

Data-driven control with Koopman Operators

Hakan Girgin

L2C course

10.06.2021



EPFL

Koopman Operators: Theory

Dynamical systems can be expressed in discrete-time as $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$ where $\mathbf{x}_t, \mathbf{x}_{t+1} \in \mathbb{R}^{d_x}$ are respectively the current and next states, $\mathbf{u}_t \in \mathbb{R}^{d_u}$ is the current control command and the function $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x}$ describes the evolution of the dynamical system

Koopman theory asserts that any nonlinear dynamical system can be expressed linearly in infinite dimensional Hilbert space. We define a discrete-time Koopman operator \mathcal{K} as the linear operator which advances observation functions $\phi : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_l}$ of the state in time as in

$$\phi(\mathbf{x}_{t+1}) = \mathcal{K} \circ \phi(\mathbf{x}_t)$$

where d_l is the dimension of the output of the observation function.

Finite approximation:

$$\mathbf{K} \in \mathbb{R}^{d_l \times d_l}$$

Koopman Operators: Example of slow manifold dynamics

An example in continuous time dynamics:

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \underbrace{\begin{bmatrix} \mu x_1 \\ \lambda(x_2 - x_1^2) \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} 0 \\ 1 \end{bmatrix}}_{\tilde{B}} u \implies \frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \end{bmatrix} = \underbrace{\begin{bmatrix} \mu & 0 & 0 \\ 0 & \lambda & -\lambda \\ 0 & 0 & 2\mu \end{bmatrix}}_A \underbrace{\begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \end{bmatrix}}_{\phi(\mathbf{x})} + \underbrace{\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}}_B u$$

Koopman Operators: Related Work I

- Some methods assume linear observables $\phi(\mathbf{x}) = \mathbf{x}$, while others use predefined observation function dictionary for such as polynomial, radial, Fourier basis functions.
- Considering an autonomous nonlinear dynamical system:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t)$$

we can find such an approximation of the Koopman operator by setting a least-squares problem as

$$\min_{\mathbf{K}} \sum_{t=1}^T \|\phi(\mathbf{x}_{t+1}) - \mathbf{K}\phi(\mathbf{x}_t)\|^2$$

- Some methods parameterize the observation functions with an unknown parameter θ which is learned together with the Koopman operator at the same time.

$$\min_{\mathbf{K}, \theta} \sum_{t=1}^T \|\phi_{\theta}(\mathbf{x}_{t+1}) - \mathbf{K}\phi_{\theta}(\mathbf{x}_t)\|^2$$

Challenges

1. What are the observables?
2. What are the A and B matrices?
3. How to learn the dynamics model?
4. How to learn the dynamics model for control?

Koopman Operators: Related Work II

- Many of them include also an exogenous control input to transform the problem to a more general version as in

$$\min_{\mathbf{K}, \boldsymbol{\theta}} \sum_{t=1}^T \left\| \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_{t+1}, \mathbf{u}_{t+1}) - \mathbf{K} \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_t, \mathbf{u}_t) \right\|^2$$

- Although this formulation in helps to determine a better approximation for more complicated dynamics models, it does not admit a linear form that the linear control theory works on. So many of methods investigated:

$$\min_{\mathbf{A}, \mathbf{B}, \boldsymbol{\theta}} \sum_{t=1}^T \left\| \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_{t+1}) - \mathbf{A} \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_t) - \mathbf{B} \mathbf{u}_t \right\|^2$$

- Some work investigate the multi-step prediction errors as in the following as opposed to the previous single-step prediction errors minimization

$$\min_{\mathbf{A}, \mathbf{B}, \boldsymbol{\theta}} \sum_{t=1}^T \left\| \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_{t+1}) - \mathbf{A}^{t+1} \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x}_0) - \sum_{\tau=0}^t \mathbf{A}^{\tau} \mathbf{B} \mathbf{u}_{t-\tau} \right\|^2.$$

Challenges

1. What are the observables?
2. What are the A and B matrices?
3. How to learn the dynamics model?
4. How to learn the dynamics model for control?

Koopman Operators: Related Work III

- The linear approximation of the dynamics is appealing as it allows us to use well-established model-based linear control techniques such as LQR.
- If the dynamics model is given as:

$$\phi_{\theta}(\mathbf{x}_{t+1}) = \mathbf{A}\phi_{\theta}(\mathbf{x}_t) + \mathbf{B}u_t$$

and if the cost is:

$$c = \phi_{\theta}(\mathbf{x}_t)^{\top} \mathbf{Q}\phi_{\theta}(\mathbf{x}_t) + u_t^{\top} \mathbf{R}u_t$$

then this is an LQR problem.


However: how to design a cost function with respect to observables instead of original variables (state)?

Answer: Include the state itself in the new observables: $\psi(\mathbf{x}) = \begin{bmatrix} \mathbf{x} \\ \phi(\mathbf{x}) \end{bmatrix}$

Challenges

1. What are the observables?
2. What are the A and B matrices?
3. How to learn the dynamics model?
4. How to learn the dynamics model for control?

Koopman Operators: Approximation of the dynamics



PROBLEMS

Control affine dynamics: (e.g. manipulator dynamics)

$$\dot{x} = f(x) + g(x)u$$

Multiplicative terms of x and u are most of the time ignored.

$$z_{\theta} = \phi_{\theta}(x)$$

Basis functions are chosen **either** manually requiring a large amount of choices or prior knowledge, **or** using deep networks requiring a large amount of data and a good regression cost.

If the regression cost is **one-step ahead prediction squared error**, the resulting model may diverge as we perform trajectory rollouts, if it is **multi-step ahead prediction squared error**, the problem is insanely nonconvex.

Example: nD robot arm dynamics

Rigid body dynamics of a robot arm is written as:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) + \mathbf{f}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{u}$$

where $\mathbf{M}(\mathbf{q})$ is the mass matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ is the centrifugal and coriolis matrix, $\mathbf{g}(\mathbf{q})$ is the gravity vector and $\mathbf{f}(\mathbf{q}, \dot{\mathbf{q}})$ are the other non-modeled nonlinearities of the system.

The state space representation is given as:

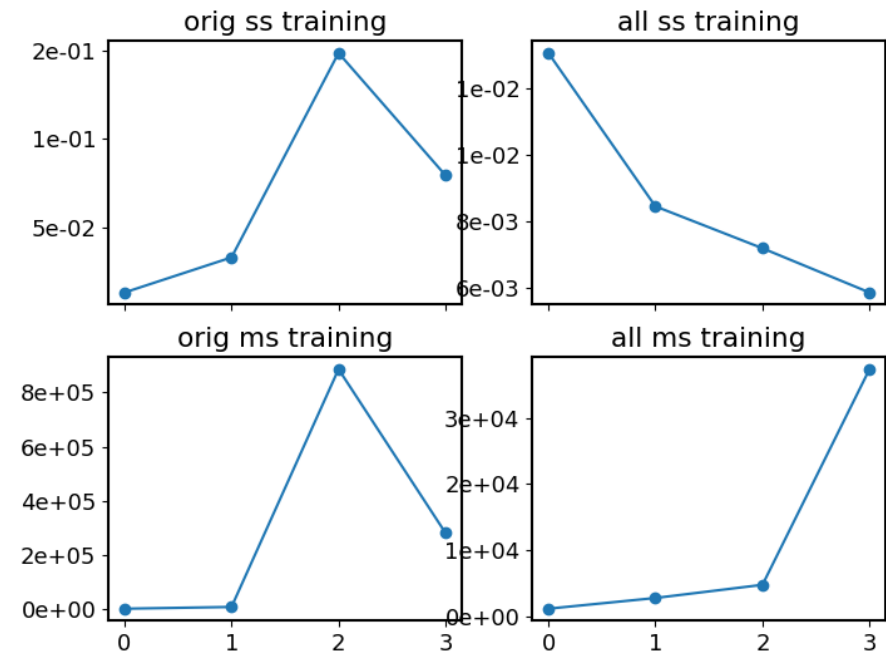
$$\frac{d}{dt} \underbrace{\begin{bmatrix} \mathbf{q} \\ \dot{\mathbf{q}} \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} \dot{\mathbf{q}} \\ -\mathbf{M}^{-1}(\mathbf{C}\dot{\mathbf{q}} + \mathbf{g} + \mathbf{f}) \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1} \end{bmatrix}}_{\tilde{\mathbf{B}}(\mathbf{x})} \mathbf{u}$$

For a simple pendulum \mathbf{M} is a constant \rightarrow we can learn a good Koopman representation!

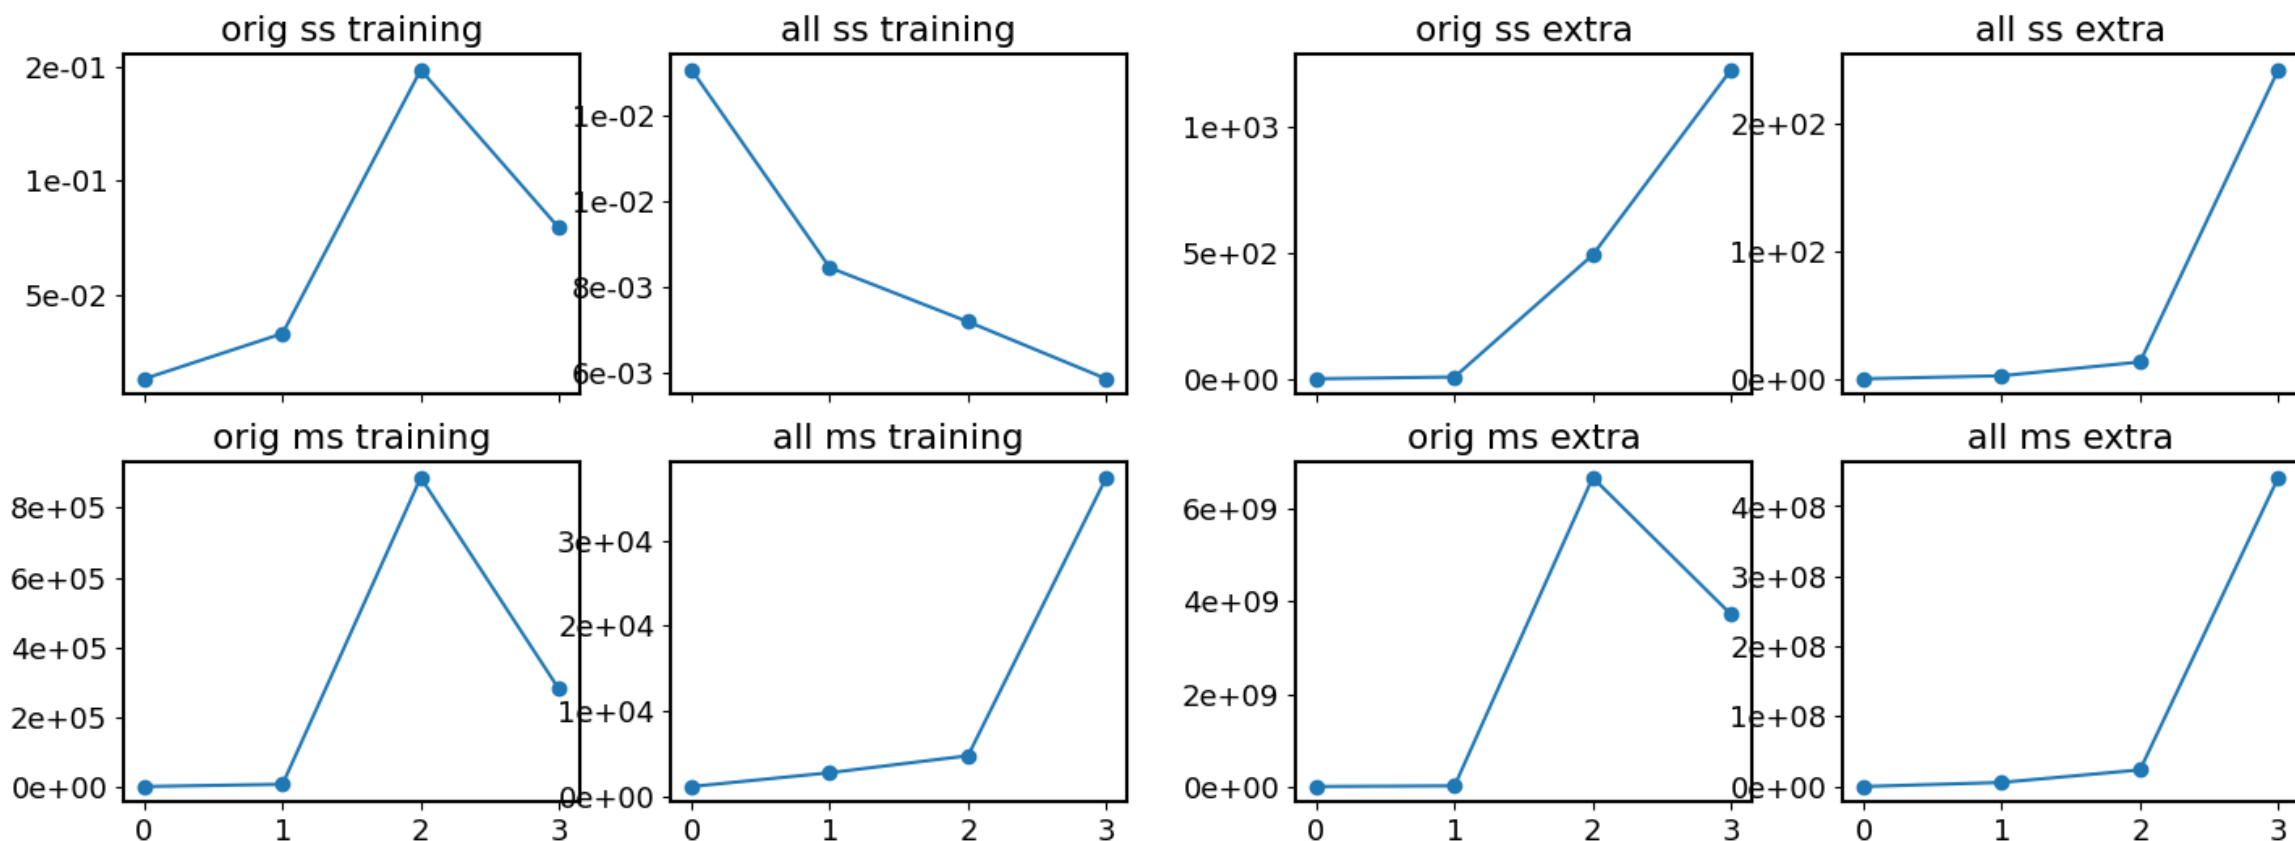
Example: 3D robot arm dynamics

- Experiments done with known lifting functions (observables) chosen arbitrarily with increasing complexity.
- Since the observables are known, this is just a least-squares problem.
- The error is computed as first the single-step prediction errors on the training data (called *training*) and on the testing data (called *extra*), and then as the multi-step prediction errors over the whole horizon.
- *Orig* refers to the error only on the original state variables and *all* refers to the error on the overall lifted vector


$$\begin{aligned}\phi_0(\mathbf{x}) &= [\mathbf{x}] & \phi_2(\mathbf{x}) &= \begin{bmatrix} \mathbf{x} \\ \mathbf{x}^2 \end{bmatrix} \\ \phi_1(\mathbf{x}) &= \begin{bmatrix} \mathbf{x} \\ \cos(\mathbf{x}) \\ \sin(\mathbf{x}) \end{bmatrix} & \phi_3(\mathbf{x}) &= \begin{bmatrix} \mathbf{x} \\ \cos(\mathbf{x}) \\ \sin(\mathbf{x}) \\ \mathbf{x}^2 \end{bmatrix}\end{aligned}$$



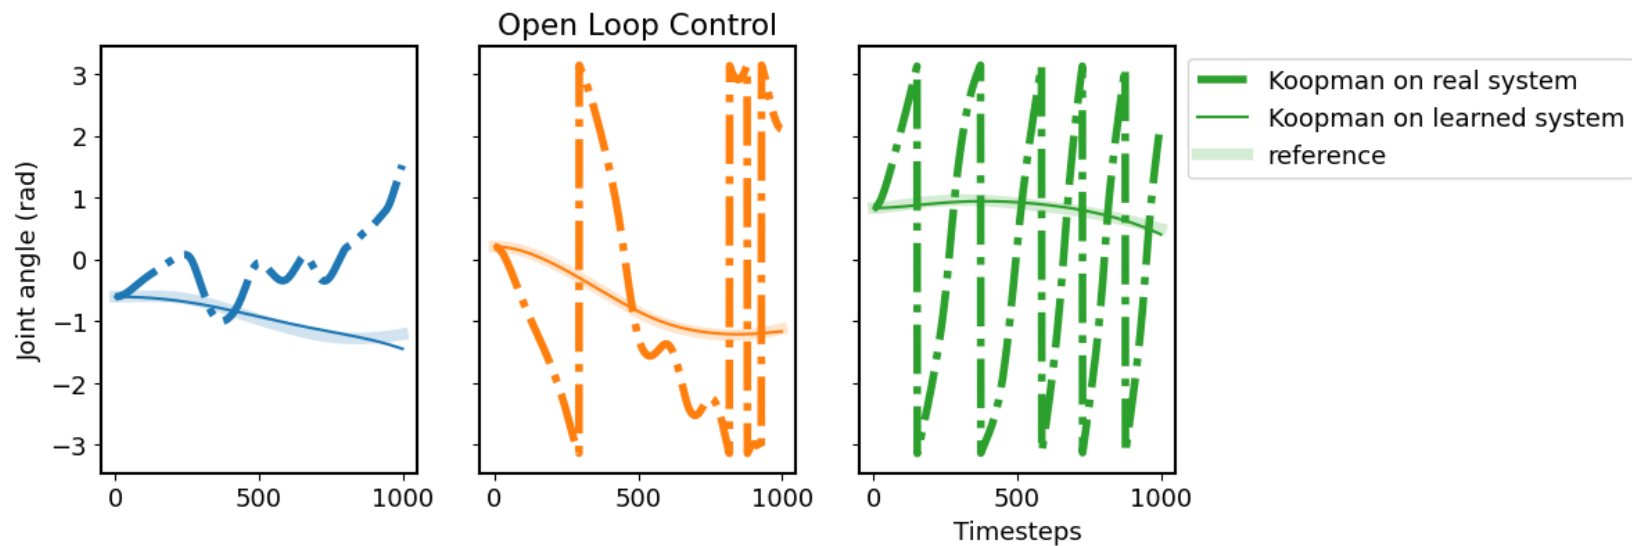
Example: 3D robot arm dynamics - results



LQT applied to the learned model - 1




	Tracking Cost	Control Cost
Feedback	6.40E+00	8.34E+05
Open-Loop	9.90E+03	1.00E+06

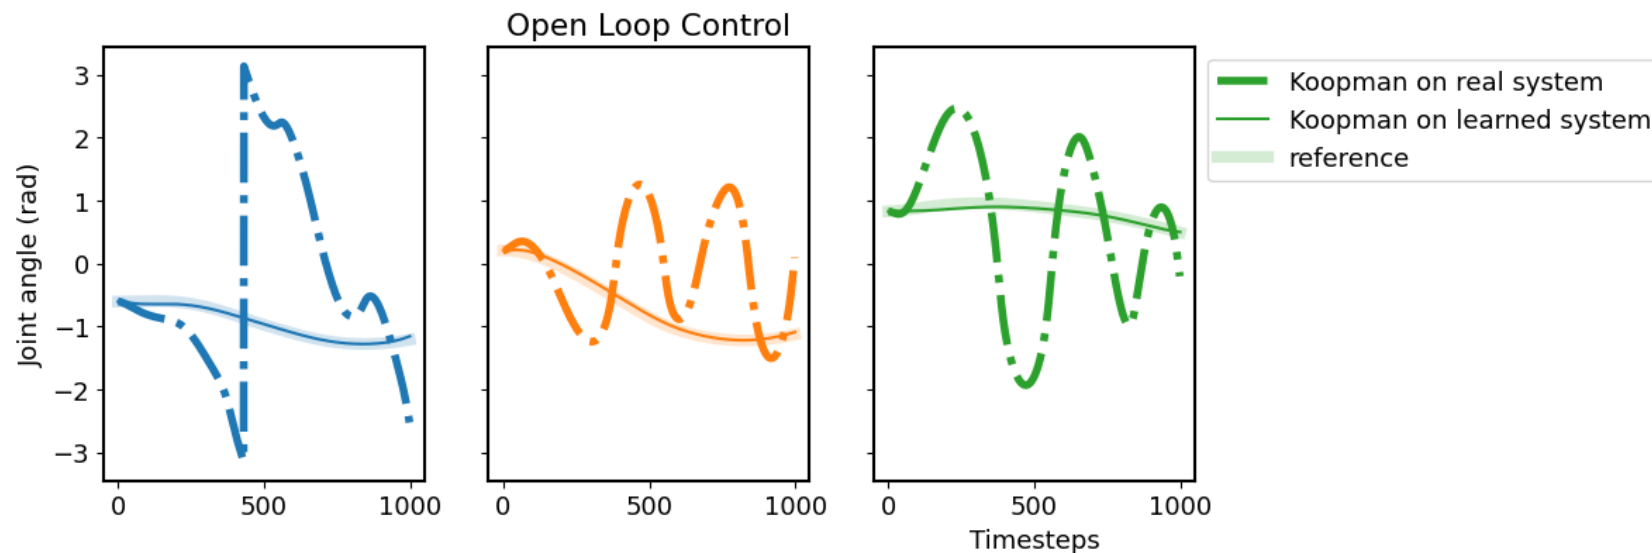


$$\phi_0(\mathbf{x}) = [\mathbf{x}]$$

LQT applied to the learned model - 2



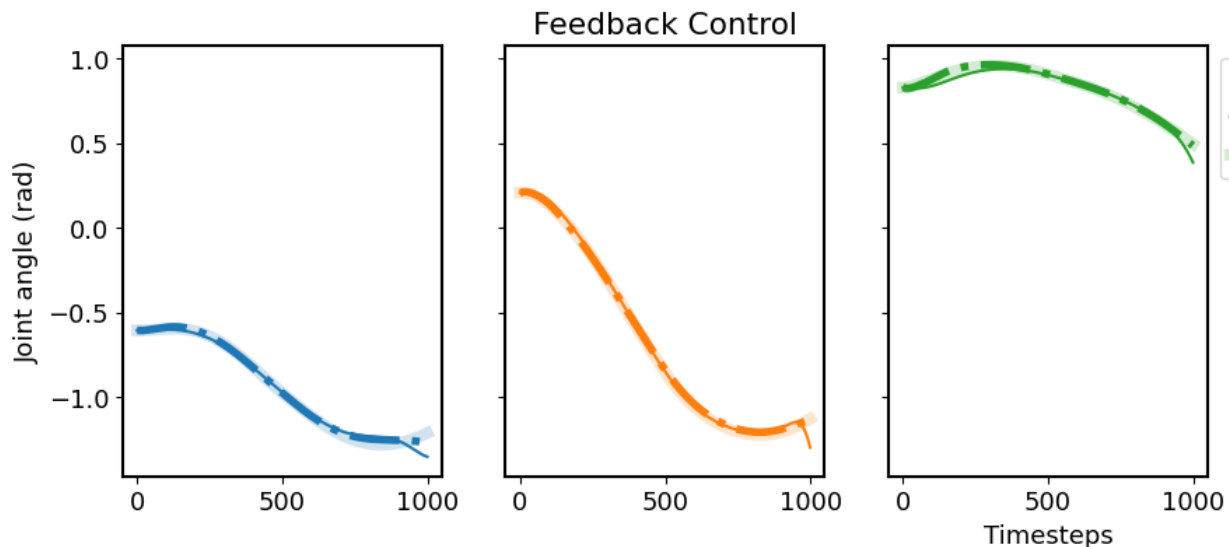
	Tracking Cost	Control Cost
Feedback	1.35E+01	1.06E+06
Open-Loop	5.99E+03	2.77E+06



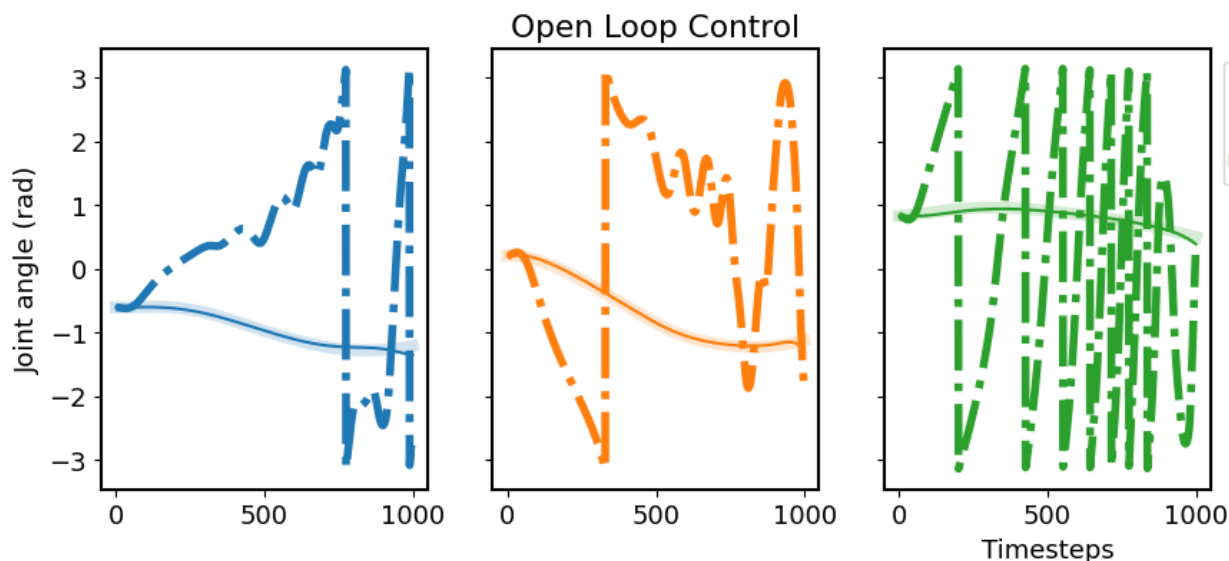
$$\phi_1(\mathbf{x}) = \begin{bmatrix} \mathbf{x} \\ \cos(\mathbf{x}) \\ \sin(\mathbf{x}) \end{bmatrix}$$

Others \rightarrow divergence or oscillations ..

Learning the lifting functions with MLP and apply LQT



	Tracking Cost	Control Cost
Feedback	4.11E-01	1.66E+06
Open-Loop	1.22E+04	1.13E+07



$$\phi_{\theta}(\mathbf{x}) = \begin{bmatrix} \mathbf{x} \\ \text{MLP}(\mathbf{x}) \end{bmatrix}$$

Comparison of these costs for each of three cases

	Tracking Cost	Control Cost
Feedback	6.40E+00	8.34E+05
Open-Loop	9.90E+03	1.00E+06

$$\phi_0(\mathbf{x}) = [\mathbf{x}]$$

	Tracking Cost	Control Cost
Feedback	1.35E+01	1.06E+06
Open-Loop	5.99E+03	2.77E+06

$$\phi_1(\mathbf{x}) = \begin{bmatrix} \mathbf{x} \\ \cos(\mathbf{x}) \\ \sin(\mathbf{x}) \end{bmatrix}$$

	Tracking Cost	Control Cost
Feedback	4.11E-01	1.66E+06
Open-Loop	1.22E+04	1.13E+07

$$\phi_\theta(\mathbf{x}) = \begin{bmatrix} \mathbf{x} \\ \text{MLP}(\mathbf{x}) \end{bmatrix}$$

Dynamical consistency

Control affine dynamics:
(e.g. manipulator dynamics)

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t) + \mathbf{g}(\mathbf{x}_t)\mathbf{u}_t$$

Problem was: Multiplicative terms of \mathbf{x} and \mathbf{u} are most of the time ignored.

Consider the following Koopman structure:

$$\dot{\psi}_t = \mathbf{A}\psi_t + \mathbf{B}(\psi_t)\mathbf{u}_t$$



$$\begin{aligned}\dot{\psi}(\mathbf{x}_t) &= \frac{\partial \psi_t(\mathbf{x}_t)}{\partial \mathbf{x}_t} \dot{\mathbf{x}}_t \\ &= \frac{\partial \psi_t(\mathbf{x}_t)}{\partial \mathbf{x}_t} \mathbf{f}(\mathbf{x}_t) + \frac{\partial \psi_t(\mathbf{x}_t)}{\partial \mathbf{x}_t} \mathbf{g}(\mathbf{x}_t)\mathbf{u}_t \\ &= \mathbf{A}\psi_t + \mathbf{B}(\psi_t)\mathbf{u}_t\end{aligned}$$

Then the constraints for dynamical consistency are:

$$\mathbf{A}\psi_t = \frac{\partial \psi_t(\mathbf{x}_t)}{\partial \mathbf{x}_t} \mathbf{f}(\mathbf{x}_t),$$

$$\mathbf{B}(\psi_t)\mathbf{u}_t = \frac{\partial \psi_t(\mathbf{x}_t)}{\partial \mathbf{x}_t} \mathbf{g}(\mathbf{x}_t)\mathbf{u}_t$$

$$\psi_t = \psi(\mathbf{x}_t) = \begin{bmatrix} \mathbf{x}_t \\ \phi(\mathbf{x}_t) \end{bmatrix}$$

Dynamical consistency –a proposed solution

Control affine dynamics:
(e.g. manipulator dynamics)

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t) + \mathbf{g}(\mathbf{x}_t)\mathbf{u}_t$$

Problem was: Multiplicative terms of \mathbf{x} and \mathbf{u} are most of the time ignored.

$$\dot{\boldsymbol{\psi}}_t = \mathbf{A}\boldsymbol{\psi}_t + \mathbf{B}(\boldsymbol{\psi}_t)\mathbf{u}_t$$

$$\boldsymbol{\psi}_t = \boldsymbol{\psi}(\mathbf{x}_t) = \begin{bmatrix} \mathbf{x}_t \\ \boldsymbol{\phi}(\mathbf{x}_t) \end{bmatrix}$$

For the sake of simplicity, we choose the control matrix to be a linear function of the lifted state as

$$\mathbf{B}(\boldsymbol{\psi}_t) = \mathbf{W}\boldsymbol{\psi}_t\mathbf{p}^\top + \mathbf{C}, \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{d \times d}$, $\mathbf{p} \in \mathbb{R}^{d_u}$ and $\mathbf{C} \in \mathbb{R}^{d \times d_u}$ are the unknown parameters of the control matrix to be learned.

Dynamical consistency - control

For the sake of simplicity, we choose the control matrix to be a linear function of the lifted state as

$$\mathbf{B}(\boldsymbol{\psi}_t) = \mathbf{W}\boldsymbol{\psi}_t\mathbf{p}^\top + \mathbf{C}, \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{d \times d}$, $\mathbf{p} \in \mathbb{R}^{d_u}$ and $\mathbf{C} \in \mathbb{R}^{d \times d_u}$ are the unknown parameters of the control matrix to be learned.

Our Koopman dynamics model assumes the following final form:

$$\dot{\boldsymbol{\psi}}_t = f(\boldsymbol{\psi}_t, \mathbf{u}_t) = \mathbf{A}\boldsymbol{\psi}_t + (\mathbf{W}\boldsymbol{\psi}_t\mathbf{p}^\top + \mathbf{C})\mathbf{u}_t.$$

$$\frac{\partial f(\boldsymbol{\psi}_t, \mathbf{u}_t)}{\partial \boldsymbol{\psi}_t} = \mathbf{A} + \nabla \mathbf{B}\mathbf{u}_t,$$

$$\frac{\partial f(\boldsymbol{\psi}_t, \mathbf{u}_t)}{\partial \mathbf{u}_t} = (\mathbf{W}\boldsymbol{\psi}_t\mathbf{p}^\top + \mathbf{C}),$$

where $\nabla \mathbf{B}_{ijk} = \mathbf{W}_{ik}\mathbf{p}_j$. The derivative of the forward dynamics with respect to the lifted state is therefore a linear function of the control, whereas with respect to the control is a linear function of the lifted state.

Open problems

- Gathering meaningful experimental data so that the system can learn as fast as possible. → ideas around active learning, and iterative data gathering to learn better and faster with each collected trajectory.
- What should be done so that we extrapolate well even though we have small amount of data → ideas around a new structure of MLP that can help us, can also discuss this.
- What should be the training cost? Should there be equations related to dynamic consistency as constraints? Some papers deal with finding stable A and B matrices, but we need stabilizable A and B matrices, right?
- When we learn the Koopman dynamics model, we add another layer of complexity to predict also the evolution of the observables. However, they are functions of the state, therefore their evolution is related to the evolution of the state itself. Do we do something redundant?