Galerkin solution (from Eq. (2.8))

$$\begin{aligned}\mathbf{V}_j^{(k)T}\mathbf{A}_j\mathbf{V}_j^{(k)}\hat{\mathbf{x}}_j^{(k)} &= \mathbf{V}_j^{(k)T}\mathbf{r}_j(\mathbf{x}_j^{(0)}),\\ \mathbf{x}_j^{(k)} &= \mathbf{x}_j^{(0)} + \mathbf{V}_j^{(k)}\hat{\mathbf{x}}_j^{(k)}, \quad k = 1, \ldots, k_j,\end{aligned}$$

where $\mathbf{V}_j^{(k)} \in \mathbb{R}_\star^{n\times k}$ constitutes a basis for the subspace $K_j^{(k)}(\mathbf{x}_j^{(0)})$ and $\mathbf{r}_j : \mathbf{x} \mapsto \mathbf{b}_j - \mathbf{A}_j\mathbf{x}$ denotes the residual for linear system j; and 4) as the solution yielding residual orthogonality

$$\mathbf{V}_j^{(k)T}\mathbf{r}_j(\mathbf{x}_j^{(k)}) = 0.$$

2.2. Augmented conjugate-gradient method. Krylov-subspace recycling aims to reduce the computational burden of solving linear system j of Eqs. (2.1) by exploiting a previously computed 'augmenting' subspace $\mathcal{Y}_j \subseteq \mathbb{R}^n$ of dimension $y_j \leq n$ spanned by a basis $\mathbf{Y}_j \in \mathbb{R}_\star^{n\times y_j}$. To achieve this, recycling strategies employ augmented Krylov-subspace methods as a primary tool; rather than perform the typical sequence of Krylov iterations (e.g., Eqs. (2.3)), these methods first compute a solution in the augmenting subspace, and then compute an increment in a newly generated Krylov subspace. Critically, these methods ensure that the final solution minimizes the error over the sum of augmenting and Krylov subspaces; this generally requires maintaining orthogonality of new Krylov vectors (or $\mathbf{A}_j$-orthogonality of new search directions) to the augmenting subspace.

In the context of the conjugate-gradient method, these methods first solve a minimization problem in the affine subspace $\bar{\mathbf{x}}_j + \mathcal{Y}_j$, where $\bar{\mathbf{x}}_j$ is an initial guess:

$$\mathbf{x}_j^{\mathcal{Y}_j} = \arg\min_{\mathbf{x}\in\bar{\mathbf{x}}_j+\mathcal{Y}_j} \left\|\mathbf{x}_j^\star - \mathbf{x}\right\|_{\mathbf{A}_j} = \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right). \tag{2.10}$$

As before, the solution can be computed via Galerkin projection

$$\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j\hat{\mathbf{y}}_j = \mathbf{Y}_j^T\left(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j\right), \quad \mathbf{x}_j^{\mathcal{Y}_j} = \bar{\mathbf{x}}_j + \mathbf{Y}_j\hat{\mathbf{y}}_j. \tag{2.11}$$

Subsequently, these methods solve a minimization problem in the 'augmented' Krylov subspace

$$\mathbf{x}_j^{(k)} = \arg \min_{\mathbf{x} \in \bar{\mathbf{x}}_j + \mathcal{Y}_j + K_j^{(k)}(\mathbf{x}_j^{\mathcal{Y}_j})} \left\| \mathbf{x}_j^\star - \mathbf{x} \right\|_{\mathbf{A}_j}, \quad k = 1, \ldots, k_j.$$

The final solution can then be expressed as

$$\mathbf{V}_j^T \mathbf{A}_j \mathbf{V}_j \hat{\mathbf{v}}_j = \mathbf{V}_j^T \left(\mathbf{b}_j - \mathbf{A}_j \bar{\mathbf{x}}_j \right), \quad \mathbf{x}_j = \bar{\mathbf{x}}_j + \mathbf{Y}_j \hat{\mathbf{y}}_j + \mathbf{V}_j \hat{\mathbf{v}}_j, \tag{2.12}$$

where $\mathbf{V}_j := \mathbf{V}_j^{(k_j)}$ denotes a basis that satisfies

$$\mathbf{V}_j^T \mathbf{A}_j \mathbf{V}_j = \mathbf{\Gamma}_j, \quad \mathbf{Y}_j^T \mathbf{A}_j \mathbf{V}_j = \mathbf{0}, \tag{2.13}$$

$$\mathcal{Y}_j \oplus \mathcal{V}_j = \mathcal{Y}_j + K_j^{(k)}(\mathbf{x}_j^{\mathcal{Y}_j}), \tag{2.14}$$

where $\mathbf{\Gamma}_j$ is diagonal. Critically, $\mathbf{A}_j$-orthogonality of the augmented CG search directions to the augmenting basis (2.13) ensures the 'direct sum' property (2.14). This allows the solution to be expressed as the (unique) sum of components in $\mathcal{Y}_j$ and $\mathcal{V}_j$ by decoupling the Galerkin problems (2.11) and (2.12). The solution increment $\mathbf{V}_j \hat{\mathbf{v}}_j$ in Eq. (2.12) can be computed via the augmented PCG algorithm 1 as

$$(k_j, \hat{\mathbf{v}}_j, \mathbf{V}_j, \mathbf{\Gamma}_j) = \texttt{augmented_pcg}(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j \bar{\mathbf{x}}_j, \hat{\mathbf{y}}_j, \mathbf{Y}_j, \mathbf{M}_j, \epsilon_j).$$

Here, steps 4–5, and 13–14 enforce $\mathbf{A}_j$-orthogonality property (2.13).

Algorithm 1 `augmented_pcg`

Input: $\mathbf{A}$, $\mathbf{b}$, $\hat{\mathbf{x}}^{(0)}$, $\mathbf{Y}$, $\mathbf{M}$, ϵ
Output: k, $\hat{\mathbf{v}}$, $\mathbf{V}$, $\mathbf{\Gamma}$

1: $\mathbf{x}^{(0)} = \mathbf{Y}\hat{\mathbf{x}}^{(0)}$
2: $\mathbf{r}^{(0)} = \mathbf{b} - \mathbf{A}\mathbf{x}^{(0)}$
3: $\mathbf{z}^{(0)} = \mathbf{M}^{-1}\mathbf{r}^{(0)}$
4: Solve $\mathbf{Y}^T\mathbf{A}\mathbf{Y}\boldsymbol{\mu}^{(0)} = \mathbf{Y}^T\mathbf{A}\mathbf{z}^{(0)}$.
5: $\mathbf{p}^{(0)} = \mathbf{z}^{(0)} - \mathbf{Y}\boldsymbol{\mu}^{(0)}$
6: **for** $k = 0, 1, \ldots$ **do**
7: $\quad \gamma^{(k)} = (\mathbf{A}\mathbf{p}^{(k)}, \mathbf{p}^{(k)})$
8: $\quad \alpha^{(k)} = (\mathbf{r}^{(k)}, \mathbf{z}^{(k)})/\gamma^{(k)}$
9: $\quad \mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha^{(k)}\mathbf{p}^{(k)}$
10: $\quad \mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} - \alpha^{(k)}\mathbf{A}\mathbf{p}^{(k)}$
11: $\quad \mathbf{z}^{(k+1)} = \mathbf{M}^{-1}\mathbf{r}^{(k+1)}$
12: $\quad \beta^{(k+1)} = \frac{(\mathbf{r}^{(k+1)}, \mathbf{z}^{(k+1)})}{(\mathbf{r}^{(k)}, \mathbf{z}^{(k)})}$
13: $\quad$ Solve $\mathbf{Y}^T\mathbf{A}\mathbf{Y}\boldsymbol{\mu}^{(k+1)} = \mathbf{Y}^T\mathbf{A}\mathbf{z}^{(k+1)}$.
14: $\quad \mathbf{p}^{(k+1)} = \mathbf{z}^{(k+1)} + \beta^{(k+1)}\mathbf{p}^{(k)} - \mathbf{Y}\boldsymbol{\mu}^{(k+1)}$
15: $\quad$ **if** $\|\mathbf{r}^{(k+1)}\| \leq \epsilon$ **then**
16: $\quad\quad$ Exit.
17: $\quad$ **end if**
18: **end for**
19: $k = k + 1$, $\hat{\mathbf{v}} = \left[\alpha^{(0)} \ \cdots \ \alpha^{(k-1)}\right]^T$, $\mathbf{V} = \left[\mathbf{p}^{(0)} \ \cdots \ \mathbf{p}^{(k-1)}\right]$, $\mathbf{\Gamma} = \mathrm{diag}(\gamma^{(0)}, \ldots, \gamma^{(k-1)})$

Several strategies exist for selecting the augmenting basis $\mathbf{Y}_j$. Typically, the columns of this matrix consist of *all* Krylov vectors generated during the solution of previous linear systems; this provides the interpretation of 'recycling' Krylov subspaces. In this case, we have

$$\mathbf{Y}_j = \left[\mathbf{V}_1 \ \cdots \ \mathbf{V}_{j-1}\right].$$

However, after a modest number of linear systems has been solved, it becomes memory- and computation-intensive to retain and orthogonalize against this complete set of Krylov vectors. Therefore, truncation techniques have been devised to retain only the subspace that is 'most important' in some sense. In particular, after solving linear system $j-1$, these methods compute $\mathbf{Y}_j$ such that $\mathcal{Y}_j \subseteq \mathcal{Z}_j$, where $\mathbf{Z}_j := [\mathbf{Y}_{j-1},\ \mathbf{V}_{j-1}] \in \mathbb{R}^{n \times z_j}$ denotes the matrix of all preserved vectors accumulated over the first $j-1$ linear solves, with $z_j = y_{j-1} + k_{j-1}$. Note that $\mathbf{Z}_j = [\mathbf{V}_1\ \cdots\ \mathbf{V}_{j-1}]$ before truncation first occurs.

2.3. Deflation with harmonic Ritz vectors. Deflation techniques aim to retain the subspace associated with eigenvectors that tend to hamper convergence, i.e., those with eigenvalues close to zero. Such techniques have been developed for multiple linear systems with an invariant (e.g., restarting) [7, 38] and varying [32, 29] matrix.

To accomplish this, these techniques compute the harmonic Ritz vectors—which approximate the eigenvectors of $\mathbf{A}_j$ with the smallest eigenvalues—by solving the following problem: Find $\bar{\mathbf{y}}_i \in \text{range}(\mathbf{A}_{j-1}\mathbf{Z}_j)$ and $\lambda_i \in \mathbb{R}$, $i = 1, \ldots, y_j$ with $y_j \leq z_j$ such that

$$\left(\boldsymbol{w}, \mathbf{A}_{j-1}^{-1}\bar{\mathbf{y}}_i - \bar{\mathbf{y}}_i\lambda_i\right) = 0, \quad \forall \boldsymbol{w} \in \text{range}(\mathbf{A}_{j-1}\mathbf{Z}_j).$$

An equivalent problem statement is the following: Find $\mathbf{y}_i \in \mathcal{Z}_j$ and $\theta_i \in \mathbb{R}$, $i = 1, \ldots, y_j$ with $y_j \leq z_j$ such that

$$(\boldsymbol{w}, \mathbf{A}_{j-1}\mathbf{y}_i - \mathbf{y}_i\theta_i) = 0, \quad \forall \boldsymbol{w} \in \text{range}(\mathbf{A}_{j-1}\mathbf{Z}_j),$$

where $\theta_i = \lambda_i^{-1}$ and $\bar{\mathbf{y}}_i = \mathbf{A}_{j-1}\mathbf{y}_i$, $i = 1, \ldots, y_{j+1}$. Because vector $\mathbf{y}_i$ is equivalent to eigenvector $\bar{\mathbf{y}}_i$ with one step of inverse iteration, vectors $\mathbf{y}_i$ are typically employed as eigenvector approximations [25]. Algebraically, this corresponds to solving the generalized eigenvalue problem

$$\mathbf{Z}_j^T\mathbf{A}_{j-1}^T\mathbf{A}_{j-1}\mathbf{Z}_j\mathbf{G} = \mathbf{Z}_j^T\mathbf{A}_{j-1}\mathbf{Z}_j\mathbf{G}\boldsymbol{\Lambda}^{-1} \tag{2.15}$$

with $\boldsymbol{\Lambda} = \text{diag}\left(\lambda_1, \ldots, \lambda_{z_j}\right)$ and subsequently setting $\mathbf{Y}_j \leftarrow \left[\mathbf{Z}_j\mathbf{g}_1\ \cdots\ \mathbf{Z}_j\mathbf{g}_{y_j}\right]$, where the eigenvalues λ_i are ordered in decreasing magnitude.

Deflation has been pursued in the context of GMRES with deflated restarting (GMRES-DR [7, 32, 33, 26]), GCR with orthogonalization and deflated restarting (GCRO-DR) [29], and the deflated conjugate gradient method [7]. This approach is effective primarily in cases where the matrix is characterized by a small number of eigenvalues close to zero that hamper convergence. Further, because this approach aims to promote convergence to the 'exact' solution, it is not tailored for the efficient computation of *inexact* solutions, which is the focus of this work. To this end, we propose a novel truncation strategy inspired by model reduction.

3. POD-augmented CG. Because this work focuses on efficiently computing *inexact* solutions, we aim to compute an 'ideal' low-dimensional subspace $\mathcal{Y}_j^\star$ that directly minimizes the error in the augmenting-subspace solution appearing in Eq. (2.10), i.e.,

$$\mathcal{Y}_j^\star = \arg\min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \left\| \mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j + \mathcal{Y}_j}\left(\mathbf{x}_j^\star\right) \right\|_{\mathbf{A}_j}, \tag{3.1}$$

where $\mathcal{G}(m, n)$ with $m \leq n$ denotes the set of m-dimensional subspaces of $\mathbb{R}^n$ (the Grassmannian). Alternatively, we may be interested in computing inexact solutions that most accurately represent output quantities of interest. In this case, we can also consider an ideal output-oriented subspace

$$\bar{\mathcal{Y}}_j^\star = \arg\min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j + \mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right)\|_2. \tag{3.2}$$

Clearly, these ideal subspaces are computable if the exact solution $\mathbf{x}_j^\star$ is known; in this case, we can enforce $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j \in \mathcal{Y}_j^\star$ (resp. $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j \in \bar{\mathcal{Y}}_j^\star$)—which yields a zero objective-function value in Eq. (3.1) (resp. Eq. (3.2))—for any $y_j \geq 1$; if $y_j = 1$, then $\mathcal{Y}_j^\star = \text{span}\{\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\}$. However, the exact solution $\mathbf{x}_j^\star$ is not known before the jth linear system is solved. In this case, we aim to compute a subspace that *approximately* solves minimization problem (3.1) (resp. (3.2)). For this purpose, we employ goal-oriented POD with carefully chosen snapshots, weights, and metrics [2, 4].

3.1. POD. The POD method [18, 30] generates a basis that optimally represents a set of vectors (or 'snapshots') in a certain sense. The technique was developed in the context of model reduction, and has been applied to reduced-order modeling in structural dynamics [22, 23, 6] and fluid dynamics [42, 40, 27, 43, 39, 5] among others; it is also closely related to the singular value decomposition, principal component analysis, and the Karhunen–Loève expansion. We aim to select goal-oriented POD ingredients (i.e., snapshots, weights, and metric) [2, 4] that align the POD optimality property with the objective function in Eq. (3.1).

Given a matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$ whose columns represent s snapshots, a vector of weights $\boldsymbol{\gamma} \in \mathbb{R}^s$, and pseudometric associated with matrix $\boldsymbol{\Theta} \in \text{SPSD}(n)$, where $\text{SPSD}(n)$ denotes the set of $n \times n$ symmetric-positive-semidefinite matrices, the POD method computes a subspace of dimension $y \leq s$ that minimizes the sum of squared projection errors, i.e.,

$$\mathcal{U}_y^{\boldsymbol{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) = \arg\min_{\mathcal{A}\in\mathcal{G}(y,n)} \sum_{i=1}^{s} \| \left(\mathbf{I} - \mathbf{P}_{\boldsymbol{\Theta}}^{\mathcal{A}}\right)(\gamma_i \mathbf{s}_i) \|_{\boldsymbol{\Theta}}^2, \quad y = 1, \ldots, s. \tag{3.3}$$

A basis $\mathbf{U}_y^{\boldsymbol{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) \in \mathbb{R}_\star^{n\times y}$ for this POD subspace can be computed by solving the eigenvalue problem

$$\text{diag}(\gamma_1, \ldots, \gamma_s)\, \mathbf{S}^T \boldsymbol{\Theta} \mathbf{S} \text{diag}(\gamma_1, \ldots, \gamma_s)\, \mathbf{V} = \mathbf{V}\boldsymbol{\Sigma}^2$$

with $\boldsymbol{\Sigma} := \text{diag}(\sigma_1, \ldots, \sigma_s)$ and $\mathbf{V} \in \mathcal{S}_s(\mathbb{R}^s)$ and setting

$$\mathbf{U}_y^{\boldsymbol{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) = \mathbf{S}\text{diag}(\gamma_1, \ldots, \gamma_s) \left[\frac{1}{\sigma_1}\mathbf{v}_1 \ \ldots \ \frac{1}{\sigma_y}\mathbf{v}_y\right], \tag{3.4}$$

where the eigenvalues σ_i^2 are ordered in decreasing magnitude. Here, $\mathcal{S}_m(\mathbb{R}^n)$ denotes the Stiefel manifold: the set of orthonormal m-frames in $\mathbb{R}^n$. Alternatively, the POD basis can be computed via the singular value decomposition; Appendix A describes both algorithms. The resulting basis is nested

$$\mathbf{U}_{i+1}^{\boldsymbol{\Theta}} = [\mathbf{U}_i^{\boldsymbol{\Theta}} \ \mathbf{u}_{i+1}^{\boldsymbol{\Theta}}], \quad i = 1, \ldots, s-1 \tag{3.5}$$

with $\mathbf{u}_i^{\boldsymbol{\Theta}} \in \mathbb{R}^n$, $i = 1, \ldots, s$ and $\mathbf{U}_1^{\boldsymbol{\Theta}} = \mathbf{u}_1^{\boldsymbol{\Theta}}$ and exhibits $\boldsymbol{\Theta}$-orthogonality, i.e.,

$$\left(\mathbf{u}_i^{\boldsymbol{\Theta}}, \mathbf{u}_j^{\boldsymbol{\Theta}}\right)_{\boldsymbol{\Theta}} = \delta_{ij}, \tag{3.6}$$

where δ_{ij} denotes the Kronecker delta. We note that the algebraic definition (2.6) is not valid when considering pseudometrics, as $\mathbf{V}^T\boldsymbol{\Theta}\mathbf{V}$ may not be invertible if $\boldsymbol{\Theta}$ is semidefinite. In this case, definition (2.7) is appropriate, as $\boldsymbol{\Theta}^{1/2}\mathbf{V}$ always has a pseudoinverse.

REMARK 1. Note that Eqs. (3.3)–(3.5) imply that the POD basis vectors are optimally ordered, i.e., the first y POD basis vectors span the optimal y-dimensional subspace in the sense of Eq. (3.3).

3.2. POD and the augmented conjugate gradient method. We propose employing goal-oriented POD to define the augmenting subspace, i.e.,

$$\mathcal{Y}_j = \mathcal{U}_y^{\mathbf{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) \tag{3.7}$$

for specific choices of the snapshots $\mathbf{S}$, weights $\boldsymbol{\gamma}$, and metric $\mathbf{\Theta}$. The following results provide guidance toward this end. We first show that POD is related to minimizing $\left\|\mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right\|_{\mathbf{A}_j}$, which is the error minimized by ideal subspace $\mathcal{Y}_j^\star$ in problem (3.1).

THEOREM 3.1. *The POD subspace* $\mathcal{U}_y^{\mathbf{A}_j}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)$ *with*

$$\boldsymbol{\eta}_j^\star := \left(\mathbf{Z}_j^T\mathbf{A}_j\mathbf{Z}_j\right)^{-1}\mathbf{Z}_j^T\left(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j\right) \tag{3.8}$$

minimizes an upper bound for $\left\|\mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right\|_{\mathbf{A}_j}$ *over all* y*-dimensional subspaces of* $\mathcal{Z}_j \subseteq \mathbb{R}^n$*, where* $y \leq z_j$*.* See Appendix B for the proof (and all subsequent proofs).

We now show that other POD ingredients can be selected to align POD with minimizing $\|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right)\|_2$, which is the error minimized by ideal subspace $\bar{\mathcal{Y}}_j^\star$ in problem (3.2).

THEOREM 3.2. *The POD subspace* $\mathcal{U}_y^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \bar{\boldsymbol{\eta}}_j^\star\right)$ *with*

$$\bar{\boldsymbol{\eta}}_j^\star := \left(\mathbf{C}\mathbf{Z}_j\right)^+\mathbf{C}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right) \tag{3.9}$$

minimizes an upper bound for the output error $\|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right)\|_2$ *over all* y*-dimensional subspaces of* $\mathcal{Z}_j \subseteq \mathbb{R}^n$*, where* $y \leq z_j \leq n$*.*

3.3. Goal-oriented POD ingredients. In light of these theoretical results, we now propose several practical choices for POD ingredients that align goal-oriented POD with augmented CG.

3.3.1. Snapshots. Theorems 3.1 and 3.2 show that POD minimizes an upper bound for errors of interest if snapshots are set to vectors accumulated over the first $j-1$ linear solves. We therefore employ $\mathbf{S} = \mathbf{Z}_j$ in Eq. (3.7).

3.3.2. Weights. In practice, the augmenting subspace $\mathcal{Y}_j$ is computed after linear system $j-1$ is solved. As such, we cannot compute the 'ideal weights' $\bar{\boldsymbol{\eta}}_j^\star$ defined in (3.9), which requires knowledge of $\mathbf{x}_j^\star$. While $\boldsymbol{\eta}_j^\star$ in Eq. (3.8) can be computed by solving $\mathbf{Z}_j^T\mathbf{A}_j\mathbf{Z}_j\boldsymbol{\eta}_j^\star = \mathbf{Z}_j^T\left(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j\right)$, this is not practical: the computational cost of doing so is equivalent to solving Eq. (2.11) with $\mathbf{Y}_j = \mathbf{Z}_j$, i.e., employing an augmenting subspace that has not been truncated. As we aim to employ a truncated augmented space for the jth linear-system solve, we consider two approximations to these ideal weights.

1. **Previous weights**. This approach employs weights of

$$\boldsymbol{\eta}_j^{\text{prev}} = \left(\mathbf{Z}_j^T\mathbf{A}_{j-1}\mathbf{Z}_j\right)^{-1}\mathbf{Z}_j^T\left(\mathbf{b}_{j-1} - \mathbf{A}_{j-1}\bar{\mathbf{x}}_{j-1}\right). \tag{3.10}$$

Comparing Eqs. (3.8) and (3.10) reveals that $\boldsymbol{\eta}_j^{\text{prev}}$ is 'close' to $\boldsymbol{\eta}_j^\star$, but employs readily available data, as these weights are equal to the coefficient in the expansion of the solution at the previous time step, i.e., $\mathbf{x}_{j-1} = \bar{\mathbf{x}}_{j-1} + \mathbf{Z}_j\boldsymbol{\eta}_j^{\text{prev}}$. We now provide bounds for the difference between these computable weights and the ideal weights.

THEOREM 3.3. *The difference between the previous weights and ideal weights can be bounded as*

$$\|\boldsymbol{\eta}_j^\star - \boldsymbol{\eta}_j^{prev}\| \leq \frac{1}{\sigma_{min}} \left(\|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j} - \mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}\| \|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\| + \sigma_1 \|(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - (\mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1})\| \right) \tag{3.11}$$

and

$$\|\bar{\boldsymbol{\eta}}_j^\star - \boldsymbol{\eta}_j^{prev}\| \leq \frac{1}{\sigma_{min}} \left(\|\mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Z}_j} - \mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}\| \|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\| + \sigma_1 \|(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - (\mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1})\| \right), \tag{3.12}$$

where σ_{min} denotes the smallest singular value of $\mathbf{Z}_j$ and σ_1 denotes the largest singular value of $\mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}$. We now provide conditions under which the computable weights are equal to the ideal weights.

COROLLARY 3.4. *If $\mathbf{A}_{j-1} = \mathbf{A}_j$ and $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j = \mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1}$, then $\boldsymbol{\eta}_j^\star = \boldsymbol{\eta}_j^{prev}$. Alternatively, if $\mathbf{C}^T\mathbf{C} = \mathbf{A}_j$ and $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j = \mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1}$, then $\bar{\boldsymbol{\eta}}_j^\star - \boldsymbol{\eta}_j^{prev}$.*

Proof. The result follows trivially from Eqs. (3.11) *and* (3.12). □

2. **Radial-basis-function weights**. Noting that $\mathbf{Z}_j \boldsymbol{\eta}_j^\star = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j} \left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j \right)$, we can approximate ideal weights $\boldsymbol{\eta}_j^\star$ by approximating the component of $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j$ in $\mathcal{Z}_j$. To achieve this, we assume the jth solution can be approximated as a linear combination of previous solutions, i.e.,

$$\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j} \left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j \right) \approx \sum_{i=1}^{\omega} \rho^{\mathrm{RBF}}(j, j-i) (\mathbf{P}_{\mathbf{A}_{j-i}}^{\mathcal{Z}_{j-i}} \left(\mathbf{x}_{j-i}^\star - \bar{\mathbf{x}}_{j-i} \right)),$$

where $\omega \in \mathbb{N}(j-1)$ denotes the number of previous solutions to include and $\rho^{\mathrm{RBF}}(i,j)$ denotes a radial basis function. This implies weights of

$$\begin{aligned} \boldsymbol{\eta}_j^{\mathrm{RBF}}(\omega) &= \sum_{i=1}^{\omega} \rho^{\mathrm{RBF}}(j, j-i) \left(\mathbf{Z}_{j+1-i}^T \mathbf{A}_{j-i} \mathbf{Z}_{j+1-i} \right)^{-1} \mathbf{Z}_{j+1-i}^T \mathbf{A}_{j-i} \left(\mathbf{x}_{j-i}^\star - \bar{\mathbf{x}}_{j-i} \right) \\ &= \sum_{i=1}^{\omega} \rho^{\mathrm{RBF}}(j, j-i) \boldsymbol{\eta}_{j+1-i}^{\mathrm{prev}}. \end{aligned} \tag{3.13}$$

In practice, we employ a inverse-distance-weight radial basis function of the form $\rho^{\mathrm{RBF}}(i,j) = \rho^{\mathrm{IDW}}(|i-j|)$ with $\rho^{\mathrm{IDW}} : r \mapsto 1/2^{r-1}$; we set ω to be the number of linear systems since the most recent truncation.

3.3.3. Metric. Theorems 3.1 and 3.2 demonstrate that POD minimizes an upper bound for the errors $\left\| \mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j + \mathcal{Y}_j} \left(\mathbf{x}_j^\star \right) \right\|_{\mathbf{A}_j}$ and $\|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j + \mathcal{Y}_j} \left(\mathbf{x}_j^\star \right) \right)\|_2$ only if POD metrics of $\boldsymbol{\Theta} = \mathbf{A}_j$ and $\boldsymbol{\Theta} = \mathbf{C}^T\mathbf{C}$ are used, respectively. However, as previously discussed, we aim to avoid using $\mathbf{A}_j$ to truncate the basis $\mathbf{Z}_j$, as it entails reduced computations with the jth linear system (i.e., solving Eqs. (2.11) with $\mathbf{Y}_j = \mathbf{Z}_j$). Therefore, we employ two practical choices for metrics in Eq. (3.7):

1. $\boldsymbol{\Theta} = \mathbf{A}_{j-1}$. This approach aligns the truncation with minimizing $\left\| \mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j + \mathcal{Y}_j} \left(\mathbf{x}_j^\star \right) \right\|_{\mathbf{A}_j}$. From Eq. (3.6), the resulting basis is $\mathbf{A}_{j-1}$-orthogonal. Here, computing the POD basis via Algorithm 4 is appropriate, as a symmetric factorization of $\mathbf{A}_{j-1}$ is not readily available.
2. $\boldsymbol{\Theta} = \mathbf{C}^T\mathbf{C}$. This is an output-oriented approach associated with minimizing the output error. From Eq. (3.6), the resulting basis is $\mathbf{C}^T\mathbf{C}$-orthogonal. Computing this POD basis via Algorithm 5 is appropriate, as a symmetric factor of the (pseudo)metric is readily available as $\mathbf{C}$.

3.3.4. Computable goal-oriented POD methods. In summary, the four (computable) goal-oriented POD truncation methods we propose are

1. $\mathbf{Y}_j = \mathbf{U}_{y_j}^{\mathbf{A}_{j-1}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}\right)$, computable by $(\mathbf{Y}_j) = \texttt{pod_evd}(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}, \mathbf{A}_{j-1}, \varepsilon_y)$,
2. $\mathbf{Y}_j = \mathbf{U}_{y_j}^{\mathbf{A}_{j-1}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega)\right)$, computable by $(\mathbf{Y}_j) = \texttt{pod_evd}(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega), \mathbf{A}_{j-1}, \varepsilon_y)$,
3. $\mathbf{Y}_j = \mathbf{U}_{y_j}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}\right)$, computable by $(\mathbf{Y}_j) = \texttt{pod_svd}(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}, \mathbf{C}, \varepsilon_y)$,
4. $\mathbf{Y}_j = \mathbf{U}_{y_j}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega)\right)$, computable by $(\mathbf{Y}_j) = \texttt{pod_svd}(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega), \mathbf{C}, \varepsilon_y)$,

where $\varepsilon_y \in [0,1]$ is a statistical 'energy criterion' used to truncate the POD basis, and `pod_evd` and `pod_svd` are described in Algorithms 4 and 5, respectively. We now provide a result that allows us to bound the distance between the four computable goal-oriented POD subspaces above and their ideal counterparts, which correspond to $\mathcal{U}_{y_j}^{\mathbf{A}_j}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)$ for 1–2 and $\mathcal{U}_{y_j}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \bar{\boldsymbol{\eta}}_j^\star\right)$ for 3–4 above.

THEOREM 3.5 (Distance between goal-oriented POD subspaces). *The distance between two goal-oriented POD subspaces $\mathcal{U}^a := \mathcal{U}_y^{\boldsymbol{\Theta}^a}(\mathbf{S}, \boldsymbol{\gamma}^a)$ and $\mathcal{U}^b := \mathcal{U}_y^{\boldsymbol{\Theta}^b}(\mathbf{S}, \boldsymbol{\gamma}^b)$—which are characterized by different metrics $\boldsymbol{\Theta}^a \in SPSD(n)$ and $\boldsymbol{\Theta}^b \in SPSD(n)$ with $\boldsymbol{\Theta}^a = \boldsymbol{\Theta}^b + \boldsymbol{\Delta}$ and different snapshot weights $\boldsymbol{\Gamma}^a := diag(\gamma_1^a, \ldots, \gamma_s^a)$ and $\boldsymbol{\Gamma}^b := diag(\gamma_1^b, \ldots, \gamma_s^b)$ associated with snapshots $\mathbf{S} \in \mathbb{R}^{n\times s}$—can be bounded as*

$$d\left(\mathcal{U}^a, \mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b)\Big(\|\left(\boldsymbol{\Gamma}^b + \boldsymbol{\Gamma}^a\right)\left(\boldsymbol{\Gamma}^b - \boldsymbol{\Gamma}^a\right)[\boldsymbol{\Gamma}^a]^{-1}\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\boldsymbol{\Gamma}^a\|_2 + \|\boldsymbol{\Delta}\|_2\|\mathbf{S}\boldsymbol{\Gamma}^a\|_2^2\Big)/absSep(\boldsymbol{\Lambda}_\perp^a, \boldsymbol{\Lambda}^b) \tag{3.14}$$

if the matrix $\mathbf{S}\boldsymbol{\Gamma}^a$ has full column rank. Here, $d(\mathcal{U}, \mathcal{V}) := \max_{u\in\mathcal{U}, \|u\|=1}\min_{v\in\mathcal{V}}\|u - v\| = \sin(\theta_{max})$, where θ_{max} is the largest principal angle between $\mathcal{U}$ and $\mathcal{V}$, $\kappa(\mathbf{A})$ denotes the condition number of matrix $\mathbf{A}$, and $absSep(\boldsymbol{\Lambda}_1, \boldsymbol{\Lambda}_2) := \min_{\|\boldsymbol{Z}\|_2=1}\|\boldsymbol{\Lambda}_1\boldsymbol{Z} - \boldsymbol{Z}\boldsymbol{\Lambda}_2\|_2$ denotes the absolute separation between two spectra. In addition, $\mathbf{V}^b \in \mathcal{S}_y(\mathbb{R}^s)$ and $\boldsymbol{\Lambda}^b \in \mathbb{R}^{y\times y}$ denote the first y eigenvectors and (diagonal matrix of) eigenvalues, respectively, of the matrix $\boldsymbol{\Gamma}^b\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\boldsymbol{\Gamma}^b$, assuming eigenvalues are ordered in descending magnitude. Similarly, $\boldsymbol{\Lambda}_\perp^a$ denotes the diagonal matrix of eigenvalues associated with the orthogonal complement to the invariant subspace spanned by the first y eigenvectors of $\boldsymbol{\Gamma}^a\mathbf{S}^T\boldsymbol{\Theta}^a\mathbf{S}\boldsymbol{\Gamma}^a$.

This theorem allows us to bound the distances between the four computable POD subspaces and their ideal counterparts by substituting the appropriate quantities into Theorem 3.5. That is, the theorem allows us to bound the following quantities:

1. $d\left(\mathcal{U}_y^{\mathbf{A}_{j-1}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}\right), \mathcal{U}_y^{\mathbf{A}_j}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)\right)$,
2. $d\left(\mathcal{U}_y^{\mathbf{A}_{j-1}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega)\right), \mathcal{U}_y^{\mathbf{A}_j}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)\right)$,
3. $d\left(\mathcal{U}_y^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{prev}}\right), \mathcal{U}_y^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)\right)$, and
4. $d\left(\mathcal{U}_y^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^{\text{RBF}}(\omega)\right), \mathcal{U}_y^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_j, \boldsymbol{\eta}_j^\star\right)\right)$.

REMARK 2. The subspace-distance bound (3.14) simplifies under certain conditions.
Fixed snapshot weights. First, note that when the snapshot weights are identical but the metric differes, i.e., $\boldsymbol{\Gamma}^b = \boldsymbol{\Gamma}^a$, the first term within parentheses vanishes and $\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b) = 1$ due to symmetry of $\mathbf{B}^b$, which yields

$$d\left(\mathcal{U}^a, \mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\|\boldsymbol{\Delta}\|_2\|\mathbf{S}\boldsymbol{\Gamma}^a\|_2^2/\text{absSep}(\boldsymbol{\Lambda}_\perp^a, \boldsymbol{\Lambda}^b).$$

Thus, the subspace distance depends linearly on the norm of the metric perturbation in this case.

Fixed metric. If instead the snapshot weights differ while $\boldsymbol{\Delta} = \mathbf{0}$, then the second term within parentheses vanishes and we have

$$d\left(\mathcal{U}^a,\mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b)\|\left(\boldsymbol{\Gamma}^b+\boldsymbol{\Gamma}^a\right)\left(\boldsymbol{\Gamma}^b-\boldsymbol{\Gamma}^a\right)[\boldsymbol{\Gamma}^a]^{-1}\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\boldsymbol{\Gamma}^a\|_2/\text{abssep}(\boldsymbol{\Lambda}^a_\perp,\boldsymbol{\Lambda}^b)$$

In practice, this situation corresponds to cases 3–4 above (where $\boldsymbol{\Theta}^b = \boldsymbol{\Theta}^a = \mathbf{C}^T\mathbf{C}$) or to cases 1–2 when $\mathbf{A}_j = \mathbf{A}_{j-1}$. Notice that the relative weight perturbation $(\boldsymbol{\Gamma}^b - \boldsymbol{\Gamma}^a)[\boldsymbol{\Gamma}^a]^{-1}$ appears within the remaining term. When the relative weight perturbation can be bounded as

$$\max_k (\boldsymbol{\gamma}^b_j - \boldsymbol{\gamma}^a)_k/(\boldsymbol{\gamma}^a)_k \leq \hat{\delta},$$

then it can be shown that the bound simplifies to

$$d\left(\mathcal{U}^a,\mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b)\hat{\delta}(2+\hat{\delta})\|\boldsymbol{\Gamma}^a\mathbf{S}^T\boldsymbol{\Theta}^a\mathbf{S}\boldsymbol{\Gamma}^a\|_2/\text{abssep}(\boldsymbol{\Lambda}^a_\perp,\boldsymbol{\Lambda}^b).$$

Further, when $\boldsymbol{\Gamma}^a$ commutes with $\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}$, then the bound simplifies to

$$d\left(\mathcal{U}^a,\mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b)\|\left(\boldsymbol{\Gamma}^b+\boldsymbol{\Gamma}^a\right)\left(\boldsymbol{\Gamma}^b-\boldsymbol{\Gamma}^a\right)\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\|_2/\text{abssep}(\boldsymbol{\Lambda}^a_\perp,\boldsymbol{\Lambda}^b),$$

which now depends on the absolute weight perturbation, $\boldsymbol{\Gamma}^b - \boldsymbol{\Gamma}^a$, rather than a relative perturbation. This occurs for cases 1–2 above when $\mathbf{A}_1 = \cdots = \mathbf{A}_j = \mathbf{A}$ (i.e., an invariant matrix), as $\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}$ is diagonal due to enforced $\mathbf{A}$-orthogonality of $\mathbf{S}$.

Strongly separated eigenvalues. If the eigenvalues associated with $\boldsymbol{\Lambda}^b$ and $\boldsymbol{\Lambda}^a_\perp$ in Theorem 3.5 are strongly separated, i.e.,

$$0 < \delta_a := \max\{\min_{\lambda\in\boldsymbol{\Lambda}^b}|\lambda| - \max_{\hat{\lambda}\in\boldsymbol{\Lambda}^a_\perp}|\hat{\lambda}|,\ \min_{\hat{\lambda}\in\boldsymbol{\Lambda}^a_\perp}|\hat{\lambda}| - \max_{\lambda\in\boldsymbol{\Lambda}^b}|\lambda|\}$$

then the expression for the subspace-distance bound becomes [19, Theorem 5.3]

$$d\left(\mathcal{U}^a,\mathcal{U}^b\right) \leq \kappa(\mathbf{S}\boldsymbol{\Gamma}^a)\kappa([\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b\mathbf{V}^b)\Big(\|\left(\boldsymbol{\Gamma}^b+\boldsymbol{\Gamma}^a\right)\left(\boldsymbol{\Gamma}^b-\boldsymbol{\Gamma}^a\right)[\boldsymbol{\Gamma}^a]^{-1}\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\boldsymbol{\Gamma}^a\|_2 + \|\boldsymbol{\Delta}\|_2\|\mathbf{S}\boldsymbol{\Gamma}^a\|_2^2\Big)/\delta_a.$$

Finally, we note that we can apply the results in Eq. (3.11) from Theorem 3.3 to those of Theorem 3.5 via the triangle inequality to obtain a more explicit subspace-distance bound for case 1 above, where $\boldsymbol{\Gamma}^a = \text{diag}\left([\boldsymbol{\eta}^\star_j]_1,\ldots,[\boldsymbol{\eta}^\star_j]_s\right)$ and $\boldsymbol{\Gamma}^b = \text{diag}\left([\boldsymbol{\eta}^{\text{prev}}_j]_1,\ldots,[\boldsymbol{\eta}^{\text{prev}}_j]_s\right)$.

4. Three-stage algorithm. From Section 2.2, we know that augmented PCG relies on first computing a solution over the augmenting subspace $\bar{\mathbf{x}}_j + \mathcal{Y}_j$. While this can be accomplished by solving the Galerkin system (2.11) directly, this requires assembling the reduced matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$, which incurs y_j matrix–vector products and $\mathcal{O}((\tau_j + y_j)y_j n)$ flops, where τ_j denotes the average number of nonzeros per row of $\mathbf{A}_j$ This cost can be significant when the dimension of the augmenting space y_j becomes large.

To mitigate this effect, we propose a novel three-stage algorithm for augmented PCG. The algorithm leverages an augmenting basis that is optimally ordered and yields a well-conditioned reduced matrix. In particular, for linear system j, we assume that an augmenting basis $\mathbf{Y}_j \in \mathbb{R}^{n\times y_j}_\star$ is available that 1) contains a low-dimensional basis $\mathbf{W}_j \in \mathbb{R}^{n\times w_j}_\star$ with $\mathcal{W}_j \subseteq \mathcal{Y}_j \subseteq \mathbb{R}^n$ and $w_j \leq y_j \leq n$ that can capture an accurate solution, and 2) yields a well-conditioned reduced

matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$. Given this basis, stage 1 first computes an accurate solution over the subspace $\bar{\mathbf{x}}_j + \mathcal{W}_j$; a direct approach is suitable due to its assumed low dimensionality. Stage 2 iteratively computes a solution over $\bar{\mathbf{x}}_j + \mathcal{Y}_j$. This should be efficient, as $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$ need never be formed: the well-conditioned assumption implies that a preconditioner is not required for fast convergence and matrix–vector products of the form $\mathbf{Y}^T\mathbf{A}\mathbf{Y}\mathbf{p}$ can be computed using a single matrix–vector product as $\mathbf{Y}^T(\mathbf{A}(\mathbf{Y}\mathbf{p}))$. Finally, stage 3 applies augmented PCG in the full space to compute a solution satisfying the specified tolerance ϵ_j.

Goal-oriented POD as proposed in Section 3.3 fits naturally into this framework. First, the basis vectors are optimally ordered (Remark 1): the first few POD basis vectors span an optimal subspace in the sense of minimizing the objective function in Eq. (3.3), which is an upper bound for the $\mathbf{A}_j$-norm and $\mathbf{C}^T\mathbf{C}$-norm of the error, respectively (Theorems 3.1 and 3.2). This implies that the first few POD vectors could be employed as $\mathbf{W}_j$. Second, it automatically yields a well-conditioned reduced matrix if $\mathbf{\Theta} = \mathbf{A}_j$; this will be further discussed in Section 4.4.1. However, other truncation methods that satisfy these properties can also be considered within the proposed three-stage algorithm.

4.1. Stage 1. The objective of Stage 1 is to 'jump start' the algorithm by computing an accurate solution at very low cost. To achieve this, this stage computes $\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j = \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{W}_j}\left(\mathbf{x}_j^\star\right)$ by solving

$$\mathbf{W}_j^T\mathbf{A}_j\mathbf{W}_j\hat{\mathbf{w}}_j = \mathbf{W}_j^T(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j)$$

directly. The assumptions placed on $\mathbf{W}_j$ imply that the solution $\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j$ will be accurate and a direct solve will be inexpensive. Computing this solution can be executed via Algorithm 2 as

$$(\hat{\mathbf{w}}_j, \hat{\mathbf{R}}_j) = \texttt{direct_reduced_solve}(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \mathbf{W}_j).$$

Algorithm 2 `direct_reduced_solve`

Input: $\mathbf{A}$, $\mathbf{b}$, $\mathbf{W}$
Output: $\hat{\mathbf{w}}$, $\hat{\mathbf{R}}$
1: Compute $\hat{\mathbf{A}} = \mathbf{W}^T\mathbf{A}\mathbf{W}$ and $\hat{\mathbf{b}} = \mathbf{W}^T\mathbf{b}$
2: Solve $\hat{\mathbf{A}}\hat{\mathbf{w}} = \hat{\mathbf{b}}$ by Cholesky factorization $\hat{\mathbf{A}} = \hat{\mathbf{R}}^T\hat{\mathbf{R}}$

4.2. Stage 2. The objective of Stage 2 is to improve upon the stage-1 solution $\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j$ by efficiently solving over the entire augmenting subspace $\mathcal{Y}_j$, whose dimension y_j may be large, while avoiding the cost of explicitly forming the reduced matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$. To achieve this, this stage solves

$$\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j\hat{\mathbf{y}}_j^\star = \mathbf{Y}_j^T\left(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j - \mathbf{A}_j\mathbf{W}_j\hat{\mathbf{w}}_j\right)$$

iteratively to tolerance $\hat{\epsilon}_j$ via augmented CG with augmenting basis $\hat{\mathbf{W}}_j \in \mathbb{R}_\star^{y_j\times w_j}$, which represents the stage-1 basis in augmenting-subspace coordinates, i.e., $\mathbf{W}_j = \mathbf{Y}_j\hat{\mathbf{W}}_j$. This approach is promising for two reasons. First, it precludes the need to explicitly compute the reduced matrix, as matrix–vector products of the form $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j\mathbf{p}$ can be computed via $\mathbf{Y}_j^T(\mathbf{A}_j(\mathbf{Y}_j\mathbf{p}))$ using a single matrix–vector product and $\mathcal{O}((2y_j + \tau_j)n)$ flops. Second, a small number of iterations will be needed for convergence if the reduced matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$ is well conditioned, as has been assumed.

Therefore, this stage amounts to executing Algorithm 1 as

$$(\hat{k}_j, \hat{\mathbf{x}}_j, \hat{\mathbf{X}}_j, \hat{\mathbf{\Gamma}}_j) = \texttt{augmented_pcg}(\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j, \mathbf{Y}_j^T(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j), \hat{\mathbf{w}}_j, \hat{\mathbf{W}}_j, \mathbf{I}, \hat{\epsilon}_j)$$

with an implementation optimization: the solves in steps 4 and 13 of Algorithm 1 can be performed directly in $\mathcal{O}(w_j^2)$ flops, as the Cholesky factorization of $\mathbf{W}_j^T\mathbf{A}_j\mathbf{W}_j = \hat{\mathbf{R}}_j^T\hat{\mathbf{R}}_j$ can be reused from stage 1.

The (inexact) solution increment $\mathbf{Y}_j\hat{\mathbf{y}}_j = \mathbf{X}_j\hat{\mathbf{x}}_j$ computed by stage 2 lies in the range of a (reduced) Krylov basis $\mathbf{X}_j = \mathbf{Y}_j\hat{\mathbf{X}}_j \in \mathbb{R}^{n\times\hat{k}_j}$ with $\mathcal{X}_j \subseteq \mathcal{Y}_j \subseteq \mathbb{R}^n$ and $\hat{\mathbf{X}}_j \in \mathbb{R}^{y_j\times\hat{k}_j}$, where $\hat{k}_j \leq y_j$ denotes the number of stage-2 iterations. This basis satisfies

$$\mathbf{X}_j^T\mathbf{A}_j\mathbf{X}_j = \hat{\mathbf{\Gamma}}_j, \quad \mathbf{X}_j^T\mathbf{A}_j\mathbf{W}_j = 0,$$
$$\hat{\mathcal{W}}_j \oplus \hat{\mathcal{X}}_j = \hat{\mathcal{W}}_j + K^{(\hat{k}_j)}\left(\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j, \mathbf{Y}_j^T(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j - \mathbf{A}_j\mathbf{W}_j\hat{\mathbf{w}}_j)\right),$$

while the resulting solution satisfies

$$\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j = \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{W}_j\oplus\mathcal{X}_j}\left(\mathbf{x}_j^\star\right)$$

Ref. [33] proposed a similar idea referred to as the iterative reuse of Krylov subspaces (IRKS). This approach did not employ stage-1 direct solve and supported using either $\mathbf{Y}_j = [\mathbf{V}_1\ \cdots\ \mathbf{V}_{j-1}]$ or $\mathbf{Y}_j = \mathbf{V}_{j-1}$ (i.e., no truncation).

4.3. Stage 3. The objective of stage 3 is to continue iterating in the full space until the specified tolerance ϵ_j is satisfied. This stage therefore solves

$$\mathbf{A}_j\delta\mathbf{x}_j^\star = \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j - \mathbf{A}_j\mathbf{W}_j\hat{\mathbf{w}}_j - \mathbf{A}_j\mathbf{X}_j\hat{\mathbf{x}}_j$$

via augmented PCG with augmenting basis $[\mathbf{W}_j,\ \mathbf{X}_j]$. Critically, note that the augmenting subspace used in stage 3 satisfies range$([\mathbf{W}_j,\ \mathbf{X}_j]) \subseteq \mathcal{Y}_j$. This choice is made because $\mathbf{A}_j[\mathbf{W}_j,\ \mathbf{X}_j]$ and $[\mathbf{W}_j,\ \mathbf{X}_j]^T\mathbf{A}_j[\mathbf{W}_j,\ \mathbf{X}_j]$—which are required for orthogonalization in step 13 of Algorithm 1—are available from stages 1 and 2. Treating $\mathcal{Y}_j$ as the augmenting subspace in stage 3 would necessitate computing $\mathbf{A}_j\mathbf{Y}_j$ and $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$, which would eliminate any cost savings realized during stages 1 and 2, as computing and factorizing the reduced matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$ incurs $\mathcal{O}((\tau_j + {y_j}^2)n + {y_j}^3)$ flops: the same cost as executing stage 1 with $\mathbf{W}_j = \mathbf{Y}_j$. Thus, stage 3 executes Algorithm 1 as

$$(k_j, \hat{\mathbf{v}}_j, \mathbf{V}_j, \mathbf{\Gamma}_j) = \texttt{augmented_pcg}(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \left[\hat{\mathbf{w}}_j^T,\ \hat{\mathbf{x}}_j^T\right]^T, [\mathbf{W}_j,\ \mathbf{X}_j], \mathbf{M}_j, \epsilon_j)$$

with the following optimization: the solves in steps 4 and 13 can be performed directly in $\mathcal{O}(w_j^2 + \hat{k}_j)$ flops, as the Cholesky factorization $[\mathbf{W}_j,\ \mathbf{X}_j]^T\mathbf{A}[\mathbf{W}_j,\ \mathbf{X}_j] = \mathbf{R}_j^T\mathbf{R}_j$ is readily available from stages 1 and 2 as

$$\mathbf{R}_j = \begin{bmatrix} \hat{\mathbf{R}}_j & \mathbf{0} \\ \mathbf{0} & \sqrt{\hat{\mathbf{\Gamma}}_j} \end{bmatrix}.$$

The solution increment $\delta\mathbf{x}_j = \mathbf{V}_j\hat{\mathbf{v}}_j$ computed by stage 3 lies in the range of a Krylov basis $\mathbf{V}_j \in \mathbb{R}_\star^{n\times k_j}$, where $k_j \leq n$ denotes the number of stage-3 iterations. This basis satisfies

$$\mathbf{V}_j^T\mathbf{A}_j\mathbf{V}_j = \mathbf{\Gamma}_j, \quad \mathbf{V}_j^T\mathbf{A}_j[\mathbf{W}_j,\ \mathbf{X}_j] = \mathbf{0},$$

$$\mathcal{W}_j \oplus \mathcal{X}_j \oplus \mathcal{V}_j = \mathcal{W}_j + \mathbf{Y}_j K^{(\hat{k}_j)}\left(\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j, \mathbf{Y}_j^T(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j - \mathbf{A}_j\mathbf{W}_j\hat{\mathbf{w}}_j)\right) + K_j^{(k_j)}(\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j)$$

with $\mathbf{\Gamma}_j \in \mathbb{R}^{k_j\times k_j}$ diagonal. The resulting solution satisfies

$$\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j + \mathbf{V}_j\hat{\mathbf{v}}_j = \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{W}_j\oplus\mathcal{X}_j\oplus\mathcal{V}_j}\left(\mathbf{x}_j^\star\right).$$

4.3.1. Accounting for entire augmenting subspace. To preserve computational-cost savings, the stage-3 implementation described above augments with the subspace range($[\mathbf{W}_j,\ \mathbf{X}_j]$) $\subseteq \mathcal{Y}_j$. However, we know that the stage-2 solution $\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j$ satisfies

$$\|\mathbf{Y}_j^T\mathbf{r}_j(\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j)\| \leq \hat{\epsilon}_j,$$

where $\hat{\epsilon}_j$ is the stage-2 tolerance. If this tolerance is small, then orthogonality condition (2.9) implies that this solution is *nearly* optimal over the entire augmenting subspace, i.e.,

$$\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j = \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{W}_j\oplus\mathcal{X}_j}\left(\mathbf{x}_j^\star\right) \approx \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right).$$

Therefore, employing $\mathbf{Y} \leftarrow \mathbf{Y}_j$ in stage 3 may reduce the number of iterations to convergence, as this ensures that new search directions remain $\mathbf{A}_j$-orthogonal to the full augmenting subspace $\mathcal{Y}_j$ over which the solution has already been computed to tolerance $\hat{\epsilon}_j$. However, for this approach to be practical, the reduced solves in steps 4 and 13 must be performed efficiently, i.e., without assembling and factorizing the matrix $\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j$. We therefore propose performing these solves *iteratively*, i.e, by executing stage 2 within stage 3. In particular, this stage-3 variant executes Algorithm 1 as

$$(k_j, \hat{\mathbf{v}}_j, \mathbf{V}_j, \mathbf{\Gamma}_j) = \texttt{augmented_pcg}(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \hat{\mathbf{W}}_j\hat{\mathbf{w}}_j + \hat{\mathbf{X}}_j\hat{\mathbf{x}}_j, \mathbf{Y}_j, \mathbf{M}_j, \epsilon_j)$$

with the following optimizations:

- At iteration k, the solves in steps 4 and 13 can be performed iteratively by executing Algorithm 1 as

  $$(\bar{k}_j^{(k)}\ \hat{\mathbf{x}}_j^{(k)}\ \bar{\mathbf{X}}_j^{(k)}\ \bar{\mathbf{\Gamma}}_j^{(k)}) = \texttt{augmented_pcg}(\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j, \mathbf{Y}_j^T(\mathbf{A}_j\mathbf{z}_j^{(k+1)}), \mathbf{0}, [\hat{\mathbf{W}}_j,\ \hat{\mathbf{X}}_j,\ \bar{\mathbf{X}}_j^{(0)},\ \cdots,\ \bar{\mathbf{X}}_j^{(k-1)}], \mathbf{I}, \bar{\epsilon}_j)$$

  with the following optimizations:
  - The solves in steps 4 and 13 can be performed directly in $\mathcal{O}(w_j^2 + \hat{k}_j + \sum_{\ell=0}^{k-1}\bar{k}_j^{(\ell)})$ flops, as the Cholesky factorization $\mathbf{Y}^T\mathbf{A}\mathbf{Y} = \bar{\mathbf{R}}_j^{(k)}\bar{\mathbf{R}}_j^{(k)}$ with

    $$\bar{\mathbf{R}}_j^{(k)} = \begin{bmatrix} \hat{\mathbf{R}}_j & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \sqrt{\hat{\mathbf{\Gamma}}_j} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \sqrt{\bar{\mathbf{\Gamma}}_j^{(0)}} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \sqrt{\bar{\mathbf{\Gamma}}_j^{(k-1)}} \end{bmatrix}$$

    can be reused from stage 1, stage 2, and earlier stage-3 iterations.

This solution increment $\delta\mathbf{x}_j = \mathbf{V}_j\hat{\mathbf{v}}_j$ also lies in the range of a Krylov basis $\mathbf{V}_j \in \mathbb{R}_\star^{n\times k_j}$; however, the basis now satisfies

$$\mathbf{V}_j^T\mathbf{A}_j\mathbf{V}_j = \mathbf{\Gamma}_j, \quad \mathbf{V}_j^T\mathbf{A}_j[\mathbf{W}_j,\ \mathbf{X}_j,\ \mathbf{Y}_j\bar{\mathbf{X}}_j^{(0)},\ \cdots,\ \mathbf{Y}_j\bar{\mathbf{X}}_j^{(k_j-1)}] = \mathbf{0},$$
$$\mathcal{Y}_j \oplus \mathcal{V}_j \approx \mathcal{Y}_j + K_j^{(k_j)}(\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j),$$

while the resulting solution satisfies

$$\bar{\mathbf{x}}_j + \mathbf{W}_j\hat{\mathbf{w}}_j + \mathbf{X}_j\hat{\mathbf{x}}_j + \mathbf{V}_j\hat{\mathbf{v}}_j \approx \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right).$$

Note that the solution does not associate with an exact orthogonal projection.

4.4. Overall algorithm. Algorithm 3 reports the proposed three-stage algorithm. Here, $\bar{y}$ denotes the maximum number of accumulated vectors to preserve before truncation. The variable $\varrho \in [0, 1]$ is a threshold for determining which new vectors should be included in the stage 1 basis. Finally, φ is a boolean variable whose value is 1 if stage 3 enforcing orthogonality with the entire augmenting subspace (not just stage 1 and 2 directions) as discussed in Section 4.3.1.

Algorithm 3 `three_stage_algorithm`

Input: $\{\mathbf{A}_j\}_{j=1}^p$, $\{\mathbf{b}_j\}_{j=1}^p$, $\{\bar{\mathbf{x}}_j\}_{j=1}^p$, forcing sequence $\{\epsilon_j\}_{j=1}^p$, storage threshold $\bar{y}$, stage-1 option ϱ, stage-3 option φ

Output:

1: $\mathbf{W}_1 \leftarrow \emptyset$, $\mathbf{X}_1 \leftarrow \emptyset$, $\bar{j} \leftarrow 0$
2: **for** $j = 1, \ldots, p$ **do**
3: **if** $\mathbf{W}_j \neq \emptyset$ **then**
4: *Stage 1*: $(\hat{\mathbf{w}}_j, \hat{\mathbf{R}}_j) =$ `direct_reduced_solve`$(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \mathbf{W}_j)$
5: **end if**
6: **if** $\mathbf{X}_j \neq \emptyset$ **then**
7: *Stage 2*: $(\hat{k}_j, \hat{\mathbf{x}}_j, \hat{\mathbf{X}}_j, \hat{\mathbf{\Gamma}}_j) =$ `augmented_pcg`$(\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j, \mathbf{Y}_j^T(\mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j), \hat{\mathbf{w}}_j, \hat{\mathbf{W}}_j, \mathbf{I}, \hat{\epsilon}_j)$. Note optimizations discussed in Section 4.2
8: **end if**
9: **if** $\varphi = 1$ **then** {orthogonalize against entire $\mathbf{Y}_j$}
10: *Stage 3*: $(k_j, \hat{\mathbf{v}}_j, \mathbf{V}_j, \mathbf{\Gamma}_j) =$ `augmented_pcg`$(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \hat{\mathbf{W}}_j\hat{\mathbf{w}}_j + \hat{\mathbf{X}}_j\hat{\mathbf{x}}_j, \mathbf{Y}_j, \mathbf{M}_j,$ and $\epsilon_j)$. Note optimizations discussed in Section 4.3.1
11: **else**
12: *Stage 3*: $(k_j, \hat{\mathbf{v}}_j, \mathbf{V}_j, \mathbf{\Gamma}_j) =$ `augmented_pcg`$(\mathbf{A}_j, \mathbf{b}_j - \mathbf{A}_j\bar{\mathbf{x}}_j, \left[\hat{\mathbf{w}}_j^T\ \hat{\mathbf{x}}_j^T\right]^T, [\mathbf{W}_j\ \mathbf{X}_j], \mathbf{M}_j, \epsilon_j)$. Note optimizations discussed in Section 4.3
13: **end if**
14: $\mathbf{W}_{j+1} \leftarrow \mathbf{W}_j$
15: $\mathcal{K}_j \leftarrow \{i \mid [\mathbf{\Gamma}_j]_{ii} / \sum_{k=1}^{k_j} [\mathbf{\Gamma}_j]_{kk} > \varrho\}$
16: **for** $k \in \mathcal{K}_j$ **do**
17: $\mathbf{W}_{j+1} \leftarrow [\mathbf{W}_{j+1},\ [\mathbf{V}_j]_k\sqrt{[\mathbf{\Gamma}_j]_{kk}}]$
18: **end for**
19: **if** $\varrho = 1$ **then** {include new vectors in stage-1 basis}
20: $\mathbf{W}_{j+1} \leftarrow [\mathbf{W}_j,\ \mathbf{V}_j\sqrt{\mathbf{\Gamma}_j}]$
21: **else**
22: $\mathbf{W}_{j+1} \leftarrow \mathbf{W}_j$
23: **end if**
24: **if** $y_{j+1} > \bar{y}$ **then** {truncate}
25: $\mathbf{Y}_{j+1} =$ `compression`$([\mathbf{Y}_j, \mathbf{V}_j])$ (note that $\mathbf{Z}_{j+1} = [\mathbf{Y}_j, \mathbf{V}_j]$)
26: Enforce $\mathbf{A}_j$-orthogonality: $\mathbf{Y}_{j+1}^T\mathbf{A}_j\mathbf{Y}_{j+1} = \mathbf{L}_j\mathbf{L}_j^T$ (Cholesky factorization), $\mathbf{Y}_{j+1} \leftarrow \mathbf{Y}_{j+1}\mathbf{L}^{-T}$
27: Determine $\mathbf{W}_{j+1}$ with range$(\mathbf{W}_{j+1}) \subseteq$ range$(\mathbf{Y}_{j+1})$
28: Update lastest-truncation index $\bar{j} \leftarrow j$
29: **else**
30: $\mathbf{Y}_{j+1} \leftarrow \left[\mathbf{Y}_j,\ \mathbf{V}_j\sqrt{\mathbf{\Gamma}_j}\right]$
31: **end if**
32: **end for**

Step 26 of Algorithm 3 ensures that $\mathbf{Y}_{\bar{j}+1}^T\mathbf{A}_{\bar{j}}\mathbf{Y}_{\bar{j}+1} = \mathbf{I}$. We now show that this step is critical for enabling fast stage-2 convergence: it leads to well-conditioned reduced matrices for slowly varying matrix sequences.

THEOREM 4.1. *In Algorithm 3, if either 1)* $\varrho = 1$ *and* $\mathbf{W}_{\bar{j}+1} = \mathbf{Y}_{\bar{j}+1}$ *in step 27, or 2)* $\varphi = 1$

and $\hat{\epsilon}_k = \bar{\epsilon}_k = 0$, $k = \bar{j}+1, \ldots, j-1$, *then*

$$\|\mathbf{Y}_j \mathbf{A}_j \mathbf{Y}_j - \mathbf{I}\| \leq \sum_{k=\bar{j}+1}^{j} \|\mathbf{Y}_k\|^2 \|\mathbf{A}_k - \mathbf{A}_{k-1}\|. \tag{4.1}$$

4.4.1. Integration with goal-oriented POD. Section 3.3.4 described four sets of goal-oriented POD ingredients (with $\omega = j - \bar{j}$) that could be employed as `compression` in step 25 of Algorithm 3. In all cases, step 27 of Algorithm 3 can be executed by setting $\mathbf{W}_{j+1}$ equal to the first w_j vectors of $\mathbf{Y}_{j+1}$, where w_j is determined from the appropriate POD algorithm (Algorithm 4 or 5) with a modest statistical energy criterion $\varepsilon \leftarrow \varepsilon_w$ with $\varepsilon_w \leq \varepsilon_y$.

The first two options $\mathbf{Y}_{j+1} = \mathbf{U}_{y_{j+1}}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\text{prev}}\right)$ and $\mathbf{Y}_{j+1} = \mathbf{U}_{y_{j+1}}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\text{RBF}}(\omega)\right)$ expose two implementation optimizations. First, during any truncation iteration $\bar{j}$, the algorithm should employ $\mathbf{W}_{\bar{j}} = \mathbf{Y}_{\bar{j}}$. The reason is that the truncation step requires computing $\mathbf{Y}_{\bar{j}}^T \mathbf{A}_{\bar{j}} \mathbf{Y}_{\bar{j}}$ explicitly; because this is the same matrix used in stage one with $\mathbf{W}_{\bar{j}} = \mathbf{Y}_{\bar{j}}$, employing this choice can lead to faster stage-3 convergence at no additional computational cost. Second, orthogonalization step 26 in Algorithm 3 requires *no operations*, as the basis is automatically $\mathbf{A}_{\bar{j}}$-orthogonal (see Eq. (3.6)).

5. Numerical experiments.

5.1. Problem description. We now assess the proposed methodology using model problems from Sierra/SolidMechanics [24], which is a Lagrangian, three-dimensional code for finite element analysis of solids and structures.

5.1.1. Problem 1: Pancake problem. We first consider computing the quasistatic response of the 'pancake' domain pictured in Figure 1. The material is steel, which is characterized by a Young's modulus of $E = 2.0 \times 10^8 \ \frac{\text{N}}{\text{mm} \cdot \text{s}^2}$, Poisson's ratio of $\nu = 0.3$, and density of $\rho = 7.86 \times 10^{-6}$ kg/mm^3. The logarithmic thermal strain of steel is linearly dependent on temperature $\epsilon_{\text{thermal}} = (11.7 \times 10^{-6})\Delta T \ \frac{1}{\text{K}}$, where ΔT denotes the change in temperature (in Kelvin) from the reference temperature.

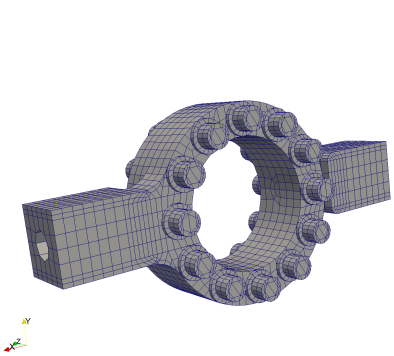

(a) Finite-element mesh

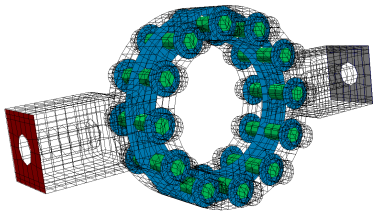

(b) Pressure-loaded surface (red), contact surfaces (blue), prescribed temperature (green), Dirichlet boundary condition (gray).

Fig. 1: Pancake problem.

The x-, y-, and z-displacements of the rightmost surface (gray in Figure 1(b)) are set to zero. The x- and y- displacements of the leftmost surface (red in Figure 1(b)) are set to zero; this surface is also subjected to the time-dependent pressure load depicted in Figure 2a.

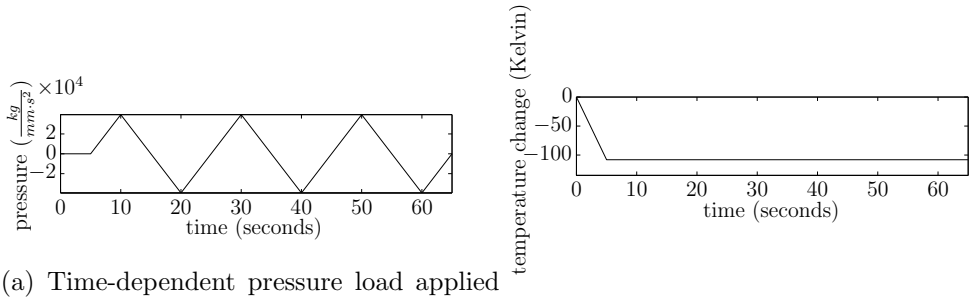

(a) Time-dependent pressure load applied to leftmost surface (extrema are $\pm 3.94 \times 10^4$ $\frac{\text{kg}}{\text{mm}\cdot\text{s}^2}$).

(b) Time-dependent temperature load applied to bolts.

Fig. 2: Time-dependent loadings for the pancake problem.

The time-dependent thermal load depicted in Figure 2b is applied to the bolts (green components in Figure 1(b)) to emulate a pre-loading condition. Contact surfaces are shown in blue in Figure 1(b); they are enforced by an augmented Lagrangian approach with DASH search using a penalty factor of 1.25 and a friction coefficient of 0.5.

The problem is discretized by the finite-element method using a mesh generated by the SIERRA toolkit [11]. The mesh consists of 9108 nodes and 4719 hexahedral elements. At each node, there are three degrees of freedom (the x-, y-, and z-displacements), which leads to $27,324$ total degrees of freedom for the finite-element model.

As the time scales of the load application are relatively large, we neglect inertial effects and consider solving the quasi-static equations

$$\boldsymbol{f}^{\text{int}}\left(\boldsymbol{u}_i\right) + \boldsymbol{f}^{\text{contact}}\left(\boldsymbol{u}_i, \lambda^l\right) = \boldsymbol{f}^{\text{ext}}(t_i), \quad l = 1, \ldots, L, \quad i = 1, \ldots, T, \tag{5.1}$$

where T denotes the number of time steps, $\boldsymbol{f}^{\text{int}} : \mathbb{R}^n \to \mathbb{R}^n$ is a nonlinear operator representing the internal forces, $\boldsymbol{f}^{\text{contact}} : \mathbb{R}^n \times \mathbb{R}^+ \to \mathbb{R}^n$ is nonlinear in its first argument and represents contact forces, $\boldsymbol{f}^{\text{ext}} : [0, 65] \to \mathbb{R}^n$ is the external-force vector, and $\boldsymbol{u}_i \in \mathbb{R}^n$ denotes the displacement at time t_i. As the contact constraints are enforced using an augmented Lagrangian approach within a continuation loop, $\lambda^l \in \mathbb{R}^+$ with $\lambda^l \leq \lambda^{l+1}$, $l \in \mathbb{N}(L)$ denotes the penalty factor at continuation iteration l. Note that these equations also include equality constraints arising from the Dirichlet boundary conditions.

Solving Eqs. (5.1) is mathematically equivalent to solving

$$\underset{\mathbf{z} \in \mathbb{R}^n}{\text{minimize}} \quad g_i^l(\mathbf{z}) \tag{5.2}$$

with $-\nabla g_i^l : \mathbf{z} \mapsto \boldsymbol{f}^{\text{ext}}(t_i) - \boldsymbol{f}^{\text{int}}\left(\mathbf{z}\right) - \boldsymbol{f}^{\text{contact}}\left(\mathbf{z}, \lambda^l\right)$ for a local minimum. Sierra/SolidMechanics solves problem (5.2) using the nonlinear conjugate gradient method with a displacement-dependent preconditioner $\bar{\mathbf{M}}^l : \bar{\boldsymbol{u}} \mapsto \nabla_{\boldsymbol{u}} \boldsymbol{f}^{\text{int}}\left(\bar{\boldsymbol{u}}\right) + \nabla_{\boldsymbol{u}} \boldsymbol{f}^{\text{contact}}\left(\bar{\boldsymbol{u}}, \lambda^l\right) \in \text{SPD}(n)$. This results in a sequence of linear systems of the original form in Eqs. (2.1): one at each nonlinear conjugate-gradient iteration. Here, $\mathbf{A}_j = \bar{\mathbf{M}}^l(\boldsymbol{u}_i^{l(k)})$, $\mathbf{b}_j = \nabla g_i^l(\boldsymbol{u}_i^{l(k)})$, $\boldsymbol{u}_i^{l(k)} \in \mathbb{R}^n$ denotes the displacement at time step

i, continuation iteration l, and nonlinear conjugate gradient iteration k. The mapping between timestep i, continuation iteration l, nonlinear conjugate-gradient iteration k and the index of the linear system j is provided by $j : (i, l, k) \mapsto k + \sum_{\mathrm{i}=1}^{i} \sum_{\mathrm{l}=1}^{l} K(\mathrm{i}, \mathrm{l})$, where $K(i, l)$ denotes the number of conjugate-gradient iterations needed to solve problem (5.2).

For this problem, the total number of linear systems we consider is $p = 47$. Each of these linear systems is preconditioned using a three-level algebraic multigrid (AMG) preconditioner $\mathbf{M}_j$ with incomplete Cholesky smoothing for both pre-smoothing and post-smoothing. This preconditioner tends to be expensive to apply, especially when compared with the cost of a matrix–vector product for this system.

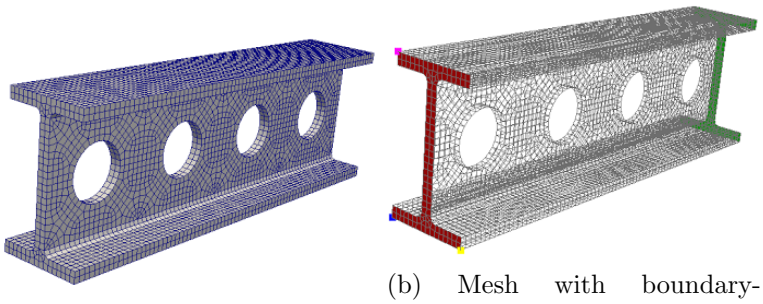

(a) Finite-element mesh

(b) Mesh with boundary-condition information.

Fig. 3: I-beam problem.

5.1.2. Problem 2: I-beam problem. We next consider computing the quasistatic response (neglecting thermal effects) of an I-beam with holes in the web section; the domain is pictured in Figure 3a. The material is steel 304L, which is characterized by a Young's modulus of $E = 2.1 \times 10^8 \ \frac{\mathrm{N}}{\mathrm{mm} \cdot \mathrm{s}^2}$, Poisson's ratio of $\nu = 0.33$, and density of $\rho = 7.8 \times 10^{-6}$ kg/mm^3. The x-, y-, and z-displacement of the bottom-left point (blue point in Figure 3b) are set to zero. The x- and y- displacement of the bottom-right point (yellow point in Figure 3b) are set to zero. The x-displacement of the top-left point (magenta point in Figure 3b) is set to zero. Finally, a torsional traction is applied the end surfaces (red and green in Figure 3); Ref. [1, Eq. (11)] reports details on the tractional loading, where the scale factor for the current problem is 0.01. The mesh consists of 13,137 nodes and 8,576 hexahedral elements. Each node is characterized by three degrees of freedom (the x-, y-, and z-displacements), which leads to $39,411$ total degrees of freedom in the finite-element model.

We again neglect inertial effects and consider solving the quasi-static equations

$$\boldsymbol{f}^{\mathrm{int}}(\boldsymbol{u}_i) = \boldsymbol{f}^{\mathrm{ext}}(t_i), \quad i = 1, \ldots, T. \tag{5.3}$$

These equations also include equality constraints arising from the Dirichlet boundary conditions; note that the continuation loop for computing contact forces does not appear in this problem. Solving Eqs. (5.3) is equivalent to solving

$$\underset{\mathbf{z} \in \mathbb{R}^n}{\text{minimize}} \quad g_i(\mathbf{z}) \tag{5.4}$$

with $-\nabla g_i : \mathbf{z} \mapsto \boldsymbol{f}^{\mathrm{ext}}(t_i) - \boldsymbol{f}^{\mathrm{int}}(\mathbf{z})$. We again use the nonlinear conjugate gradient method to solve problem (5.4); we also again employ a displacement-dependent preconditioner $\bar{\mathbf{M}} : \boldsymbol{u} \mapsto$

$\nabla_{\boldsymbol{u}} \boldsymbol{f}^{\text{int}}(\boldsymbol{u}) + \nabla_{\boldsymbol{u}} \boldsymbol{f}^{\text{contact}}\left(\boldsymbol{u}, \lambda^l\right) \in \text{SPD}(n)$. This results in a sequence of linear systems of the original form in Eqs. (2.1): one at each nonlinear conjugate-gradient iteration. Here, $\mathbf{A}_j = \bar{\mathbf{M}}(\boldsymbol{u}_i^{(k)})$ and $\mathbf{b}_j = \nabla g_i(\boldsymbol{u}_i^{(k)})$, where the mapping between time step and nonlinear conjugate-gradient iteration is provided by $j : (i,k) \mapsto k + \sum_{\mathrm{i}=1}^{i} K(\mathrm{i})$. Here, $K(i)$ denotes the number of conjugate-gradient iterations needed to solve problem (5.4). For this problem, the total number of linear systems we consider is $p = 49$. Each of these linear systems is preconditioned using the same multigrid preconditioner $\mathbf{M}_j$ as described in Section 5.1.1; the only modification is that four levels of multigrid are used in this case due to the larger number of degrees of freedom.

5.2. Experimental setup. We implemented our proposed method in Matlab and ran experiments on a Macbook Pro with Intel 2.7 GHz i5 processor and 8 GB of RAM; the implementation performs the linear-system solves after reading in the linear systems generated by Sierra/SolidMechanics as described in Section 5.1.

For all problems, we test our framework using a full-orthogonalization method (FOM) rather than the conjugate-gradient recurrence; this amounts to replacing Step 14 in Algorithm 1 with the following [37]:

$\mathbf{p}^{(k+1)} = \mathbf{z}^{(k+1)} - \mathbf{Y}\boldsymbol{\mu}^{(k+1)}$
for $i = 1, \ldots, k$ **do**
  $\beta^{(k+1),i} = \frac{(\mathbf{r}^{(k+1)})^T \mathbf{z}^{(k+1)}}{(\mathbf{r}^{(i)})^T \mathbf{z}^{(i)}}$
  $\mathbf{p}^{(k+1)} = \mathbf{p}^{(k+1)} + \beta^{(k+1),i} \mathbf{p}^{(i)}$
end for

In exact arithmetic, this modification does not affect the solution. However, in finite precision, this modification ensures that that the basis $\mathbf{V}_j$ is full rank; this is sometimes necessary to ensure nonsingular systems during stage-1 solves. Removing the effect of possible rank deficiency from the numerical experiments also simplifies interpretation of the results.

Iterative-solver performance can be measured in three primary ways: the number of incurred matrix–vector products, the number of stage-3 iterations (which is equal to the number of preconditioner applications), and the wall time per linear-system solve. While the wall time is the most important metric in practice, we report all three metrics (in terms of their averages over all linear systems) to provide a more complete picture of performance, as the specific linear system and choice of preconditioner can have a strong effect on the relative cost of operations.

5.3. Method comparison: Pancake problem. This section compares the following methods:

1. *FOM.* This approach solves each linear system independently without recycling using the full orthogonalization method.
2. *No truncation.* This approach does not perform truncation, and employs $\bar{y} = \infty$, $\varrho = 1$, and $\varphi = 0$ in Algorithm 3. Because it employs all Krylov vectors as the augmenting-subspace basis, it requires the fewest stage-3 iterations and, thus, the fewest number of preconditioner applications; however, the memory and orthogonalization costs are the largest for this method.
3. *DF(100,0).* This is the standard approach for deflation-based truncation, which places all augmenting-subspace basis vectors in the stage-1 basis. Here, Algorithm 3 parameters are $\bar{y} = 200$, $\varrho = 1$, $\varphi = 0$, and $\mathbf{W}_{j+1} = \left[\left[\mathbf{Y}_{j+1}\right]_1 \ \cdots \ \left[\mathbf{Y}_{j+1}\right]_{100}\right]$ in step 27. In step 25, the augmenting space is computed by solving Eq. (2.15) for $j \leftarrow j+1$ and setting $\mathbf{Y}_{j+1} \leftarrow [\mathbf{Z}_{j+1}\mathbf{g}_1 \ \cdots \ \mathbf{Z}_{j+1}\mathbf{g}_{100}]$.
4. *POD(100,0).* This approach employs POD truncation and places augmenting-subspace

basis vectors in the stage-1 basis. Algorithm 3 parameters are $\bar{y} = 200$, $\varrho = 1$, $\varphi = 0$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_{100}\right]$ in step 27.

5. *POD(5,95)it stg1.* This approach employs POD truncation, places only the dominant POD modes in the stage-1 basis, and places all post-truncation Krylov vectors in the stage-1 basis. Algorithm 3 parameters are $\bar{y} = 200$, $\varrho = 1$, $\varphi = 1$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_5\right]$ in step 27. Note that $\hat{\epsilon}_j = 10^{-4}\epsilon_j$ for $\epsilon_j \geq 10^{-3}$, and $\hat{\epsilon}_j = 10^{-5}\epsilon_j$ otherwise, while $\bar{\epsilon}_j = 10^{-2}\epsilon_j$ for all tolerances.
6. *POD(5,95)it mixed.* This approach employs POD truncation, places only the dominant POD modes in the stage-1 basis, and places only the dominant post-truncation Krylov vectors in the stage-1 basis. Algorithm 3 parameters are employs $\bar{y} = 200$, $\varrho = 1 \times 10^{-3}$, $\varphi = 1$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_5\right]$ in step 27. Note $\hat{\epsilon}_j = 10^{-4}\epsilon_j$ for $\epsilon_j \geq 10^{-2}$, and $\hat{\epsilon}_j = 10^{-6}\epsilon_j$ otherwise and $\bar{\epsilon}_j = 10^{-2}\epsilon_j$ for all ϵ_j.
7. *POD(5,95)it stg2.* This approach employs POD truncation, places only the dominant POD modes in the stage-1 basis, and places none of the post-truncation Krylov vectors in the stage-1 basis. Algorithm 3 parameters are employs $\bar{y} = 200$, $\varrho = 0$, $\varphi = 1$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_5\right]$ in step 27. Note $\hat{\epsilon}_j = 10^{-4}\epsilon_j$ for $\epsilon_j \geq 10^{-2}$, and $\hat{\epsilon}_j = 10^{-7}\epsilon_j$ otherwise. $\bar{\epsilon}_j = 10^{-2}\epsilon_j$ for $\epsilon_j \geq 10^{-3}$ and $\bar{\epsilon}_j = 10^{-3}\epsilon_j$ otherwise.

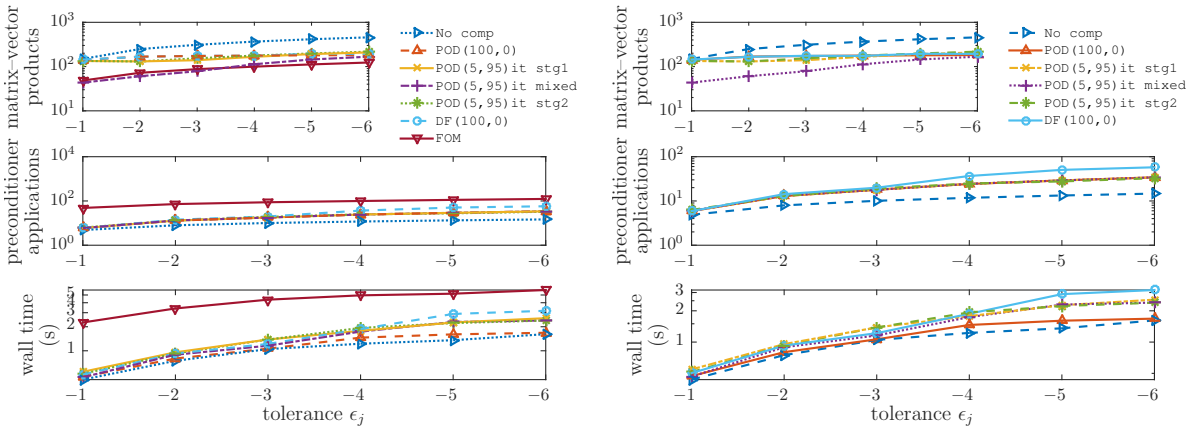

(a) All methods

(b) Recycling methods only

Fig. 4: Pancake problem: average number of matrix–vector products, preconditioner applications, and wall time to compute solutions within tolerances $\epsilon_j = 10^{-1}$ through $\epsilon_j = 10^{-6}$.

Figure 4a reports results for all tested methods. First, we note that applying the AMG preconditioner is computationally expensive for this problem, especially relative to matrix–vector products. Therefore, there is a strong relationship between the wall-time performance and the preconditions-application performance of iterative methods for this example. Next, we note that recycling provides a significant benefit, as applying FOM without recycling is the slowest method for all tested tolerances. To more clearly distinguish the differences between recycling methods, Figure 4b reports

the same results with the FOM performance removed. Here, we see that the no-truncation case yields the best performance as measured in both wall time and preconditioner applications; however, it performs the worst in matrix–vector products. This arises from two primary effects: 1) the preconditioner application is the dominant cost for this problem, so minimizing the number of stage-3 iterations—which will always occur by not truncating the augmenting subspace—yields the best wall time performance, and 2) the problem is small scale, so there is not a significant penalty to retaining all Krylov vectors. We also note that the POD methods (especially the POD(100,0) method) perform similarly to the no truncation method; this suggests that POD truncation effectively captures the most important subspace from the set of available vectors.

Next, we note that the 'inner' iterative method described in Section 4.3.1—which is used by the three POD(5,95) methods—produces the same number of preconditioner applications as POD(100,0). This illustrates that the inner iterative method has successfully orthogonalized against the entire augmenting subspace. Figure 5 illustrates this point further by comparing three methods: the POD(5,95)it stg1 method, the POD(5,95) stg1 method (which is identical to POD(5,95)it stg1 except that it employs $\varphi = 0$ in Algorithm 3), and the POD(100,0) method. While POD(5,95)it stg1 matches the preconditioning applications of POD(100,0), this comes at the cost of additional matrix–vector products for stricter tolerances; further these matrix–vector products are generally not applied as a block as in POD(100,0), which makes them more expensive on average. In addition, the cost of repeatedly multiplying vectors by the full augmenting-subspace basis $\mathbf{Y}_j$ within the inner iterative method is not counted toward matrix–vector products, even though this operation incurs a non-negligible cost due to the density of $\mathbf{Y}_j$. As a result, POD(100,0) yields the lowest wall time; we might expect POD(5,95)it stg1 to produce a lower wall time when the dimension of the augmenting subspace is larger.

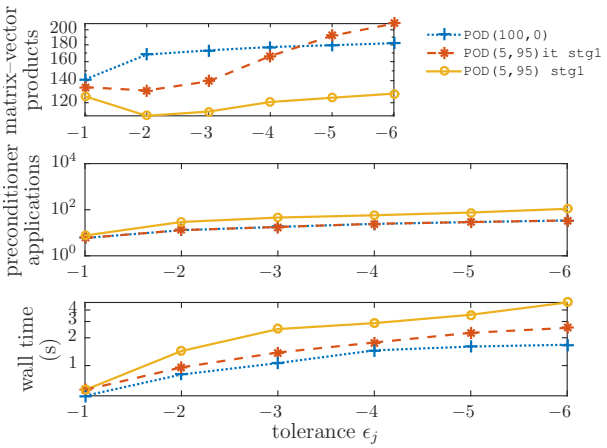

Fig. 5: Comparison of iterative and non-iterative Stage 3 for the problem 1.

As was shown in Theorem 4.1, the reduced matrix $\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j$ should be well conditioned if the system matrices do not vary significantly. Figure 6 illustrates this effect: the condition number of the reduced system is close to one for all linear systems. This implies fast convergence of stage 2. Note that this effect comes 'for free' when employing a POD metric of $\mathbf{A}_{\bar{j}}$, as the basis is automatically $\mathbf{A}_{\bar{j}}$-orthogonal; this precludes the need for orthogonalization step 26 in Algorithm 3 as was noted in Section 4.4.1.

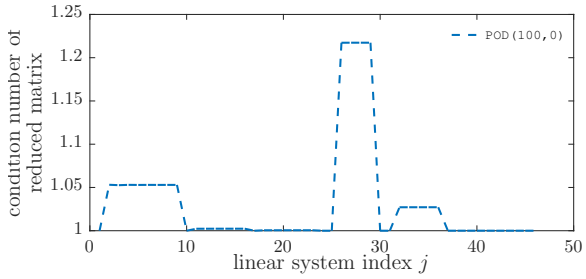

Fig. 6: Condition number for reduced linear systems.

5.4. Method comparison: I-beam problem. We compare the same methods as in Section 5.3; the only modification is that all POD(5,95) methods employ parameters $\hat{\epsilon}_j = 10^{-4}\epsilon_j$, $\bar{\epsilon}_j = 10^{-2}\epsilon_j$ in Algorithm 3.

Figure 7a reports results for all tested methods. Again, we first note that the use of recycling leads to significant improvements, as the FOM method (without recycling) produces the largest wall time and requires the largest number of preconditioner applications. To more easily distinguish the relative merits of the recycling methods, Figure 7b reports the results without FOM. This figure illustrates the need for truncation within recycling. As for the pancake problem, the 'no truncation' case minimizes the number of preconditioner applications; however, the lower matrix–vector multiplication cost and the lower overhead in the stage-3 orthogonalization steps lead to lower overall costs for several truncation methods. Figure 7b also illustrates the benefit of using the hybrid direct/iterative approach to solve over the augmenting subspace, as POD(5,95)it stg2—which employs this approach—produces the lowest wall time and number of matrix–vector products for all tested tolerances. We also note that all POD-based truncation methods outperform deflation in terms of preconditioner applications; all POD methods except for POD(5,95)it mixed also outperform deflation in terms of matrix–vector multiplications and wall time.

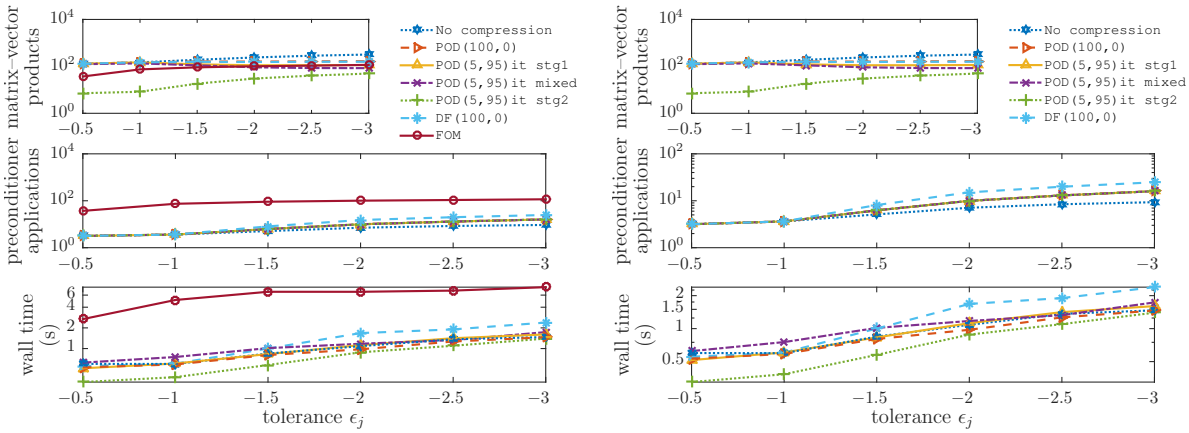

(a) All methods

(b) Recycling methods only

Fig. 7: I-beam problem: average number of matrix–vector products, preconditioner applications, and wall time to compute solutions within tolerances $\epsilon_j = 10^{-0.5}$ through $\epsilon_j = 10^{-3}$.

Finally, Figure 8 assess the performance of the 'inner' iterative method proposed in Section 4.3.1, which aims to orthogonalize against the entire augmenting subspace within stage 3 using an iterative method. These data show that the iterative method in POD(5,95)it stg1 successfully matches the number of stage-3 iterations (i.e., preconditioner applications) as POD(100,0), yet it incurs far fewer matrix–vector products. In fact, the number of matrix–vector products is close to that realized by POD(5,95) stg1 without the inner iterative method.

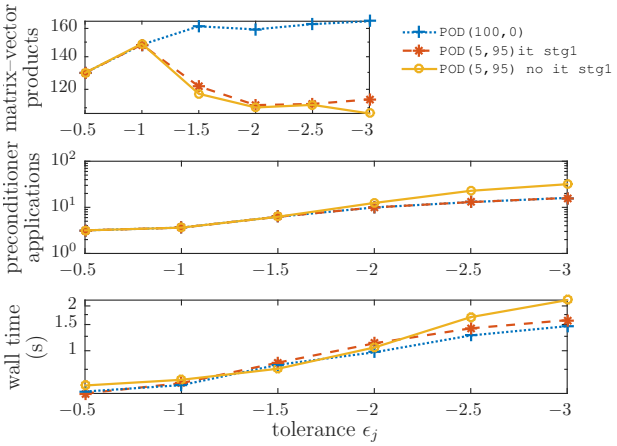

Fig. 8: I-beam problem: POD-method performance in terms of average number of matrix–vector products, preconditioner applications, and wall time to compute solutions within tolerances $\epsilon_j = 10^{-0.5}$ through $\epsilon_j = 10^{-3}$.

5.5. POD-weights experiments. This section compares the performance of POD-based truncation methods when various POD weights are employed (see Section 3.3.2). To assess this, POD truncation is performed after the solution of the first ten linear systems, and performance of the resulting truncated augmented subspace is assessed for the eleventh linear system. That is, we propose computing the truncated augmenting subspace according to

$$\mathcal{Y}_{11} = \mathcal{U}_{y_{11}}^{\mathbf{A}_{10}}(\mathbf{Z}_{10}, \boldsymbol{\gamma})$$

for various choices of $\boldsymbol{\gamma}$. We compare the following POD weights proposed in Section 3.3.2:

1. *Ideal weights.* This case employs $\boldsymbol{\gamma} = \boldsymbol{\eta}_{11}^{\star}$ as defined in Eq. (3.8). Although this weighting scheme is not practical, it illustrates the best possible choice.
2. *Previous weights.* This case employs $\boldsymbol{\gamma} = \boldsymbol{\eta}_{11}^{\text{prev}}$ as defined in Eq. (3.10).
3. *Radial-basis-function weights.* This case employs $\boldsymbol{\gamma} = \boldsymbol{\eta}_{11}^{\text{RBF}}(11)$ as defined in Eq. (3.13). This is the choice used by the previous experiments in this section.

Figure 9a illustrates that the ideal weights minimize the residual after the reduced system is solved, which implies that the ideal weights lead to a better estimate of the solution in the augmenting subspace. Note that the radial-basis-function weights (which was employed for POD results in previous sections) yields results that are close to the ideal case. Figure 9b shows that the ideal weights also minimize the number of stage-3 iterations, while the two other methods produce a similar number of stage-3 iterations as the ideal-weights case. This suggests that radial-basis-function weights provide a good approximation of the ideal weights for both producing an accurate solution in the augmenting subspace and yielding similar stage-3 convergence.

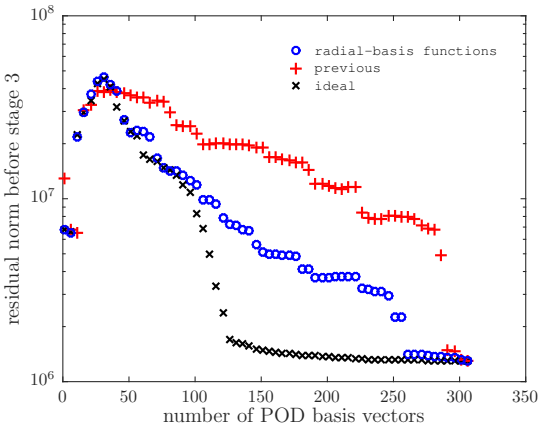

(a) The residual norm before stage 3 as a function of the number of POD vectors k.

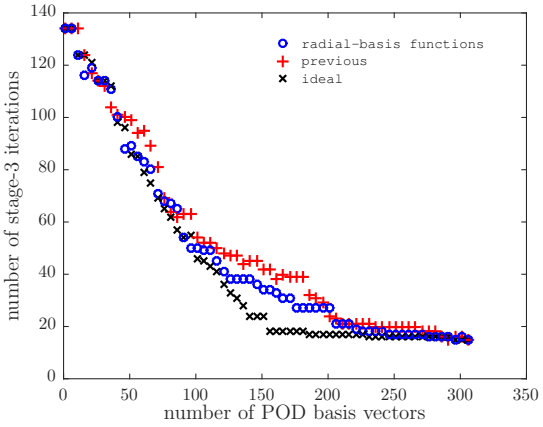

(b) Number of stage-3 iterations taken to satisfy a tolerance of $\epsilon_{11} = 10^{-6}$ as a function of the number of POD vectors k.

Fig. 9: Pancake problem: POD-weights experiments. Results correspond to solving linear system eleven after performing POD-based truncation for linear system ten using radial-basis-function weights, previous weights, and ideal weights.

5.6. Output-oriented POD experiments. This section assesses the performance of output-oriented POD, i.e., when the metric is set to $\mathbf{\Theta} = \mathbf{C}^T\mathbf{C}$ as proposed in Section 3.3.3. For this purpose, we consider a set of $q = 100$ output quantities of interest that are random linear functionals of the solution; as such $\mathbf{C} \in [0,1]^{100\times n}$ with entries drawn from a uniform distribution in the interval $[0,1]$. In addition to assessing the performance of the No truncation, DF(100,0), and POD(100,0) as described in Section 5.4, we also compare the following methods:

1. *POD(5,95).* This approach employs POD truncation, places only the dominant POD modes in the stage-1 basis, and places all post-truncation Krylov vectors in the stage-1 basis. Algorithm 3 parameters are $\bar{y} = 200$, $\varrho = 1$, $\varphi = 0$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_5\right]$ in step 27.
2. *POD(5,95)it.* This approach employs POD truncation, places only the dominant POD modes in the stage-1 basis, and places all post-truncation Krylov vectors in the stage-1 basis. Algorithm 3 parameters are $\bar{y} = 200$, $\varrho = 1$, $\varphi = 1$, $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{A}_j}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$ in step 25, and $\mathbf{W}_{j+1} = \left[[\mathbf{Y}_{j+1}]_1 \ \cdots \ [\mathbf{Y}_{j+1}]_5\right]$ in step 27. Note that $\bar{\epsilon}_j = 10^{-2}\epsilon_j$.
3. *POD-C^TC(100,0).* Identical to POD(100,0) except for $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$.
4. *POD-C^TC(5,95).* Identical to POD(5.95) except for $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$.
5. *POD-C^TC(5,95)it.* Identical to POD(5.95) except for $\mathbf{Y}_{j+1} = \mathbf{U}_{100}^{\mathbf{C}^T\mathbf{C}}\left(\mathbf{Z}_{j+1}, \boldsymbol{\eta}_j^{\mathrm{RBF}}(j-\bar{j})\right)$.

To assess the performance of output-oriented POD truncation—which aims to accurately represent the output quantity of interest—we monitor the error of the solution in the output-oriented norm $||\mathbf{x}_j^\star - \mathbf{x}_j^{(k)}||_{\mathbf{C}^T\mathbf{C}}$. For the pancake problem, we employ $\epsilon_j = 10^{-6}$, while we use $\epsilon_j = 10^{-3}$ for the I-beam problem. During the execution of the stage-3 algorithm, we track the output-oriented error and report the average number of matrix–vector products, preconditioner applications, and

wall time for the error to satisfy $||\mathbf{x}_j^\star - \mathbf{x}_j^{(k)}||_{\mathbf{C}^T\mathbf{C}} < \tau$ for a variety of output-oriented tolerances τ.

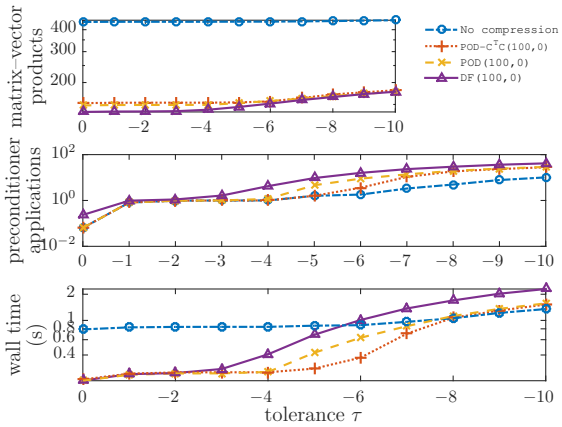

(a) Pancake problem: stage-1 methods only.

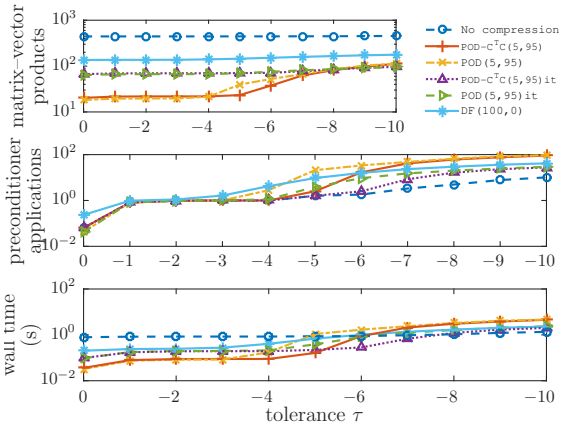

(b) Pancake problem: stage-1 and stage-2 methods.

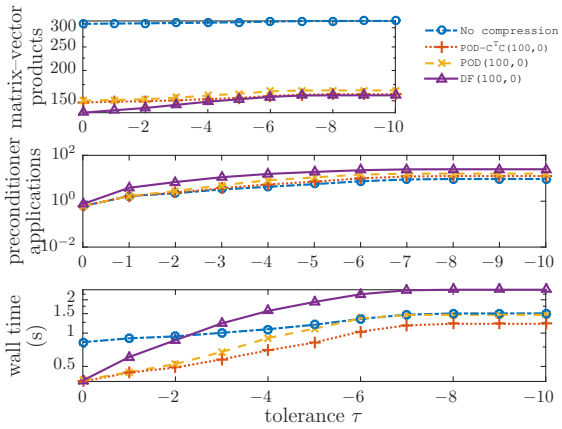

(c) I-beam problem: stage-1 methods only.

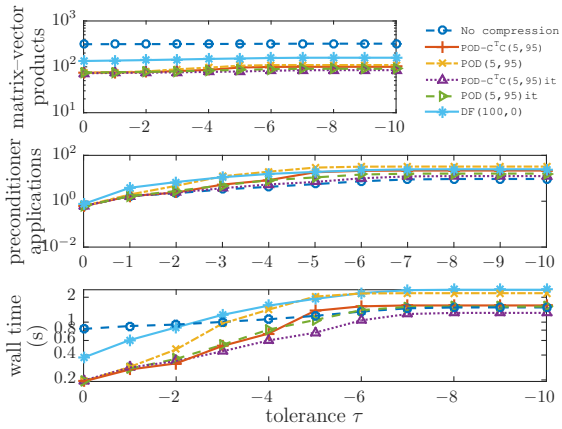

(d) I-beam problem: stage-1 and stage-2 methods.

Fig. 10: Output-oriented truncation experiments: average number of matrix–vector products, preconditioner applications, and wall time to compute solutions as a function of the output-oriented tolerance τ. Note that (average) preconditioner applications below one indicate that the solution satisfied the specified tolerance before entering stage 3.

Figures 10a–10d compare the performances of the assessed methods as a function of the output-oriented tolerance τ. First, note that Figures 10a–10b show that POD-based truncation methods produce the lowest wall times for *inexact output-oriented tolerances* for the pancake problem. These plots also illustrate the benefit of output-oriented metric over the $\mathbf{A}_j$-metric: Figure 10a shows the superior performance of the former for modest output-oriented tolerances, i.e., those between $\tau = 10^{-4}$ and 10^{-7}. Figure 10b shows that the POD(5,95) methods without the inner iterative method described in Section 4.3.1 are faster for inexact output-oriented tolerances; this occurs because very few stage-3 iterations are required for convergence in the output-oriented norm to these tolerances. In contrast, for stricter tolerances $\tau \leq 10^{-6}$, the inner iterative modification (or no truncation) yields superior performance.

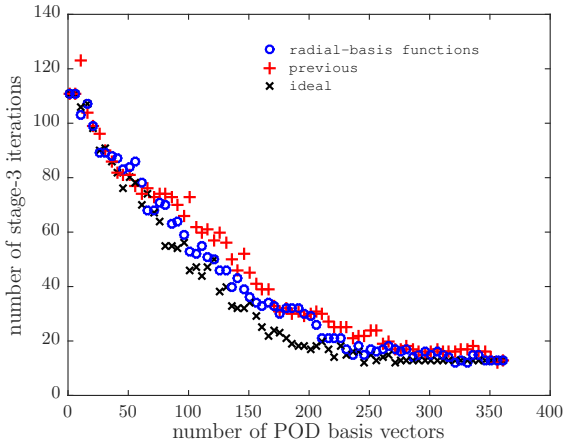

Fig. 11: Number of stage-3 iterations taken to satisfy an output-oriented tolerance of $\tau = 10^{-6}$ for the pancake problem as a function of the augmenting-subspace dimension k for radial-basis-function weights, previous weights, and ideal weights.

Figures 10c–10d present results for the I-beam problem. These results illustrate the advantage to using the output-oriented metric for most values of the output-oriented tolerance τ, as POD-C^TC(100,0) produces the lowest wall time in Figure 10c and the POD-C^TC(5,95)it method produces the lowest wall time for $\tau \leq 10^{-3}$.

Finally, we repeat the POD-weights experiments discussed in Section 5.5 for the output-oriented POD-based truncation method applied to the pancake problem. The experiments measure the number of stage-3 iterations required to converge to a output-oriented tolerance of $\tau = 10^{-6}$ as a function of the dimension of the augmenting subspace. We employ a stage-3 tolerance of $\epsilon_j = 10^{-8}$ to accrue the search directions for the first 10 linear systems and subsequently perform truncation using the output-oriented POD metric with different POD weights. The resulting stage-3 iterations are for the eleventh linear system. As we observed for the $\mathbf{A}_j$-metric, Figure 11 shows that the ideal weights yield the best performance, as they minimizing the number of stage-3 iterations required for convergence in the output-oriented norm, while the other weighting schemes closely follow the ideal weights.

6. Conclusions. This work has proposed a novel strategy for Krylov-subpace recycling inspired by goal-oriented proper orthogonal decomposition (POD). We performed analyses that expose the close connection between model reduction and Krylov-subspace recycling, proposed specific goal-oriented POD ingredients for truncating previous Krylov vectors, and developed a new 'three-stage' algorithm that employs a hybrid direct/iterative approach for efficiently solving over the sum of augmenting and Krylov subspaces. Results on several solid-mechanics problems highlighted the benefits of the new contributions, especially for efficiently computing quantities of interest to modest tolerances.

Acknowledgments. We thank Paul Tsuji for his contributions to obtaining the linear systems used in the numerical experiments, as well as Kendall Pierson for his assistance in using Adagio. We also thank Howard Elman for both insightful conversations and his support in establishing this collaboration. We are also grateful to Eric de Sturler and Michael Parks for fruitful conversations an invaluable suggestions. Sandia National Laboratories is a multi-program laboratory managed

and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. The content of this publication does not necessarily reflect the position or policy of any of these institutions, and no official endorsement should be inferred.

Appendix A. POD computation.

This section describes two techniques for computing a POD basis using given snapshots, weights, and a pseudometric (see Ref. [4] for additional details). Algorithm 4 describes the first technique, which is based on the eigenvalue decomposition and is equivalent to the well-known "method of snapshots" [41]. Algorithm 5 reports the second case, which is based on the singular value decomposition (SVD) and is more appropriate when the symmetric factorization $\mathbf{\Theta} = (\mathbf{\Theta}^{1/2})^T\mathbf{\Theta}^{1/2}$ is available, where $\mathbf{\Theta}^{1/2}$ need not be upper triangular; this approach leads to a more well-conditioned linear system. Note that Algorithms 4 and 5 produce equivalent POD bases (in exact arithmetic) and differ only in their first two steps.

Algorithm 4 `pod_evd`

Input: snapshot matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$, weights $\boldsymbol{\gamma} \in \mathbb{R}^s$, pseudometric $\mathbf{\Theta} \in \mathrm{SPSD}(n)$, and energy criterion $\varepsilon \in [0,1]$

Output: POD basis $\mathbf{U}_y^{\mathbf{\Theta}}(\mathbf{S}, (\gamma_1, \ldots, \gamma_s))$

1: $\bar{\mathbf{\Theta}} = \mathrm{diag}(\gamma_1, \ldots, \gamma_s)\,\mathbf{S}^T\mathbf{\Theta}\mathbf{S}\mathrm{diag}(\gamma_1, \ldots, \gamma_s)$

2: Compute symmetric eigenvalue decomposition $\bar{\mathbf{\Theta}} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^T$

3: Choose dimension of truncated basis $y = \min_{i\in A} i$ with $A := \{i \in \mathbb{N}(\mathrm{rank}(\bar{\mathbf{\Theta}})) \mid \sum_{k=1}^{i}\sigma_k^2 / \sum_{\ell=1}^{s}\sigma_\ell^2 \geq \varepsilon\}$

4: $\mathbf{U}_y^{\mathbf{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) = \mathbf{S}\mathrm{diag}(\gamma_1, \ldots, \gamma_s)\left[\frac{1}{\sigma_1}\mathbf{v}_1 \;\ldots\; \frac{1}{\sigma_y}\mathbf{v}_y\right]$, where $\mathbf{\Sigma} := \mathrm{diag}(\sigma_1, \ldots, \sigma_s)$

Algorithm 5 `pod_svd`

Input: snapshot matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$, weights $\boldsymbol{\gamma} \in \mathbb{R}^s$, pseudometric factor $\mathbf{\Theta}^{1/2}$ such that $\mathbf{\Theta} = (\mathbf{\Theta}^{1/2})^T\mathbf{\Theta}^{1/2} \in \mathrm{SPSD}(n)$, and energy criterion $\varepsilon \in [0,1]$

Output: POD basis $\mathbf{U}_y^{\mathbf{\Theta}}(\mathbf{S}, (\gamma_1, \ldots, \gamma_s))$

1: $\bar{\mathbf{S}} = \mathbf{\Theta}^{1/2}\mathbf{S}\mathrm{diag}(\gamma_1, \ldots, \gamma_s)$

2: Compute thin singular value decomposition $\bar{\mathbf{S}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$

3: Choose dimension of truncated basis $y = \min_{i\in A} i$ with $A := \{i \in \mathbb{N}(\mathrm{rank}(\bar{\mathbf{S}})) \mid \sum_{k=1}^{i}\sigma_k^2 / \sum_{\ell=1}^{s}\sigma_\ell^2 \geq \varepsilon\}$

4: $\mathbf{U}_y^{\mathbf{\Theta}}(\mathbf{S}, \boldsymbol{\gamma}) = \mathbf{S}\mathrm{diag}(\gamma_1, \ldots, \gamma_s)\left[\frac{1}{\sigma_1}\mathbf{v}_1 \;\ldots\; \frac{1}{\sigma_y}\mathbf{v}_y\right]$, where $\mathbf{\Sigma} := \mathrm{diag}(\sigma_1, \ldots, \sigma_s)$

Appendix B. Proofs.

This section provides proofs for the theoretical results presented in this paper.

Proof of Theorem 3.1 We begin by decomposing the centered exact solution as $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j = \mathbf{x}_j^{\|} + \mathbf{x}_j^{\perp}$, where $\mathbf{x}_j^{\|} = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right) = \mathbf{Z}_j\boldsymbol{\eta}_j^\star$ (from Eqs. (2.1) and (2.6)) and $\mathbf{x}_j^{\perp} = \left(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j}\right)\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)$.

We can then bound the error as

$$\begin{aligned}\left\|\mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right\|_{\mathbf{A}_j} &= \left\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)\right\|_{\mathbf{A}_j} = \left\|\mathbf{x}_j^{\|} + \mathbf{x}_j^{\perp} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\left(\mathbf{x}_j^{\|}\right)\right\|_{\mathbf{A}_j} \\ &= \left\|\left(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\right)\left(\mathbf{Z}_j\boldsymbol{\eta}_j^\star\right)\right\|_{\mathbf{A}_j} + \left\|\mathbf{x}_j^{\perp}\right\|_{\mathbf{A}_j} \\ &\leq z_j^{1/2}\sqrt{\sum_{i=1}^{z_j}\left\|\left(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\right)\left([\boldsymbol{\eta}_j^\star]_i\mathbf{z}_i\right)\right\|_{\mathbf{A}_j}^2} + \underbrace{\left\|\mathbf{x}_j^{\perp}\right\|_{\mathbf{A}_j}}_{\text{(I)}}, \end{aligned} \tag{B.1}$$

where we have used $\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\left(\mathbf{x}_j^{\perp}\right) = 0$ (because $\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j}\left(\mathbf{x}_j^{\perp}\right) = 0$ and $\mathcal{Y}_j \subseteq \mathcal{Z}_j$), the triangle inquality, and the norm-equivalence relation $\|\mathbf{x}\|_1 \leq n^{1/2}\|\mathbf{x}\|_2$. By comparing Eqs. (3.3) and (B.1), using monotinicity of the square root function, and noting that term (I) is independent of the subspace $\mathcal{Y}_j$, it is clear that the POD subspace minimizes (over all y-dimensional subspaces of $\mathcal{Z}_j$) an upper bound for $\left\|\mathbf{x}_j^\star - \mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right\|_{\mathbf{A}_j}$ if $\boldsymbol{\Theta} = \mathbf{A}_j$, $\mathbf{s}_i = \mathbf{z}_i$ and $\gamma_i = [\boldsymbol{\eta}_j^\star]_i$ for $i = 1, \ldots, s$ with $s = z_j$. □

Proof of Theorem 3.2 We again decompose the centered exact solution $\mathbf{x}_j^\star - \bar{\mathbf{x}}_j$ according to

$$\mathbf{x}_j^\star - \bar{\mathbf{x}}_j = \mathbf{x}_j^{\mathbf{C}} + \mathbf{x}_j^{\mathbf{C}^\perp} = \underline{\mathbf{x}_j}^{\|} + \underline{\mathbf{x}_j}^{\perp} + \mathbf{x}_j^{\mathbf{C}^\perp}.$$

Here, $\mathbf{x}_j^{\mathbf{C}} = \mathbf{P}_{\mathbf{I}}^{\text{range}(\mathbf{C}^T)}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)$, $\mathbf{x}_j^{\mathbf{C}^\perp} = \left(\mathbf{I} - \mathbf{P}_{\mathbf{I}}^{\text{range}(\mathbf{C}^T)}\right)\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)$, $\underline{\mathbf{x}_j}^{\|} = \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Z}_j}\left(\mathbf{x}_j^{\mathbf{C}}\right) = \mathbf{Z}_j\bar{\boldsymbol{\eta}}_j^\star$ (from Eq. (2.7)), and $\underline{\mathbf{x}_j}^{\perp} = \left(\mathbf{I} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Z}_j}\right)\left(\mathbf{x}_j^{\mathbf{C}}\right)$. We can then bound the error as

$$\begin{aligned}\|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right)\|_2 &= \|\mathbf{C}(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - \mathbf{C}(\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right))\|_2 = \left\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)\right\|_{\mathbf{C}^T\mathbf{C}} \\ &\leq \left\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)\right\|_{\mathbf{C}^T\mathbf{C}} + \left\|(\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j})(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j)\right\|_{\mathbf{C}^T\mathbf{C}} \\ &\leq \left\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right)\right\|_{\mathbf{C}^T\mathbf{C}} + \left\|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\right\|_{\mathbf{C}^T\mathbf{C}}\left\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right\|_{\mathbf{C}^T\mathbf{C}} \\ &= \left\|\underline{\mathbf{x}_j}^{\|} + \underline{\mathbf{x}_j}^{\perp} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\left(\underline{\mathbf{x}_j}^{\|}\right)\right\|_{\mathbf{C}^T\mathbf{C}} + \left\|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\right\|_{\mathbf{C}^T\mathbf{C}}\left\|\underline{\mathbf{x}_j}^{\|} + \underline{\mathbf{x}_j}^{\perp}\right\|_{\mathbf{C}^T\mathbf{C}} \\ &= \left\|\left(\mathbf{I} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\right)\left(\underline{\mathbf{x}_j}^{\|}\right)\right\|_{\mathbf{C}^T\mathbf{C}} + \left\|\underline{\mathbf{x}_j}^{\perp}\right\|_{\mathbf{C}^T\mathbf{C}} + \left\|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\right\|_{\mathbf{C}^T\mathbf{C}}\left\|\underline{\mathbf{x}_j}^{\|} + \underline{\mathbf{x}_j}^{\perp}\right\|_{\mathbf{C}^T\mathbf{C}} \\ &\leq z_j^{1/2}\sqrt{\sum_{i=1}^{z_j}\left\|\left(\mathbf{I} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\right)\left([\bar{\boldsymbol{\eta}}_j^\star]_i\mathbf{z}_i\right)\right\|_{\mathbf{C}^T\mathbf{C}}^2} + \underbrace{\left\|\underline{\mathbf{x}_j}^{\perp}\right\|_{\mathbf{C}^T\mathbf{C}} + \alpha\left\|\mathbf{Z}_j\bar{\boldsymbol{\eta}}_j^\star\right\|_{\mathbf{C}^T\mathbf{C}}}_{\text{(I)}} \end{aligned} \tag{B.2}$$

Here, we have used $\mathbf{C}\mathbf{x}_j^{\mathbf{C}^\perp} = 0$ and $\mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}_j}\left(\underline{\mathbf{x}_j}^{\perp}\right) = 0$. Also, $\|\mathbf{A}\|_{\mathbf{C}^T\mathbf{C}} = \sup_{\mathbf{x}\neq 0}\|\mathbf{A}\mathbf{x}\|_{\mathbf{C}^T\mathbf{C}} / \|\mathbf{x}\|_{\mathbf{C}^T\mathbf{C}}$ with $\mathbf{A} \in \mathbb{R}^{n\times n}$ is an induced matrix norm and $\alpha := \sup_{\mathcal{Y}\in\mathcal{G}(y,n)}\left\|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}} - \mathbf{P}_{\mathbf{C}^T\mathbf{C}}^{\mathcal{Y}}\right\|_{\mathbf{C}^T\mathbf{C}}$. Within inequality (B.2), term (I) is independent of the subspace $\mathcal{Y}$; as such, a comparison of Eqs. (3.3) and (B.2) reveals that the POD subspace minimizes (over all subspaces $\mathcal{Y}$) an upper bound for $\|\mathbf{q}\left(\mathbf{x}_j^\star\right) - \mathbf{q}\left(\mathbf{P}_{\mathbf{A}_j}^{\bar{\mathbf{x}}_j+\mathcal{Y}_j}\left(\mathbf{x}_j^\star\right)\right)\|_2$ if $\boldsymbol{\Theta} = \mathbf{C}^T\mathbf{C}$, $\mathbf{s}_i = \mathbf{z}_i$ and $\gamma_i = [\bar{\boldsymbol{\eta}}_j^\star]_i$ for $i = 1, \ldots, s$ with $s = z_j$ under the stated assumptions. □

Proof of Theorem 3.3 We prove the result in Eq. (3.11); Eq. (3.12) follows trivially. First note that $\mathbf{Z}_j\left(\boldsymbol{\eta}_j^\star - \boldsymbol{\eta}_j^{\text{prev}}\right) = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j}(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - \mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}(\mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1})$ from Eqs. (2.6), (3.8), and (3.10). Then, we have

$$\begin{aligned}\|\mathbf{Z}_j\left(\boldsymbol{\eta}_j^\star - \boldsymbol{\eta}_j^{\text{prev}}\right)\| \leq& \|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j}(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - \mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j)\| + \|\mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}\left(\left(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\right) - \left(\mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1}\right)\right)\| \\ \leq& \|\mathbf{P}_{\mathbf{A}_j}^{\mathcal{Z}_j} - \mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}\|\|\mathbf{x}_j^\star - \bar{\mathbf{x}}_j\| + \sigma_1\|(\mathbf{x}_j^\star - \bar{\mathbf{x}}_j) - (\mathbf{x}_{j-1}^\star - \bar{\mathbf{x}}_{j-1})\|\end{aligned}$$

where we have applied the triangle inequality and used $\sigma_1 = \max_{\mathbf{x}\neq 0}\|\mathbf{P}_{\mathbf{A}_{j-1}}^{\mathcal{Z}_j}\mathbf{x}\|/\|\mathbf{x}\|$. Eq. (3.11) follows from applying $\sigma_{\min} = \min_{\mathbf{x}\neq 0}\|\mathbf{Z}_j\mathbf{x}\|/\|\mathbf{x}\|$. □

Proof of Theorem 3.5 We proceed by leveraging perturbation bounds for invariant subspaces. From the algebraic definition of a POD basis provided by Eq. (3.4), we can write a basis for $\mathcal{U}^a$ as

$$\mathbf{U}^a := \mathbf{U}_y^{\boldsymbol{\Theta}^a}(\mathbf{S}, \boldsymbol{\gamma}^a) = \mathbf{S}\boldsymbol{\Gamma}^a\mathbf{V}^a[\boldsymbol{\Lambda}^a]^{-\frac{1}{2}}.$$

The matrices $\mathbf{V}^a \in \mathcal{S}_y(\mathbb{R}^s)$ and $\boldsymbol{\Lambda}^a \in \mathbb{R}^{y\times y}$ are obtained by truncation, i.e., by retaining the first y columns of eigenvector matrix $\bar{\mathbf{V}}^a \in \mathcal{S}_s(\mathbb{R}^s)$ and the first y rows and columns of the eigenvalue matrix $\bar{\boldsymbol{\Lambda}}^a \in \mathbb{R}^{s\times s}$, which are associated with the eigenvalue decomposition

$$\boldsymbol{\Gamma}^a\mathbf{S}^T\boldsymbol{\Theta}^a\mathbf{S}\boldsymbol{\Gamma}^a = \bar{\mathbf{V}}^a\bar{\boldsymbol{\Lambda}}^a[\bar{\mathbf{V}}^a]^T. \tag{B.3}$$

Similarly, the second POD subspace $\mathcal{U}^b$ is spanned by a basis

$$\mathbf{U}^b := \mathbf{U}_y^{\boldsymbol{\Theta}^b}\left(\mathbf{S}, \boldsymbol{\gamma}^b\right) = \mathbf{S}\boldsymbol{\Gamma}^b\mathbf{V}^b[\boldsymbol{\Lambda}^b]^{-\frac{1}{2}}$$

where again $\mathbf{V}^b$ and $\boldsymbol{\Lambda}^b \in \mathbb{R}^{y\times y}$ are the truncated counterparts of the eigenvectors $\bar{\mathbf{V}}^b \in \mathcal{S}_s(\mathbb{R}^s)$ and eigenvalues $\bar{\boldsymbol{\Lambda}}^b \in \mathbb{R}^{s\times s}$ associated with eigenvalue problem

$$\boldsymbol{\Gamma}^b\mathbf{S}^T\boldsymbol{\Theta}^b\mathbf{S}\boldsymbol{\Gamma}^b = \bar{\mathbf{V}}^b\bar{\boldsymbol{\Lambda}}^b[\bar{\mathbf{V}}^b]^T. \tag{B.4}$$

Noting that a diagonal column scaling of a basis does not affect the associated subspace, we can write

$$\mathcal{U}^a = \text{range}(\mathbf{S}\boldsymbol{\Gamma}^a\mathbf{V}^a) \quad \text{and} \quad \mathcal{U}^b = \text{range}(\mathbf{S}\boldsymbol{\Gamma}^b\mathbf{V}^b) = \text{range}(\mathbf{S}\boldsymbol{\Gamma}^a\mathbf{D}\mathbf{V}^b),$$

with $\mathbf{D} := [\boldsymbol{\Gamma}^a]^{-1}\boldsymbol{\Gamma}^b$. Thus, we can compute the distance between $\mathcal{U}^a$ and $\mathcal{U}^b$ in two steps: 1) compute the distance between subspaces spanned by bases $\mathbf{V}^a$ and $\mathbf{D}\mathbf{V}^b$, and 2) apply a change in coordinates by pre-multiplying each basis by $\mathbf{S}\boldsymbol{\Gamma}^a$.

Step 1. We employ eigenvector perturbation theory [8, 9, 21, 13, 20] to bound the distance between range($\mathbf{V}^a$) and range($\mathbf{D}\mathbf{V}^b$). First, we recall a general eigenvector perturbation bound that characterizes the distance between y-dimensional invariant subspaces $\mathcal{S}^a \in \mathcal{G}(y, s)$ and $\mathcal{S}^b \in \mathcal{G}(y, s)$ associated with matrices $\mathbf{B}^a \in \mathbb{R}^{s\times s}$ and $\mathbf{B}^b \in \mathbb{R}^{s\times s}$, respectively.

From Ref. [20], we have the following bound when both $\mathbf{B}^b$ and $\mathbf{B}^a$ are diagonalizable:

$$d\left(\mathcal{S}^a, \mathcal{S}^b\right) \leq \kappa^a\kappa^b\|\mathbf{E}\|_2/.\text{abssep}(\boldsymbol{\Lambda}_{B,\perp}^a, \boldsymbol{\Lambda}_B^b) \tag{B.5}$$

Here, the difference matrix is denoted by $\mathbf{E} := \mathbf{B}^b - \mathbf{B}^a$, κ^b denote the condition number of the matrix whose columns comprise the first y eigenvectors of $\mathbf{B}^b$, and κ^a denote condition numbers of the matrix whose columns comprise the last $n - y$ eigenvectors of $\mathbf{B}^a$; these condition numbers evaluate to one when the associated matrix is normal. Also, $\boldsymbol{\Lambda}_B^b$ denotes the diagonal matrix

of eigenvalues associated with the invariant subspace spanned by the first y eigenvectors of the matrix $\mathbf{B}^b$, while $\mathbf{\Lambda}^a_{B,\perp}$ denotes the diagonal matrix of eigenvalues associated with the orthogonal complement to the invariant subspace spanned by the first y eigenvectors of $\mathbf{B}^a$.

To apply bound (B.5) in the present context, we take

$$\mathbf{B}^a = \mathbf{\Gamma}^a\mathbf{S}^T\mathbf{\Theta}^a\mathbf{S}\mathbf{\Gamma}^a = \mathbf{\Gamma}^a\mathbf{S}^T(\mathbf{\Theta}^b + \mathbf{\Delta})\mathbf{S}\mathbf{\Gamma}^a = \bar{\mathbf{V}}^a\bar{\mathbf{\Lambda}}^a[\bar{\mathbf{V}}^a]^T$$
$$\mathbf{B}^b = \mathbf{D}\mathbf{\Gamma}^b\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^b\mathbf{D}^{-1},$$

where the latter satisfies

$$\mathbf{B}^b\mathbf{D}\bar{\mathbf{V}}^b = \mathbf{D}\bar{\mathbf{V}}^b\bar{\mathbf{\Lambda}}^b.$$

Here, we have used Eqs. (B.3), (B.4), and $[\mathbf{D}\bar{\mathbf{V}}^b]^{-1} = [\bar{\mathbf{V}}^b]^T\mathbf{D}^{-1}$. Thus, we have $\mathbf{\Lambda}^b_B = \mathbf{\Lambda}^b$ because $\mathbf{B}^b$ is similar to $\mathbf{\Gamma}^b\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^b$; we also can write $\mathbf{\Lambda}^a_{B,\perp} = \mathbf{\Lambda}^a_\perp$. We also have $\mathcal{S}^a = \text{range}(\mathbf{V}^a)$ and $\mathcal{S}^b = \text{range}(\mathbf{D}\mathbf{V}^b)$. Further, note that $\kappa^a = 1$ because $\mathbf{B}^a$ is normal, while $\kappa^b = \kappa(\mathbf{D}\mathbf{V}^b)$ due to non-normality of $\mathbf{B}^b$.

To complete bound (B.5), we bound $\|\mathbf{E}\|_2$ using commutativity of diagonal matrices and the relations $[\mathbf{\Gamma}^b]^2[\mathbf{\Gamma}^a]^{-1} = \mathbf{D}\mathbf{\Gamma}^b$ and $\mathbf{\Gamma}^a = \mathbf{\Gamma}^b\mathbf{D}^{-1}$ as

$$\begin{aligned}
\|\mathbf{E}\|_2 &= \|\mathbf{B}^b - \mathbf{B}^a\|_2 \\
&= \|\mathbf{D}\mathbf{\Gamma}^b\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^b\mathbf{D}^{-1} \quad - \mathbf{\Gamma}^a\mathbf{S}^T\left(\mathbf{\Theta}^b + \mathbf{\Delta}\right)\mathbf{S}\mathbf{\Gamma}^a\|_2 \\
&= \|\mathbf{D}\mathbf{\Gamma}^b\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a \quad - \mathbf{\Gamma}^a\mathbf{S}^T\left(\mathbf{\Theta}^b + \mathbf{\Delta}\right)\mathbf{S}\mathbf{\Gamma}^a\|_2 \\
&= \|\left(\mathbf{D}\mathbf{\Gamma}^b - \mathbf{\Gamma}^a\right)\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a \quad - \mathbf{\Gamma}^a\mathbf{S}^T\mathbf{\Delta}\mathbf{S}\mathbf{\Gamma}^a\|_2 \\
&= \|\left([\mathbf{\Gamma}^b]^2[\mathbf{\Gamma}^a]^{-1} - \mathbf{\Gamma}^a\right)\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a \quad - \mathbf{\Gamma}^a\mathbf{S}^T\mathbf{\Delta}\mathbf{S}\mathbf{\Gamma}^a\|_2 \\
&= \|(\mathbf{\Gamma}^b + \mathbf{\Gamma}^a)(\mathbf{\Gamma}^b - \mathbf{\Gamma}^a)[\mathbf{\Gamma}^a]^{-1}\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a - \mathbf{\Gamma}^a\mathbf{S}^T\mathbf{\Delta}\mathbf{S}\mathbf{\Gamma}^a\|_2 \\
&\leq \|(\mathbf{\Gamma}^b + \mathbf{\Gamma}^a)(\mathbf{\Gamma}^b - \mathbf{\Gamma}^a)[\mathbf{\Gamma}^a]^{-1}\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a\|_2 + \|\mathbf{\Delta}\|_2\|\mathbf{S}\mathbf{\Gamma}^a\|_2^2. \qquad \text{(B.6)}
\end{aligned}$$

Substituting inequality (B.6) and the values of the condition numbers in bound (B.5) leads to

$$\begin{aligned}
d\left(\mathcal{S}^a, \mathcal{S}^b\right) \leq \kappa(\mathbf{D}\mathbf{V}^b)\Big(&\|\left(\mathbf{\Gamma}^b + \mathbf{\Gamma}^a\right)\left(\mathbf{\Gamma}^b - \mathbf{\Gamma}^a\right)[\mathbf{\Gamma}^a]^{-1}\mathbf{S}^T\mathbf{\Theta}^b\mathbf{S}\mathbf{\Gamma}^a\|_2 + \\
&\|\mathbf{\Delta}\|_2\|\mathbf{S}\mathbf{\Gamma}^a\|_2^2\Big)/\text{abssep}(\mathbf{\Lambda}^a_\perp, \mathbf{\Lambda}^b) \qquad \text{(B.7)}
\end{aligned}$$

Step 2. We next use the result in Reference [44, Proposition 2.6.14] that allows for a change of coordinates using a nonsingular square transformation matrix. In particular, if $[\mathbf{S}\mathbf{\Gamma}^a]$ has full column rank, then $[\mathbf{S}\mathbf{\Gamma}^a,\ \mathbf{Z}_\perp]$ with $\mathbf{Z}_\perp \in \mathcal{S}_{n-s}(\mathbb{R}^n)$ and $\text{range}(\mathbf{Z}_\perp) = \text{range}(\mathbf{S}\mathbf{\Gamma}^a)^\perp$ is nonsingular. Then, we have

$$d\left(\mathcal{U}^a, \mathcal{U}^b\right) \leq \kappa(\mathbf{S}\mathbf{\Gamma}^a)d\left(\mathcal{S}^a, \mathcal{S}^b\right). \qquad \text{(B.8)}$$

Here, we have used $\kappa([\mathbf{S}\mathbf{\Gamma}^a,\ \mathbf{Z}_\perp]) = \kappa([\mathbf{S}\mathbf{\Gamma}^a])$. Combining bounds (B.7) and (B.8) yields the desired result. □

Proof of Theorem 4.1

$$
\begin{aligned}
\|\mathbf{Y}_j^T\mathbf{A}_j\mathbf{Y}_j - \mathbf{I}\| &= \|\mathbf{Y}_j^T\mathbf{A}_{j-1}\mathbf{Y}_j - \mathbf{I} + \mathbf{Y}_j^T(\mathbf{A}_j - \mathbf{A}_{j-1})\mathbf{Y}_j\| \\
&\leq \| \left[\mathbf{Y}_{j-1},\ \mathbf{V}_{j-1}\sqrt{\mathbf{\Gamma}_{j-1}}\right]^T \mathbf{A}_{j-1} \left[\mathbf{Y}_{j-1},\ \mathbf{V}_{j-1}\sqrt{\mathbf{\Gamma}_{j-1}}\right] - \mathbf{I}\| + \|\mathbf{Y}_j\|^2\|\mathbf{A}_j - \mathbf{A}_{j-1}\| \\
&= \| \begin{bmatrix} \mathbf{Y}_{j-1}\mathbf{A}_{j-1}\mathbf{Y}_{j-1} - \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \| + \|\mathbf{Y}_j\|^2\|\mathbf{A}_j - \mathbf{A}_{j-1}\|, \\
&= \|\mathbf{Y}_{j-1}\mathbf{A}_{j-1}\mathbf{Y}_{j-1} - \mathbf{I}\| + \|\mathbf{Y}_j\|^2\|\mathbf{A}_j - \mathbf{A}_{j-1}\|,
\end{aligned}
$$

where we have used $\mathbf{Y}_k^T\mathbf{A}_k\mathbf{V}_k = \mathbf{0}$, $k = 1, \ldots, j-1$ (which holds under the stated assumptions) and $\sqrt{\mathbf{\Gamma}_k}^T\mathbf{V}_k^T\mathbf{A}_k\mathbf{V}_k\sqrt{\mathbf{\Gamma}_k} = \mathbf{I}$, $k = 1, \ldots, j-1$. By induction, we arrive at inequality (4.1), where we have used $\mathbf{Y}_{\bar{j}+1}\mathbf{A}_{\bar{j}}\mathbf{Y}_{\bar{j}+1} = \mathbf{I}$. □

REFERENCES

[1] J. Bishop, J. Emery, R. Field, C. Weinberger, and D. Littlewood, *Direct numerical simulations in solid mechanics for understanding the macroscale effects of microscale material variability*, Computer Methods in Applied Mechanics and Engineering, 287 (2015), pp. 262–289.

[2] K. Carlberg and C. Farhat, *A compact proper orthogonal decomposition basis for optimization-oriented reduced-order models*, AIAA Paper 2008-5964, 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Victoria, Canada, (2008).

[3] ———, *An adaptive POD-Krylov reduced-order model for structural optimization*, 8th World Congress on Structural and Multidisciplinary Optimization, Lisbon, Portugal, (2009).

[4] ———, *A low-cost, goal-oriented 'compact proper orthogonal decomposition' basis for model reduction of static systems*, International Journal for Numerical Methods in Engineering, 86 (2011), pp. 381–402.

[5] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem, *The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows*, Journal of Computational Physics, 242 (2013), pp. 623–647.

[6] K. Carlberg, R. Tuminaro, and P. Boggs, *Preserving Lagrangian structure in nonlinear model reduction with application to structural dynamics*, SIAM J. Sci. Comput., 37 (2015), pp. B153—B184.

[7] A. Chapman and Y. Saad, *Deflated and augmented Krylov subspace techniques*, Numerical Linear Algebra with Applications, 4 (1997), pp. 43–66.

[8] C. Davis and W. Kahan, *Some new bounds on perturbation of subspaces*, Bull. Amer. Math. Soc., 75 (1969), pp. 863–868.

[9] ———, *The rotation of eigenvectors by a perturbation. III*, SIAM Journal on Numerical Analysis, 7 (1970), pp. 1–46.

[10] E. De Sturler, *Truncation strategies for optimal Krylov subspace methods*, SIAM Journal of Numerical Analysis, 36 (1999), pp. 864–889.

[11] H. Edwards, A. Williams, G. Sjaardema, D. Baur, and W. Cochran, *SIERRA toolkit computational mesh conceptual model*, Sandia National Laboratories SAND Series, SAND, 1192 (2010), p. 2010.

[12] S.C. Eisenstat and H.F. Walker, *Choosing the forcing terms in an inexact Newton method*, SIAM J. Sci. Comput, (1996).

[13] Stanley C. Eisenstat and Ilse C. F. Ipsen, *Relative perturbation techniques for singular value problems*, SIAM Journal on Numerical Analysis, 32 (1995), pp. 1972–1988.

[14] J. Erhel and F. Guyomarc'h, *An augmented conjugate gradient method for solving consecutive symmetric positive definite linear systems*, SIAM Journal on Matrix Analysis and Applications, 21 (2000), pp. 1279–1299.

[15] C. Farhat, *Optimizing substructuring methods for repeated right hand sides, scalable parallel coarse solvers, and global/local analysis*, Domain-Based Parallelism and Problem Decomposition Methods in Computational Science and Engineering, (1995), pp. 141–160.

[16] C. Farhat, L. Crivelli, and F.X. Roux, *Extending substructure based iterative solvers to multiple load and repeated analyses*, Computer Methods in Applied Mechanics and Engineering, 117 (1994), pp. 195–209.

[17] C. Farhat, K. Pierson, and M. Lesoinne, *The second generation FETI methods and their application to the parallel solution of large-scale linear and geometrically non-linear structural analysis problems*, Computer Methods in Applied Mechanics and Engineering, 184 (2000), pp. 333–374.
[18] P. Holmes, J. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, 1996.
[19] Ilse CF Ipsen, *Absolute and relative perturbation bounds for invariant subspaces of matrices*, Linear Algebra and its Applications, 309 (2000), pp. 45–56.
[20] Ilse C.F. Ipsen, *An overview of relative sin Θ theorems for invariant subspaces of complex matrices*, Journal of Computational and Applied Mathematics, 123 (2000), pp. 131–153.
[21] ———, *A note on unifying absolute and relative perturbation bounds*, Linear Algebra and its Applications, 358 (2003), pp. 239 – 253.
[22] S. Lall, P. Krysl, and J.E. Marsden, *Structure-preserving model reduction for mechanical systems*, Physica D: Nonlinear Phenomena, 184 (2003), pp. 304–318.
[23] B.M. Lewis, K. Ito, and H.T. Tran, *Galerkin proper orthogonal decomposition error estimates for a class of structural models*, Journal of Computational and Applied Mathematics, (2012).
[24] J. A. Mitchell, A. S. Gullerud, W. M. Scherzinger, R. Koteras, and V. L. Porter, *Adagio: non-linear quasi-static structural response using the SIERRA framework*, (2001).
[25] R.B. Morgan, *Computing interior eigenvalues of large matrices*, Linear Algebra and its Applications, 154 (1991), pp. 289–309.
[26] ———, *GMRES with deflated restarting*, SIAM Journal on Scientific Computing, 24 (2003), pp. 20–37.
[27] B. Noack, P. Papas, and P. Monkewitz, *The need for a pressure-term representation in empirical galerkin models of incompressible shear flows*, Journal of Fluid Mechanics, 523 (2005), pp. 339–365.
[28] D. P. O'Leary, *The block conjugate gradient algorithm and related methods*, Linear Algebra and its Applications, 29 (1980), pp. 293–322.
[29] M.L. Parks, E. De Sturler, G. Mackey, D.D. Johnson, and S. Maiti, *Recycling Krylov subspaces for sequences of linear systems*, SIAM Journal on Scientific Computing, 28 (2007), pp. 1651–1674.
[30] M. Rathinam and L. R. Petzold, *A new look at proper orthogonal decomposition*, SIAM Journal on Numerical Analysis, 41 (2003), pp. 1893–1925.
[31] C. Rey, *Developpement d'algorithmes paralleles de resolution en calcul non-lineaire de structures heterogenes: application au cas d'une butee acier-elastomere*, PhD thesis, Laboratoire de Modelisation et Mecanique des Structures, University of Paris VI, December 1994.
[32] C. Rey and F. Risler, *A Rayleigh–Ritz preconditioner for the iterative solution to large scale nonlinear problems*, Numerical Algorithms, 17 (1998), pp. 279–311.
[33] F. Risler and C. Rey, *Iterative accelerating algorithms with Krylov subspaces for the solution to large-scale nonlinear problems*, Numerical Algorithms, 23 (2000), pp. 1–30.
[34] F.X. Roux, *Parallel implementation of a domain decomposition method for non-linear elasticity*, SIAM, 1995, pp. 161–175.
[35] Y. Saad, *On the Lanczos method for solving symmetric linear systems with several right-hand sides*, Mathematics of Computation, (1987), pp. 651–662.
[36] ———, *Analysis of augmented Krylov subspace methods*, SIAM Journal on Matrix Analysis and Applications, 18 (1997), pp. 435–449.
[37] ———, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
[38] Y. Saad, M. Yeung, J. Erhel, and F. Guyomarc'h, *A deflated version of the conjugate gradient algorithm*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1909–1926.
[39] O. San and T. Iliescu, *Proper orthogonal decomposition closure models for fluid flows: Burgers equation*, arXiv preprint arXiv:1308.3276, (2013).
[40] S. Sirisup and G.E. Karniadakis, *A spectral viscosity method for correcting the long-term behavior of pod models*, Journal of Computational Physics, 194 (2004), pp. 92–116.
[41] L. Sirovich, *Turbulence and the dynamics of coherent structures. I: Coherent structures*, Quarterly of Applied Mathematics, 45 (1987), pp. 561–571.
[42] ———, *Turbulence and the dynamics of coherent structures. III: dynamics and scaling*, Quarterly of Applied Mathematics, 45 (1987), pp. 583–590.
[43] Z. Wang, I. Akhtar, J. Borggaard, and T. Iliescu, *Proper orthogonal decomposition closure models for turbulent flows: a numerical comparison*, Computer Methods in Applied Mechanics and Engineering, 237 (2012), pp. 10–26.
[44] D. Watkins, *The matrix eigenvalue problem: GR and Krylov subspace methods*, SIAM, 2007.