Data-driven optimal control of affine systems: A linear programming perspective

A. Martinelli, M. Gargiani, M. Draskovic, J. Lygeros
Automatic Control Laboratory, ETH Zurich

61st IEEE Conference on Decision and Control (CDC), Cancun, Mexico 2022

Outline

Introduction to data-driven optimal control via linear programming

Estimation of Bellman inequalities from data

Willems' Fundamental Lemma for affine systems

Conclusion and future work

Problem setup: stochastic optimal control

Ingredients:

- A discrete-time stochastic system $x^+ = f(x, u, \psi)$ with possibly infinite state & action spaces
- A stage-cost function $\ell : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_+$

Problem setup: stochastic optimal control

Ingredients:

- A discrete-time stochastic system $x^+ = f(x, u, \psi)$ with possibly infinite state & action spaces
- A stage-cost function $\ell : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_+$

The γ-discounted ∞**-horizon cost** associated to a stationary feedback policy $\pi : \mathbb{X} \to \mathbb{U}$ is

$$v_\pi(x) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k \ell(x_k, \pi(x_k)) \,\middle|\, x_0 = x\right]$$

Problem setup: stochastic optimal control

Ingredients:

- A discrete-time stochastic system $x^+ = f(x, u, \psi)$ with possibly infinite state & action spaces
- A stage-cost function $\ell : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_+$

The γ-discounted ∞**-horizon cost** associated to a stationary feedback policy $\pi : \mathbb{X} \to \mathbb{U}$ is

$$v_\pi(x) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k \ell(x_k, \pi(x_k)) \,\middle|\, x_0 = x\right]$$

Objective: find an optimal policy π^* such that $v_{\pi^*}(x) = \inf_\pi v_\pi(x) = v^*(x)$

Problem setup: stochastic optimal control

Ingredients:

- A discrete-time stochastic system $x^+ = f(x, u, \psi)$ with possibly infinite state & action spaces
- A stage-cost function $\ell : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_+$

The γ-discounted ∞-**horizon cost** associated to a stationary feedback policy $\pi : \mathbb{X} \to \mathbb{U}$ is

$$v_\pi(x) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k \ell(x_k, \pi(x_k)) \,\middle|\, x_0 = x\right]$$

Objective: find an optimal policy π^* such that $v_{\pi^*}(x) = \inf_\pi v_\pi(x) = v^*(x)$

- The value function admits a recursive definition – the **Bellman equation**

$$v^*(x) = \inf_{u \in \mathbb{U}} \left\{ \ell(x, u) + \gamma \mathbb{E}\left[v^*(f(x, u, \psi))\right] \right\}$$

The Q-function

- If one is able to obtain v^*, then

$$\pi^*(x) = \arg\min_{u \in \mathbb{U}} \left\{ \ell(x,u) + \gamma \mathbb{E}\left[v^*(f(x,u,\psi)) \right] \right\}$$

Problem: *policy extraction is in general not possible if f or ℓ are not known*

The Q-function

- If one is able to obtain v^*, then

$$\pi^*(x) = \arg\min_{u \in \mathbb{U}} \left\{ \ell(x,u) + \gamma \mathbb{E}\left[v^*(f(x,u,\psi)) \right] \right\}$$

Problem: *policy extraction is in general not possible if f or ℓ are not known*

- Introducing Q-functions [1]

$$q^*(x,u) = \underbrace{\ell(x,u) + \gamma \mathbb{E}\left[\inf_{w \in \mathbb{U}} q^*(f(x,u,\psi), w) \right]}_{\textbf{Q-Bellman operator } (\mathcal{F}q^*)(x,u)}$$

[1]Watkins and Dayan, *Machine Learning*, 1992

The Q-function

- If one is able to obtain v^*, then

$$\pi^*(x) = \arg\min_{u \in \mathbb{U}} \left\{ \ell(x,u) + \gamma \mathbb{E}\left[v^*(f(x,u,\psi)) \right] \right\}$$

Problem: *policy extraction is in general not possible if f or ℓ are not known*

- Introducing Q-functions [1]

$$q^*(x,u) = \underbrace{\ell(x,u) + \gamma \mathbb{E}\left[\inf_{w \in \mathbb{U}} q^*(f(x,u,\psi), w) \right]}_{\textbf{\textit{Q}-Bellman operator } (\mathcal{F} q^*)(x,u)}$$

- Since $v^*(x) = \min_{u \in \mathbb{U}} q^*(x,u)$, now policy extraction is **model-free**:

$$\pi^*(x) = \arg\min_{u \in \mathbb{U}} q^*(x,u)$$

[1] Watkins and Dayan, *Machine Learning*, 1992

The linear programming formulation

- $\mathcal{F}$ is a monotone contraction mapping with a unique fixed point (q^*), *i.e.*,

$$q \leq \mathcal{F}q \implies q \leq q^* \quad \text{and} \quad q^* = \mathcal{F}q^*$$

The linear programming formulation

- $\mathcal{F}$ is a monotone contraction mapping with a unique fixed point (q^*), *i.e.*,

$$q \leq \mathcal{F}q \implies q \leq q^* \quad \text{and} \quad q^* = \mathcal{F}q^*$$

- Natural to define q^* as the greatest function that satisfies the *Bellman inequalities*,

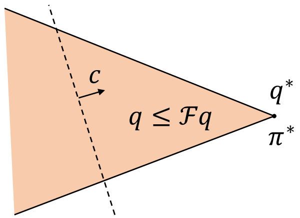

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X}\times\mathbb{U}} q(x,u)c(dx,du)$$
$$\text{s.t.} \;\; q(x,u) \leq (\mathcal{F}q)(x,u) \quad \forall x,u$$

Problem: **nonlinear** and ∞-dimensional

The linear programming formulation

- $\mathcal{F}$ is a monotone contraction mapping with a unique fixed point (q^*), *i.e.*,

$$q \leq \mathcal{F}q \implies q \leq q^* \quad \text{and} \quad q^* = \mathcal{F}q^*$$

- Natural to define q^* as the greatest function that satisfies the *Bellman inequalities*,

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X}\times\mathbb{U}} q(x,u)c(dx,du)$$
$$\text{s.t.} \;\; q(x,u) \leq (\mathcal{F}q)(x,u) \quad \forall x,u$$

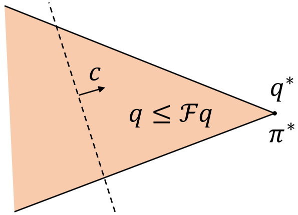

Problem: **nonlinear** and ∞-dimensional

- Introducing the **relaxed *Q*-Bellman operator** [2]

$$(\hat{\mathcal{F}}q)(x,u) = \ell(x,u) + \gamma \inf_{w \in \mathbb{U}} \mathbb{E}\big[q(f(x,u,\psi),w)\big]$$

[2] Martinelli *et al.*, *Automatica*, 2022

The relaxed Bellman operator

Lemma (Properties of the relaxed operator)

(i) *$\hat{\mathcal{F}}$ is a monotone contraction mapping with a unique fixed point ($\hat{q}$)*

(ii) $q^* \leq \hat{q}$

(iii) *For linear (affine) systems, it holds* $\arg\min_{u \in \mathbb{U}} q^* = \arg\min_{u \in \mathbb{U}} \hat{q}$ ***(policy preservation)***

The relaxed Bellman operator

Lemma (Properties of the relaxed operator)

(i) *$\hat{\mathcal{F}}$ is a monotone contraction mapping with a unique fixed point ($\hat{q}$)*

(ii) $q^* \leq \hat{q}$

(iii) *For linear (affine) systems, it holds* $\arg\min_{u \in \mathbb{U}} q^* = \arg\min_{u \in \mathbb{U}} \hat{q}$ ***(policy preservation)***

- One can find $\hat{q}$ via

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X} \times \mathbb{U}} q(x,u) c(dx, du)$$

$$\text{s.t.} \quad q(x,u) \leq (\hat{\mathcal{F}} q)(x,u) \quad \forall x, u$$

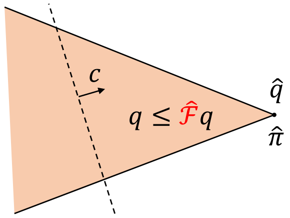

The relaxed Bellman operator

Lemma (Properties of the relaxed operator)

(i) *$\hat{\mathcal{F}}$ is a monotone contraction mapping with a unique fixed point ($\hat{q}$)*

(ii) $q^* \leq \hat{q}$

(iii) *For linear (affine) systems, it holds* $\arg\min_{u \in \mathbb{U}} q^* = \arg\min_{u \in \mathbb{U}} \hat{q}$ ***(policy preservation)***

- One can find $\hat{q}$ via

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X} \times \mathbb{U}} q(x,u) c(dx,du)$$
$$\text{s.t.} \quad q(x,u) \leq (\hat{\mathcal{F}} q)(x,u) \quad \forall x,u$$

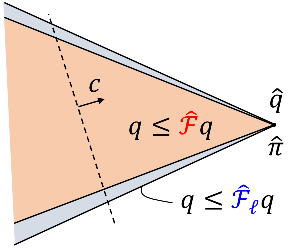

- Equivalently with $(\hat{\mathcal{F}}_\ell q)(x,u) = \ell(x,u) + \gamma \mathbb{E}\big[q(f(x,u,\psi),w)\big] \quad \forall (x,u,w)$

Finite and model-free LP formulation

LPs can be solved efficiently in general, but here several **sources of intractability** arise:

- q is an optimization variable in the ∞-dimensional space $\mathbb{Q}$
- ∞ number of constraints

$$\begin{aligned} &\sup_{q \in \mathbb{Q}} \int_{\mathbb{X}\times\mathbb{U}} q(x,u)c(dx,du) \\ &\text{s.t.} \quad q(x,u) \le \ell(x,u) + \gamma \mathbb{E}\big[q(f(x,u,\psi),w)\big] \quad \forall (x,u,w) \end{aligned}$$

Finite and model-free LP formulation

LPs can be solved efficiently in general, but here several **sources of intractability** arise:

- q is an optimization variable in the ∞-dimensional space $\mathbb{Q}$
- ∞ number of constraints

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X}\times\mathbb{U}} q(x,u) c(dx,du)$$

$$\text{s.t.} \quad q(x,u) \leq \ell(x,u) + \gamma \mathbb{E}\left[q(f(x,u,\psi),w)\right] \quad \forall (x,u,w)$$

We can substitute $\sum_i \varphi_i \phi_i(x,u)$ and sample a finite subset of constraints

$$q(x_i,u_i) \leq \ell(x_i,u_i) + \gamma \mathbb{E}\left[q(x_i^+,w)\right] \quad \forall (x_i,u_i,w_i) \in \mathcal{D}$$

Finite and model-free LP formulation

LPs can be solved efficiently in general, but here several **sources of intractability** arise:

- q is an optimization variable in the ∞-dimensional space $\mathbb{Q}$
- ∞ number of constraints

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X} \times \mathbb{U}} q(x,u) c(dx, du)$$

$$\text{s.t.} \quad q(x,u) \leq \ell(x,u) + \gamma \mathbb{E}\big[q(f(x,u,\psi), w)\big] \quad \forall (x,u,w)$$

We can substitute $\sum_i \varphi_i \phi_i(x,u)$ and sample a finite subset of constraints

$$q(x_i, u_i) \leq \ell(x_i, u_i) + \gamma \mathbb{E}\big[q(x_i^+, w)\big] \quad \forall (x_i, u_i, w_i) \in \mathcal{D}$$

Model-free/RL framework: construct one constraint for each observation

$$\{x_i, u_i, \ell(x_i, u_i), x_i^+\}_{i=0}^{d-1}$$

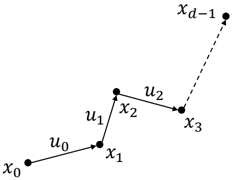

Finite and model-free LP formulation

LPs can be solved efficiently in general, but here several **sources of intractability** arise:

- q is an optimization variable in the ∞-dimensional space $\mathbb{Q}$
- ∞ number of constraints

$$\sup_{q \in \mathbb{Q}} \int_{\mathbb{X}\times\mathbb{U}} q(x,u)c(dx,du)$$

$$\text{s.t.} \quad q(x,u) \leq \ell(x,u) + \gamma \mathbb{E}\big[q(f(x,u,\psi),w)\big] \quad \forall (x,u,w)$$

We can substitute $\sum_i \varphi_i \phi_i(x,u)$ and sample a finite subset of constraints

$$q(x_i,u_i) \leq \ell(x_i,u_i) + \gamma \mathbb{E}\big[q(x_i^+,w)\big] \quad \forall (x_i,u_i,w_i) \in \mathcal{D}$$

Model-free/RL framework: construct one constraint for each observation

$$\{x_i, u_i, \ell(x_i,u_i), x_i^+\}_{i=0}^{d-1}$$

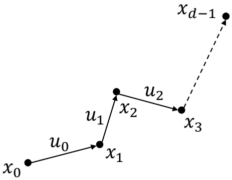

Outline

Introduction to data-driven optimal control via linear programming

Estimation of Bellman inequalities from data

Willems' Fundamental Lemma for affine systems

Conclusion and future work

Estimation of Bellman inequalities

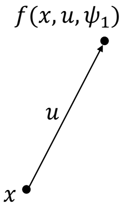

Estimation of Bellman inequalities

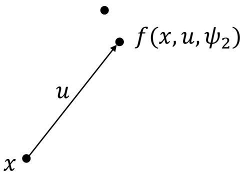

Estimation of Bellman inequalities

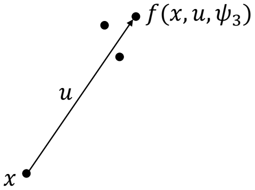

Estimation of Bellman inequalities

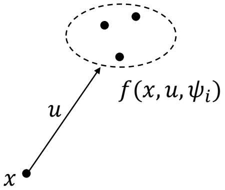

Estimation of Bellman inequalities

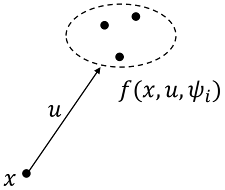

Unbiased estimator (reinitialization)

$$\hat{\theta}_{x,u} = \frac{1}{d} \sum_{i=0}^{d-1} v(x_i^+)$$

- $\mathbb{E}[\hat{\theta}_{x,u}] = \mathbb{E}\left[v(f(x, u, \psi))\right]$
- $\mathrm{Var}(\hat{\theta}_{x,u}) = \frac{1}{d}\mathrm{Var}(v(f(x, u, \psi)))$

Estimation of Bellman inequalities

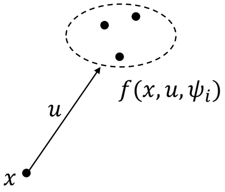

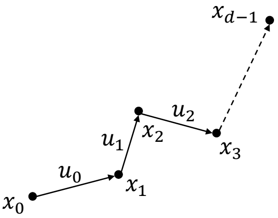

$$
\begin{aligned}
X &= [x_0 \; \cdots \; x_{d-1}] \\
U &= [u_0 \; \cdots \; u_{d-1}] \\
X^+ &= \left[x_0^+ \; \cdots \; x_{d-1}^+\right] \\
\Psi &= [\psi_0 \; \cdots \; \psi_{d-1}]
\end{aligned}
$$

Unbiased estimator (reinitialization)

$$\hat{\theta}_{x,u} = \frac{1}{d}\sum_{i=0}^{d-1} v(x_i^+)$$

- $\mathbb{E}[\hat{\theta}_{x,u}] = \mathbb{E}\left[v(f(x,u,\psi))\right]$
- $\mathrm{VAR}(\hat{\theta}_{x,u}) = \frac{1}{d}\mathrm{VAR}(v(f(x,u,\psi)))$

Estimation of Bellman inequalities

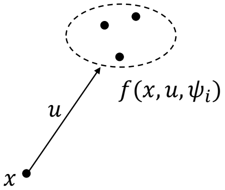

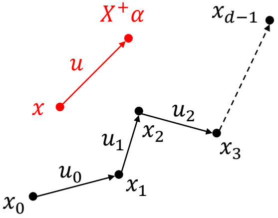

$$
\begin{aligned}
X &= \begin{bmatrix} x_0 & \cdots & x_{d-1} \end{bmatrix} \\
U &= \begin{bmatrix} u_0 & \cdots & u_{d-1} \end{bmatrix} \\
X^+ &= \begin{bmatrix} x_0^+ & \cdots & x_{d-1}^+ \end{bmatrix} \\
\Psi &= \begin{bmatrix} \psi_0 & \cdots & \psi_{d-1} \end{bmatrix}
\end{aligned}
$$

Unbiased estimator (reinitialization)

$$\hat{\theta}_{x,u} = \frac{1}{d} \sum_{i=0}^{d-1} v(x_i^+)$$

- $\mathbb{E}[\hat{\theta}_{x,u}] = \mathbb{E}\left[v(f(x,u,\psi))\right]$
- $\mathrm{VAR}(\hat{\theta}_{x,u}) = \frac{1}{d}\mathrm{VAR}(v(f(x,u,\psi)))$

Biased estimator (w/o reinitialization)

Select $\alpha \in \mathbb{R}^d \; : \; \begin{bmatrix} X \\ U \end{bmatrix} \alpha = \begin{bmatrix} x \\ u \end{bmatrix}$ and compute

$$\bar{\theta}_{x,u} = v(X^+\alpha)$$

- $\mathbb{E}[\bar{\theta}_{x,u}] \neq \mathbb{E}\left[v(f(x,u,\psi))\right]$
- What if $X^+\alpha = f(x,u,\Psi\alpha)$?

Estimation for affine systems (1/2)

Fact

Consider $f(x, u, \psi) = Ax + Bu + c + \psi$ *and* $\alpha \in \mathbb{R}^d \; : \; \left[\begin{smallmatrix} X \\ U \end{smallmatrix}\right]\alpha = \left[\begin{smallmatrix} x \\ u \end{smallmatrix}\right]$.
If we impose $\mathbf{1}^\top \alpha = 1$, *then*

$$X^+ \alpha = f(x, u, \Psi\alpha)$$

Estimation for affine systems (1/2)

Fact

Consider $f(x,u,\psi) = Ax + Bu + c + \psi$ *and* $\alpha \in \mathbb{R}^d \;:\; \begin{bmatrix} X \\ U \end{bmatrix}\alpha = \begin{bmatrix} x \\ u \end{bmatrix}$.
If we impose $\mathbf{1}^\top \alpha = 1$, *then*

$$X^+\alpha = f(x,u,\Psi\alpha)$$

- Is the estimator **unbiased** in this case? *i.e.*, is

$$\mathbb{E}[\bar{\theta}_{x,u}] = \mathbb{E}[v(X^+\alpha)] = \mathbb{E}\left[v(f(x,u,\psi))\right]?$$

Estimation for affine systems (1/2)

Fact

Consider $f(x,u,\psi) = Ax + Bu + c + \psi$ *and* $\alpha \in \mathbb{R}^d$: $\left[\begin{smallmatrix} X \\ U \end{smallmatrix}\right]\alpha = \left[\begin{smallmatrix} x \\ u \end{smallmatrix}\right]$.
If we impose $\mathbf{1}^\top \alpha = 1$*, then*

$$X^+\alpha = f(x,u,\Psi\alpha)$$

- Is the estimator **unbiased** in this case? *i.e.*, is

$$\mathbb{E}[\bar{\theta}_{x,u}] = \mathbb{E}[v(X^+\alpha)] = \mathbb{E}\left[v(f(x,u,\psi))\right]? \quad \textbf{No!}$$

Proposition

For affine systems, if $\begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix}\alpha = \begin{bmatrix} x \\ u \\ 1 \end{bmatrix}$*, then*

(i) $\mathbb{E}[\bar{\theta}_{x,u}] = \mathbb{E}\left[v(f(x,u,\bar{\psi}))\right]$, *where* $\mathbb{E}[\bar{\psi}] = \mathbb{E}[\psi]$ *and* $\mathrm{VAR}(\bar{\psi}) = \|\alpha\|_2^2 \mathrm{VAR}(\psi)$

(ii) $\mathrm{VAR}(\bar{\theta}_{x,u}) = \|\alpha\|_2^2 \mathrm{VAR}(v(f(x,u,\psi)))$

Estimation for affine systems (2/2)

Lemma

The LP with biased constraints is policy-preserving. Indeed,

$$\bar{q}(x,u) = \hat{q}(x,u) - \frac{\gamma}{1-\gamma}\textsc{Tr}\left((1 - \|\alpha\|_2^2) q^* \textsc{Var}(\psi)\right)$$

Estimation for affine systems (2/2)

Lemma

The LP with biased constraints is policy-preserving. Indeed,

$$\bar{q}(x,u) = \hat{q}(x,u) - \tfrac{\gamma}{1-\gamma}\,\textsc{Tr}\left((1-\|\alpha\|_2^2)q^*\textsc{Var}(\psi)\right)$$

Fact

If the data matrix $\begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix}$ *is full row-rank then we can construct (biased) constraints for all* (x,u)

- Can we meet such condition with a **persistently exciting** input?

Outline

Introduction to data-driven optimal control via linear programming

Estimation of Bellman inequalities from data

Willems' Fundamental Lemma for affine systems

Conclusion and future work

The Fundamental Lemma for **linear** systems

- A sequence $S = \begin{bmatrix} S_1 & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{m \times d}$ is **persistently exciting** of order K if the associated *Hankel* matrix of depth K,

$$\mathcal{H}_K(S) = \begin{bmatrix} S_1 & S_2 & \cdots & S_{d-K+1} \\ S_2 & S_3 & \cdots & S_{d-K+2} \\ \vdots & \vdots & & \vdots \\ S_K & S_{K+1} & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{mK \times (d-K+1)},$$

is full row-rank, *i.e.* $\textsc{Rank}\, \mathcal{H}_K(S) = mK$

The Fundamental Lemma for **linear** systems

- A sequence $S = \begin{bmatrix} S_1 & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{m \times d}$ is **persistently exciting** of order K if the associated *Hankel* matrix of depth K,

$$\mathcal{H}_K(S) = \begin{bmatrix} S_1 & S_2 & \cdots & S_{d-K+1} \\ S_2 & S_3 & \cdots & S_{d-K+2} \\ \vdots & \vdots & & \vdots \\ S_K & S_{K+1} & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{mK \times (d-K+1)},$$

  is full row-rank, *i.e.* $\textsc{Rank}\, \mathcal{H}_K(S) = mK$

- The information contained in a **sufficiently rich and long trajectory** of a *linear* system is enough to describe any other trajectory of appropriate length that the system can generate [3]

[3]Willems' *et al.*, *Syst. Control Lett.*, 2005

The Fundamental Lemma for **linear** systems

- A sequence $S = \begin{bmatrix} S_1 & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{m \times d}$ is **persistently exciting** of order K if the associated *Hankel* matrix of depth K,

$$\mathcal{H}_K(S) = \begin{bmatrix} S_1 & S_2 & \cdots & S_{d-K+1} \\ S_2 & S_3 & \cdots & S_{d-K+2} \\ \vdots & \vdots & & \vdots \\ S_K & S_{K+1} & \cdots & S_d \end{bmatrix} \in \mathbb{R}^{mK \times (d-K+1)},$$

  is full row-rank, *i.e.* $\textsc{Rank}\, \mathcal{H}_K(S) = mK$

- The information contained in a **sufficiently rich and long trajectory** of a *linear* system is enough to describe any other trajectory of appropriate length that the system can generate [3]

- Sufficient conditions are that
  (i) the system is controllable, and
  (ii) the input is persistently exciting of sufficient order

[3] Willems' *et al.*, *Syst. Control Lett.*, 2005

Extending the Fundamental Lemma to affine systems

- **Objective:** Guarantee $\textsc{Rank}\begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix} = n + m + 1$ when data are generated by affine dynamics $x^+ = Ax + Bu + c$.

Extending the Fundamental Lemma to affine systems

- **Objective:** Guarantee $\textsc{Rank}\begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix} = n + m + 1$ when data are generated by affine dynamics $x^+ = Ax + Bu + c$.
- Nontrivial extension since
  - Trajectories do not form a linear subspace anymore
  - We must guarantee $\mathbf{1}^\top \notin \textsc{Rowsp}\left[\begin{smallmatrix} X \\ U \end{smallmatrix}\right]$

Extending the Fundamental Lemma to affine systems

- **Objective:** Guarantee $\textsc{Rank}\begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix} = n + m + 1$ when data are generated by affine dynamics $x^+ = Ax + Bu + c$.
- Nontrivial extension since
  - Trajectories do not form a linear subspace anymore
  - We must guarantee $\mathbf{1}^\top \notin \textsc{Rowsp}\begin{bmatrix} X \\ U \end{bmatrix}$
- Transformations that don't work:
  - $x^+ = Ax + \begin{bmatrix} B & I \end{bmatrix}\begin{bmatrix} u \\ c \end{bmatrix} \quad \longrightarrow \quad (A, \begin{bmatrix} B & I \end{bmatrix})$ is controllable but $\begin{bmatrix} u \\ c \end{bmatrix}$ is not *p.e.*
  - $\begin{bmatrix} x^+ \\ y^+ \end{bmatrix} = \begin{bmatrix} A & c \\ 0 & 1 \end{bmatrix}\begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} u \quad \longrightarrow \quad u$ is *p.e.* but $(\begin{bmatrix} A & c \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} B \\ 0 \end{bmatrix})$ is not controllable

Extending the Fundamental Lemma to affine systems

- **Objective:** Guarantee $\textsc{Rank} \begin{bmatrix} X \\ U \\ \mathbf{1}^\top \end{bmatrix} = n + m + 1$ when data are generated by affine dynamics $x^+ = Ax + Bu + c$.
- Nontrivial extension since
  - Trajectories do not form a linear subspace anymore
  - We must guarantee $\mathbf{1}^\top \notin \textsc{Rowsp} \begin{bmatrix} X \\ U \end{bmatrix}$
- Transformations that don't work:
  - $x^+ = Ax + \begin{bmatrix} B & I \end{bmatrix} \begin{bmatrix} u \\ c \end{bmatrix} \quad \longrightarrow \quad (A, \begin{bmatrix} B & I \end{bmatrix})$ is controllable but $\begin{bmatrix} u \\ c \end{bmatrix}$ is not *p.e.*
  - $\begin{bmatrix} x^+ \\ y^+ \end{bmatrix} = \begin{bmatrix} A & c \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} u \quad \longrightarrow \quad u$ is *p.e.* but $(\begin{bmatrix} A & c \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} B \\ 0 \end{bmatrix})$ is not controllable

Proposition

If S is persistently exciting of order K, then for all $K' < K$ it holds that

$$\mathbf{1}^\top \notin \textsc{Rowsp}\, \mathcal{H}_{K'}(S)$$

Willems' Fundamental Lemma for **affine** systems

Theorem

Consider a dataset (X, U, X^+, Y) *of length* d *with* $X^+ = AX + BU + c\mathbf{1}^\top, \; Y = CX + DU + r\mathbf{1}^\top$.

If U *is persistently exciting of order* $n + L + 1$ *and* (A, B) *is a controllable pair, then*

(i) $\textsc{Rank}\begin{bmatrix} \mathcal{H}_1(X_{1:d-L+1}) \\ \mathcal{H}_L(U) \\ \mathbf{1}^\top \end{bmatrix} = n + mL + 1$,

(ii) $(\tilde{X}, \tilde{U}, \tilde{X}^+, \tilde{Y})$ *is a dataset of length* $L \iff$ *there exists* $g \in \mathbb{R}^{d-L+1}$ *such that*

$$\begin{bmatrix} \textsc{Vec}\,\tilde{U} \\ \textsc{Vec}\,\tilde{Y} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathcal{H}_L(U) \\ \mathcal{H}_L(Y) \\ \mathbf{1}^\top \end{bmatrix} g.$$

Willems' Fundamental Lemma for **affine** systems

Theorem

Consider a dataset (X, U, X^+, Y) *of length* d *with* $X^+ = AX + BU + c\mathbf{1}^\top$, $Y = CX + DU + r\mathbf{1}^\top$.

If U *is persistently exciting of order* $n + L + 1$ *and* (A, B) *is a controllable pair, then*

(i) $\textsc{Rank}\begin{bmatrix} \mathcal{H}_1(X_{1:d-L+1}) \\ \mathcal{H}_L(U) \\ \mathbf{1}^\top \end{bmatrix} = n + mL + 1$,

(ii) $(\tilde{X}, \tilde{U}, \tilde{X}^+, \tilde{Y})$ *is a dataset of length* L $\iff$ *there exists* $g \in \mathbb{R}^{d-L+1}$ *such that*

$$\begin{bmatrix} \textsc{Vec}\,\tilde{U} \\ \textsc{Vec}\,\tilde{Y} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathcal{H}_L(U) \\ \mathcal{H}_L(Y) \\ \mathbf{1}^\top \end{bmatrix} g.$$

- (i) $\implies$ (ii) was proven in [4]
- Inspired by state-space proof in [5], we show that higher order *p.e.* + controllability $\implies$ (i)

---

[4] Berberich *et al.*, *IEEE TAC*, 2022
[5] van Waarde *et al.*, *IEEE L-CSS*, 2020

Outline

Conclusion and future work

Takeaway message: The LP approach shows lots of **potential** and is relatively **underexplored**

- Can handle nonlinear, stochastic, data-driven problems
- Flexibility and integration with other methods

Conclusion and future work

Takeaway message: The LP approach shows lots of **potential** and is relatively **underexplored**

- Can handle nonlinear, stochastic, data-driven problems
- Flexibility and integration with other methods

Lot to be done: Nested **approximation architectures** are hard to study

- Performance bounds due to function approximation and constraint sampling
- Exploration logic to guarantee bounded solutions

Thank you!

Andrea Martinelli
Automatic Control Laboratory
ETH Zurich
andremar@ethz.ch