

PROCEEDINGS A

rspa.royalsocietypublishing.org

Research



Article submitted to journal

Subject Areas:

applied mathematics, mathematical modelling, robotics, control

Keywords:

Koopman operator theory, linear matrix inequalities, regularization, linear systems theory, system norms, asymptotic stability

Author for correspondence:

Steven Dahdah

e-mail:

steven.dahdah@mail.mcgill.ca

System Norm Regularization Methods for Koopman Operator Approximation

Steven Dahdah¹ and James R. Forbes¹¹Department of Mechanical Engineering, McGill University, Montreal QC H3A 0C3, Canada

Approximating the Koopman operator from data is numerically challenging when many lifting functions are considered. Even low-dimensional systems can yield unstable or ill-conditioned results in a high-dimensional lifted space. In this paper, Extended Dynamic Mode Decomposition (DMD) and DMD with control, two methods for approximating the Koopman operator, are reformulated as convex optimization problems with linear matrix inequality constraints. Asymptotic stability constraints and system norm regularizers are then incorporated as methods to improve the numerical conditioning of the Koopman operator. Specifically, the $\mathcal{H}_\infty$ norm is used to penalize the input-output gain of the Koopman system. Weighting functions are then applied to penalize the system gain at specific frequencies. These constraints and regularizers introduce bilinear matrix inequality constraints to the regression problem, which are handled by solving a sequence of convex optimization problems. Experimental results using data from an aircraft fatigue structural test rig and a soft robot arm highlight the advantages of the proposed regression methods.

1. Introduction

Koopman operator theory [1–4] allows a nonlinear system to be exactly represented as a linear system in terms of an infinite set of *lifting functions*. The *Koopman operator* advances each of these lifting function to the next timestep. Thanks to recent theoretical developments [2–4] and the widespread availability of computational resources, there has been a recent resurgence of interest in using data-driven methods to approximate the Koopman operator. The Koopman operator defines a linear state-space system in the chosen lifted space, making it convenient for control system design. Koopman models have been paired with a wide variety of existing linear optimal control techniques [5–10] with great success.




Figure 1: Overview of the data-driven Koopman workflow, including the role of the proposed regression methods. (a) First, data is collected from the system to be identified. (b) Next, a set of lifting functions is chosen. (c) The lifting functions are then applied to the data and snapshot matrices are formed. (d) Finally, one of several regression methods is used to approximate the Koopman matrix. The proposed regression methods seek to address the numerical problems often present in this step by viewing it as an optimization over discrete-time linear systems.

To use the Koopman representation in practical applications, a finite-dimensional approximation of the infinite-dimensional Koopman operator must be found. First, a finite set of lifting functions is selected. These functions are often hand-picked based on known dynamics [7,8], or are combinations of sinusoids, polynomials, and other basis functions [9,11]. Time delay embeddings are also popular [5,9]. However, there is no universally agreed-upon method for selecting lifting functions. Given a set of lifting functions, linear regression is used to find the matrix approximation of the Koopman operator, also called a *Koopman matrix* [6,12].

Unfortunately, the regression problem associated with finding an approximate Koopman operator is numerically challenging, as complex lifting function choices can yield unstable or ill-conditioned Koopman models for stable systems [13]. Regularization techniques play a crucial role in obtaining usable Koopman models for prediction and control applications. Standard regularization techniques like Tikhonov regularization [14] or the lasso [15] are often used to promote well-conditioned Koopman matrices. These regularization techniques penalize different matrix norms of the Koopman matrix, without considering the fact that the Koopman matrix defines a discrete-time linear system with input, state, and output. While these methods may indirectly promote asymptotic stability in this Koopman system, their success is highly dependent on the regularization coefficient used. Furthermore, they do not consider the input-output gain of the Koopman system. This paper takes a systems view of the Koopman matrix regression problem, proposing regression methods that constrain the asymptotic stability of the

Koopman system and regularize the regression problem by penalizing its input-output gain, as represented by its $\mathcal{H}_\infty$ norm. To accomplish this, the *Extended Dynamic Mode Decomposition* (EDMD) [16] and *Dynamic Mode Decomposition with control* (DMDc) [17] methods are reformulated as convex optimization problems with linear matrix inequality (LMI) constraints. Regularizers and additional constraints are incorporated in a modular fashion as either LMI constraints, or bilinear matrix inequality (BMI) constraints. The data-driven Koopman workflow and the proposed regression methods are summarized in Figure 1.

(a) Related work

Convex optimization and LMIs have previously been used to synthesize controllers for Koopman models [10], but they have not yet been leveraged to regularize the Koopman matrix regression problem. A related optimization problem is posed in [18], where both the Koopman matrix and lifting functions are treated as unknowns. While this problem is NP-hard, a convex relaxation allows both to be found by solving two semidefinite programs. The $\mathcal{H}_\infty$ norm of the Koopman operator has previously been considered in [19], however it is in the form of a hard constraint on the system's dissipativity, rather than as a regularizer in the cost function.

In [20], the problem of learning Lyapunov stable and asymptotically stable linear systems is explored in the context of subspace identification, where Lyapunov inequalities are used to enforce the corresponding stability conditions. The related problem of learning positive real and strictly positive real systems using constrained subspace identification is discussed in [21]. A convex relaxation of the Lyapunov inequality is considered in [22], where linear constraints are added incrementally to enforce the Lyapunov stability of a system. In [13], a gradient-descent method called SOC [23] is applied to find locally optimal Lyapunov stable or asymptotically stable Koopman matrices. The method relies on a parameterization of the Koopman matrix that guarantees Lyapunov stability or asymptotic stability [24]. While addressing the asymptotic stability problem, this formulation lacks the modularity of the proposed approach.

(b) Contribution

The core contributions of this paper are solving the EDMD and DMDc problems with asymptotic stability constraints and with system norm regularizers. Of particular focus is the use of the $\mathcal{H}_\infty$ norm as a regularizer, which penalizes the worst-case gain of the Koopman system over all frequencies. The BMI formulation of the EDMD problem with asymptotic stability constraints and $\mathcal{H}_\infty$ norm regularization were previously explored by the authors in [25]. LMI formulations for Tikhonov regularization, matrix two-norm regularization, and nuclear norm regularization were also presented in [25]. This paper expands on [25] to include an LMI formulation of the DMDc problem, and discusses the corresponding asymptotic stability constraint and $\mathcal{H}_\infty$ norm regularizer. As with EDMD, these modifications add BMI constraints to the DMDc problem. Furthermore, weighted $\mathcal{H}_\infty$ norm regularization is explored, which allows the Koopman system's gain to be penalized in a specific frequency band, where experimental measurements may be less reliable, or where system dynamics may be irrelevant. Finally, the proposed regression methods are evaluated using two experimental datasets, one from a fatigue structural testing platform, and the other from a soft robot arm. The significance of this work is the use of a system norm to regularize the Koopman regression problem, which is viewed as a regression problem over discrete-time linear systems, resulting in a numerically better conditioned data-driven model.

2. Background

(a) Koopman operator theory

Consider the discrete-time nonlinear process

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k), \quad (2.1)$$

where $\mathbf{x}_k \in \mathcal{M}$ evolves on a manifold $\mathcal{M} \subseteq \mathbb{R}^{m \times 1}$. Let $\psi : \mathcal{M} \rightarrow \mathbb{R}$ be a *lifting function*. Any scalar function of the state $\mathbf{x}_k$ qualifies as a lifting function. The lifting functions therefore form an infinite-dimensional Hilbert space $\mathcal{H}$. The *Koopman operator* $\mathcal{U} : \mathcal{H} \rightarrow \mathcal{H}$ is a linear operator that advances all scalar-valued lifting functions $\psi \in \mathcal{H}$ in time by one timestep. That is [12, §3.2],

$$(\mathcal{U}\psi)(\cdot) = (\psi \circ \mathbf{f})(\cdot). \quad (2.2)$$

Using (2.2), the dynamics of (2.1) can be rewritten linearly in terms of ψ as

$$\psi(\mathbf{x}_{k+1}) = (\mathcal{U}\psi)(\mathbf{x}_k). \quad (2.3)$$

In finite dimensions, (2.3) is approximated by

$$\psi(\mathbf{x}_{k+1}) = \mathbf{U}\psi(\mathbf{x}_k) + \mathbf{r}_k, \quad (2.4)$$

where $\psi : \mathcal{M} \rightarrow \mathbb{R}^{p \times 1}$, $\mathbf{U} \in \mathbb{R}^{p \times p}$, and $\mathbf{r}_k$ is the residual error. Each element of the *vector-valued lifting function* ψ is a lifting function in $\mathcal{H}$. The *Koopman matrix* $\mathbf{U}$ is a matrix approximation of the Koopman operator.

(b) Koopman operator theory with inputs

If a discrete-time nonlinear process with exogenous inputs is considered, the definitions of the lifting functions and Koopman operator must be modified. Consider

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k), \quad (2.5)$$

where $\mathbf{x}_k \in \mathcal{M} \subseteq \mathbb{R}^{m \times 1}$ and $\mathbf{u}_k \in \mathcal{N} \subseteq \mathbb{R}^{n \times 1}$. In this case, the lifting functions become $\psi : \mathcal{M} \times \mathcal{N} \rightarrow \mathbb{R}$ and the Koopman operator $\mathcal{U} : \mathcal{H} \rightarrow \mathcal{H}$ is instead defined so that

$$(\mathcal{U}\psi)(\mathbf{x}_k, \mathbf{u}_k) = \psi(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k), \star), \quad (2.6)$$

where $\star = \mathbf{u}_k$ if the input has state-dependent dynamics, or $\star = \mathbf{0}$ if the input has no dynamics [12, §6.5]. If the input is computed by a controller, it is often considered to have state-dependent dynamics. Let the vector-valued lifting function $\psi : \mathcal{M} \times \mathcal{N} \rightarrow \mathbb{R}^{p \times 1}$ be partitioned as

$$\psi(\mathbf{x}_k, \mathbf{u}_k) = \begin{bmatrix} \vartheta(\mathbf{x}_k) \\ \mathbf{v}(\mathbf{x}_k, \mathbf{u}_k) \end{bmatrix}, \quad (2.7)$$

where $\vartheta : \mathcal{M} \rightarrow \mathbb{R}^{p_\vartheta \times 1}$, $\mathbf{v} : \mathcal{M} \times \mathcal{N} \rightarrow \mathbb{R}^{p_v \times 1}$, and $p_\vartheta + p_v = p$. When the input is exogenous, (2.6) has the form [12, §6.5.1]

$$\vartheta(\mathbf{x}_{k+1}) = \mathbf{U}\vartheta(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{r}_k, \quad (2.8)$$

where $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$. Expanding (2.8) yields the familiar linear state-space form,

$$\vartheta(\mathbf{x}_{k+1}) = \mathbf{A}\vartheta(\mathbf{x}_k) + \mathbf{B}\mathbf{v}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{r}_k. \quad (2.9)$$

When identifying a Koopman model for control, the input is often left unlifted, that is, $\mathbf{v}(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{u}_k$ [5]. However, recent work demonstrates that this choice of lifting functions is insufficient for describing control affine systems, which are ubiquitous in real-world applications [26]. An alternative choice of input-dependent lifting functions proposed in [26] is

$$\mathbf{v}(\mathbf{x}_k, \mathbf{u}_k) = \begin{bmatrix} \mathbf{u}_k \otimes \vartheta(\mathbf{x}_k) \\ \mathbf{u}_k \end{bmatrix}, \quad (2.10)$$

where $\otimes$ denotes the Kronecker product. These bilinear input-dependent lifting functions are capable of representing all control affine systems, and therefore present an interesting alternative to leaving the input unlifted [26].

(c) Approximating the Koopman operator from data

To approximate the Koopman matrix from a dataset $\mathcal{D} = \{\mathbf{x}_k, \mathbf{u}_k\}_{k=0}^q$, consider the lifted snapshot matrices

$$\Psi = [\psi_0 \quad \psi_1 \quad \cdots \quad \psi_{q-1}] \in \mathbb{R}^{p \times q}, \quad (2.11)$$

$$\Theta_+ = [\vartheta_1 \quad \vartheta_2 \quad \cdots \quad \vartheta_q] \in \mathbb{R}^{p_\vartheta \times q}, \quad (2.12)$$

where $\psi_k = \psi(\mathbf{x}_k, \mathbf{u}_k)$ and $\vartheta_k = \vartheta(\mathbf{x}_k)$. The Koopman matrix that minimizes

$$J(\mathbf{U}) = \|\Theta_+ - \mathbf{U}\Psi\|_F^2 \quad (2.13)$$

is [12, §1.2.1]

$$\mathbf{U} = \Theta_+ \Psi^\dagger, \quad (2.14)$$

where $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse.

3. Reformulating the Koopman Operator Regression Problem

(a) Extended DMD

The direct least-squares method of approximating the Koopman operator in (2.14) is fraught with numerical and performance issues. Namely, computing the pseudoinverse of Ψ is costly when the dataset contains many snapshots. Extended Dynamic Mode Decomposition (EDMD) [16] reduces the dimension of the pseudoinverse required to compute (2.14) when the number of snapshots is much larger than the dimension of the lifted state (i.e., $p \ll q$) [12, §10.3].

Extended DMD consists of computing (2.14) using

$$\mathbf{U} = \Theta_+ (\Psi^\top \Psi^{\top\dagger}) \Psi^\dagger = (\Theta_+ \Psi^\top) (\Psi \Psi^\top)^\dagger = \mathbf{G} \mathbf{H}^\dagger, \quad (3.1)$$

where

$$\mathbf{G} = \frac{1}{q} \Theta_+ \Psi^\top \in \mathbb{R}^{p_\vartheta \times p}, \quad \mathbf{H} = \frac{1}{q} \Psi \Psi^\top \in \mathbb{R}^{p \times p}. \quad (3.2)$$

Now, only a $p \times p$ pseudoinverse is required, rather than a $p \times q$ pseudoinverse. To improve numerical conditioning, $\mathbf{G}$ and $\mathbf{H}$ are often scaled by the number of snapshots q , as in (3.2). Note that $\mathbf{H} = \mathbf{H}^\top > 0$ if the columns of Ψ are linearly independent.

(b) LMI reformulation of EDMD

To incorporate regularizers and constraints in a modular fashion, the Koopman operator regression problem is reformulated as a convex optimization problem with LMI constraints. Recall that the Koopman matrix $\mathbf{U}$ minimizes (2.13). It therefore also minimizes

$$J(\mathbf{U}) = \frac{1}{q} \|\Theta_+ - \mathbf{U}\Psi\|_F^2 = \frac{1}{q} \text{tr}((\Theta_+ - \mathbf{U}\Psi)(\Theta_+ - \mathbf{U}\Psi)^\top) \quad (3.3)$$

$$= \text{tr}\left(\frac{1}{q} \Theta_+ \Theta_+^\top - \text{He}\{\mathbf{U}\mathbf{G}^\top\} + \mathbf{U}\mathbf{H}\mathbf{U}^\top\right) \quad (3.4)$$

$$= c - 2\text{tr}(\mathbf{U}\mathbf{G}^\top) + \text{tr}(\mathbf{U}\mathbf{H}\mathbf{U}^\top), \quad (3.5)$$

where $c = \frac{1}{q} \Theta_+ \Theta_+^\top$ is a scalar constant, $\mathbf{G}$ and $\mathbf{H}$ are defined in (3.2), and $\text{He}\{\cdot\} = (\cdot) + (\cdot)^\top$. The minimization of (3.5) is equivalent to the minimization of

$$J(\mathbf{U}, \nu, \mathbf{W}) = c - 2\text{tr}(\mathbf{U}\mathbf{G}^\top) + \nu \quad (3.6)$$

subject to

$$\text{tr}(\mathbf{W}) < \nu, \quad \mathbf{W} > 0, \quad \mathbf{U}\mathbf{H}\mathbf{U}^\top < \mathbf{W}, \quad (3.7)$$

where ν and $\mathbf{W}$ are slack variables that allow the cost function to be rewritten using LMIs [27, §2.15.1]. To rewrite the quadratic term in (3.7) as an LMI, consider the matrix decomposition $\mathbf{H} = \mathbf{L}\mathbf{L}^T$. The matrix $\mathbf{L}$ can be found using a Cholesky factorization or eigendecomposition of $\mathbf{H}$, or a singular value decomposition of Ψ . Assuming the decomposition has been performed, the quadratic term in the optimization problem becomes

$$\mathbf{W} - \mathbf{U}\mathbf{H}\mathbf{U}^T = \mathbf{W} - \mathbf{U}\mathbf{L}\mathbf{L}^T\mathbf{U}^T = \mathbf{W} - (\mathbf{U}\mathbf{L})\mathbf{1}(\mathbf{U}\mathbf{L})^T, \quad (3.8)$$

where $\mathbf{1}$ is the identity matrix. Applying the Schur complement [27, §2.3.1] to (3.8) yields

$$\begin{bmatrix} \mathbf{W} & \mathbf{U}\mathbf{L} \\ \mathbf{L}^T\mathbf{U}^T & \mathbf{1} \end{bmatrix} > 0. \quad (3.9)$$

The LMI form of the optimization problem is therefore

$$\min J(\mathbf{U}, \nu, \mathbf{W}) = c - 2\text{tr}(\mathbf{U}\mathbf{G}^T) + \nu \quad (3.10)$$

$$\text{s.t. } \text{tr}(\mathbf{W}) < \nu, \quad \mathbf{W} > 0, \quad \begin{bmatrix} \mathbf{W} & \mathbf{U}\mathbf{L} \\ \mathbf{L}^T\mathbf{U}^T & \mathbf{1} \end{bmatrix} > 0, \quad (3.11)$$

where $\mathbf{H} = \mathbf{L}\mathbf{L}^T$.

As previously mentioned, the decomposition $\mathbf{H} = \mathbf{L}\mathbf{L}^T$ can be realized via a Cholesky factorization, eigendecomposition, or singular value decomposition. Using a Cholesky factorization directly gives $\mathbf{L}$. When using an eigendecomposition, $\mathbf{H} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, it follows that $\mathbf{L} = \mathbf{V}\sqrt{\mathbf{\Lambda}}$. Alternatively, using a singular value decomposition, $\Psi = \mathbf{Q}\mathbf{\Sigma}\mathbf{Z}^T$, and substituting it into the definition of $\mathbf{H}$ in (3.2) yields $\mathbf{H} = \frac{1}{q}\mathbf{Q}\mathbf{\Sigma}^2\mathbf{Q}^T$. It follows that $\mathbf{L} = \frac{1}{\sqrt{q}}\mathbf{Q}\mathbf{\Sigma}$. A singular value decomposition is used to compute $\mathbf{L}$ in the experiments presented in this paper.

4. Asymptotic Stability Constraint

Since many systems of interest have asymptotically stable dynamics, it is desirable to identify Koopman systems that share this property. In [13], it is proven that an asymptotically stable nonlinear system can only be represented accurately by an asymptotically stable Koopman system, thus highlighting the importance of enforcing the asymptotic stability property during the regression process.

However, in practice, it is possible to identify an unstable Koopman system from measurements of an asymptotically stable system [13]. Even if the identified Koopman system is asymptotically stable in theory, the eigenvalues of its $\mathbf{A}$ matrix may be so close to the unit circle that it is effectively unstable in practice. One solution, presented in this section, is to constrain the largest eigenvalue of $\mathbf{A}$ to be strictly less than one in magnitude, within a desired tolerance.

(a) Constraint formulation

To ensure that the system defined by the Koopman matrix $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$ is asymptotically stable, the eigenvalues of $\mathbf{A}$ must be constrained to have magnitude strictly less than one. A modified Lyapunov constraint [28, §1.4.4]

$$\mathbf{P} > 0, \quad (4.1)$$

$$\mathbf{A}^T\mathbf{P}\mathbf{A} - \bar{\rho}^2\mathbf{P} < 0, \quad (4.2)$$

can be added to ensure that the magnitude of the largest eigenvalue of $\mathbf{A}$ is no larger than $0 < \bar{\rho} < 1$. Applying the Schur complement to (4.2) yields

$$\mathbf{A}^\top \mathbf{P} \mathbf{A} - \bar{\rho}^2 \mathbf{P} < 0 \iff (\mathbf{A}^\top \mathbf{P}) \mathbf{P}^{-1} (\mathbf{A}^\top \mathbf{P})^\top - \bar{\rho}^2 \mathbf{P} < 0 \quad (4.3)$$

$$\iff -\bar{\rho} \mathbf{P} - (-\mathbf{A}^\top \mathbf{P}) (-\bar{\rho} \mathbf{P})^{-1} (-\mathbf{A}^\top \mathbf{P})^\top < 0 \quad (4.4)$$

$$\iff \begin{bmatrix} -\bar{\rho} \mathbf{P} & -\mathbf{A}^\top \mathbf{P} \\ -\mathbf{P}^\top \mathbf{A} & -\bar{\rho} \mathbf{P} \end{bmatrix} < 0, \quad -\bar{\rho} \mathbf{P} < 0 \quad (4.5)$$

$$\iff \begin{bmatrix} \bar{\rho} \mathbf{P} & \mathbf{A}^\top \mathbf{P} \\ \mathbf{P}^\top \mathbf{A} & \bar{\rho} \mathbf{P} \end{bmatrix} > 0, \quad \mathbf{P} > 0. \quad (4.6)$$

The full optimization problem with asymptotic stability constraint is therefore

$$\min J(\mathbf{U}, \nu, \mathbf{W}, \mathbf{P}; \bar{\rho}) = c - 2\text{tr}(\mathbf{U} \mathbf{G}^\top) + \nu \quad (4.7)$$

$$\text{s.t. } \text{tr}(\mathbf{W}) < \nu, \quad \mathbf{W} > 0, \quad \begin{bmatrix} \mathbf{W} & \mathbf{U} \mathbf{L} \\ \mathbf{L}^\top \mathbf{U}^\top & \mathbf{1} \end{bmatrix} > 0, \quad \mathbf{P} > 0, \quad \begin{bmatrix} \bar{\rho} \mathbf{P} & \mathbf{A}^\top \mathbf{P} \\ \mathbf{P}^\top \mathbf{A} & \bar{\rho} \mathbf{P} \end{bmatrix} > 0, \quad (4.8)$$

where $\mathbf{H} = \mathbf{L} \mathbf{L}^\top$ and $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$.

Since both $\mathbf{A}$ and $\mathbf{P}$ are unknown, (4.6) includes a BMI constraint. The optimization problem is therefore nonconvex and NP-hard. One method to find a locally optimal solution is outlined in [29]. First, assume $\mathbf{P}$ is constant with an initial guess of $\mathbf{P} = \mathbf{1}$ and solve (4.7)–(4.8) as an LMI problem. Then, hold the remaining variables constant using the solution from that optimization, and solve a feasibility problem for $\mathbf{P}$. Repeat this process until the cost function stops changing significantly. Although this approach is rather simple, it does result in an asymptotically stable Koopman system with reasonable prediction error.

(b) Experimental results

The effectiveness of the proposed asymptotic stability constraint is demonstrated using experimental data collected from the Fatigue Structural Testing Equipment Research (FASTER) platform at the National Research Council of Canada (NRC) [30], which is used for aircraft fatigue testing research. In the fatigue structural testing dataset used in this paper, the structure under test is an aluminum-composite beam. During a test, the FASTER platform applies a force to this beam using a hydraulic actuator, which is controlled by a voltage. This input, the applied force, and the structure's deflection are recorded at 128 Hz. The actuator voltage is determined by a linear controller designed to track a reference force profile. In this case, the dynamics of the controller are neglected and the actuator voltage is considered to be exogenous. All states and inputs in the FASTER dataset are normalized. A photograph of this experimental setup can be found in [30].

The Koopman lifting functions chosen for the FASTER platform are first- and second-order monomials. The full lifting procedure consists of several steps to improve numerical conditioning. First, all states and inputs are normalized so that they do not grow when passed through the monomial lifting functions. Then, all first- and second-order monomials of the state and input are computed. Finally, the lifted states are standardized to ensure that they are evenly weighted in the regression. To standardize the lifted states, their means are subtracted and they are rescaled to have unit variance [31, §4.6.6]. When using the identified Koopman matrices for prediction, the state and input are lifted, then multiplied by the Koopman matrix. The state is then recovered and re-lifted with the input for the next timestep. Using this prediction method leads to the local error definition presented in [13], which is used throughout this paper.

Given that the true FASTER system is asymptotically stable, it is crucial to ensure that the identified Koopman system shares this property. To demonstrate the impact of the asymptotic stability constraint, Koopman matrices with maximum spectral radii of $\bar{\rho} = 1.00$ and $\bar{\rho} = 0.99$ are




Figure 2: (a) Eigenvalues and spectral radius constraints of Koopman $\mathbf{A}$ matrices approximated from the FASTER dataset. The eigenvalues of $\mathbf{A}$ satisfy their respective spectral radius constraints. Additionally, lowering the spectral radius constraint from $\bar{\rho} = 1.00$ to $\bar{\rho} = 0.99$ does not significantly alter the eigenvalues. (b) Multi-step prediction error of Koopman systems approximated from the FASTER dataset. All units are normalized. States are recovered and re-lifted between prediction timesteps. Although both systems are asymptotically stable, only the system with $\bar{\rho} = 0.99$ is usable in practice, as the other system's response diverges due to the accumulation of numerical error.

computed and compared to an unconstrained Koopman matrix computed with standard EDMD. Figure 2a shows the eigenvalues of the two constrained Koopman $\mathbf{A}$ matrices. In both cases, the eigenvalues indicate that the systems obey their respective maximum spectral radii. However, the multi-step prediction errors of the two Koopman systems in Figure 2b show that only the system with the spectral radius constraint of $\bar{\rho} = 0.99$ is usable in practice. The other system produces prediction errors that diverge to infinity due to the accumulation of numerical error. The Koopman matrix without asymptotic stability constraint behaves identically to the matrix with maximum spectral radius $\bar{\rho} = 1.00$, and is therefore not shown Figure 2. By introducing a small amount of conservatism in the spectral radius constraint, the identified Koopman system is rendered asymptotically stable in the face of accumulated numerical error.

5. System Norm Regularization

While constraining the asymptotic stability of the identified Koopman system helps ensure that the system's predictions are usable, it does not consider the input-output properties of the system. A system norm like the $\mathcal{H}_2$ norm or the $\mathcal{H}_\infty$ norm can be used to regularize the system's gain, while also ensuring asymptotic stability for any regularization coefficient. Specifically, the existence of a finite $\mathcal{H}_2$ or $\mathcal{H}_\infty$ norm guarantees the asymptotic stability of the resulting linear time-invariant (LTI) system [32]. Regularizing using the $\mathcal{H}_2$ norm can be thought of as penalizing the average system gain over all frequencies, while using the $\mathcal{H}_\infty$ norm penalizes the worst-case system gain. As such, the use of the $\mathcal{H}_\infty$ norm as a regularizer is explored next. Weighted

$\mathcal{H}_\infty$ norms are also considered as regularizers, which allow the regularization problem to be tuned in the frequency domain with weighting functions.

(a) $\mathcal{H}_\infty$ norm regularization

Since approximating the Koopman matrix amounts to identifying a discrete-time LTI system, it is natural to consider a system norm as a regularizer rather than a matrix norm. The Koopman representation of a nonlinear ODE can be thought of as a discrete-time LTI system $\mathcal{G} : \ell_{2e} \rightarrow \ell_{2e}$, where ℓ_{2e} is the extended inner product sequence space [33, §B.1.1] and $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$. Consider the lifted output equation,

$$\zeta_k = \mathbf{C}\vartheta_k + \mathbf{D}v_k, \quad (5.1)$$

where $\mathbf{C} \in \mathbb{R}^{p_\zeta \times p_\vartheta}$ and $\mathbf{D} \in \mathbb{R}^{p_\zeta \times p_v}$. In the simplest formulation, $\mathbf{C} = \mathbf{1}$, and $\mathbf{D} = \mathbf{0}$. The Koopman system is then

$$\mathcal{G} \stackrel{\min}{\sim} \left[\begin{array}{c|c} \mathbf{A} & \mathbf{B} \\ \hline \mathbf{C} & \mathbf{D} \end{array} \right], \quad (5.2)$$

where $\stackrel{\min}{\sim}$ denotes a minimal state-space realization [33, §3.2.1]. The Koopman system $\mathcal{G}$ has a corresponding discrete-time transfer matrix [32, §3.7], $\mathbf{G}(z) = \mathbf{C}(z\mathbf{1} - \mathbf{A})^{-1}\mathbf{B} + \mathbf{D}$.

The $\mathcal{H}_\infty$ norm of $\mathcal{G}$ is the worst-case gain from $\|v\|_2$ to $\|\mathcal{G}v\|_2$. That is [33, §B.1.1]

$$\|\mathcal{G}\|_\infty = \sup_{v \in \ell_2, v \neq \mathbf{0}} \frac{\|\mathcal{G}v\|_2}{\|v\|_2}, \quad (5.3)$$

where ℓ_2 is the inner product sequence space [33, §B.1.1]. In the frequency domain, this definition is equivalent to [33, §B.1.1]

$$\|\mathcal{G}\|_\infty = \sup_{\theta \in (-\pi, \pi]} \bar{\sigma}(\mathbf{G}(e^{j\theta})), \quad (5.4)$$

where $\bar{\sigma}(\cdot)$ denotes the maximum singular value of a matrix. In (5.4), the transfer function is evaluated at $z = e^{j\theta}$, where $\theta = 2\pi\Delta t f$ is the discrete-time frequency, Δt is the sampling timestep, and f is the continuous-time frequency.

With $\mathcal{H}_\infty$ norm regularization, the cost function associated with the regression problem is

$$J(\mathbf{U}; \beta) = \|\boldsymbol{\Theta}_+ - \mathbf{U}\boldsymbol{\Psi}\|_F^2 + \beta\|\mathcal{G}\|_\infty, \quad (5.5)$$

where β is the regularization coefficient. To integrate the $\mathcal{H}_\infty$ norm into the regression problem, its LMI formulation must be considered. The inequality $\|\mathcal{G}\|_\infty < \gamma$ holds if and only if [27, §3.2.2]

$$\mathbf{P} > 0, \quad \begin{bmatrix} \mathbf{P} & \mathbf{A}\mathbf{P} & \mathbf{B} & \mathbf{0} \\ \mathbf{P}^\top \mathbf{A}^\top & \mathbf{P} & \mathbf{0} & \mathbf{P}\mathbf{C}^\top \\ \mathbf{B}^\top & \mathbf{0} & \gamma\mathbf{1} & \mathbf{D}^\top \\ \mathbf{0} & \mathbf{C}\mathbf{P}^\top & \mathbf{D} & \gamma\mathbf{1} \end{bmatrix} > 0. \quad (5.6)$$

The full optimization problem with $\mathcal{H}_\infty$ regularization is

$$\min J(\mathbf{U}, \nu, \mathbf{W}, \gamma, \mathbf{P}; \beta) = c - 2\text{tr}(\mathbf{U}\mathbf{G}^\top) + \nu + \frac{\beta}{q}\gamma \quad (5.7)$$

$$\text{s.t. } \text{tr}(\mathbf{W}) < \nu, \quad \mathbf{W} > 0, \quad \begin{bmatrix} \mathbf{W} & \mathbf{U}\mathbf{L} \\ \mathbf{L}^\top \mathbf{U}^\top & \mathbf{1} \end{bmatrix} > 0, \quad (5.8)$$

$$\mathbf{P} > 0, \quad \begin{bmatrix} \mathbf{P} & \mathbf{A}\mathbf{P} & \mathbf{B} & \mathbf{0} \\ \mathbf{P}^\top \mathbf{A}^\top & \mathbf{P} & \mathbf{0} & \mathbf{P}\mathbf{C}^\top \\ \mathbf{B}^\top & \mathbf{0} & \gamma\mathbf{1} & \mathbf{D}^\top \\ \mathbf{0} & \mathbf{C}\mathbf{P}^\top & \mathbf{D} & \gamma\mathbf{1} \end{bmatrix} > 0, \quad (5.9)$$

where $\mathbf{H} = \mathbf{L}\mathbf{L}^\top$ and $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$.

Like the asymptotic stability constraint proposed in Section 4, (5.9) includes a BMI constraint in terms of the unknowns $\mathbf{P}$ and $\mathbf{A}$. As such, the optimization problem in (5.7)–(5.9) is nonconvex and NP-hard. However, it can be solved using the same iterative procedure described in Section 4 [29].

(b) Weighted $\mathcal{H}_\infty$ norm regularization

The $\mathcal{H}_\infty$ norm used in (5.7) can be weighted by cascading $\mathcal{G}$ with another LTI system, $\mathcal{G}^w$. For example, choosing $\mathcal{G}^w$ to be a highpass filter penalizes system gains at high frequencies. Weights can be cascaded before or after $\mathcal{G}$ to weight the inputs, outputs, or both. Recall that multi-input multi-output LTI systems do not commute.

Consider the weight

$$\mathcal{G}^w \stackrel{\min}{\sim} \left[\begin{array}{c|c} \mathbf{A}^w & \mathbf{B}^w \\ \hline \mathbf{C}^w & \mathbf{D}^w \end{array} \right], \quad (5.10)$$

with state ϑ^w , input v^w , and output ζ^w . Cascading $\mathcal{G}^w$ after $\mathcal{G}$ yields the augmented state-space system

$$\begin{bmatrix} \vartheta_{k+1} \\ \vartheta_{k+1}^w \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{B}^w \mathbf{C} & \mathbf{A}^w \end{bmatrix} \begin{bmatrix} \vartheta_k \\ \vartheta_k^w \end{bmatrix} + \begin{bmatrix} \mathbf{B} \\ \mathbf{B}^w \mathbf{D} \end{bmatrix} v_k, \quad (5.11)$$

$$\zeta_k^w = \begin{bmatrix} \mathbf{D}^w \mathbf{C} & \mathbf{C}^w \end{bmatrix} \begin{bmatrix} \vartheta_k \\ \vartheta_k^w \end{bmatrix} + \mathbf{D}^w \mathbf{D} v_k. \quad (5.12)$$

Minimizing the $\mathcal{H}_\infty$ norm of the augmented system

$$\mathcal{G}^w \mathcal{G} \stackrel{\min}{\sim} \left[\begin{array}{cc|c} \mathbf{A} & \mathbf{0} & \mathbf{B} \\ \mathbf{B}^w \mathbf{C} & \mathbf{A}^w & \mathbf{B}^w \mathbf{D} \\ \hline \mathbf{D}^w \mathbf{C} & \mathbf{C}^w & \mathbf{D}^w \mathbf{D} \end{array} \right] \quad (5.13)$$

is equivalent to minimizing the weighted $\mathcal{H}_\infty$ norm of the original system. The choice of weighting function used can be viewed as another hyperparameter in the regression problem.

Weighting the regression problem in the frequency domain comes at the cost of increasing the dimension of the optimization problem. When cascading the weight before $\mathcal{G}$, the dimension of ϑ^w scales with the dimension of v . When cascading after $\mathcal{G}$, the dimension of ϑ^w scales with the dimension of ζ . In the regression problems considered here, only post-weighting is considered, since ζ has a much smaller dimension.

(c) Experimental results

The unique advantages of the $\mathcal{H}_\infty$ norm regularizer are demonstrated using the soft robot dataset published alongside [34] and [9]. Unregularized EDMD and Tikhonov-regularized EDMD [14,25], two standard Koopman matrix approximation methods, are compared with the asymptotic stability constraint from Section 4 and the $\mathcal{H}_\infty$ norm regularizer presented in this section. The Koopman systems identified using these regression methods are analyzed in terms of their prediction errors, system properties, and numerical conditioning.

Identifying a Koopman representation of a soft robot arm is a particularly interesting problem, as its dynamics are not easily modelled from first principles. The soft robot under consideration consists of two flexible segments with a laser pointer mounted at the end. The laser pointer projects a dot onto a board positioned below the robot. The two states of the system are the Cartesian coordinates of the dot on the board, as measured by a camera. The soft robot arm is actuated by three pressure regulators, each controlled by a voltage. The dot position and control voltages are recorded at 12 Hz. Thirteen training episodes and four test episodes were recorded in this manner. The third test episode is shown in Figure 3. Photographs of this experimental setup, along with additional details, can be found in [34] and [9].




Figure 3: Multi-step prediction error and trajectory plot of the third test episode for Koopman systems approximated from the soft robot dataset. States are recovered and re-lifted between prediction timesteps. All Koopman systems have comparable prediction errors, with the exception of two large error spikes in the system with the asymptotic stability constraint.




Figure 4: Eigenvalues of Koopman A matrices approximated from the soft robot dataset. EDMD without regularization and EDMD with Tikhonov regularization both identify unstable systems, while the asymptotic stability constraint and $\mathcal{H}_\infty$ norm regularizer yield asymptotically stable Koopman systems.




Figure 5: Singular values of Koopman **A** and **B** matrices approximated from the soft robot dataset plotted on a logarithmic scale. The **A** and **B** matrices computed using unregularized EDMD both have large singular values. EDMD with Tikhonov regularization decreases the singular values of both matrices, but does not yield an asymptotically stable system. With an asymptotic stability constraint, the singular values of **A** are reduced, but the singular values of **B** are unaffected and remain large. Using an $\mathcal{H}_\infty$ norm regularizer reduces the singular values of both **A** and **B** significantly, yielding a better-conditioned Koopman matrix.

The lifting functions chosen for the soft robot system consist of a time delay step, followed by a third-order monomial transformation. Although time delays do not, strictly speaking, meet the definition of a lifting function outlined in Section 2, they are commonly used in the lifted states of Koopman identification problems [5,9]. As with the FASTER dataset in Section 4, the states and inputs are first normalized. Then, the states and inputs are augmented with their delayed versions, where the delay period is one timestep. Next, all first-, second-, and third-order monomials are computed. Finally, the lifted states are standardized. Since the time delay step occurs before the monomial lifting step, cross-terms including delayed and non-delayed states and inputs occur in the lifted state.

Unregularized EDMD and Tikhonov-regularized EDMD are used as baselines for comparison with the proposed regression methods. Tikhonov regularization improves the numerical conditioning of **U** by penalizing its squared Frobenius norm [14,25]. The regularizers used in this section have coefficients of $\beta = 7.5 \times 10^{-3}$, while the asymptotic stability constraint has a maximum spectral radius of $\bar{\rho} = 0.999$. Figure 3 shows the multi-step prediction errors of the four Koopman systems for the third test episode of the dataset. The prediction errors are comparable for all four Koopman systems, aside from two large error spikes produced by the system with the asymptotic stability constraint.

One way to compare the resulting Koopman systems is to analyze the eigenvalues of their **A** matrices. Figure 4 shows that unregularized EDMD and Tikhonov-regularized EDMD produce unstable Koopman systems, even though, as shown in Figure 7, the multi-step prediction errors do not happen to diverge in any test episodes. As expected, EDMD with an $\mathcal{H}_\infty$ norm regularizer and EDMD with an asymptotic stability constraint both yield asymptotically stable Koopman systems.

While the Koopman system identified with an asymptotic stability constraint is indeed asymptotically stable, the system is not well-conditioned. To see this, consider Figure 5, which shows the singular values of the Koopman **A** and **B** matrices. These singular values indicate the




Figure 6: (a) Bode plot of Koopman systems approximated from the soft robot dataset. The Koopman system identified with unregularized EDMD has high gain, with a resonant peak at a high frequency. EDMD with Tikhonov regularization reduces the system's gain, but identifies an unstable system and retains an undesirable high-frequency resonant peak in the frequency response. Constraining the asymptotic stability of the system does not significantly reduce the gain compared to the unregularized system. Penalizing the $\mathcal{H}_\infty$ norm of the Koopman system reduces its gain at all frequencies without compromising prediction error. (b) Bode plot of unweighted and weighted Koopman systems approximated from the soft robot dataset, along with weighting function. The dashed line representing the weighting function uses the right axis, while the solid lines use the left axis.

sizes of the entries in each matrix. With unregularized EDMD, both matrices have singular values on the order of 10^3 . Using Tikhonov regularization decreases the singular values of both $\mathbf{A}$ and $\mathbf{B}$, though it still yields an unstable system. Constraining the spectral radius of $\mathbf{A}$ greatly reduces the singular values of $\mathbf{A}$ but increases the singular values of $\mathbf{B}$. The numerical conditioning of this Koopman system is arguably worse, as the $\mathbf{A}$ and $\mathbf{B}$ matrices contain entries of drastically different scales. Regularizing using the $\mathcal{H}_\infty$ norm resolves this problem, as it reduces the singular values in both matrices, yielding a better-conditioned, asymptotically stable Koopman system with similar prediction error. The key takeaway is that constraining the spectral radius of $\mathbf{A}$ is not sufficient to guarantee a well-conditioned Koopman matrix, as the constraint does not directly impact $\mathbf{B}$. Using the $\mathcal{H}_\infty$ norm as a regularizer considers the system as a whole, thus impacting both $\mathbf{A}$ and $\mathbf{B}$, and reducing their entries to reasonable sizes.

Another way to compare the identified Koopman systems is by looking at their frequency responses, which can be found by plotting the maximum singular value of the transfer matrix at each frequency. Figure 6a shows the magnitude responses of the four Koopman systems, and paints a similar picture to Figure 5. Unregularized EDMD yields a Koopman system with very high gain and a resonant peak in the upper frequency range. Incorporating Tikhonov regularization reduces the system's gain, but retains the resonant peak at a high frequency. Constraining the asymptotic stability of the system does not significantly impact the system's gain, which, given the large singular values of $\mathbf{B}$ in Figure 5, is not surprising. However, regularizing using the $\mathcal{H}_\infty$ norm directly penalizes the peak of the Bode plot in Figure 6a, yielding a system with significantly lower gain and similar prediction error.

The $\mathcal{H}_\infty$ norm regularizer can be expanded upon by weighting the $\mathcal{H}_\infty$ norm using a highpass filter. This weighting function penalizes high gains at high frequencies while allowing higher gains at low frequencies. Penalizing gain at high frequencies is desirable because relevant system dynamics typically occupy low frequencies, while high frequencies are corrupted by measurement noise. Furthermore, causal physical systems have frequency responses that roll off as frequency grows, since it is unrealistic for a system to have infinite gain at infinitely high frequencies. Figure 6b demonstrates the impact of weighting the $\mathcal{H}_\infty$ norm regularizer with a highpass filter that has a zero at 4 Hz and a pole slightly below 6 Hz. This weighted regularizer yields a Koopman system with high gain at low frequencies and decreasing gain at high frequencies.

Note that Figure 6 shows the frequency response of the Koopman system in the lifted space, which is not the same as the “frequency response of the nonlinear system.” Since the ultimate goal is to design linear controllers in the lifted space, only the frequency response of the Koopman system in the lifted space and the corresponding $\mathcal{H}_\infty$ norm are relevant.




Figure 7: RMS Euclidean errors of Koopman systems approximated from the soft robot dataset. Error bars indicate mean and standard deviation of RMS error over the four test episodes. In terms of mean error, all identified Koopman systems perform similarly well. However, the system identified using $\mathcal{H}_\infty$ norm regularizer performed more consistently throughout the test set.

In Figure 7, the multi-step prediction errors of the five identified Koopman systems are compared across each episode in the test set. Given that the laser pointer dot projected by the soft robot moves within a circle of radius 10 cm, it's practically meaningless to distinguish each method based on mean prediction error alone. However, the distribution of each identified system's prediction errors across the test set clearly highlights the importance of regularization. The Tikhonov regularizer, $\mathcal{H}_\infty$ norm regularizer, and weighted $\mathcal{H}_\infty$ norm regularizer lead to systems with smaller standard deviations in the RMS error over the test set. In contrast, EDMD without regularization and EDMD with an asymptotic stability constraint lead to systems that perform inconsistently over the test set. Comparing standard deviations indicates that the regularized systems generalize better to previously unseen data, with the $\mathcal{H}_\infty$ norm regularizer performing most consistently over the four test episodes.

The results in this section highlight the desirable properties of the proposed $\mathcal{H}_\infty$ norm regularizers. While all systems have comparable RMS prediction errors on the test set, unregularized EDMD results in an unstable, poorly conditioned Koopman system with large gain. While Tikhonov regularization improves numerical conditioning, the resulting system is still unstable. Conversely, constraining the asymptotic stability of the system does not improve its numerical properties. Only the $\mathcal{H}_\infty$ norm regularizers guarantee asymptotic stability while improving the numerical conditioning of the system. These key results are summarized in Table 1, which also highlights the difference between the unweighted and weighted $\mathcal{H}_\infty$ norm regularizers. Weighting the $\mathcal{H}_\infty$ norm regularizer further improves the condition numbers of $\mathbf{A}$ and $\mathbf{B}$.

Table 1: Comparison of regression methods through the condition numbers of their Koopman matrices and asymptotic stability guarantees. Only the $\mathcal{H}_\infty$ regularizers guarantee asymptotic stability while significantly improving the condition number of the Koopman matrices. In this case, weighting the $\mathcal{H}_\infty$ norm further improves $\text{cond}(\mathbf{A})$ and $\text{cond}(\mathbf{B})$.

regression method	$\text{cond}(\mathbf{A})$	$\text{cond}(\mathbf{B})$	asymptotic stability
no regularization	5.77×10^7	3.40×10^7	no
Tikhonov regularization	4.39×10^5	2.90×10^3	no
asymptotic stability constraint	7.32×10^4	4.87×10^3	yes
$\mathcal{H}_\infty$ regularization	3.87×10^4	2.14×10^2	yes
weighted $\mathcal{H}_\infty$ regularization	1.69×10^3	5.43×10^1	yes

6. Reducing the Size of the Regression Problem

As the number of lifting functions required grows, so does the size of the optimization problem. With several hundred lifting functions, finding a solution can take days and consume an intractable amount of memory. To combat this limitation, an approach reminiscent of DMDc is now presented.

(a) DMD with control

DMD with control [17] reduces the dimension of the Koopman matrix regression problem when the dataset contains many more lifted states than time snapshots (i.e., $p \gg q$) [12, §10.3]. In DMDc, the Koopman matrix is projected onto the left singular vectors of Θ_+ . The size of the problem is then controlled by retaining only the $\hat{r}$ largest singular values in the SVD. Consider the truncated singular value decomposition $\Theta_+ \approx \hat{\mathbf{Q}}\hat{\mathbf{\Sigma}}\hat{\mathbf{Z}}^T$, where $\hat{\mathbf{Q}} \in \mathbb{R}^{p_\theta \times \hat{r}}$, $\hat{\mathbf{\Sigma}} \in \mathbb{R}^{\hat{r} \times \hat{r}}$, and $\hat{\mathbf{Z}} \in \mathbb{R}^{q \times \hat{r}}$. Instead of solving for $\mathbf{U} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix}$, the regression problem is written in terms of [17]

$$\hat{\mathbf{U}} = \begin{bmatrix} \hat{\mathbf{Q}}^T \mathbf{A} \hat{\mathbf{Q}} & \hat{\mathbf{Q}}^T \mathbf{B} \end{bmatrix} = \hat{\mathbf{Q}}^T \mathbf{U} \begin{bmatrix} \hat{\mathbf{Q}} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix}, \quad (6.1)$$

where $\hat{\mathbf{U}} \in \mathbb{R}^{\hat{r} \times \hat{r} + p_v}$ is significantly smaller than $\mathbf{U} \in \mathbb{R}^{p_\theta \times p}$. The least-squares solution to the Koopman matrix (2.14) can be written as

$$\mathbf{U} = \Theta_+ \tilde{\mathbf{Z}} \tilde{\mathbf{\Sigma}}^\dagger \tilde{\mathbf{Q}}^T, \quad (6.2)$$

where $\Psi \approx \tilde{\mathbf{Q}} \tilde{\mathbf{\Sigma}} \tilde{\mathbf{Z}}^T$. The number of singular values retained in this SVD is denoted $\tilde{r}$. Thus $\tilde{\mathbf{Q}} \in \mathbb{R}^{p \times \tilde{r}}$, $\tilde{\mathbf{\Sigma}} \in \mathbb{R}^{\tilde{r} \times \tilde{r}}$, and $\tilde{\mathbf{Z}} \in \mathbb{R}^{q \times \tilde{r}}$. The standard solution to the DMDc problem is obtained by

substituting (6.2) into (6.1), yielding

$$\hat{\mathbf{U}} = \hat{\mathbf{Q}}^T \boldsymbol{\Theta}_+ \tilde{\mathbf{Z}} \tilde{\boldsymbol{\Sigma}}^\dagger \tilde{\mathbf{Q}}^T \begin{bmatrix} \hat{\mathbf{Q}} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix}. \quad (6.3)$$

The SVD dimensions $\hat{r}$ and $\tilde{r}$ are a design choice. A common approach is the hard-thresholding algorithm described in [35]. Note that typically $\hat{r} < \tilde{r}$ [12, §6.1.3].

(b) LMI reformulation of DMDc

To reformulate DMDc as a convex optimization problem with LMI constraints, the cost function is rewritten in terms of $\hat{\mathbf{U}}$ instead of $\mathbf{U}$. Recall that the rescaled Koopman cost function (3.3) can be written as

$$\min J(\mathbf{U}) = \frac{1}{q} \text{tr} \left(\boldsymbol{\Theta}_+ \boldsymbol{\Theta}_+^T - \text{He} \left\{ \mathbf{U} \boldsymbol{\Psi} \boldsymbol{\Theta}_+^T \right\} + \mathbf{U} \boldsymbol{\Psi} \boldsymbol{\Psi}^T \mathbf{U}^T \right). \quad (6.4)$$

With the introduction of a slack variable, (6.4) becomes [27, §2.15.1]

$$\min J(\mathbf{U}, \mathbf{W}) = \frac{1}{q} \text{tr}(\mathbf{W}) \quad (6.5)$$

$$\text{s.t. } \mathbf{W} > 0, \quad \boldsymbol{\Theta}_+ \boldsymbol{\Theta}_+^T - \text{He} \left\{ \mathbf{U} \boldsymbol{\Psi} \boldsymbol{\Theta}_+^T \right\} + \mathbf{U} \boldsymbol{\Psi} \boldsymbol{\Psi}^T \mathbf{U}^T < \mathbf{W}. \quad (6.6)$$

Substituting the SVDs of $\boldsymbol{\Theta}_+$ and $\boldsymbol{\Psi}$ into the optimization problem yields

$$\min J(\mathbf{U}, \mathbf{W}) = \frac{1}{q} \text{tr}(\mathbf{W}) \quad (6.7)$$

$$\text{s.t. } \mathbf{W} > 0, \quad \hat{\mathbf{Q}} \hat{\boldsymbol{\Sigma}}^2 \hat{\mathbf{Q}}^T - \text{He} \left\{ \mathbf{U} \tilde{\mathbf{Q}} \tilde{\boldsymbol{\Sigma}} \tilde{\mathbf{Z}}^T \hat{\mathbf{Z}} \hat{\boldsymbol{\Sigma}} \hat{\mathbf{Q}}^T \right\} + \mathbf{U} \tilde{\mathbf{Q}} \tilde{\boldsymbol{\Sigma}}^2 \tilde{\mathbf{Q}}^T \mathbf{U}^T < \mathbf{W}. \quad (6.8)$$

Next, consider the projection of (6.7) and (6.8) onto the column space of $\hat{\mathbf{Q}}$, denoted $\mathcal{R}(\hat{\mathbf{Q}})$. Recall that $\mathbf{W} > 0$ is equivalent to

$$\boldsymbol{\vartheta}^T \mathbf{W} \boldsymbol{\vartheta} > 0, \quad \forall \boldsymbol{\vartheta} \neq \mathbf{0} \in \mathbb{R}^{p_\vartheta \times 1}. \quad (6.9)$$

Since (6.9) holds over all of $\mathbb{R}^{p_\vartheta \times 1}$, it must also hold over the subspace $\mathcal{R}(\hat{\mathbf{Q}})$. Let the vectors in $\mathcal{R}(\hat{\mathbf{Q}})$ be parameterized by

$$\boldsymbol{\vartheta} = \hat{\mathbf{Q}} \hat{\boldsymbol{\vartheta}}, \quad (6.10)$$

where $\hat{\boldsymbol{\vartheta}} \in \mathbb{R}^{\hat{r} \times 1}$. Substituting (6.10) into (6.9) yields

$$\hat{\boldsymbol{\vartheta}}^T \hat{\mathbf{Q}}^T \mathbf{W} \hat{\mathbf{Q}} \hat{\boldsymbol{\vartheta}} > 0, \quad \forall \hat{\boldsymbol{\vartheta}} \neq \mathbf{0} \in \mathbb{R}^{\hat{r} \times 1}, \quad (6.11)$$

which is equivalent to

$$\hat{\mathbf{W}} = \hat{\mathbf{Q}}^T \mathbf{W} \hat{\mathbf{Q}} > 0 \quad (6.12)$$

over $\mathcal{R}(\hat{\mathbf{Q}})$. Applying the same logic to (6.8) yields

$$\hat{\boldsymbol{\Sigma}}^2 - \text{He} \left\{ \hat{\mathbf{Q}}^T \mathbf{U} \tilde{\mathbf{Q}} \tilde{\boldsymbol{\Sigma}} \tilde{\mathbf{Z}}^T \hat{\mathbf{Z}} \hat{\boldsymbol{\Sigma}} \right\} + \hat{\mathbf{Q}}^T \mathbf{U} \tilde{\mathbf{Q}} \tilde{\boldsymbol{\Sigma}}^2 \tilde{\mathbf{Q}}^T \mathbf{U}^T \hat{\mathbf{Q}} < \hat{\mathbf{W}}, \quad (6.13)$$

where the fact that $\hat{\mathbf{Q}}^T \hat{\mathbf{Q}} = \mathbf{1}$ has been used.

To further simplify the problem, it is advantageous to rewrite (6.7) in terms of $\hat{\mathbf{W}}$. To accomplish this, first recall that the trace of a matrix is equal to the sum of its eigenvalues. The eigenvalue problem for $\mathbf{W}$ is

$$\mathbf{W} \mathbf{v}_i = \lambda_i \mathbf{v}_i. \quad (6.14)$$

Projecting (6.14) onto $\mathcal{R}(\hat{\mathbf{Q}})$ by substituting $\mathbf{v}_i = \hat{\mathbf{Q}} \hat{\mathbf{v}}_i$, then premultiplying the result by $\hat{\mathbf{Q}}^T$, yields

$$\hat{\mathbf{Q}}^T \mathbf{W} \hat{\mathbf{Q}} \hat{\mathbf{v}}_i = \lambda_i \hat{\mathbf{v}}_i. \quad (6.15)$$

Thus, $\mathbf{W}$ and $\hat{\mathbf{W}}$ share the same eigenvalues for eigenvectors in $\mathcal{R}(\hat{\mathbf{Q}})$, which indicates that minimizing $\text{tr}(\hat{\mathbf{W}})$ is equivalent to minimizing $\text{tr}(\mathbf{W})$ in that subspace.

The full regression problem projected onto $\mathcal{R}(\hat{\mathbf{Q}})$ is therefore

$$\min J(\mathbf{U}, \hat{\mathbf{W}}) = \frac{1}{q} \text{tr}(\hat{\mathbf{W}}) \quad (6.16)$$

$$\text{s.t. } \hat{\mathbf{W}} > 0, \quad \hat{\Sigma}^2 - \text{He}\{\hat{\mathbf{Q}}^\top \mathbf{U} \tilde{\mathbf{Q}} \tilde{\Sigma} \tilde{\mathbf{Z}}^\top \hat{\mathbf{Z}} \hat{\Sigma}\} + \hat{\mathbf{Q}}^\top \mathbf{U} \tilde{\mathbf{Q}} \tilde{\Sigma}^2 \tilde{\mathbf{Q}}^\top \mathbf{U}^\top \hat{\mathbf{Q}} < \hat{\mathbf{W}}. \quad (6.17)$$

Substituting $\hat{\mathbf{U}}$ from (6.1) into the optimization problem yields

$$\min J(\hat{\mathbf{U}}, \hat{\mathbf{W}}) = \frac{1}{q} \text{tr}(\hat{\mathbf{W}}) \quad (6.18)$$

$$\text{s.t. } \hat{\mathbf{W}} > 0, \quad \hat{\Sigma}^2 - \text{He}\{\hat{\mathbf{U}} \tilde{\mathbf{Q}} \tilde{\Sigma} \tilde{\mathbf{Z}}^\top \hat{\mathbf{Z}} \hat{\Sigma}\} + \hat{\mathbf{U}} \tilde{\mathbf{Q}} \tilde{\Sigma}^2 \tilde{\mathbf{Q}}^\top \hat{\mathbf{U}}^\top < \hat{\mathbf{W}}, \quad (6.19)$$

where

$$\tilde{\mathbf{Q}} = \begin{bmatrix} \hat{\mathbf{Q}} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix}^\top \hat{\mathbf{Q}}. \quad (6.20)$$

Applying the Schur complement to (6.19) yields the LMI formulation of DMDc,

$$\min J(\hat{\mathbf{U}}, \hat{\mathbf{W}}) = \frac{1}{q} \text{tr}(\hat{\mathbf{W}}) \quad (6.21)$$

$$\text{s.t. } \hat{\mathbf{W}} > 0, \quad \begin{bmatrix} -\hat{\mathbf{W}} + \hat{\Sigma}^2 - \text{He}\{\hat{\mathbf{U}} \tilde{\mathbf{Q}} \tilde{\Sigma} \tilde{\mathbf{Z}}^\top \hat{\mathbf{Z}} \hat{\Sigma}\} & \hat{\mathbf{U}} \tilde{\mathbf{Q}} \tilde{\Sigma} \\ \tilde{\Sigma} \tilde{\mathbf{Q}}^\top \hat{\mathbf{U}}^\top & -\mathbf{1} \end{bmatrix} < 0. \quad (6.22)$$

This is now a significantly smaller optimization problem, as its size is controlled by the truncation of the SVD of Θ_+ . Reducing the first dimension of $\hat{\mathbf{U}}$ also reduces the dimension of the slack variable $\hat{\mathbf{W}}$.

(c) Constraints and regularization

The projection in (6.10) defines a new Koopman system,

$$\hat{\boldsymbol{\vartheta}}_{k+1} = \hat{\mathbf{A}} \hat{\boldsymbol{\vartheta}}_k + \hat{\mathbf{B}} \mathbf{v}_k, \quad (6.23)$$

$$\zeta_k = \hat{\mathbf{C}} \hat{\boldsymbol{\vartheta}}_k + \mathbf{D} \mathbf{v}_k, \quad (6.24)$$

where $\hat{\mathbf{C}} = \mathbf{C} \hat{\mathbf{Q}}$. The asymptotic stability constraint discussed in Section 4 and the $\mathcal{H}_\infty$ norm regularizers discussed in Section 5 can be equally applied to the projected system in (6.23) and (6.24) to ensure that the projected system of smaller dimension has the desired stability and frequency response characteristics.

(d) Experimental results

The properties of the asymptotic stability constraint from Section 4 and the $\mathcal{H}_\infty$ norm regularizer from Section 5 are now compared when applied to the EDMD and DMDc regression problems. The same dataset and experimental setup as Section 5 is used here.

An important decision in the DMDc algorithm is the choice of singular value truncation method. Optimal hard singular value truncation [35] is used to determine $\hat{r}$, while $\tilde{r}$ is left at full rank. For the soft robot dataset, the optimal hard truncation algorithm retains only 14 of the 34 singular values of $\mathbf{A}$. Figure 8 demonstrates that the DMDc methods indeed reduce the dimensionality of the problem, while also showing that the remaining singular values are close to their EDMD counterparts. Figure 9a shows that the frequency responses of the original Koopman systems are preserved by the DMDc methods. In spite of their reduced dimensionality, the Koopman systems identified with DMDc retain their frequency domain properties.

In Figure 9b, the RMS Euclidean errors of the EDMD and DMDc methods with asymptotic stability constraints and $\mathcal{H}_\infty$ norm regularization are summarized. Since the Koopman systems identified by the DMDc methods are of a lower order, their mean prediction error is higher




Figure 8: Singular values of Koopman **A** and **B** matrices approximated from the soft robot dataset using EDMD and DMDc regressors. Singular values smaller than 10^{-12} are not shown. Note the logarithmic scale. While the EDMD methods retain all 34 singular values, the DMDc methods truncate all but the first 14. The singular values retained by the DMDc methods are close to the corresponding singular values computed by the EDMD methods.




Figure 9: (a) Bode plots of Koopman systems approximated from the soft robot dataset using EDMD and DMDc regressors. The DMDc methods preserve the frequency responses of the corresponding systems identified with the EDMD methods. (b) RMS Euclidean errors of Koopman systems approximated from the soft robot dataset using EDMD and DMDc regressors. Error bars indicate standard deviation of RMS error over the four test episodes. Since the DMDc methods identify reduced-order Koopman models of the system, they have larger mean errors. However, the $\mathcal{H}_\infty$ norm regularizer still significantly reduces the standard deviation.

than that of the EDMD methods. However, the $\mathcal{H}_\infty$ norm regularizer retains its benefit of improving prediction consistency. Furthermore, as demonstrated by Figure 9a, the frequency response properties of the EDMD regression methods are preserved by the reduced-order DMDc models.

Note that, in one case, the Koopman system identified using stability-constrained DMDc diverged due to poor numerical conditioning. This short segment of the dataset was omitted throughout the paper to allow for a more fair comparison between regression methods. This finding highlights the advantages of the $\mathcal{H}_\infty$ regularization method in identifying numerically well-conditioned Koopman matrices.




Figure 10: (a) Execution time per iteration and (b) peak memory consumption of EDMD and DMDc regression methods using the soft robot dataset. The DMDc methods run significantly faster and consume less memory than the EDMD methods. Tests were run on a PC with an Intel Core i7-10700K processor using the MOSEK solver.

The most important advantage of the DMDc regression methods is their computational savings when many lifting functions are required. The long execution times of the EDMD methods, along with their high memory consumption, make cross-validation impractical. Figure 10 demonstrates the significant resource savings provided by the DMDc methods in both execution time and peak memory consumption. Memory consumption is of particular importance when running multiple instances of a regressor in a multi-process cross-validation scheme. The DMDc regression methods presented provide significant computational savings while still retaining the frequency-domain characteristics of their EDMD counterparts. In spite of their higher mean prediction error, it is often worthwhile to leverage them for Koopman operator identification, particularly when hyperparameter optimization is a priority.

7. Conclusion

Approximating the Koopman matrix using linear regression proves challenging as lifting function complexity increases. Even small problems can become ill-conditioned when many lifting functions are required for an accurate fit. Viewing the problem from a systems perspective, where system inputs pass through dynamics and lead to outputs, provides multiple avenues to

enforce asymptotic stability and penalize large input-output gains in the system, thus ensuring improved numerical conditioning. In particular, regularizing the regression problem with the $\mathcal{H}_\infty$ norm provides the opportunity to tune the regularization process in the frequency domain using weighting functions. The significant performance savings presented by the DMDc-based regression methods allow the $\mathcal{H}_\infty$ norm regularizer to be applied to much larger systems while still remaining tractable.

The nonconvex optimization problems required to use the asymptotic stability constraint and $\mathcal{H}_\infty$ norm regularizers limit their applicability to practical problems. Future research will address this limitation by making use of more efficient BMI solution methods, including Iterative Convex Overbounding [36] and branch-and-bound methods [37]. Although the use of the $\mathcal{H}_\infty$ norm [27, §3.2] as a regularizer is explored in this paper, any system norm, like the $\mathcal{H}_2$ norm [27, §3.3] or a mixed $\mathcal{H}_2$ norm [27, §3.5], can be used. The unique properties of system norms prove useful in addressing the numerical challenges associated with approximating the Koopman operator from data, and will be explored further in future work.

Data Accessibility. The methods presented in this paper and its predecessor [25] are implemented in release v1.0.4 of `pykoop`, the authors' open source Koopman operator identification library [38]. The code required to reproduce the plots in this paper is available in a companion repository at https://github.com/decarsg/system_norm_koopman, release v1.0.3.

Authors' Contributions. Both authors conceived of the research and performed the formal analysis. S.D. wrote the software, performed the experiments, and prepared the visualizations. J.R.F. supervised the research, formulated its high-level objectives, and assisted with detailed derivations.

Competing Interests. The authors declare that they have no competing interests.

Funding. This work was supported by the Mecaademic Inc. through the Mitacs Accelerate program, and by the Natural Sciences and Engineering Research Council of Canada (NSERC), the National Research Council of Canada (NRC), the Toyota Research Institute, the National Science Foundation Career Award [grant number 1751093], and the Office of Naval Research [grant number N00014-18-1-2575].

Acknowledgements. The authors thank Daniel Bruder, Xun Fu, and Ram Vasudevan for graciously providing the soft robot dataset used in this research. The authors also acknowledge Doug Shi-Dong, Robyn Fortune, Shaowu Pan, Karthik Duraisamy, and Matthew M. Peet for productive discussions about regularization techniques, the Koopman operator, and methods for handling BMI constraints.

References

1. Koopman BO. 1931 Hamiltonian systems and transformations in Hilbert space. *Proc. Nat. Acad. Sci.* **17**, 315–318.
2. Mezić I. 2019 Spectrum of the Koopman Operator, Spectral Expansions in Functional Spaces, and State-Space Geometry. *J. Nonlinear Sci.* **30**, 2091–2145.
3. Budišić M, Mohr R, Mezić I. 2012 Applied Koopmanism. *Chaos* **22**.
4. Mauroy A, Mezić I, Susuki Y, editors. 2020 *The Koopman Operator in Systems and Control*. Cham, Switzerland: Springer.
5. Korda M, Mezić I. 2018 Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica* **93**, 149–160.
6. Otto SE, Rowley CW. 2021 Koopman Operators for Estimation and Control of Dynamical Systems. *Annu. Rev. Control, Robot., Auton. Syst.* **4**, 59–87.
7. Abraham I, Murphey TD. 2019 Active Learning of Dynamics for Data-Driven Control Using Koopman Operators. *IEEE Trans. Robot.* **35**, 1071–1083.
8. Mamakoukas G, Castano M, Tan X, Murphey T. 2019 Local Koopman Operators for Data-Driven Control of Robotic Systems. In *Proc. Robot.: Sci. Syst. XV* Freiburg im Breisgau, Germany.
9. Bruder D, Gillespie B, Remy CD, Vasudevan R. 2019 Modeling and Control of Soft Robots Using the Koopman Operator and Model Predictive Control. In *Proc. Robot.: Sci. Syst. XV* Freiburg im Breisgau, Germany.

10. Uchida D, Yamashita A, Asama H. 2021 Data-Driven Koopman Controller Synthesis Based on the Extended $\mathcal{H}_2$ Norm Characterization. *IEEE Contr. Syst. Lett.* **5**, 1795–1800.
11. Abraham I, de la Torre G, Murphey T. 2017 Model-Based Control Using Koopman Operators. In *Proc. Robot.: Sci. Syst. XIII* Cambridge, MA.
12. Kutz NJ, Brunton SL, Brunton BW, Proctor JL. 2016 *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. Philadelphia, PA: SIAM.
13. Mamakoukas G, Abraham I, Murphey TD. 2020 Learning Data-Driven Stable Koopman Operators. *arXiv:2005.04291v1 [cs.RO]*.
14. Tikhonov AN, Goncharsky A, Stepanov VV, Yagola AG. 1995 *Numerical Methods for the Solution of Ill-Posed Problems*. Dordrecht, Netherlands: Springer.
15. Tibshirani R. 1996 Regression Shrinkage and Selection Via the Lasso. *J. Roy. Statistical Soc.: Ser. B* **58**, 267–288.
16. Williams MO, Kevrekidis IG, Rowley CW. 2015 A Data-Driven Approximation of the Koopman Operator: Extending Dynamic Mode Decomposition. *J. Nonlinear Sci.* **25**, 1307–1346.
17. Proctor JL, Brunton SL, Kutz JN. 2014 Dynamic mode decomposition with control. *arXiv:1409.6358v1 [math.OC]*.
18. Sznaiier M. 2021 A Convex Optimization Approach to Learning Koopman Operators. *arXiv:2102.03934v1 [eess.SY]*.
19. Hara K, Inoue M, Sebe N. 2020 Learning Koopman Operator under Dissipativity Constraints. *IFAC-PapersOnLine* **53**, 1169–1174.
20. Lacy SL, Bernstein DS. 2002 Subspace identification with guaranteed stability using constrained optimization. In *Proc. 2002 Amer. Control Conf.* Anchorage, AK.
21. Hoagg JB, Lacy SL, Erwin RS, Bernstein DS. 2004 First-order-hold sampling of positive real systems and subspace identification of positive real models. In *Proc. 2004 Amer. Control Conf.* Boston, MA.
22. Siddiqi SM, Boots B, Gordon GJ. 2008 A Constraint Generation Approach to Learning Stable Linear Dynamical Systems. Technical report Defense Technical Information Center.
23. Mamakoukas G, Xherija O, Murphey T. 2020 Memory-Efficient Learning of Stable Linear Dynamical Systems for Prediction and Control. *Advances Neural Inf. Process. Syst.* **33**, 13527–13538.
24. Gillis N, Karow M, Sharma P. 2020 A note on approximating the nearest stable discrete-time descriptor systems with fixed rank. *Appl. Numer. Math.* **148**, 131–139.
25. Dahdah S, Forbes JR. 2021 Linear Matrix Inequality Approaches to Koopman Operator Approximation. *arXiv:2102.03613v2 [eess.SY]*.
26. Bruder D, Fu X, Vasudevan R. 2021 Advantages of Bilinear Koopman Realizations for the Modeling and Control of Systems With Unknown Dynamics. *IEEE Trans. Robot. Autom.* **6**, 4369–4376.
27. Caverly RJ, Forbes JR. 2019 LMI Properties and Applications in Systems, Stability, and Control Theory. *arXiv:1903.08599v3 [cs.SY]*.
28. El Ghaoui L, Niculescu SI. 2000 *Advances in Linear Matrix Inequality Methods in Control*. SIAM.
29. Doroudchi A, Shivakumar S, Fisher RE, Marvi H, Aukes D, He X, Berman S, Peet MM. 2018 Decentralized Control of Distributed Actuation in a Segmented Soft Robot Arm. In *2018 IEEE Conf. Decision Control* Miami Beach, FL.
30. Fortune R, Beltempo CA, Forbes JR. 2019 System Identification and Feedforward Control of a Fatigue Structural Testing Rig: The Single Actuator Case. *IFAC-PapersOnLine* **52**, 382–387.
31. James G, Witten D, Hastie T, Tibshirani R. 2013 *An Introduction to Statistical Learning*. New York, NY: Springer.
32. Zhou K, Doyle JC, Glover K. 1995 *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice Hall.
33. Green M, Limebeer DJN. 1994 *Linear Robust Control*. London, England: Prentice Hall.
34. Bruder D, Fu X, Gillespie RB, Remy CD, Vasudevan R. 2021 Data-Driven Control of Soft Robots Using Koopman Operator Theory. *IEEE Trans. Robot.* **37**, 948–961.
35. Gavish M, Donoho DL. 2014 The Optimal Hard Threshold for Singular Values is $4/\sqrt{3}$. *IEEE Trans. Inf. Theory* **60**, 5040–5053.
36. Warner E, Scruggs J. 2017 Iterative Convex Overbounding Algorithms for BMI Optimization Problems. *IFAC-PapersOnLine* **50**, 10449–10455.
37. VanAntwerp JG, Braatz RD. 2000 A tutorial on linear and bilinear matrix inequalities. *J. Process Control* **10**, 363–385.
38. Dahdah S, Forbes JR. 2021 decarsg/pykoop. <https://github.com/decarsg/pykoop>.