

Optimal Sensor and Actuator Selection using Balanced Model Reduction

Krithika Manohar, *Member, IEEE*, J. Nathan Kutz, *Member, IEEE*, and Steven L. Brunton, *Member, IEEE*

Abstract—Optimal sensor and actuator selection is a central challenge in high-dimensional estimation and control. Nearly all subsequent control decisions are affected by these sensor and actuator locations. In this work, we exploit balanced model reduction and greedy optimization to efficiently determine sensor and actuator selections that optimize observability and controllability. In particular, we determine locations that optimize scalar measures of observability and controllability via greedy matrix QR pivoting on the dominant modes of the direct and adjoint balancing transformations. Pivoting runtime scales linearly with the state dimension, making this method tractable for high-dimensional systems. The results are demonstrated on the linearized Ginzburg-Landau system, for which our algorithm approximates known optimal placements computed using costly gradient descent methods.

Index Terms—optimal control, balanced truncation, sensor selection, actuator selection, observability, controllability.

I. INTRODUCTION

Optimizing the selection of sensors and actuators is one of the foremost challenges in feedback control [1]. For high-dimensional systems it is impractical to monitor or actuate every state, hence a few sensors and actuators must be carefully positioned for effective estimation and control. Determining optimal selections with respect to a desired objective is an NP-hard selection problem, and in general can only be solved by enumerating all possible configurations. This combinatorial growth in complexity is intractable; therefore, the placement of sensors and actuators are typically chosen according to heuristics and intuition. In this paper, we propose a greedy algorithm for sensor and actuator selection based on jointly maximizing observability and controllability in linear time-invariant systems. Our approach (see Fig. 1) exploits low-rank transformations that balance the observability and controllability gramians to bypass the combinatorial search, enabling favorable scaling for high-dimensional systems.

To understand the challenges of sensor and actuator placement for estimation and control, we will first consider optimal sensor placement, which has mostly been used to reconstruct static signals. The primary challenge of sensor selection is that given n possible locations and a budget of r sensors, there are combinatorially many, $\binom{n}{r}$, configurations to evaluate in a brute-force search. Fortunately, there are heuristics that employ greedy selection of sensors based on maximizing mutual information [2] and information theoretic criteria [3]. Another popular approach relaxes sensor selection to a weighted convex combination of possible sensors [4], [5], [6], typically solved using semidefinite programming. Both heuristic approaches optimize submodular objective functions [7], which bound the distance between heuristic and optimal placement. Some objectives, such as those based on the quality of a Kalman filter, are not submodular [8]. Alternatively, sparsity-promoting optimization can be used to determine sensors and actua-

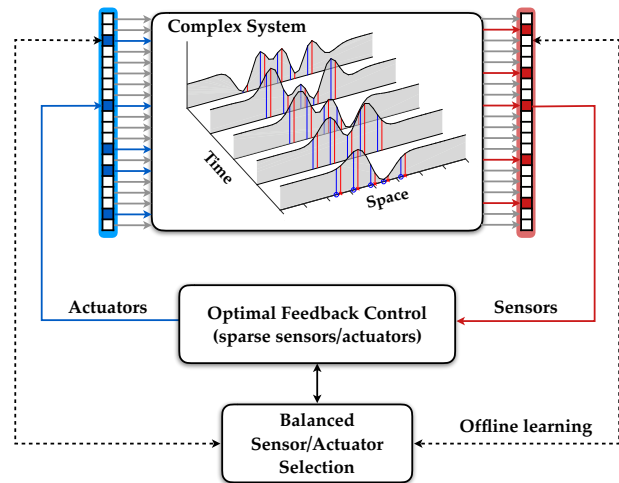


Fig. 1: Schematic of balanced sensor and actuator selection for the optimal control of a high-dimensional system.

tors [9], [10], [11], although non-differentiability of sparsity promoting terms motivates other optimization techniques [12].

Even such heuristics cannot accommodate the large dimension of many physical models, such as in fluid dynamics. Fortunately, high-dimensional systems often evolve according to relatively few intrinsic degrees of freedom. Thus, it is possible to leverage dimensionality reduction to strategically select sensors. Such sensor placement approaches [13], [14] build upon discrete empirical interpolation methods (DEIM) [15], which select the optimal measurement locations to interpolate low-rank proper orthogonal decomposition (POD) modes.

For systems with actuation, it is necessary to simultaneously consider the placement of sensors and actuators, since the most observable and most controllable subspaces are often different. Sensors and actuators for optimal feedback control are often placed along the most observable and controllable directions, respectively [16], [17], [18], [19], [7], using objective functions based on the associated observability or controllability gramians. Standard metrics for evaluating a certain sensor/actuator configuration include the H_2 norm [20], [16], a measure of the average impulse response, and the H_∞ norm to measure the worst case performance. A chief drawback is the need to recompute the controller with each new configuration of sensors and actuators given by either the gradient minimization computation or brute-force searches. Moreover, these methods do not exploit the state-of-the-art in model reduction to optimize sensor and actuator placement.

Contribution. This work develops a scalable sensor and actuator selection algorithm based on balanced truncation [21], in which modes are hierarchically ordered by their observability and controllability. We use empirical interpolation of the low-rank balanced representation to find maximally observable and controllable states. The resulting locations correspond to near-

optimal point sensor and actuator configurations. The quality of our optimized configurations are evaluated using the H_2 norm of the resulting system, which is an average measure of its output energy. Given a specific H_2 cost function and controller weight matrices, the closed loop H_2 norm, which measures input-output energy, is a more appropriate measure of control performance than open loop metrics. Our approach, when used to optimize the open loop H_2 norm, is neutral to specific choices of controller weight matrices, and instead maximizes the input-output energy of the reduced order model. We also show that it is possible to apply our framework to closed loop systems, demonstrating near optimal sensor and actuator selection in comparison with more expensive iterative closed loop H_2 optimization. The runtime scales linearly with the number of state variables, after a one-time offline computation of the balancing transformation, which is less expensive than iterative alternatives. The resulting sensor and actuator configurations reproduce known optimal locations at a fraction of the cost associated with gradient descent.

II. PROBLEM SETUP

Consider the following linear time-invariant system with a given state-space realization

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \quad \mathbf{x} \in \mathbb{R}^n, \mathbf{u} \in \mathbb{R}^q \quad (1a)$$

$$\mathbf{y} = \mathbf{C}\mathbf{x}, \quad \mathbf{y} \in \mathbb{R}^p, \quad (1b)$$

with large state dimension, i.e., $n \gg 1$. It is assumed that the system is stable, and \mathbf{B} and \mathbf{C} are linear *actuation* and *measurement* operators under which the system is observable and controllable. Our objective is to choose a minimal subset of these sensors and actuators to obtain a system that is most jointly controllable and observable. For example, the case $\mathbf{B} = \mathbf{C} = \mathbb{I}$ corresponds to pointwise sensing and actuation of each state. We consider the more general subset selection problem for arbitrary actuation and measurement operators \mathbf{B} and \mathbf{C} . This subset selection corresponds to multiplying inputs and outputs by the selection matrices

$$\mathbb{S}_C = [\mathbf{e}_{\gamma_1} \quad \mathbf{e}_{\gamma_2} \quad \dots \quad \mathbf{e}_{\gamma_r}]^T \quad (2a)$$

$$\mathbb{S}_B = [\mathbf{e}_{\beta_1} \quad \mathbf{e}_{\beta_2} \quad \dots \quad \mathbf{e}_{\beta_r}]. \quad (2b)$$

Here \mathbf{e}_j are the canonical basis vectors for \mathbb{R}^p or \mathbb{R}^q with a unit entry at the selected index j and zeros elsewhere, where $\gamma = \{\gamma_1, \dots, \gamma_r\} \subset \{1, \dots, p\}$ denotes the index set of sensor locations with r members. Similarly, actuator selection indices are given by $\beta = \{\beta_1, \dots, \beta_r\}$. The new measurement and actuation operators are $\hat{\mathbf{C}} = \mathbb{S}_C \mathbf{C}$ and $\hat{\mathbf{B}} = \mathbf{B} \mathbb{S}_B$ respectively. The new outputs, $\hat{\mathbf{y}} = \hat{\mathbf{C}}\mathbf{x}$, consist of r measurements of \mathbf{x} .

Problem statement: What are the best r -subsets of a given set of p sensors and q actuators, where $r \ll n$?

To answer this question, we first quantify the degree of observability and controllability for a given set of sensors and actuators, i.e. for a given choice of \mathbf{C} and \mathbf{B} . Optimizing over these directly involves a combinatorial search, and thus a heuristic approach is necessary for high-dimensional systems.

A. Observability and controllability

The degrees of observability and controllability for the state-space system (1) are quantified by the observability gramian \mathbf{W}_o and controllability gramian \mathbf{W}_c

$$\mathbf{W}_o = \int_0^\infty e^{\mathbf{A}^*t} \mathbf{C}^* \mathbf{C} e^{\mathbf{A}t} dt, \quad \mathbf{W}_c = \int_0^\infty e^{\mathbf{A}t} \mathbf{B} \mathbf{B}^* e^{\mathbf{A}^*t} dt, \quad (3)$$

which may be visualized as controllable and observable *ellipsoids* (Fig. 2). These depend on the actuation and measurement operators, which consist of all states reachable from a bounded initial state

$$\mathcal{E}_c = \{\mathbf{W}_c^{1/2} \mathbf{x} \mid \|\mathbf{x}\|_2 \leq 1\}, \quad (4)$$

and all states that may be observed

$$\mathcal{E}_o = \{\mathbf{W}_o^{1/2} \mathbf{x} \mid \|\mathbf{x}\|_2 \leq 1\}. \quad (5)$$

Because the gramians depend on \mathbf{B} and \mathbf{C} , they are often used to evaluate the observability/controllability of a given sensor and actuator placement. One important evaluation metric is the H_2 norm of a system. It measures the average output gain over all frequencies of the input, or the *output energy*. For the state-space system (1) with transfer function $G(s) = \mathbf{C}(s\mathbb{I} - \mathbf{A})^{-1}\mathbf{B}$, it is given by

$$\|G\|_2^2 = \frac{1}{4\pi^2} \int_0^\infty \text{tr}(G(j\omega)^* G(j\omega)) d\omega. \quad (6)$$

By the Plancherel theorem, it is also defined in the time domain by the impulse response $y_{ij}(t) = \mathbf{C}_i e^{\mathbf{A}t} \mathbf{B}_j$ - the output in component i given an impulse in input j ,

$$\|G\|_2^2 = \int_0^\infty \text{tr}(\mathbf{C} e^{\mathbf{A}t} \mathbf{B} \mathbf{B}^* e^{\mathbf{A}^*t} \mathbf{C}^*) dt = \text{tr}(\mathbf{C} \mathbf{W}_c \mathbf{C}^*) \quad (7a)$$

$$= \int_0^\infty \text{tr}(\mathbf{B}^* e^{\mathbf{A}^*t} \mathbf{C}^* \mathbf{C} e^{\mathbf{A}t} \mathbf{B}) dt = \text{tr}(\mathbf{B}^* \mathbf{W}_o \mathbf{B}) \quad (7b)$$

which explicitly relate each gramian to *both* \mathbf{B} and \mathbf{C} . An alternative to the average output energy metrics are the volumetric measures given by the log determinants

$$\log \det \mathbf{C} \mathbf{W}_c \mathbf{C}^*, \quad \log \det \mathbf{B}^* \mathbf{W}_o \mathbf{B}, \quad (8)$$

which are the logarithms of the geometric mean of the axes of the ellipsoid skewed by \mathbf{B} or \mathbf{C} ; the trace, in comparison, is the arithmetic mean. This metric is introduced by Summers et al [7] to place actuators using a greedy optimization scheme for the submodular objective function

$$\mathbf{B}_* = \underset{\mathbf{B}}{\text{argmax}} \log \det \mathbf{C} \mathbf{W}_c \mathbf{C}^*. \quad (9)$$

For H_2 optimal control it is desirable to minimize the average gain through the transfer function \hat{G} that maps disturbances \hat{w} to outputs $\hat{z}(s) = \hat{G}(s)\hat{w}(s)$, i.e., minimizing $\|\hat{G}\|_2$. Several strategies seek to build the controller and choose actuators simultaneously, using expensive gradient optimization schemes. The drawback of such closed loop metrics is having to recompute the gramians - an $O(n^3)$ operation - for every iteration that selects the next best actuator. This cubic scaling may be intractable for high-dimensional systems with large n .

There are cases where optimizing sensors and actuators using the closed loop H_2 norm is more relevant for control [20],

[16]. By contrast, our approach reverses the strategy by instead starting from a maximally actuated and sensed optimal controller, then seeks a subset of these sensors/actuators to preserve (maximize) the volumetric control measures

$$\mathbb{S}_{C^*} = \underset{\mathbb{S}_C}{\operatorname{argmax}} \log \det \mathbb{S}_C \mathbf{C} \mathbf{W}_c \mathbf{C}^* \mathbb{S}_C^*, \quad (10a)$$

$$\mathbb{S}_{B^*} = \underset{\mathbb{S}_B}{\operatorname{argmax}} \log \det \mathbb{S}_B^T \mathbf{B}^* \mathbf{W}_o \mathbf{B} \mathbb{S}_B. \quad (10b)$$

Now, the gramians no longer depend on the optimization variable and need only be computed once, and both objectives are still fundamentally linked to the H_2 norm of the system. Critically, we will extract the dominant controllable and observable subspaces from a *balanced* coordinate transformation of the gramians.

III. BALANCED MODEL REDUCTION

Many systems of interest are exceedingly high dimensional, making them difficult to characterize and limiting controller robustness due to significant computational time-delays. However, even if the ambient dimension is large, there may still be a few dominant coherent structures that characterize the system. Thus, significant effort has gone into obtaining efficient reduced-order models that capture the most relevant mechanisms for use in real-time feedback control [1].

The goal of balanced model reduction is to find a transformation \mathbf{T} from state-space (leaving inputs and outputs unchanged), $\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & 0 \end{bmatrix}$ to $\begin{bmatrix} \mathbf{TAT}^{-1} & \mathbf{TB} \\ \mathbf{CT}^{-1} & 0 \end{bmatrix}$, such that the transformed coordinates $\mathbf{a} = \mathbf{T}\mathbf{x}$ are hierarchically ordered by their joint observability and controllability. This permits an r -dimensional representation made possible by truncating the $n - r$ least observable and controllable states.

The seminal work of Moore in 1981 [21] showed it is possible to compute this change of coordinates Ψ under which the controllability and observability gramians are equal and diagonal, and it is given by the balanced system

$$\begin{aligned} \dot{\mathbf{a}} &= \Phi^* \mathbf{A} \Psi \mathbf{a} + \Phi^* \mathbf{B} \mathbf{u} & \mathbf{a} &\in \mathbb{R}^n, \mathbf{u} \in \mathbb{R}^q \\ \mathbf{y} &= \mathbf{C} \Psi \mathbf{a}. & \mathbf{y} &\in \mathbb{R}^p \end{aligned} \quad (11)$$

The desired transformation and its inverse are given by the *direct* modes Ψ and the *adjoint* modes Φ^* , respectively. The balanced state \mathbf{a} is then *truncated*, keeping only the first $r \ll n$ most jointly controllable and observable states in \mathbf{a}_r , so that $\mathbf{x} \approx \Psi_r \mathbf{a}_r$. This results in the *balanced truncation* model [21] $G_r = \begin{bmatrix} \Phi_r^* \mathbf{A} \Psi_r & \Phi_r^* \mathbf{B} \\ \mathbf{C} \Psi_r & 0 \end{bmatrix}$. Since gramians depend on the particular choice of coordinate system, they will transform under a change of coordinates. The controllability and observability gramians for the balanced truncated system are

$$\tilde{\mathbf{W}}_c = \Phi^* \mathbf{W}_c \Phi, \quad \tilde{\mathbf{W}}_o = \Psi^* \mathbf{W}_o \Psi. \quad (12)$$

The coordinate transformation Ψ that makes the controllability and observability gramians equal and diagonal,

$$\tilde{\mathbf{W}}_c = \tilde{\mathbf{W}}_o = \Sigma, \quad (13)$$

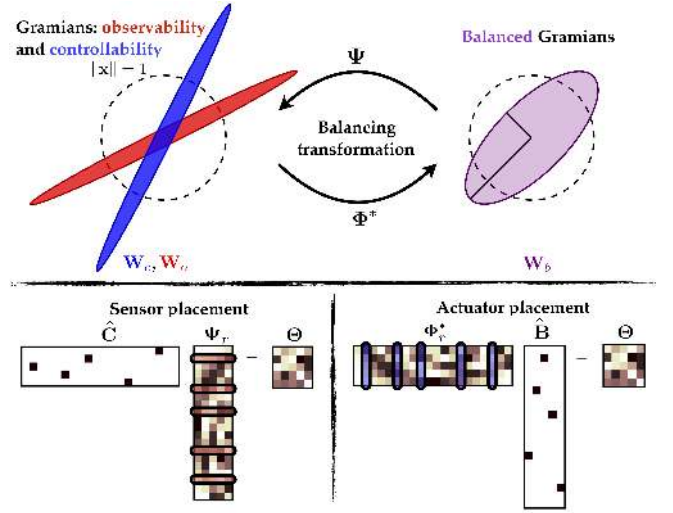


Fig. 2: (top) Illustration of the balancing transformation for gramians. The reachable set \mathcal{E}_c with unit control input is shown in blue. The corresponding observable set is shown in red. Under the balancing transformation Ψ , the gramians are equal, shown in purple. (bottom) Sensor and actuator selection based on balancing transformation.

is given by the matrix of eigenvectors of the product of the gramians $\mathbf{W}_c \mathbf{W}_o$ in the original coordinates:

$$\tilde{\mathbf{W}}_c \tilde{\mathbf{W}}_o = \Phi^* \mathbf{W}_c \mathbf{W}_o \Psi = \Sigma^2 \implies \mathbf{W}_c \mathbf{W}_o \Psi = \Psi \Sigma^2. \quad (14)$$

The H_∞ norm between the truncated system and original system is bounded by twice the sum of the neglected diagonal entries of Σ or *Hankel singular values*, $\sigma_k = \Sigma_{kk}$,

$$\|G - G_r\|_\infty \leq 2 \sum_{k=r+1}^n \sigma_k. \quad (15)$$

In practice, computing the gramians \mathbf{W}_c and \mathbf{W}_o and the eigendecomposition of the product $\mathbf{W}_c \mathbf{W}_o$ in (14) may be prohibitively expensive for high-dimensional systems. Instead, the balancing transformation may be approximated with data from impulse responses of the direct and adjoint systems, utilizing the singular value decomposition for efficient extraction of the relevant subspaces. The method of empirical gramians is quite efficient and is widely used [21], [22], [23], [24]. Moore's approach computes the entire $n \times n$ balancing transformation, which is not suitable for exceedingly high-dimensional systems. In 2002, Willcox and Peraire [23] generalized the method to high-dimensional systems, introducing a variant based on the rank- r decompositions of \mathbf{W}_c and \mathbf{W}_o obtained from snapshots of direct and adjoint simulations. It is then possible to compute the eigendecomposition of $\mathbf{W}_c \mathbf{W}_o$ using efficient eigenvalue solvers. This approach requires as many adjoint impulse-response simulations as the number of output equations, which may be prohibitively large for full-state measurements. In 2005, Rowley [24] addressed this issue by introducing output projection, which limits the number of adjoint simulations to the number of relevant POD modes in the data. It is particularly advantageous to use these data-driven methods or low-rank alternating direction methods [25]

to approximate the gramians when there are fewer than full measurements and actuation of the state.

IV. SENSOR & ACTUATOR OPTIMIZATION VIA QR PIVOTING

We now describe an efficient matrix pivoting algorithm to optimize the log determinant over the choices of sensors and actuators. The representation of the gramians in balanced truncation coordinates plays a crucial role.

A. Matrix volume objective

Recall the goal of optimizing a set of r sensors and actuators out of a fixed set p and q possible choices. The budget r determines the balancing rank truncation, which necessarily must be less than both p and q . Our sensor-actuator selection can be regarded as interpolating this rank- r representation, that is, choosing locations or interpolation points that are heavily weighted in the dominant r balanced modes.

Summers et al [7] show that it suffices to only consider controllable or observable subspaces for selecting sensors and actuators using the log determinant objective. Thus, we can substitute rank- r balanced approximation of the gramians, $\tilde{\mathbf{W}}_c$ and $\tilde{\mathbf{W}}_o$, into the log determinant objective and simplify

$$\begin{aligned} \mathbb{S}_{C^*} &\approx \underset{\mathbb{S}_C}{\operatorname{argmax}} \log \det \mathbb{S}_C \mathbf{C} \Psi_r \Sigma_r \Psi_r^* \mathbf{C}^T \mathbb{S}_C^T \\ &= \underset{\mathbb{S}_C}{\operatorname{argmax}} ((\det \mathbb{S}_C \mathbf{C} \Psi_r)^2 \cdot \det \Sigma_r) \\ &= \underset{\mathbb{S}_C}{\operatorname{argmax}} |\det \mathbb{S}_C \mathbf{C} \Psi_r|. \end{aligned} \quad (16)$$

This result follows from the monotonicity of logarithms and the product property of determinants, then omitting the term that is independent of the sensors, $\det \Sigma_r$. Likewise, in the actuator case, the objective $\log \det \tilde{\mathbf{B}}^T \tilde{\mathbf{W}}_o \tilde{\mathbf{B}}$ simplifies

$$\mathbb{S}_{B^*} \approx \underset{\mathbb{S}_B}{\operatorname{argmax}} |\det \tilde{\Phi}^* \mathbf{B} \mathbb{S}_B|. \quad (17)$$

Consider for now the case of sensor placement. The absolute determinant is a measure of matrix *volume*, and \mathbb{S}_C is a row selection matrix. The transformed objectives may be viewed as a *submatrix volume maximization* problem, which seeks the optimal r -row selection of $\mathbf{C} \Psi_r$ with the largest possible determinant. Finding this optimum is an NP-hard, intractable combinatorial search over all possible r -row submatrices of $\mathbf{C} \Psi_r$. However, it can be optimized greedily and efficiently via one-time matrix QR factorization requiring $\mathcal{O}(pr^2)$ and $\mathcal{O}(qr^2)$ operations, as described next.

B. QR pivoting algorithm

The QR factorization with column pivoting is a greedy submatrix volume optimization scheme that we will use to construct \mathbf{C} and \mathbf{B} , given Ψ_r and Φ_r . The pivoted QR factors any input matrix $\mathbf{V} \in \mathbb{R}^{r \times p}$ into a unitary matrix $\mathbf{Q} \in \mathbb{R}^{r \times r}$, upper-triangular matrix $\mathbf{R} \in \mathbb{R}^{r \times p}$, and column permutation matrix $\mathbf{P} \in \mathbb{R}^{p \times p}$ so that the permuted matrix \mathbf{VP} is better conditioned than \mathbf{V}

$$\mathbf{VP} = \mathbf{QR}. \quad (18)$$

However, we seek a well-conditioned row permutation of $\mathbf{C} \Psi_r$. Consider the input $\mathbf{V} = (\mathbf{C} \Psi_r)^*$ to the QR factorization, and the leading $r \times r$ square submatrices of the permuted input on both sides of (18),

$$\mathbf{VP}_{:,1:r} = \mathbf{QR}_{:,1:r}. \quad (19)$$

Each iteration of pivoting works by applying orthogonal projections to successive columns of \mathbf{V} to introduce subdiagonal zeros in \mathbf{R} . For our purposes, \mathbf{P} plays the crucial role: at each step \mathbf{P} stores the column ‘‘pivot’’ index of the column selected at each iteration to guarantee the following *diagonally dominant* structure in \mathbf{R}

$$|R_{ii}|^2 \geq \sum_{j=i}^k |R_{jk}|^2; \quad 1 \leq i \leq k \leq p. \quad (20)$$

Observe that the quantity of interest, the absolute determinant of the row-selected submatrix $\det |\mathbf{VP}_{:,1:r}|$ corresponding to the subset selection of measurements, now satisfies

$$|\det \mathbf{VP}_{:,1:r}| = |\det \mathbf{Q}| |\det \mathbf{R}_{:,1:r}| = \prod_{k=1}^r |R_{kk}|. \quad (21)$$

Because the determinant is the product of these diagonal entries, it can be seen that diagonal dominance property of pivoting implicitly optimizes the desired submatrix determinant. Thus \mathbb{S}_C is constructed from the first r columns of \mathbf{P}

$$\mathbb{S}_C \triangleq (\mathbf{P}_{:,1:r})^T. \quad (22)$$

Actuator selection proceeds in the same manner to construct a submatrix of r columns of $\mathbf{B}^* \Phi_r$ with maximal determinant, using one additional QR factorization

$$(\Phi_r^* \mathbf{B}) \tilde{\mathbf{P}} = \tilde{\mathbf{Q}} \tilde{\mathbf{R}}. \quad (23)$$

The solution \mathbb{S}_B is precisely the leading r columns of $\tilde{\mathbf{P}}$, $\mathbb{S}_B \triangleq \tilde{\mathbf{P}}_{:,j}$, and we denote the resulting measurement and actuation operators by

$$\hat{\mathbf{C}} = \mathbb{S}_C \mathbf{C}, \quad \hat{\mathbf{B}} = \mathbf{B} \mathbb{S}_B. \quad (24)$$

The QR pivoting routine is a standard tool in scientific computing for matrix decomposition and linear least-squares problems. We use a block accelerated implementation of classical Businger-Golub pivoting [26] in MATLAB. QR pivoting scales as $\mathcal{O}(pr^2)$ with number of candidates p and rank truncation r , making it particularly favorable in the large p and low rank setting. The computational complexity and efficient implementations have been well studied in the literature [26], [13]. Recently QR pivoting was used for reduced-order interpolation in DEIMs [13] to efficiently evaluate nonlinear terms in the model. Here, the interpolation point selection operator is analogous to our selection operator \mathbb{S}_C used with point measurements ($\mathbf{C} = \mathbb{I}$). The algorithm can be analyzed in terms of the error between the full state and the interpolant approximation using the QR selected points,

$$\hat{\mathbf{y}} = \hat{\mathbf{C}} \mathbf{x} \approx \hat{\mathbf{C}} \mathbf{U}_r \mathbf{a}_r, \quad (25)$$

where \mathbf{U}_r are the POD modes of the reduced model, and \mathbf{a}_r are modal coefficients. Recovering the state using the POD interpolant is done using standard least-squares approximation

$$\hat{\mathbf{x}} = \mathbf{U}_r(\hat{\mathbf{C}}\mathbf{U}_r)^{-1}\hat{\mathbf{y}} = \mathbf{U}_r(\hat{\mathbf{C}}\mathbf{U}_r)^{-1}\hat{\mathbf{C}}\mathbf{x}. \quad (26)$$

This can be expressed as a projection $\mathbb{P}_C \triangleq \mathbf{U}_r(\hat{\mathbf{C}}\mathbf{U}_r)^{-1}\hat{\mathbf{C}}$ of the state \mathbf{x} onto the observable subspace. The approximation error using DEIM is given by

$$\|\mathbf{x} - \mathbf{U}_r(\hat{\mathbf{C}}\mathbf{U}_r)^{-1}\hat{\mathbf{C}}\mathbf{x}\|_2. \quad (27)$$

Using a similarly defined projection onto balanced modes $\mathbb{P}_C \triangleq \Psi_r(\hat{\mathbf{C}}\Psi_r)^{-1}\hat{\mathbf{C}}$, we now use the projection error associated with balanced modes (not POD) and Hankel singular values of G to lower bound the desired submatrix determinants.

V. ANALYSIS

The best approximation to the state in the span of the direct modes is given by $\mathbf{x}_* \triangleq \Psi_r\Phi_r^*\mathbf{x}$ in the ideal measurement scenario $\mathbf{y} = \mathbf{x}$, i.e. $\mathbf{C} = \mathbb{I}$. The approximation error is bounded by twice the sum of the neglected Hankel singular values resulting from balanced truncation (15)

$$\|\mathbf{x} - \mathbf{x}_*\|_2 \leq 2(\sigma_{r+1} + \dots + \sigma_n), \quad (28)$$

which holds for all inputs of unit energy [27]. The analysis of empirical QR interpolation in the balanced modes begins with an established result for measurements selected using QR, which states that $\|(\hat{\mathbf{C}}\Psi_r)^{-1}\|_2$ at most grows as $\sqrt{p}\mathcal{O}(2^r)$.

Lemma 1 (Drmac & Gugercin [13]): For any full-rank matrix $\mathbf{M} \in \mathbb{R}^{p \times r}$, the spectral norm of $(\mathbf{S}\mathbf{M})^{-1}$, where \mathbf{S} is given by the QR pivoting algorithm (22), satisfies

$$\|(\mathbf{S}\mathbf{M})^{-1}\|_2 \leq \frac{\sqrt{p-r+1}\sqrt{4r+6r-1}}{\sigma_{\min}(\mathbf{M})}. \quad (29)$$

We generalize this result to the setting of arbitrary linear measurements and actuation, by analyzing the residual between the state and its interpolation in balanced coordinates. Note that the residual between the state and its *projection* onto balanced modes $\mathbf{v} = \mathbf{x} - \mathbf{x}_*$ satisfies

$$\mathbb{P}_C\mathbf{v} = \mathbb{P}_C\mathbf{x} - \Psi_r(\hat{\mathbf{C}}\Psi_r)^{-1}\hat{\mathbf{C}}\Psi_r\Phi_r^*\mathbf{x} = \mathbb{P}_C\mathbf{x} - \mathbf{x}_*.$$

The interpolation error from QR pivot selection satisfies

$$\begin{aligned} \|\mathbf{x} - \mathbb{P}_C\mathbf{x}\|_2 &= \|(\mathbf{v} + \mathbf{x}_*) - (\mathbb{P}_C\mathbf{v} + \mathbf{x}_*)\|_2 = \|(\mathbb{I} - \mathbb{P}_C)\mathbf{v}\|_2 \\ &\leq \|\mathbb{P}_C\|_2\|\mathbf{x} - \mathbf{x}_*\|_2 \\ &\leq \|\Psi_r\|_2\|(\hat{\mathbf{C}}\Psi_r)^{-1}\|_2\|\mathbf{C}\|_2\|\mathbf{x} - \mathbf{x}_*\|_2. \end{aligned}$$

The inequality $\|\hat{\mathbf{C}}\|_2 \leq \|\mathbf{S}\mathbf{C}\|_2\|\mathbf{C}\|_2 = \|\mathbf{C}\|_2$ is due to $\mathbf{C} \neq \mathbb{I}$. Substituting (28),(29) above yields the following result.

Theorem 2: The approximation error from interpolating QR-selected observations (22) in balanced truncated modes is controlled by the discarded Hankel singular values and the norms of the given measurements and direct modes

$$\|\mathbf{x} - \mathbb{P}_C\mathbf{x}\|_2 \leq \frac{\|\mathbf{C}\|_2\|\Psi_r\|_2}{\sigma_{\min}(\mathbf{C}\Psi_r)}\sqrt{p}\mathcal{O}(2^r)\sum_{k=r+1}^n\sigma_k. \quad (30)$$

An analogous result is obtained for actuator selection by considering the dual problem of estimating the adjoint state from actuation matrix $\hat{\mathbf{B}}$ - which is now the *measurement* operator of the adjoint system. The resulting projection operator, $\mathbb{P}_B \triangleq \Phi_r(\hat{\mathbf{B}}^*\Phi_r)^{-1}\hat{\mathbf{B}}^*$, now projects on the span of the *adjoint* modes Φ_r . Making appropriate substitutions of \mathbb{P}_B in the above results yields the following.

Corollary 1: The approximation error from interpolating QR-selected observations (24) of the adjoint state \mathbf{z} in balanced truncated modes is controlled by the discarded Hankel singular values and the norms of \mathbf{B} and Φ_r .

$$\|\mathbf{z} - \mathbb{P}_B\mathbf{z}\|_2 \leq \frac{\|\Phi_r\|_2\|\mathbf{B}\|_2}{\sigma_{\min}(\Phi_r^*\mathbf{B})}\sqrt{q}\mathcal{O}(2^r)\sum_{k=r+1}^n\sigma_k. \quad (31)$$

We now relate the approximation error bounds using QR pivot sensors and actuators to the log determinant objectives.

Theorem 3: Given direct modes Ψ_r , QR pivot sensors $\hat{\mathbf{C}}$ guarantee the following lower bound for the log determinant

$$r \log \frac{9\sigma_{\min}^2(\mathbf{C}\Psi_r)}{(p-r+1)(4^r+6r-1)} + \sum_{k=1}^r \log \sigma_k \leq \log \det \hat{\mathbf{C}}\hat{\mathbf{W}}_c\hat{\mathbf{C}}^T.$$

Proof: The absolute diagonal entries of a matrix's \mathbf{R} factor are its nondecreasing singular values, so we can lower bound the absolute determinant $|\det \hat{\mathbf{C}}\Psi_r|$

$$|\det \hat{\mathbf{C}}\Psi_r| = \prod_{k=1}^r \hat{\sigma}_k = \prod_{k=1}^r |R_{kk}| \geq |R_{rr}|^r, \quad (32)$$

where $\hat{\sigma}_k$ are the singular values of $\hat{\mathbf{C}}\Psi_r$. We obtain, upon squaring the inequality and multiplying by $\det \Sigma_r$,

$$\begin{aligned} R_{rr}^{2r} \det \Sigma_r &\leq (\det \hat{\mathbf{C}}\Psi_r)^2 \det \Sigma_r \\ &= \det \hat{\mathbf{C}}\Psi_r \Sigma_r \Psi_r^* \hat{\mathbf{C}}^T = \det \hat{\mathbf{C}}\hat{\mathbf{W}}_c\hat{\mathbf{C}}^T, \end{aligned}$$

and, taking logarithms of both sides,

$$r \log R_{rr}^2 + \sum_{k=1}^r \log \sigma_k \leq \log \det \hat{\mathbf{C}}\hat{\mathbf{W}}_c\hat{\mathbf{C}}^T.$$

Because $\|(\hat{\mathbf{C}}\Psi_r)^{-1}\|_2 = 1/|R_{rr}|$, the upper bound (29) in Lemma 2 is the inverse lower bound for $|R_{rr}|$, which can now be substituted above to obtain the final result.

An analogous lower bound can be obtained for the objective using QR pivot actuators by appropriately substituting $\hat{\mathbf{B}}, \hat{\mathbf{R}}$ and adjoint modes Φ_r in the above proof.

Corollary 2: Given adjoint modes Φ_r , $\hat{\mathbf{B}}$ satisfies the following lower bound for the log determinant

$$r \log \frac{9\sigma_{\min}^2(\Phi_r^*\mathbf{B})}{(q-r+1)(4^r+6r-1)} + \sum_{k=1}^r \log \sigma_k \leq \log \det \hat{\mathbf{B}}^T\hat{\mathbf{W}}_o\hat{\mathbf{B}}.$$

VI. RESULTS

We evaluate the selection algorithm in two settings. The first compares QR pivot selections with all possible sensor subset selections in a random state-space model of tractable size.




Fig. 3: Histograms of the log determinant and H_2 norm (trace) objectives evaluated at all possible selections of 7 sensors out of 25 show that QR pivot sensors (red) are near optimal.

Next we consider an application to closed-loop flow control using LQG control to stabilize Ginzburg-Landau dynamics. Given an LQG controller with full sensing and actuation, we approximate H_2 optimal placements computed using gradient descent [16] with our QR scheme.

A. Discrete random state space

Our first example investigates sensor and actuator selection for random state-space systems with randomized $\mathbf{A}, \mathbf{B}, \mathbf{C}$. First, we compare the results of QR sensor placement against a brute-force search across all possible $r = 7$ sensor selections for a system with $n = 25$ states and $p = q = 25$ randomized measurements. The log determinant objective (10) is evaluated for all possible choices of 7 sensors, since the system is small enough to explicitly compute the full gramian for all $\binom{n}{r} = 480,700$ choices of $\hat{\mathbf{C}}$. These results are binned in Fig. 3, and compared with the value resulting from our method (red line). The input to the QR scheme, the balancing modes, are computed only once from the full system. The sensors resulting from our method are observed to be near optimal for the log determinant, exceeding 99.99% of all others, and also good substitutes for H_2 optimal sensors. On average, our method surpasses 99.8% of possible outcomes with a standard deviation of 0.85%, over a randomly generated ensemble of 500 model realizations. Therefore, QR sensors are closer to optimal than the analysis suggests.

We now investigate performance on a larger random state-space model with $n = 100$ states, and likewise initialize the model with randomized actuation and sensing such that $p = q = 100$. Figure 4 shows the log determinant objective that is being optimized for various sensor and actuator configurations. The log determinant of the gramian volume is plotted for the truncated model with QR-optimized sensor and actuator configurations (red circles) and with random configurations (blue violin plots). The truncation level r for the balanced truncation is chosen to match the sensor and actuator budget on the x -axis. The QR-optimized configurations dramatically outperform random configurations. As more modes are retained, the chosen sensors and actuators better characterize the input-output dynamics, and their performance gap over random placement increases over all random ensembles, giving empirical validation of our approach.




Fig. 4: Sensor and actuator placement in a random state-space system. The log determinant objective is plotted for QR-optimized sensor-actuator selections (red) and an ensemble of 200 random sensor-actuator selections (blue violin plots). The truncation level r (also the sensor/actuator budget) varies on the horizontal axis.

Because the system is randomly generated and the dynamics do not evolve according to broad, non-localized features in state-space, many sensors and actuators are required to characterize the system. In particular, this is reflected in the slow decay of Hankel singular values. By contrast, the next example is generated by a physical fluid flow model, and has coherent structure that allow for a more physical interpretation of sensor and actuator placements with enhanced sparsity.

B. Linearized Ginzburg-Landau with stochastic disturbances

Consider a closed-loop, fully sensed & actuated, linearized Ginzburg-Landau model evolving velocity perturbations in a flow. The system matrix \mathbf{A}_2 is a Hermite pseudospectral discretization of the linearized Ginzburg-Landau operator $\mathcal{A} = -\nu \frac{\partial}{\partial \xi} + \mu(\xi) + \beta \frac{\partial^2}{\partial \xi^2}$ over $n = 100$ spatial gridpoints ($\xi \in \mathbb{R}^n$ in vector notation), and $\nu = 2 + .4i, \beta = 1 - i$, and $\mu(\xi) = .37 - .005\xi^2$ are complex advection, diffusion and wave amplification parameters. Sensors and actuators are spatially localized Gaussians centered at each gridpoint. For example, the k th actuator (k th column of \mathbf{B}_2) is given by $\exp(-(\xi - \xi_a^k)^2 / \sqrt{2}\sigma)$, with $\sigma = 0.4$. Plant dynamics are perturbed by white noise signals $\mathbf{w}_d \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ and $\mathbf{w}_n \sim \mathcal{N}(\mathbf{0}, 4e-8\mathbb{I})$

$$\dot{\mathbf{x}} = \mathbf{A}_2 \mathbf{x} + \mathbf{B}_2 \mathbf{u} + \mathbf{w}_d \quad (33a)$$

$$\mathbf{y} = \mathbf{C}_2 \mathbf{x} + \mathbf{w}_n, \quad (33b)$$

which is unstable because the system matrix has eigenvalues in the right half plane. We perform H_2 control using a linear quadratic Gaussian (LQG) controller that minimizes the ensemble-averaged cost function $J(t) = \langle \int_0^t [\mathbf{x}(\tau)^T \hat{\mathbf{Q}} \mathbf{x}(\tau) + \mathbf{u}(\tau)^T \hat{\mathbf{R}} \mathbf{u}(\tau)] d\tau \rangle$, where the average is taken over the noise realizations. The matrices $\hat{\mathbf{Q}} = \beta^2 \text{diag}(\mathbf{d})$ and $\hat{\mathbf{R}} = \mathbb{I}$, where \mathbf{d} are pairwise distances between the gridpoints, weight the costs of state regulation and actuation respectively, specifying the relative importance of the control objectives. We solve Riccati equations in the standard way to optimize the LQG gain matrices \mathbf{F}, \mathbf{L} , which depend on both $\hat{\mathbf{Q}}, \hat{\mathbf{R}}$ and the

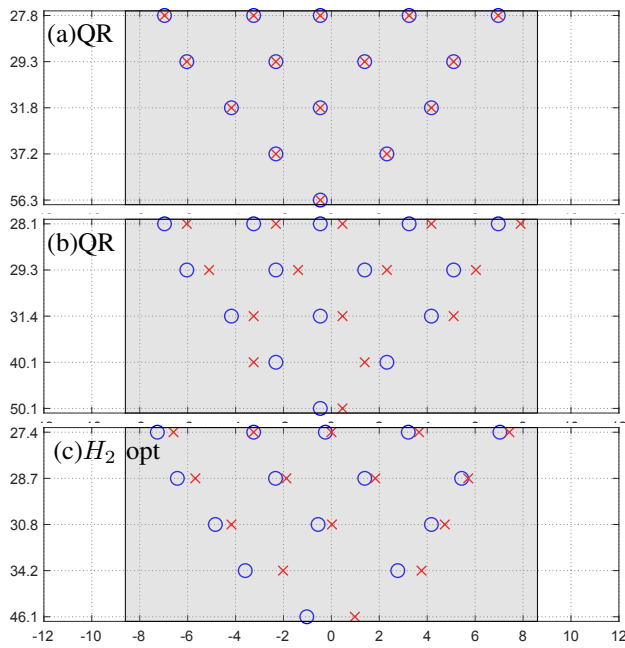


Fig. 5: Sensor (\times) and actuator (\circ) placement for linearized Ginzburg-Landau. Each row corresponds to the optimized placement for budgets of 1-5 sensors and actuators. Placements based on QR pivoting of balanced truncated modes (a) closely approximate the H_2 norms of the placements determined using gradient descent (c). The QR method can be modified to place sensors and actuators to avoid collocation (b).

noise covariances, and obtain the controller that stabilizes the plant (33)

$$\begin{bmatrix} \dot{\hat{\mathbf{x}}} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_2 - \mathbf{B}_2\mathbf{F} - \mathbf{L}\mathbf{C}_2 & \mathbf{L} \\ -\mathbf{F} & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}} \\ \mathbf{y} \end{bmatrix}. \quad (34)$$

Our algorithm generalizes optimal sensor & actuator selection for this closed loop formulation. To see why, observe that \mathbf{F}, \mathbf{L} are fixed and hence the system is structured similarly to (1), albeit with inputs and outputs swapped. Nevertheless the dominant observable and controllable subspaces of the controller may still be computed and exploited to optimize sensor-actuator selections. Thus we compute gramians and direct & adjoint modes of the controller by assigning the state-space realization $\mathbf{A} \triangleq \mathbf{A}_2 - \mathbf{B}_2\mathbf{F} - \mathbf{L}\mathbf{C}_2$, $\mathbf{B} \triangleq \mathbf{L}$, $\mathbf{C} \triangleq -\mathbf{F}$ as in (1), noting our results also hold for complex-valued systems. In this scenario, which is not always practical, we access a controller with full sensing and actuation, extract subspaces relevant to control, and proceed with sensor & actuator selection.

We compare our approach to the gradient descent scheme of Chen and Rowley [16] who simultaneously optimize the LQG controller and sensor-actuator placements. Their H_2 norm optimization scheme permits placement of sensors and actuators at locations that may not be grid points. The major drawback is that each Newton iteration requires solving $2r \ n \times \ n$ Lyapunov equations until convergence, although recent work reduces this to 2 equations per iteration [28]. Furthermore, the procedure requires an ensemble of random initial conditions




Fig. 6: LQG gain (dB) for a system with 5 sensors and actuators. Each block shows the gain from a signal $\exp(i\omega t)$ in sensor k (column) to actuator j (row), ordered upstream to downstream.

to avoid converging to a local minimum. In [16], the optimal placement is computed using conjugate gradient optimization for the same spatial discretization $n = 100$, which becomes computationally expensive as the grid resolution increases. In this case, gradient descent is more costly than balancing the fully actuated and observed system, which comes at a one-time cost of solving 2 Lyapunov equations for the gramians, and 2 Riccati equations for the LQG gain matrices ($O(n^3)$ each). Therefore, our algorithm is sensible when the grid discretization is sufficiently fine. Furthermore, our solution is a good starting point for the convergence of gradient descent, thus eliminating the need for optimization over a large ensemble of randomized starting points. QR pivoting runtime scales as $O(nr^2)$ and the deviation of the resulting placement from the H_2 optimum (fig. 5) decreases with increasing r .

Figure 5 plots sensor and actuator configurations from the QR algorithm and H_2 gradient optimization, which are compared with the H_2 optimal placements in [16]. The resulting placements for the cases $r = 1$ to $r = 5$ sensors and actuators are plotted vertically, and the horizontal axis is the spatial domain $\xi \in [-12, 12]$ with a shaded wave amplification region. For each value of r , we apply QR pivoting to the rank r truncated balanced modes. QR pivoting collocates sensors and actuators, indicating the direct and adjoint modes (first 5 plotted as shaded regions in Fig. 1) are identical up to a scaling factor. In practice, sensors are often slightly downstream to account for time delays, so we enforce via the pivoting procedure that sensors are not placed at previously chosen actuators. The H_2 norms of the resulting placement on the y -axis indicate that the QR selections closely approximate the optimal placements. The QR selection for five sensors and actuators results in an H_2 norm of 27.8, which agrees closely with the optimal H_2 norm of 27.4 [16, Fig. 4].

Figure 6 compares controller gain responses between QR pivoting and the H_2 optimum via the LQG gain of a given signal from each sensor to each actuator. The LQG gains agree closely with those produced by the H_2 optimal method of Chen and Rowley [16, Fig. 5]. Balanced truncation is applied to the closed loop system since the open loop dynamics are unstable, and it is shown in [16] that the dominant eigenmodes of the dynamics lead to vastly suboptimal placements. It is also possible to enforce non collocation of sensors and actuators as in [16]; this sensor-actuator configuration has been extensively

studied in convective flows [29].

VII. DISCUSSION AND OUTLOOK

We develop a scalable sensor and actuator selection method with runtime that scales linearly with the number of candidate locations, after a one-time computation of the balanced modes. Our approach relies on balanced model reduction [21], [23], [24], which hierarchically orders modes by their observability and controllability. We extend EIMs to interpolate the low-rank balancing modes of the system and determine maximally observable and controllable locations (sensor & actuators) in state space. The performance of this algorithm is demonstrated on random state-space systems, and optimal H_2 control of the linearized Ginzburg-Landau model. Our optimized placements vastly exceed the performance of random placements, and closely approximate H_2 optimal placements computed by costly gradient minimization schemes, but achieved at a fraction of the runtime.

Sensors and actuators are critical for feedback control of large high-dimensional complex systems. This work advocates sensor and actuator selection using QR pivots of the direct and adjoint modes of a system's balancing transformation. The resulting placement is empirically shown to preserve the dynamics of the full system. The method has deep connections to system observability, controllability, modal sampling methods and classical experimental design criteria. Furthermore, QR pivoting is more computationally efficient than leading greedy and convex optimization methods, and thus critically enlarges the search space of possible selections. This is particularly valuable in spatiotemporal models where high-resolution grids generate a large number of states, and balanced modes and QR method exploit the spatial structures.

This work opens a variety of future directions in pivoting sensor and actuator optimization. Rapid advances in data collection yield extremely large search spaces. For prohibitively large n or p , randomized linear algebra can significantly accelerate both pivoting [30] and balancing computations. In addition, it remains to study the case when only data but no model of the dynamics are given, and optimize placements based on data-driven system identification models. These models may be characterized by complex, nonlinear dynamics or sensing and actuation constraints. These computational studies and extensions are topics of future investigation.

REFERENCES

- [1] G. E. Dullerud and F. Paganini, *A course in robust control theory: A convex approach*. Springer Texts in Applied Mathematics, 2000.
- [2] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, no. Feb, pp. 235–284, 2008.
- [3] L. Paninski, "Asymptotic theory of information-theoretic experimental design," *Neural Computation*, vol. 17, no. 7, pp. 1480–1507, 2005.
- [4] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, 2009.
- [5] S. P. Chepuri and G. Leus, "Continuous sensor placement," *IEEE Signal Processing Letters*, vol. 22, no. 5, pp. 544–548, 2015.
- [6] S. Liu, S. P. Chepuri, M. Fardad, E. Masazade, G. Leus, and P. K. Varshney, "Sensor selection for estimation with correlated measurement noise," *IEEE Trans. Sig. Proc.*, vol. 64, no. 13, pp. 3509–3522, 2016.

- [7] T. H. Summers, F. L. Cortesi, and J. Lygeros, "On submodularity and controllability in complex dynamical networks," *IEEE Trans. Control of Network Systems*, vol. 3, no. 1, pp. 91–101, 2016.
- [8] H. Zhang, R. Ayoub, and S. Sundaram, "Sensor selection for Kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms," *Automatica*, vol. 78, pp. 202–210, 2017.
- [9] F. Lin, M. Fardad, and M. R. Jovanovic, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *IEEE Trans. Autom. Control*, vol. 58, no. 9, pp. 2426–2431, 2013.
- [10] U. Munz, M. Pfister, and P. Wolfrum, "Sensor and actuator placement for linear systems based on h_2 and h_∞ optimization," *IEEE Trans. Autom. Control*, vol. 59, no. 11, pp. 2984–2989, 2014.
- [11] A. Zare, N. K. Dhingra, M. R. Jovanović, and T. T. Georgiou, "Proximal algorithms for large-scale statistical modeling and optimal sensor/actuator selection," *arXiv preprint arXiv: 1807.01739*, 2018.
- [12] N. K. Dhingra, M. R. Jovanovic, and Z.-Q. Luo, "An ADMM algorithm for optimal sensor and actuator selection," *53rd IEEE Conference on Decision and Control*, pp. 4039–4044, 2014.
- [13] Z. Drmac and S. Gugercin, "A new selection operator for the discrete empirical interpolation method—improved a priori error bound and extensions," *SIAM J. Sci. Comput.*, vol. 38, no. 2, pp. A631–A648, 2016.
- [14] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, "Data-driven sparse sensor placement," *IEEE Control Syst. Mag.*, vol. 38, no. 3, pp. 63–86, 2018.
- [15] S. Chaturantabut and D. C. Sorensen, "Nonlinear model reduction via discrete empirical interpolation," *SIAM Journal on Scientific Computing*, vol. 32, no. 5, pp. 2737–2764, 2010.
- [16] K. K. Chen and C. W. Rowley, " H_2 optimal actuator and sensor placement in the linearised complex Ginzburg-Landau system," *Journal of Fluid Mechanics*, vol. 681, pp. 241–260, 2011.
- [17] T. Nestorović and M. Trajkov, "Optimal actuator and sensor placement based on balanced reduced models," *Mechanical Systems and Signal Processing*, vol. 36, no. 2, pp. 271–289, 2013.
- [18] B. T. Hinson and K. A. Morgansen, "Observability-based optimal sensor placement for flapping airfoil wake estimation," *Journal of Guidance, Control, and Dynamics*, vol. 37, no. 5, pp. 1477–1486, 2014.
- [19] D. Bhattacharjee, M. Hemati, B. Klose, and G. Jacobs, "Optimal actuator selection for airfoil separation control," in *2018 Flow Control Conference*, 2018, p. 3692.
- [20] K. Morris, M. A. Demetriou, and S. D. Yang, "Using H_2 -control performance metrics for the optimal actuator location of distributed parameter systems," *IEEE Trans. Autom. Control*, vol. 60, no. 2, pp. 450–462, 2015.
- [21] B. C. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Trans. Autom. Control*, vol. AC-26, no. 1, pp. 17–32, 1981.
- [22] S. Lall, J. E. Marsden, and S. Glavaški, "A subspace approach to balanced truncation for model reduction of nonlinear control systems," *Int. J. Robust Nonlinear Control*, vol. 12, no. 6, pp. 519–535, 2002.
- [23] K. Willcox and J. Peraire, "Balanced model reduction via the proper orthogonal decomposition," *AIAA J.*, vol. 40, no. 11, pp. 2323–2330, 2002.
- [24] C. W. Rowley, "Model reduction for fluids using balanced proper orthogonal decomposition," *Int. J. Bifurc. Chaos*, vol. 15, no. 3, pp. 997–1013, 2005.
- [25] J.-R. Li and J. White, "Low rank solution of Lyapunov equations," *SIAM J. Matrix Anal. Appl.*, vol. 24, no. 1, pp. 260–280, 2002.
- [26] P. Businger and G. H. Golub, "Linear least squares solutions by Householder transformations," *Numerische Mathematik*, vol. 7, no. 3, pp. 269–276, 1965.
- [27] A. C. Antoulas, *Approximation of large-scale dynamical systems*. SIAM, 2005.
- [28] C. Colburn, D. Zhang, and T. Bewley, "Gradient-based optimization methods for sensor & actuator placement in lti systems," *Submitted for publication, May*, 2011.
- [29] B. A. Belson, O. Semeraro, C. W. Rowley, and D. S. Henningson, "Feedback control of instabilities in the two-dimensional blasius boundary layer: the role of sensors and actuators," *Physics of fluids*, vol. 25, no. 5, p. 054106, 2013.
- [30] P.-G. Martinsson, G. Quintana Ortí, N. Heavner, and R. van de Geijn, "Householder QR factorization with randomization for column pivoting (HQRFP)," *SIAM J. Sci. Comput.*, vol. 39, no. 2, pp. C96–C115, 2017.