Optimal Data Driven Predictive Control for linear stochastic systems

The thin line between model-based and model-free

Alessandro Chiuso

University of Padova

Joint work with **V. Breschi, S. Formentin, M. Fabris**

ERNSI Workshop – September 25th, 2023

Valentina Breschi

Simone Formentin

Marco Fabris

Thanks to **Keith** and **Florian** for many discussions and new results to come…

Develop a framework for (optimal) data driven control under minimal assumptions

MAIN OUTCOMES:

- Sufficient Statistics for the control problem
- Optimal regualarization schemes found in closed form, no need for hyperparameter tuning
- Provide link with recently proposed procedures (SPC/DeePC/DDPC)

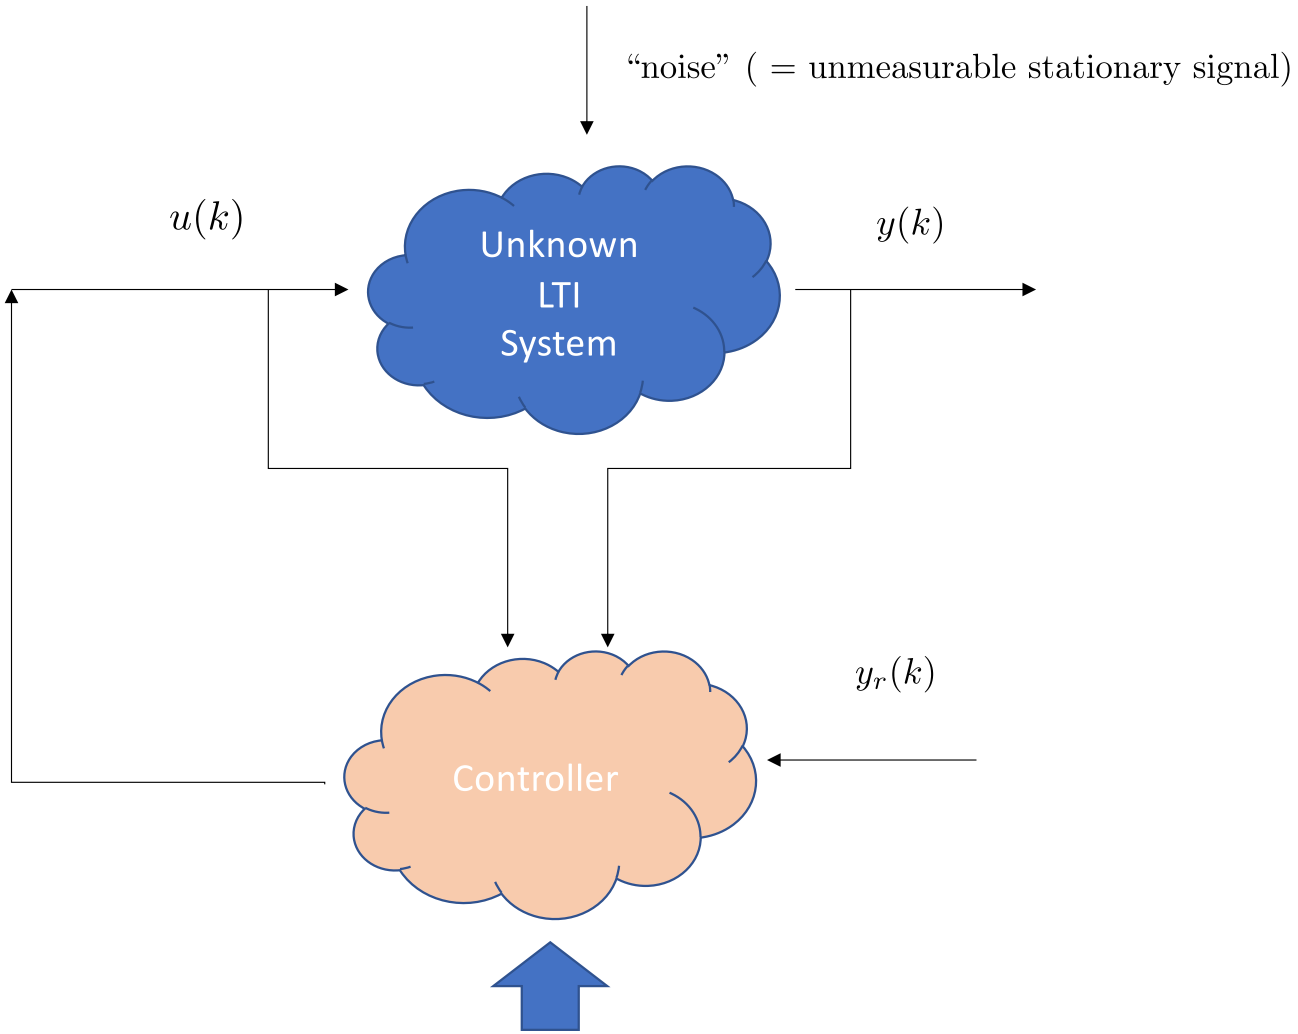

Data $\mathcal{D} := \{z(k) := [u(k); y(k)] : k = 1, .., N_s\}$ **measured historical data, possibly collected in closed-loop**

GENERAL FORMULATION

Given (at any "present" time t)

- $z_{ini} := \{u(s), y(s); s < t\}$: initial conditions
- $\mathcal{D}$: I/O data from the system
- $\mathcal{U}$, $\mathcal{Y}$: constraint sets
- $\ell_t(u_f)$: loss function

$$\ell_t(u_f) = \|y_r - y_f\|_Q^2 + \|u_r - u_f\|_R^2$$

GENERAL FORMULATION

Given (at any "present" time t)

- $z_{ini} := \{u(s), y(s); s < t\}$: initial conditions
- $\mathcal{D}$: I/O data from the system
- $\mathcal{U}$, $\mathcal{Y}$: constraint sets
- $\ell_t(u_f)$: loss function

Find Optimal u_f (future control inputs) such that

1. $\mathbb{E}[\ell_t(u_f)|\mathcal{D}]$ is minimized
2. Constraints are satisfied (possibly probabilistically)

Receding horizon framework: apply first future control and repeat optimization

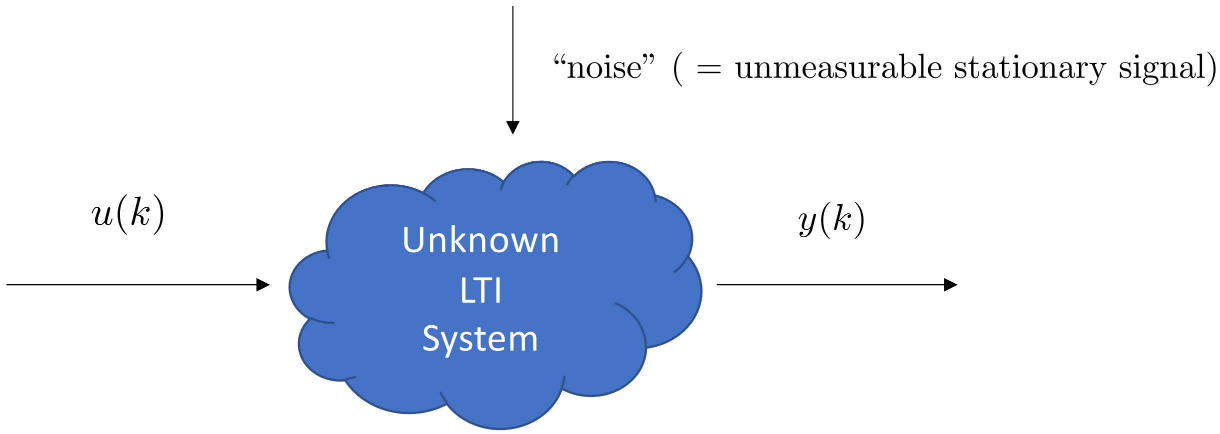

Assumptions:

1. Linear time-invariant predictor

$$y(t) = \hat{y}(t|t-1) + e(t)$$
$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k} \qquad \sum_{k=1}^{\infty} |\phi_k^u| < \infty \quad \sum_{k=1}^{\infty} |\phi_k^y| < \infty$$

2. (Conditional) Martingale difference property with constant conditional variance

$$\mathbb{E}[e(t)|Z_t^-] = 0 \qquad Var\{e(t)|Z_t^-\} = Var\{e(t)\} = \sigma^2$$

$$Z_t^- := \sigma\{z(s) := [u(s); y(s)], s < t\}$$

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

Assumptions:

- Model $\mathcal{M} = \{\phi_k^u, \phi_k^y; k = 1, .., \infty\}$ (unknown, possibly infinite dimensional)
- Prior distribution $\mathcal{M} \sim p(\mathcal{M})$

$$\underbrace{\mathbb{E}[\ell_t(u_f)|\mathcal{D}]}_{= \mathbb{E}[\underbrace{\mathbb{E}[\ell_t(u_f)|\mathcal{M}]}_{L_t(u_f) := \mathbb{E}[\ell_t(u_f)|\mathcal{M}] \propto \|y_r - f(u_f, z_{ini})\|_Q^2 + \|u_r - u_f\|_R^2}|\mathcal{D}]} \quad \Leftarrow \quad \boxed{p(\mathcal{M}|\mathcal{D})}$$

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

Assumptions:

- Model $\mathcal{M} = \{\phi_k^u, \phi_k^y; k = 1, .., \infty\}$ (unknown, possibly infinite dimensional)
- Prior distribution $\mathcal{M} \sim p(\mathcal{M})$

$$\mathbb{E}[\ell_t(u_f)|\mathcal{D}] \quad \Leftarrow \quad \boxed{p(\mathcal{M}|\mathcal{D})}$$

QUESTION: Is this system identification?

INTERMISSION

GENERAL CVaR FORMULATION

$$\hat{u}_f^\star := \arg\min_{u_f \in \mathcal{U}_f, \ell} \overbrace{\mathbb{E}[L_t(u_f) | L_t(u_f) > \ell, \mathcal{D}]}^{CVaR}$$

$$\text{s.t.} \quad \mathbb{P}[L_t(u_f) > \ell | \mathcal{D}] = \alpha.$$

If $\alpha = 1$ $CVaR = \mathbb{E}[L_t(u_f)|\mathcal{D}]$

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

Assumptions:

- Model $\mathcal{M} = \{\phi_k^u, \phi_k^y; k = 1, .., \infty\}$ (unknown, possibly infinite dimensional)
- Prior distribution $\mathcal{M} \sim p(\mathcal{M})$

Remark: Sample Complexity Model Class $>>$ $N_{data} = |\mathcal{D}|$

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

Assumptions:

- Model $\mathcal{M} = \{\phi_k^u, \phi_k^y; k = 1, .., \infty\}$ (unknown, possibly infinite dimensional)
- Prior distribution $\mathcal{M} \sim p(\mathcal{M})$

Remark: Sample Complexity Model Class $>>$ $N_{data} = |\mathcal{D}|$

Long story short (two extreme alternatives)

1. *Model free/nonparametric*: Constrain $\mathcal{M}$ so that sample complexity $\simeq$ $N_{data} = |\mathcal{D}|$ (**LONG ARX MODEL/Nonparametric stable model**)

2. *Model based*: Constrain $\mathcal{M}$ to be of "low (sample) complexity"

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

Assumptions:

- Model $\mathcal{M} = \{\phi_k^u, \phi_k^y; k = 1, .., \infty\}$ (unknown, possibly infinite dimensional)
- Prior distribution $\mathcal{M} \sim p(\mathcal{M})$

Remark: Sample Complexity Model Class $>>$ $N_{data} = |\mathcal{D}|$

Long story short (two extreme alternatives)

1. *Model free/nonparametric*: Constrain $\mathcal{M}$ so that sample complexity $\simeq$ $N_{data} = |\mathcal{D}|$ (**LONG ARX MODEL/Nonparametric stable model**)

2. *Model based*: Constrain $\mathcal{M}$ to be of "low (sample) complexity"

- Spoiler #1: Alternative 1 is basically the route followed by "modern" data-driven MPC approaches (based on Willems' Lemma)

- Spoiler #2: The model selection step in Alternative 2 may be critical

NON INFORMATIVE BAYESIAN FORMULATION

$$\hat{y}(t|t-1) = \sum_{k=1}^{\infty} \phi_k^u u_{t-k} + \phi_k^y y_{t-k}$$

BAYESIAN VARX

$$\hat{y}(t|t-1) \doteq \sum_{k=1}^{\hat{\rho}_N} \phi_k^y y(t-k) + \phi_k^u u(t-k),$$

$$\{\phi_k^u\}_{k\in\mathbb{N}} \sim \mathcal{GP}(0, K_\lambda), \quad \{\phi_k^y\}_{k\in\mathbb{N}} \sim \mathcal{GP}(0, K_\lambda)$$

$$\{\phi_k^u\}_{k\in\mathbb{N}} \perp \{\phi_k^y\}_{k\in\mathbb{N}},$$

$$K_\lambda := \lambda K, \quad \lambda \to \infty.$$

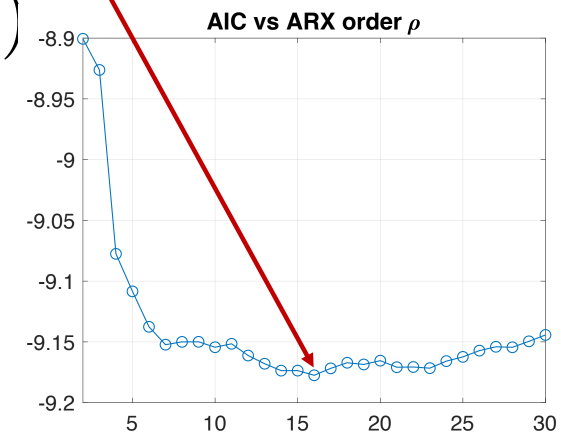

USEFUL FACTS:

$$f(u_f, z_{ini}) = \begin{bmatrix} \hat{y}(t|t-1) \\ \hat{y}(t+1|t-1) \\ \vdots \\ \hat{y}(t+T-1|t-1) \end{bmatrix}$$

$$W := I_{pT} - \Phi_y$$

$$(I_{pT} - \Phi_y) f(u_f, z_{ini}) \doteq \Phi_P z_{ini} + \Phi_u u_f, \qquad y_f - f(u_f, z_{ini}) \doteq (I_{pT} - \Phi_y)^{-1} e_f,$$

$$\Phi_P = \begin{bmatrix} \phi_\rho & \phi_{\rho-1} & \phi_{\rho-2} & \cdots & \cdots & \cdots & \phi_1 \\ 0 & \phi_\rho & \phi_{\rho-1} & \cdots & \cdots & \cdots & \phi_2 \\ 0 & 0 & \phi_\rho & \cdots & \cdots & \cdots & \phi_3 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & \phi_\rho & \cdots & \phi_T \end{bmatrix}.$$

$$\Phi_u = \begin{bmatrix} 0 & 0 & 0 \dots 0 \\ \phi_1^u & 0 & 0 \dots 0 \\ \phi_2^u & \phi_1^u & 0 \dots 0 \\ \vdots & \vdots & \ddots \ddots \vdots \\ \phi_{T-1}^u & \phi_{T-2}^u & \dots \phi_1^u \; 0 \end{bmatrix}, \quad \Phi_y = \begin{bmatrix} 0 & 0 & 0 \dots 0 \\ \phi_1^y & 0 & 0 \dots 0 \\ \phi_2^y & \phi_1^y & 0 \dots 0 \\ \vdots & \vdots & \ddots \ddots \vdots \\ \phi_{T-1}^y & \phi_{T-2}^y & \dots \phi_1^y \; 0 \end{bmatrix},$$

REMARKS:

$$W(y_f - f(z_{ini}, u_f)) = e_f$$

⟺ *Pre-multiplication of prediction error by W whitens the residual*

$$Wf(u_f, z_{ini}) \doteq \Phi_P z_{ini} + \Phi_u u_f$$

⟺ *Weighted predictor is linear in data*

⇓

Define the weighted cost:

$$L_t(u_f) = \|W(y_r - f(u_f, z_{ini}))\|_Q^2 + \|u_r - u_f\|_R^2,$$

MAIN THEOREM 1

$$\mathbb{E}[L_t(u_f)|\mathcal{D}] = J(u_f) + r(u_f),$$

Certainty equivalence cost

$$J(u_f) := \|\overline{W\Delta}_y\|_Q^2 + \|u_r - u_f\|_R^2$$
$$= \|\bar{W}y_r - (\bar{\Phi}_P z_{ini} + \bar{\Phi}_u u_f)\|_Q^2 + \|u_r - u_f\|_R^2,$$

$$\bar{\Phi}_P := \mathbb{E}[\Phi_P|\mathcal{D}], \quad \bar{\Phi}_u := \mathbb{E}[\Phi_u|\mathcal{D}], \quad \bar{W} = \mathbb{E}[W|\mathcal{D}].$$

Regularization

$$r(u_f) := Tr(Q\, Var[W\Delta_y|\mathcal{D}]),$$
$$\Delta_y(u_f) = y_r - f(u_f, z_{ini}).$$

TO BE COMPUTED

MAIN THEOREM 2

$\theta \in \mathbb{R}^{p(m+p)\rho}$ defined stacking all the $\{\phi_k^u, \phi_k^y\}_{k=1,..,\rho}$

$\lambda \to \infty$. (= NON INFORMATIVE PRIOR) *&. p = 1 (for simplicty of exposition)*

$$\hat{\theta} := \mathbb{E}[\theta|\mathcal{D}] = Y_{\rho+1} Z_P^\top (Z_P Z_P^\top)^{-1} \qquad \Sigma_\theta := Var[\theta|\mathcal{D}] = \sigma^2 (Z_P Z_P^\top)^{-1}$$

Usual Block Hankel Matrices

MAIN THEOREM 2

$\theta \in \mathbb{R}^{p(m+p)\rho}$ defined stacking all the $\{\phi_k^u, \phi_k^y\}_{k=1,..,\rho}$

$\lambda \to \infty$. (= NON INFORMATIVE PRIOR) *&. p = 1 (for simplicty of exposition)*

$$\hat{\theta} := \mathbb{E}[\theta|\mathcal{D}] = Y_{\rho+1} Z_P^\top (Z_P Z_P^\top)^{-1}$$

⇓

$$\bar{\Phi}_P := \mathbb{E}[\Phi_P|\mathcal{D}], \quad \bar{\Phi}_u := \mathbb{E}[\Phi_u|\mathcal{D}], \quad \bar{W} = \mathbb{E}[W|\mathcal{D}].$$

$$\Sigma_\theta := Var[\theta|\mathcal{D}] = \sigma^2 (Z_P Z_P^\top)^{-1}$$

⇓

$$Var[W\Delta_y|\mathcal{D}]$$

IMPORTANT REMARKS

1. (Approx) Sufficient statistic for $\mathbb{E}[L_t(u_f)|\mathcal{D}\}]$:

$$T(\mathcal{D}) := \begin{bmatrix} Y_{\rho+1}Z_P^\top & Z_PZ_P^\top \end{bmatrix}$$

= sample moments up to lag ρ

IMPORTANT REMARKS

1. (Approx) Sufficient statistic for $\mathbb{E}[L_t(u_f)|\mathcal{D}\}]$:

$$T(\mathcal{D}) := \begin{bmatrix} Y_{\rho+1} Z_P^\top & Z_P Z_P^\top \end{bmatrix}$$

= sample moments up to lag ρ

2. Same (approx) sufficient statistic for VARX model of order ρ

IMPORTANT REMARKS

1. (Approx) Sufficient statistic for $\mathbb{E}[L_t(u_f)|\mathcal{D}\}]$:

$$T(\mathcal{D}) := \begin{bmatrix} Y_{\rho+1} Z_P^\top & Z_P Z_P^\top \end{bmatrix}$$

   = sample moments up to lag ρ

2. Same (approx) sufficient statistic for VARX model of order ρ

3. To be *statistically efficient* the decision (optimal control design) should be function of $\mathcal{D}$ through a sufficient statistic $\mathcal{T} := T(\mathcal{D})$.

   = No more, no less

Connection with DeePC and co.

SPC

$$\underset{u(k),k\in[t,t+T)}{\text{minimize}} \quad J\left(\begin{bmatrix}\hat{y}_f^d \\ u_f\end{bmatrix}\right)$$

$$\text{s.t. } \alpha^\star = \begin{bmatrix}Z_P \\ U_F\end{bmatrix}^\dagger \begin{bmatrix}z_{init} \\ u_f\end{bmatrix},$$

$$\hat{y}_f^d = \hat{Y}_F \alpha^\star,$$

$$u(k) \in \mathcal{U},\ y^d(k) \in \mathcal{Y},\ k \in [t, t+T),$$

= Subspace Predictive Control (SPC)

+

Constraints

$$J\left(\begin{bmatrix}y_f^d \\ u_f\end{bmatrix}\right) = \frac{1}{2}\left[\sum_{k=t}^{t+T-1} \|y^d(k) - y_r(k)\|_Q^2 + \|u(k)\|_R^2\right],$$

DeePC

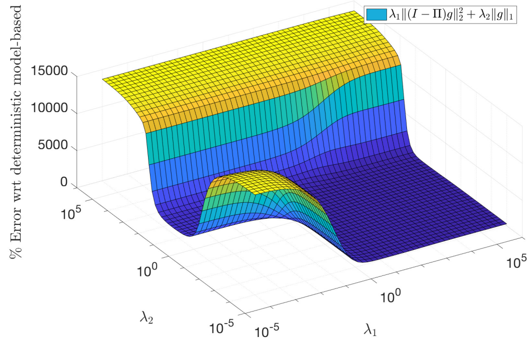

$$
\begin{aligned}
\underset{u_f, y_f^d, \alpha}{minimize} \quad & J\left(\begin{bmatrix} \hat{y}_f^d \\ u_f \end{bmatrix}\right) + \lambda_1 \|\alpha\|_1 + \lambda_2 \|(I - \Pi)\alpha\|_p \\
s.t. \quad & \begin{bmatrix} z_{init} \\ u_f \\ \hat{y}_f^d \end{bmatrix} = \begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} \alpha, \\
& u(k) \in \mathcal{U}, \ y^d(k) \in \mathcal{Y}, \ k \in [t, t+T).
\end{aligned}
$$

WHERE:

$$Y_F := \hat{Y}_F + \tilde{Y}_F$$

$$\hat{Y}_F = Y_F \Pi$$

$$\tilde{Y}_F = Y_F (I - \Pi)$$

γ-DDPC

$$\min_{\gamma_2,\gamma_3} \quad J\left(\begin{bmatrix}\hat{y}_f^d\\ u_f\end{bmatrix}\right) + \beta_2\|\gamma_2\|^2 + \beta_3\|\gamma_3\|^2$$

$$\text{s.t.} \quad \begin{bmatrix}u_f\\ \hat{y}_f^d\end{bmatrix} = \begin{bmatrix}L_{21} & L_{22} & 0\\ L_{31} & L_{32} & L_{33}\end{bmatrix}\begin{bmatrix}\gamma_1^\star\\ \gamma_2\\ \gamma_3\end{bmatrix},$$

$$u(k) \in \mathcal{U},\ \bar{y}(k) \in \mathcal{Y},\ k \in [t, t+T),$$

EXPLOITING LQ DECOMPOSITION OF HANKEL MATRICES

$$\begin{bmatrix}Z_P\\ U_F\\ Y_F\end{bmatrix} = \begin{bmatrix}L_{11} & 0 & 0\\ L_{21} & L_{22} & 0\\ L_{31} & L_{32} & L_{33}\end{bmatrix}\begin{bmatrix}Q_1\\ Q_2\\ Q_3\end{bmatrix}, \quad \hat{Y}_F = \begin{bmatrix}L_{31} & L_{32}\end{bmatrix}\begin{bmatrix}Q_1\\ Q_2\end{bmatrix}$$

DeePC

$\lambda_1 = 0 \quad \lambda_2 = \beta_3$

=

γ-DDPC

$\beta_2 = 0 \quad \beta_3 = \lambda_2$

Proposition:

one to one correspondence between LQ decomposition in γ-DDPC and bank of ARX models

$$\begin{bmatrix} Z_P \\ U_F \\ Y_F \end{bmatrix} = \begin{bmatrix} L_{11} & 0 & 0 \\ L_{21} & L_{22} & 0 \\ L_{31} & L_{32} & L_{33} \end{bmatrix} \begin{bmatrix} Q_1 \\ Q_2 \\ Q_3 \end{bmatrix},$$

$$\hat{y}_f(z_{ini}, u_f) = L_{31} L_{11}^{-1} z_{ini} + L_{32} L_{22}^{-1} \left[u_f - L_{21} L_{11}^{-1} z_{ini} \right]$$

$\Leftrightarrow$

$$Y_{\rho+i} = \Phi_{\gamma,i} \underbrace{\begin{bmatrix} Z_P \\ U_F \\ Y_{[\rho+1:\rho+i]} \end{bmatrix}}_{:= Z_{\gamma,i}} + E_{\rho+i}.$$

$$\hat{\Phi}_{\gamma,i} := Y_{\rho+i} Z_{\gamma,i}^{\top} \left[Z_{\gamma,i} Z_{\gamma,i}^{\top} \right]^{-1}$$

$(i = 1, .., T)$

$$D_{33} := \text{diag}\{L_{33}\}$$
$$\hat{W} := D_{33} L_{33}^{-1},$$

$$\hat{\Phi}_{y,\gamma} = I_{pT} - \hat{W}$$
$$\hat{\Phi}_{u,\gamma} = \hat{W} L_{32} L_{22}^{-1}$$
$$\hat{\Phi}_{P,\gamma} = \hat{W} \left[L_{31} - L_{32} L_{22}^{-1} L_{21} \right] L_{11}^{-1}$$

MAIN THEOREM 3

$$\mathbb{E}[L_t(u_f)|\hat{W}\hat{y}_f(z_{ini}, u_f), \hat{W}] \doteq J_\gamma(u_f) + r_\gamma(u_f),$$

Certainty equivalence cost

Regularization

$$J_\gamma(u_f) = \|\hat{W}(y_r - \hat{y}_f(z_{ini}, u_f))\|\|_Q^2 + \|u_r - u_f\|_R^2,$$

$$r_\gamma(u_f) = \mathrm{Tr}(QVar[\tilde{n}_f(z_{init}, u_f)]),$$

$$\tilde{n}_f(z_{ini}, u_f) := \Delta_W(y_r - \hat{y}_f(z_{ini}, u_f)) - \tilde{e}_f(z_{ini}, u_f)$$

$$\begin{aligned} \mathbb{E}[Wy_r|\mathcal{D}_\gamma] &= \hat{W}y_r \\ \mathbb{E}[Wf(z_{ini}, u_f)|\mathcal{D}_\gamma] &= \hat{W}\hat{y}_f(z_{ini}, u_f) \end{aligned}$$

TO BE COMPUTED

MAIN THEOREM 4

$$r_\gamma(u_f) = \mathrm{Tr}(QVar[\tilde{n}_f(z_{init}, u_f)]),$$

$$\tilde{n}_f(z_{ini}, u_f) := \cancel{\Delta_W(y_r - \hat{y}_f(z_{ini}, u_f))} - \tilde{e}_f(z_{ini}, u_f)$$

When the training input is white

$$r_\gamma(u_f) \xrightarrow[N\to\infty]{} \frac{qpT\sigma^2}{N}\|\gamma_2\|^2 \longleftrightarrow \beta_2\|\gamma_2\|^2$$

i.e. (sub)optimal regularization is the one used in γ-DDPC with $\beta_2 = \frac{qpT\sigma^2}{N}$, $\beta_3 = \infty$.

IMPORTANT REMARK

Theorem 1 [Optimal] *(Structured and Causal Predictor)*

$$\mathbb{E}[L_t(u_f)|\mathcal{D}] = J(u_f) + r(u_f),$$

Theorem 3 [$\mathbf{DeePC}/\gamma$**-DDPC**] *(UnStructured and NOT Causal Predictor)*

$$\mathbb{E}[L_t(u_f)|\hat{W}\hat{y}_f(z_{ini}, u_f), \hat{W}] \doteq J_\gamma(u_f) + r_\gamma(u_f),$$

Simulation results

$$(A,B) = \left(\begin{bmatrix} 1.4183 & -1.5894 & 1.3161 & -0.8864 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right);$$

$$(C,D) = \left(\begin{bmatrix} 0 & 0 & 0.2826 & 0.5067 \end{bmatrix}, 0 \right).$$

Case 1: - white noise input
- white noise measurement error

Case 2: - coloured noise input
- coloured noise measurement error

Case 1

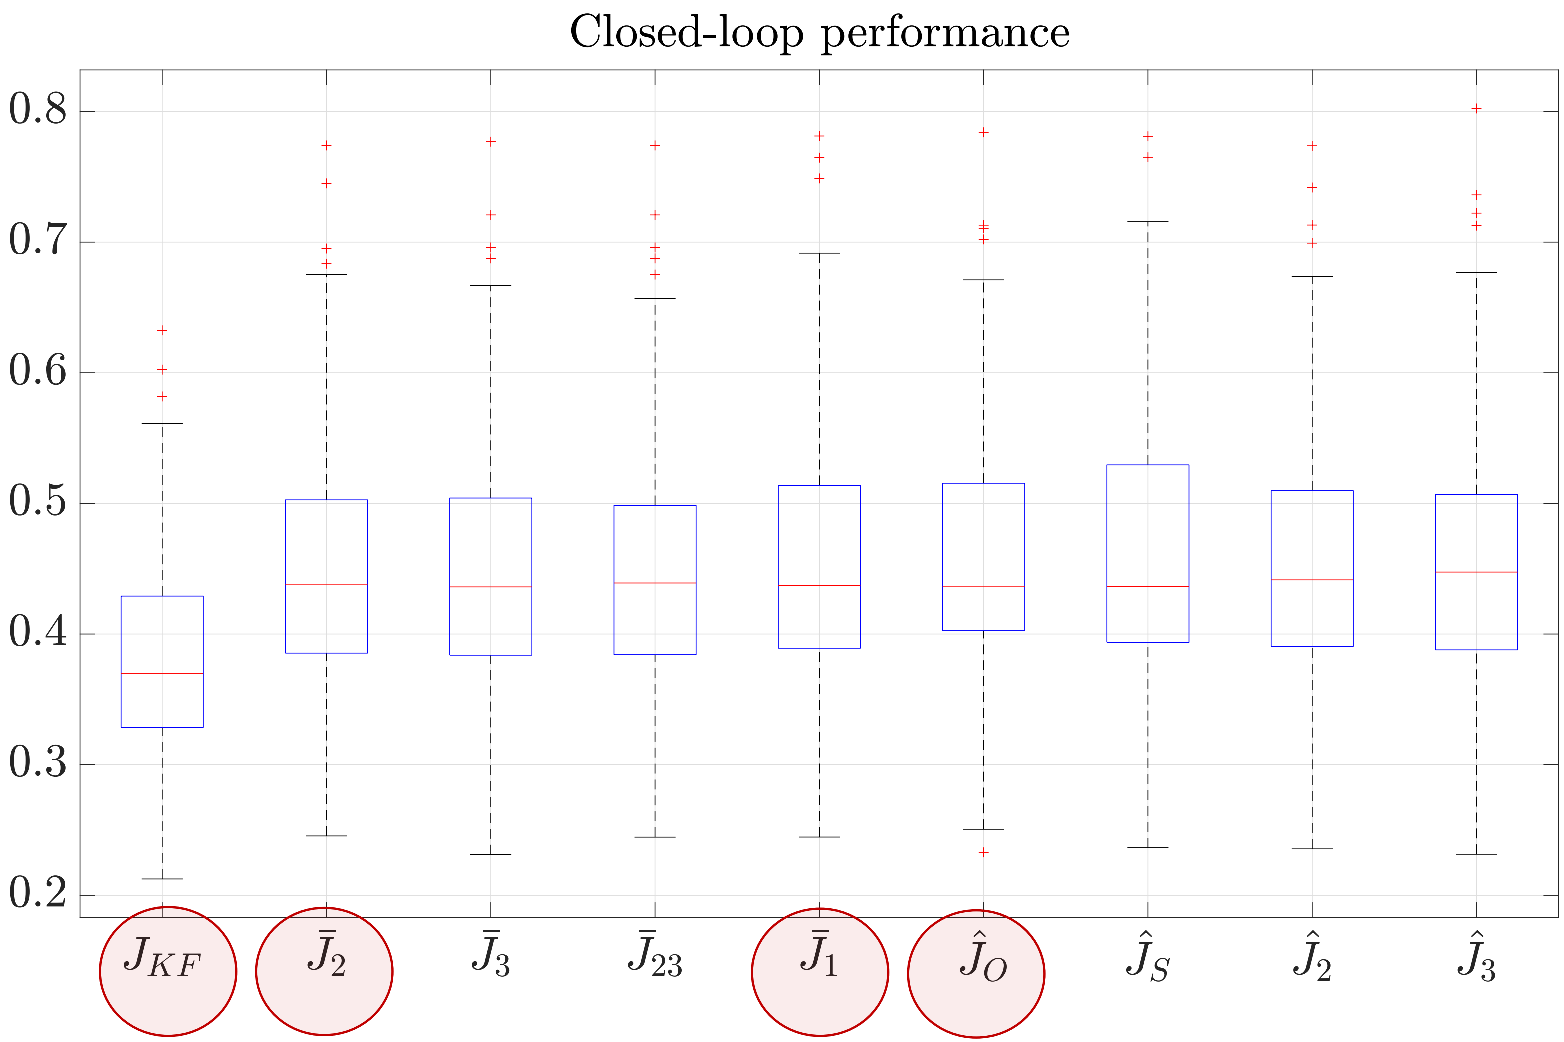

Closed-loop performance
0.8
0.7
0.6
0.5
0.4
0.3
0.2
J_{KF}
$\bar{J}_2$
$\bar{J}_3$
$\bar{J}_{23}$
$\bar{J}_1$
$\hat{J}_O$
$\hat{J}_S$
$\hat{J}_2$
$\hat{J}_3$

Case 2

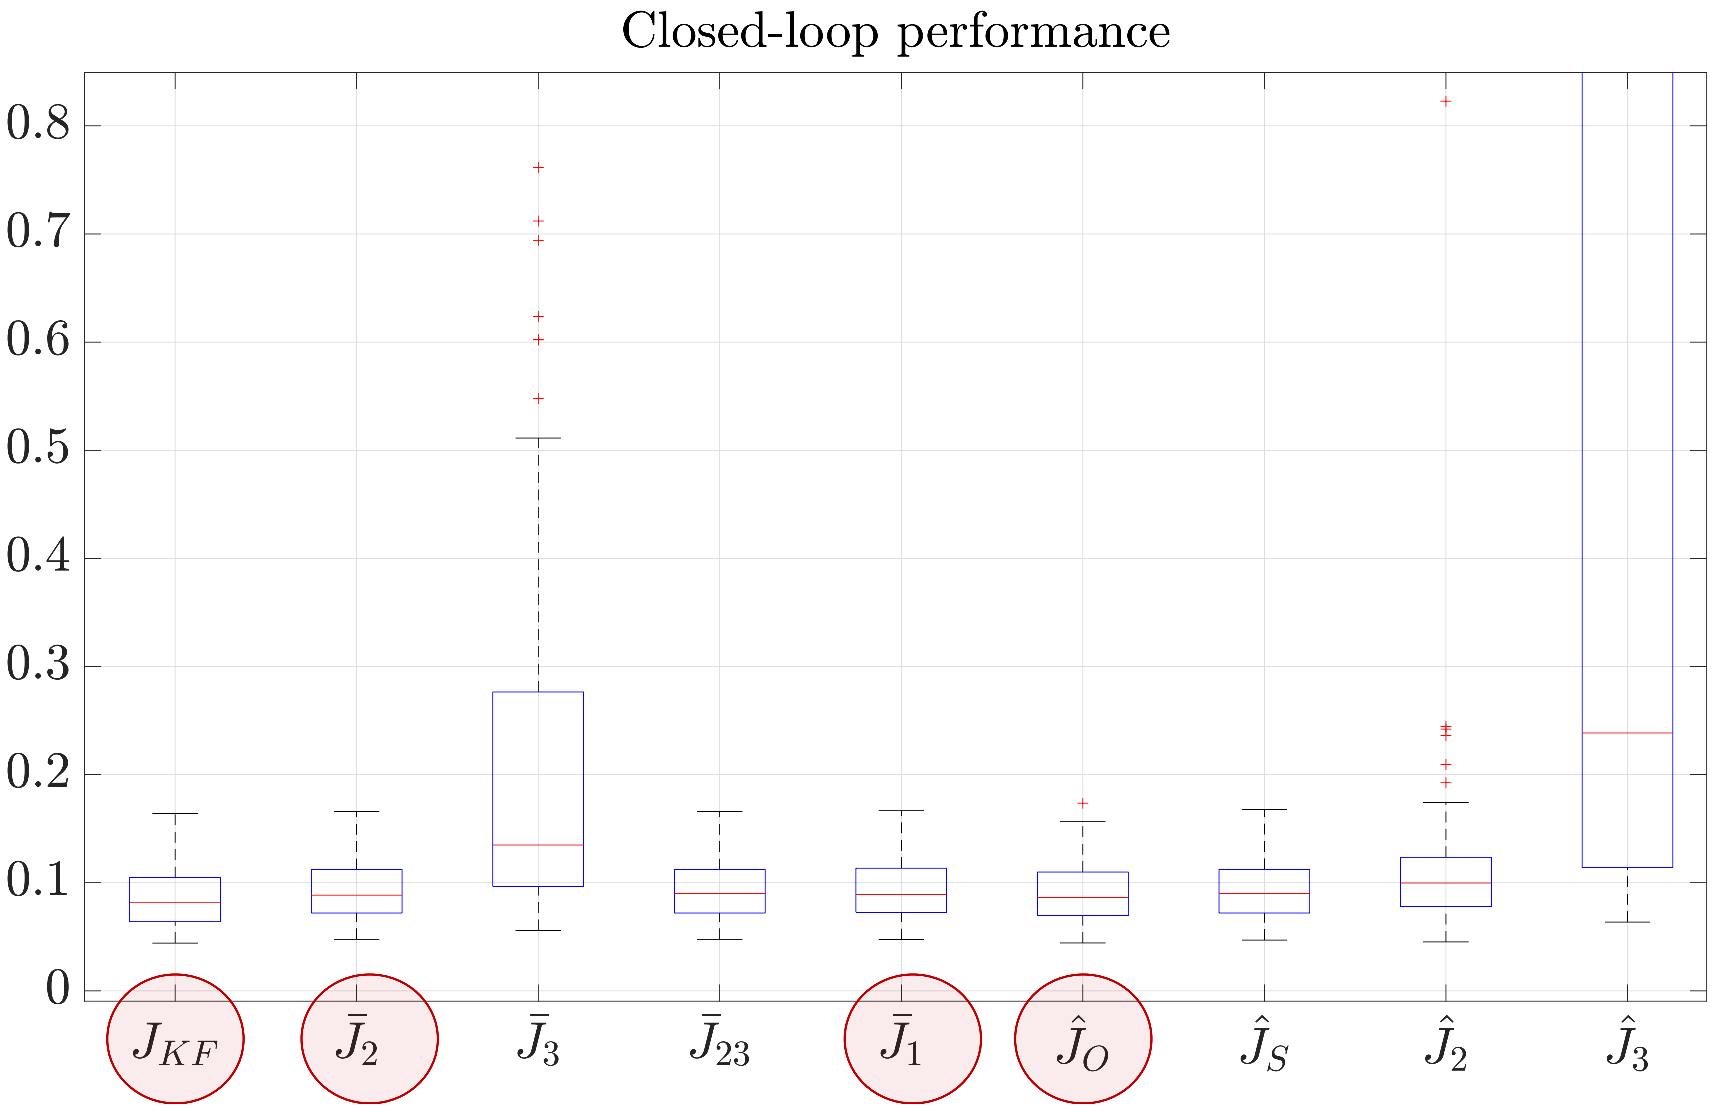

Closed-loop performance
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
J_{KF}
$\bar{J}_2$
$\bar{J}_3$
$\bar{J}_{23}$
$\bar{J}_1$
$\hat{J}_O$
$\hat{J}_S$
$\hat{J}_2$
$\hat{J}_3$

a	average execution time $\pm$ std. dev. $[ms]$
KF	5.83 ± 0.75
2, 3 (parallel search on 2 linear grids of length 202 + 2 offline runs)	492.23 ± 74.67
23 (parallel search on a 202×202 squared grid + offline run)	$(35.35 \pm 6.08) \cdot 10^3$
1 (parallel search on an 11×11 squared grid + offline run)	$(862.33 \pm 131.20) \cdot 10^3$
O (online run)	23.10 ± 6.06
S (online run)	9.10 ± 1.70
2 (online run)	344.07 ± 120.33
3 (online run)	412.62 ± 195.58

Relative frequency of estimated past length (initial condition)

Case 1

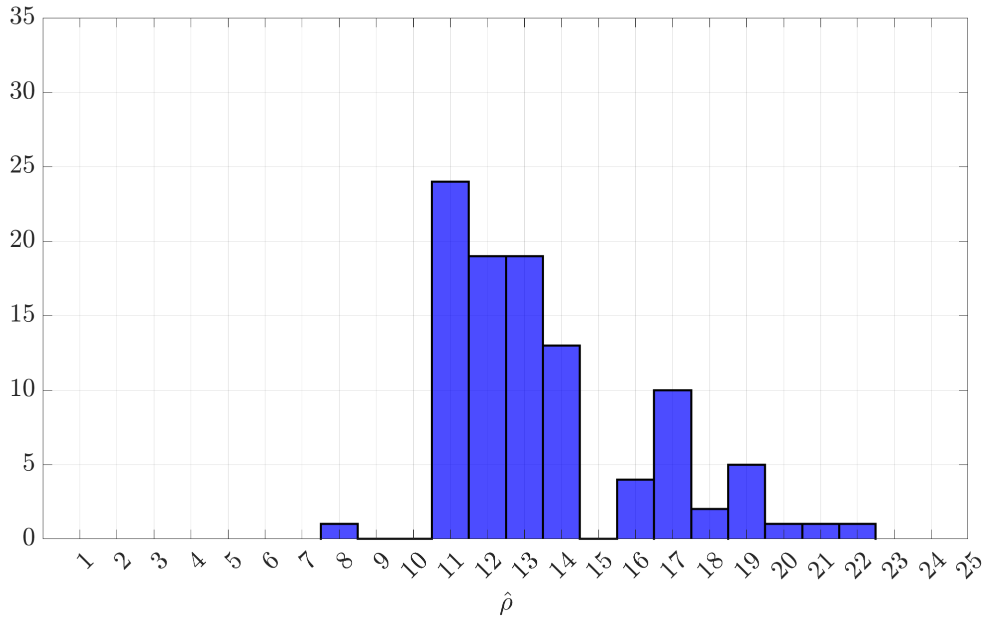

Case 2

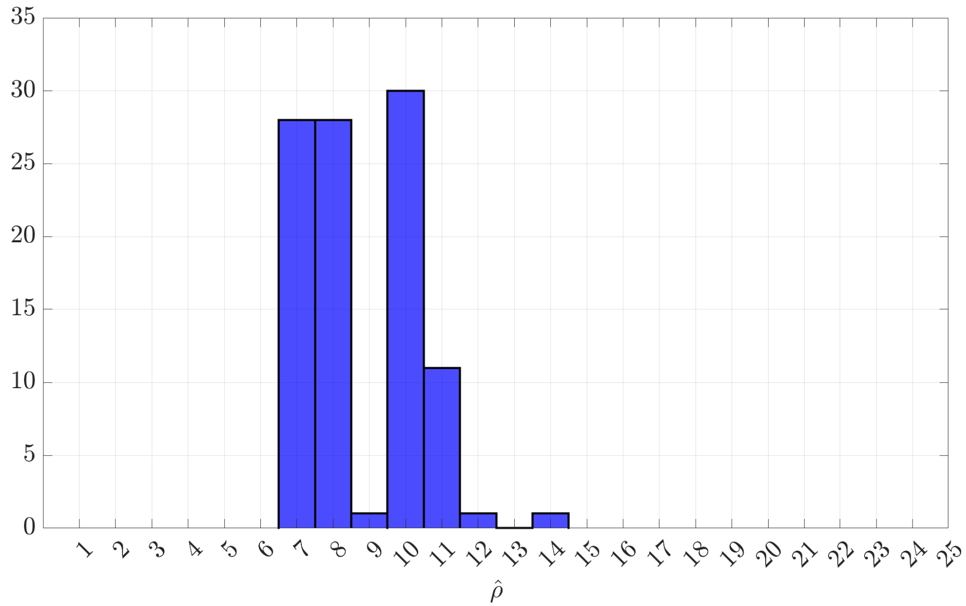

Simulation results – Closed Loop Data
[Work in progress – with K. Moffat and F. Dörfler]

True System: Double Integrator, stabilized with proportional state feedback to gather closed loop data

Simulation results – Closed Loop Data

[Work in progress – with K. Moffat and F. Dörfler]

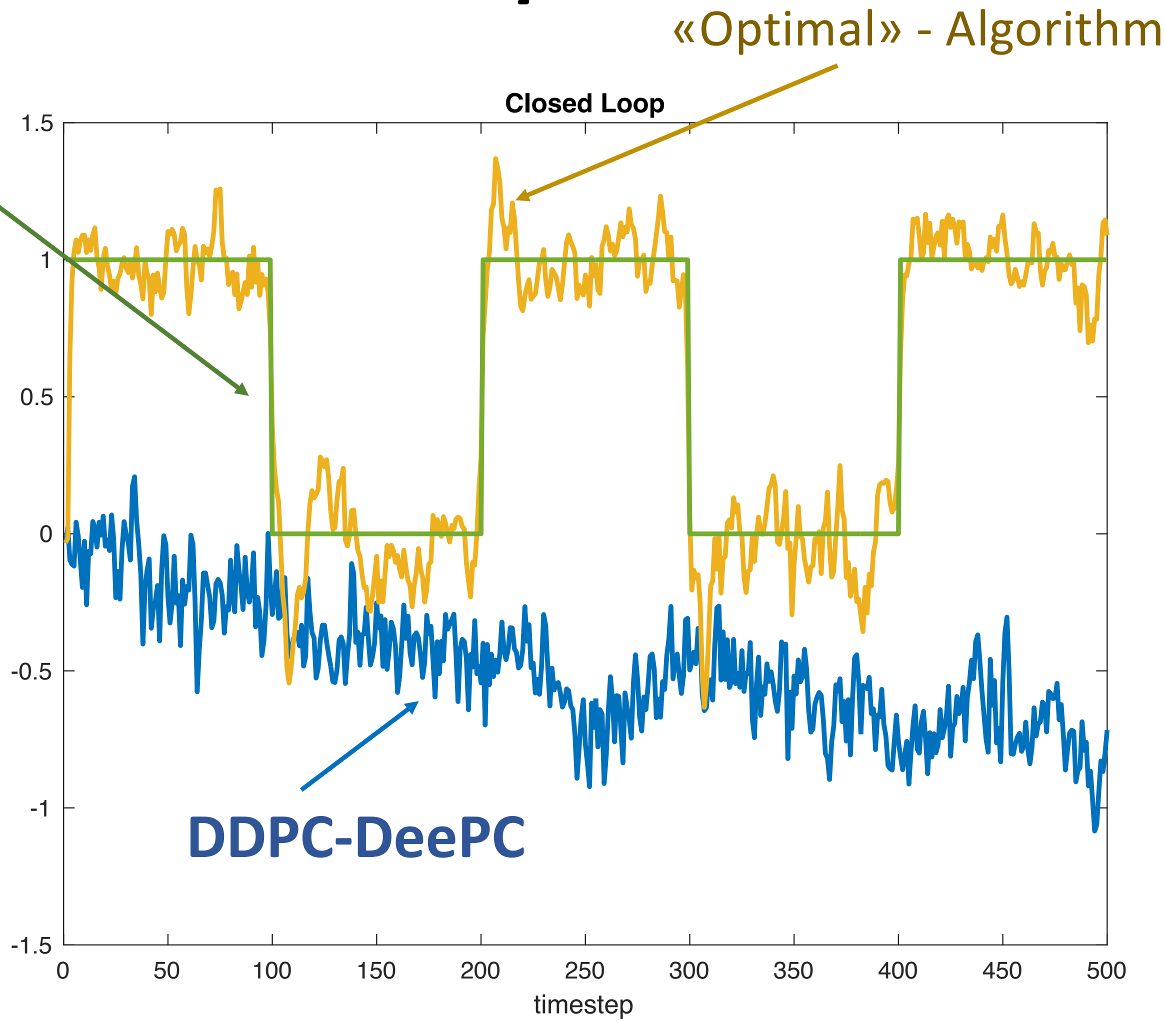

True System: Double Integrator, stabilized with proportional state feedback to gather closed loop data

TAKE HOME

- Direct Data Driven Predictive Control (Hankel Data Matrices) = VARX modeling + Reg. Optimization
- (Approx) Sufficient Statistics $Y_{\rho+1}Z_P^\top$, $Z_P Z_P^\top$
  (No reason to use Page matrices or similar)
- Closed form Optimal Regularization (no tuning): quadratic in u_f
- Bayesian framework allows to embed prior info (stability/further regularization etc.)