


# Data-Enabled Predictive Control of Autonomous Energy Systems

Florian Dörfler

ETH Zürich


DTU Summer  
School 2023



# Acknowledgements



Jeremy Coulson



Linbin Huang



John Lygeros



Ivan Markovsky

Further:

Ezzat Elokda,  
Paul Beuchat,  
Daniele Alpagò,  
Jianzhe (Trevor) Zhen,  
Claudio de Persis,  
Pietro Tesi,  
Henk van Waarde,  
Eduardo Prieto,  
Saverio Bolognani,  
Andrea Favato,  
Paolo Carlet,  
Andrea Martin,  
Luca Furieri,  
Giancarlo Ferrari-Trecate,  
Keith Moffat,

...

& many master students

# Big, deep, intelligent & so on

- **unprecedented availability** of computation, storage, & data
  - **theoretical advances** in statistics, optimization, & machine learning
  - ... and **big-data** frenzy
- increasing importance of **data-centric methods** in all of science / engineering

Make up your own opinion, but machine learning works too well to be ignored.



## From Pixels to Torques: Policy Learning with Deep Dynamical Models

Niklas Wahlström

Division of Automatic Control, Linköping University, Linköping, Sweden

NIKWA@ISY.LIU.SE

Thomas B. Schön

Department of Information Technology, Uppsala University, Sweden

THOMAS.SCHON@IT.UU.SE

Marc Peter Deisenroth

Department of Computing, Imperial College London, United Kingdom

M.DEISENROTH@IMPERIAL.AC.UK

 NVIDIA DEVELOPER

NVIDIA Developer Blog

## End-to-End Deep Learning for Self-Driving Cars

By Mariusz Bojarski, Ben Firsirot, Boxi Feng, Larry Jackel, Urs Muller, Karol Szepes and Davide Del Testa | August 17, 2016




# Thoughts on data-driven control

- **indirect data-driven control** via models:  
data  $\xrightarrow{\text{SysID}}$  model + uncertainty  $\rightarrow$  control
- growing trend: **direct data-driven control**  
by-passing models ... (again) hyped, why?

The direct approach is **viable alternative**

- for some **applications**: model-based approach is too complex to be useful  
 $\rightarrow$  too complex models, environments, sensing modalities, specifications (e.g., wind farm)
- due to (well-known) **shortcomings of ID**  
 $\rightarrow$  too cumbersome, models not identified for control, incompatible uncertainty estimates, ...
- when **brute force** data/compute available



**Central promise:** *It is often easier to learn a control policy from data rather than a model.*

**Example** 1973: autotuned PID



# Abstraction reveals pros & cons

## *indirect* (model-based) *data-driven control*

minimize control cost  $(u, x)$   
subject to  $(u, x)$  satisfy state-space model  
where  $x$  estimated from  $(u, y)$  & model  
where model identified from  $(u^d, y^d)$  data

} outer optimization } separation & certainty  
} middle opt. } equivalence  
} inner opt. } ( $\rightarrow$  LQG case)  
} no separation } ( $\rightarrow$  ID-4-control)

$\rightarrow$  nested multi-level optimization problem

## *direct* (black-box) *data-driven control*

minimize control cost  $(u, y)$   
subject to  $(u, y)$  consistent with  $(u^d, y^d)$  data


$\rightarrow$  *trade-offs*

modular vs. end-2-end  
suboptimal (?) vs. optimal  
convex vs. non-convex (?)

Additionally: account for *uncertainty* (hard to propagate in indirect approach)

# Indirect (models) vs. direct (data)

- models are useful for design & beyond
- modular → easy to debug & interpret
- id = noise filtering
- id = projection on model class
- harder to propagate uncertainty through id
- no robust separation principle → suboptimal
- ...



- some models too complex to be useful
- end-to-end → suitable for non-experts
- design handles noise
- harder to inject side info but no bias error
- transparent: no unmodeled dynamics
- possibly optimal but often less tractable
- ...

there are *no universal conclusions* & plenty of counterexamples


# A direct approach: dictionary + MPC

## ① trajectory *dictionary learning*

- motion primitives / basis functions
- theory: Koopman & Liouville  
practice: (E)DMD & particles

## ② *MPC* optimizing over dictionary span


- huge *theory vs. practice* gap
- back to basics: *impulse response*



- $\left. \begin{matrix} u_0 = s_0 \\ x_0 = 0 \end{matrix} \right\} \leadsto y(t) = \text{impulse response} =: g(t) = \{g_0, g_1, g_2, \dots\}$
- response to any other input  $u(t)$  given by convolution with  $g(t)$ :  

$$y(t) = g(t) * u(t) = \sum_{\tau=0}^{t-1} g(t-\tau) u(\tau) = g(0)u(t) + g(1)u(t-1) + \dots$$

↑ convolution



Now what if we had the impulse response recorded in our data-library?

$$[g_0 \quad g_1 \quad g_2 \quad \dots] = [y_0^d \quad y_1^d \quad y_2^d \quad \dots]$$

→ can predict any future input  $y_{\text{future}}(t)$  to any future input  $u_{\text{future}}(t)$  by convolution with  $g(t) = y^d(t)$ :  $y_{\text{future}}(t) = y_0^d u_{\text{future}}(t) + y_1^d u_{\text{future}}(t-1) + \dots$

→ dynamic matrix control

(Shell, 1970s): **predictive control from raw data**

in vector form

$$y_{\text{future}}(t) = [y_0^d \quad y_1^d \quad y_2^d \quad \dots] \cdot \begin{bmatrix} u_{\text{future}}(t) \\ u_{\text{future}}(t-1) \\ u_{\text{future}}(t-2) \\ \vdots \end{bmatrix}$$

**today**: arbitrary, finite, & corrupted data, ... stochastic & nonlinear ?

# Today's menu

1. behavioral system theory: *fundamental lemma*
2. *DeePC*: data-enabled predictive control
3. robustification via salient *regularizations*
4. cases studies from *wind* & *power systems*

*blooming literature* (2-3 ArXiv / week)


→ survey & tutorial to get started:

## DATA-DRIVEN CONTROL BASED ON BEHAVIORAL APPROACH: FROM THEORY TO APPLICATIONS IN POWER SYSTEMS

Ivan Markovsky, Linbin Huang, and Florian Dörfler

I. Markovsky is with ICREA, Pg. Lluís Companys 23, Barcelona, and CIMNE, Gran Capitán, Barcelona, Spain (e-mail: imarkovsky@cimne.upc.edu).

L. Huang and F. Dörfler are with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland (e-mails: linhuang@ethz.ch, dorfler@ethz.ch).



[[link](#)] to related publications

# Organization of the “school”

- I will **teach the basics** & provide pointers to more sophisticated research material → study cutting-edge papers yourself
- it's a school: so we will spend time on the **board** → take notes

take notes in these boxes: hello world 😊

- We teach this material also in the ETH Zürich bachelor & have plenty of **background material** + implementation experience → please reach out if you need anything
- we will take a **break** after 90 minutes → coffee ☺

# Preview


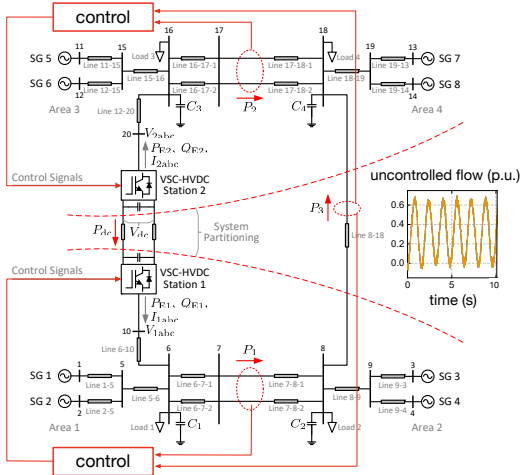
**complex** 4-area power **system**:

large ( $n=208$ ), few sensors (8),  
nonlinear, noisy, stiff, input  
constraints, & decentralized  
control specifications

**control objective**: oscillation

damping without a model

(grid has many owners, models are  
proprietary, operation in flux, ...)



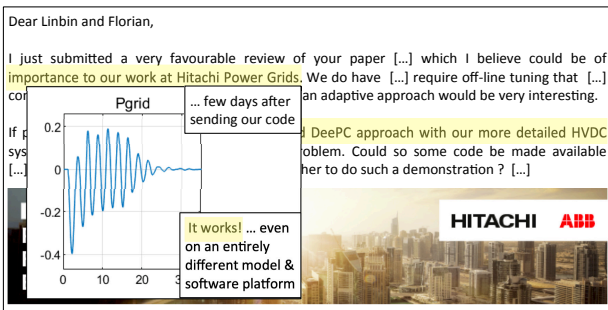
seek a method that **works reliably**, can be **efficiently** implemented, & **certifiable**

→ automating ourselves

# Reality check: black magic or hoax ?

surely, nobody would put apply such a **shaky data-driven method**

- on the **world's most complex engineered system** (the electric grid),
- using the **world's biggest actuators** (Gigawatt-sized HVDC links),
- and subject to **real-time, safety, stability, constraints** ... right?



at least someone believes that our method is practically useful ...



# LTI system representations

- <sup>exogenous</sup> ARX model:  $y(t+2) + 2y(t+1) + 3y(t) = 4u(t)$   
auto regressive
- ARX to state space:  $x(t) = \begin{bmatrix} y(t) \\ y(t+1) \end{bmatrix} \leadsto \begin{cases} \underbrace{\tilde{x}^{t+1}}_{x(t+1)} = \underbrace{\begin{bmatrix} 0 & 1 \\ -3 & -2 \end{bmatrix}}_{\substack{F \\ A}} \underbrace{x}_{x(t)} + \underbrace{\begin{bmatrix} 0 \\ 4 \end{bmatrix}}_{\substack{F \\ B}} u \\ y = \underbrace{\begin{bmatrix} 1 & 0 \end{bmatrix}}_{C} x \end{cases}$
- ARX to transfer function via z-transform  
 $y(t) \rightarrow Y(z)$   
 $y(t+1) \rightarrow z \cdot Y(z)$   
 $\leadsto Y(z) = \frac{4}{z^2 + 2z + 3} U(z)$

These are all parametric representations of an LTI system.

They are all so-called kernel representations since they multiply signals. E.g., for the ARX model, let  $z$  denote the forward time shift, e.g.,  $z \cdot y(t) = y(t+1)$ , then we can write the model as

$$\underbrace{[z^2 + 2z + 3, 4]}_{\text{"kernel"}} \cdot \underbrace{\begin{bmatrix} y(t) \\ u(t) \end{bmatrix}}_{\text{signal}} = 0$$

# Behavioral view on dynamical systems



**Definition:** A discrete-time **dynamical system** is a 3-tuple  $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$  where

- (i)  $\mathbb{Z}_{\geq 0}$  is the *discrete-time axis*,
  - (ii)  $\mathbb{W}$  is the *signal space*, &
  - (iii)  $\mathcal{B} \subseteq \mathbb{W}^{\mathbb{Z}_{\geq 0}}$  is the *behavior*.
- }  $\mathcal{B}$  is the set of all trajectories

**Definition:** The dynamical system  $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathcal{B})$  is

- (i) **linear** if  $\mathbb{W}$  is a vector space &  $\mathcal{B}$  is a subspace of  $\mathbb{W}^{\mathbb{Z}_{\geq 0}}$
- (ii) & **time-invariant** if  $\mathcal{B} \subseteq \sigma \mathcal{B}$ , where  $\sigma w_t = w_{t+1}$ .

LTI system = shift-invariant subspace of trajectory space  
→ abstract perspective suited for **data-driven control**



# Properties of the LTI trajectory space

- solution to  $\dot{x} = A x + B u$  with  $x(0) = x_{ini}$  is  
 $x(0) = x_{ini}, x(1) = A x_{ini} + B u(0), x(2) = A^2 x_{ini} + A B u(0) + B u(1), \dots$   
 $\leadsto x(t) = A^t x_{ini} + \sum_{\tau=0}^{t-1} A^{\tau} B u(t-\tau-1)$

- solution to  $\dot{x} = A x + B u, y = C x + D u, x(0) = x_{ini}$  is

$$y(t) = C A^t x_{ini} + \underbrace{\sum_{\tau=0}^{t-1} C A^{\tau} B u(t-\tau-1) + D u(t)}_{g(t) * u(t)}$$

- in vector notation for  $t \in \{0, 1, \dots, T\}$

$$\begin{bmatrix} y(0) \\ y(1) \\ y(2) \\ \vdots \\ y(T) \end{bmatrix} = \underbrace{\begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^T \end{bmatrix}}_{O_T} x_{ini} + \underbrace{\begin{bmatrix} D \\ CB \ D \\ CAB \ CB \ D \\ \ddots & \ddots & \ddots \\ CA^{T-1}B \dots CAB \ CB \ D \end{bmatrix}}_{G_T} \begin{bmatrix} u(0) \\ u(1) \\ u(2) \\ \vdots \\ u(T) \end{bmatrix}$$

$O_T$ : extended observability matrix

$G_T$ : convolution (or impulse response) matrix

- compactly:  $y = \underbrace{\mathcal{O}_T}_{L \in \mathbb{R}^{p^T \times n}} x_{ini} + \mathcal{C}_T u \quad (*)$

- observability:  $x_{ini}$  can be uniquely reconstructed from  $(*)$   
 $\Leftrightarrow n = \text{rank } \mathcal{O}_T = \text{rank} \begin{bmatrix} \mathcal{C} \\ \mathcal{C}A \\ \vdots \\ \mathcal{C}A^{T-1} \end{bmatrix}$

$\rightarrow$  the smallest  $T$  so that  $\text{rank } \mathcal{O}_T = n$  is called the lag  $\ell$  of the system. In the SISO case ( $p=1$ ), we have  $\ell = n$ ; for  $p \geq 1$ , we have  $\ell \leq n$  (MIMO case).

$\rightarrow$  given past data  $y_{ini} \in \mathbb{R}^{pT_{ini}}$ ,  $u_{ini} \in \mathbb{R}^{mT_{ini}}$  of length

$T_{ini} \geq \ell$ , we can uniquely recover  $x_{ini}$  from

$$y_{ini} = \mathcal{O}_{T_{ini}} x_{ini} + \mathcal{C}_{T_{ini}} u_{ini}$$

- dimension of the LTI trajectory space: for any  $x_{ini}$  what is the dimension  $y = \begin{bmatrix} y(0) \\ \vdots \\ y(T) \end{bmatrix}$  and  $u = \begin{bmatrix} u(0) \\ \vdots \\ u(T) \end{bmatrix}$ ?

$$\begin{array}{c} \begin{matrix} \nearrow \in \mathbb{R}^{mT} \\ \downarrow \in \mathbb{R}^{pT} \end{matrix} \begin{bmatrix} u \\ y \end{bmatrix} = \begin{bmatrix} 0 & I \\ \mathcal{A}_T & \mathcal{C}_T \end{bmatrix} \begin{bmatrix} x_{ini} \\ u \end{bmatrix} \begin{matrix} \in \mathbb{R}^n \\ \in \mathbb{R}^{mT} \end{matrix}$$

$\uparrow$  this column has full rank  $n$   
 $\uparrow$  this column has full rank  $mT$  provided  $T \geq \ell$

$\leadsto$  dimension  $\begin{bmatrix} u \\ y \end{bmatrix} = mT + n$  for  $T \geq \ell$

# LTI systems & matrix time series

foundation of state-space subspace system ID & signal recovery algorithms



$(u(t), y(t))$  satisfy recursive  
**difference equation**

$$b_0 u_t + b_1 u_{t+1} + \dots + b_n u_{t+n} + a_0 y_t + a_1 y_{t+1} + \dots + a_n y_{t+n} = 0$$

(ARX / kernel representation)



$[0 \ b_0 \ a_0 \ b_1 \ a_1 \ \dots \ b_n \ a_n \ 0]$  in left nullspace  
of **trajectory matrix** (collected data)

$$\mathcal{H} \begin{pmatrix} u^d \\ y^d \end{pmatrix} = \begin{bmatrix} \begin{pmatrix} u_{1,1}^d \\ y_{1,1}^d \end{pmatrix} & \begin{pmatrix} u_{1,2}^d \\ y_{1,2}^d \end{pmatrix} & \begin{pmatrix} u_{1,3}^d \\ y_{1,3}^d \end{pmatrix} & \dots \\ \begin{pmatrix} u_{2,1}^d \\ y_{2,1}^d \end{pmatrix} & \begin{pmatrix} u_{2,2}^d \\ y_{2,2}^d \end{pmatrix} & \begin{pmatrix} u_{2,3}^d \\ y_{2,3}^d \end{pmatrix} & \dots \\ \vdots & \vdots & \vdots & \vdots \\ \begin{pmatrix} u_{T,1}^d \\ y_{T,1}^d \end{pmatrix} & \begin{pmatrix} u_{T,2}^d \\ y_{T,2}^d \end{pmatrix} & \begin{pmatrix} u_{T,3}^d \\ y_{T,3}^d \end{pmatrix} & \dots \end{bmatrix}$$



under assumptions

where  $y_{t,i}^d$  is  $t$ th sample from  $i$ th experiment

# Fundamental Lemma



Given: data  $\begin{pmatrix} u_i^d \\ y_i^d \end{pmatrix} \in \mathbb{R}^{m+p}$  & LTI complexity parameters  $\begin{cases} \text{lag } \ell \\ \text{order } n \end{cases}$

set of all  $T$ -length trajectories =

$$\left\{ (u, y) \in \mathbb{R}^{(m+p)T} : \exists x \in \mathbb{R}^{nT} \text{ s.t.} \right.$$

$$\left. x^+ = Ax + Bu, y = Cx + Du \right\}$$

parametric state-space model

$\equiv$

$$\text{colspan} \begin{bmatrix} \begin{pmatrix} u_{1,1}^d \\ y_{1,1}^d \end{pmatrix} & \begin{pmatrix} u_{1,2}^d \\ y_{1,2}^d \end{pmatrix} & \begin{pmatrix} u_{1,3}^d \\ y_{1,3}^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_{2,1}^d \\ y_{2,1}^d \end{pmatrix} & \begin{pmatrix} u_{2,2}^d \\ y_{2,2}^d \end{pmatrix} & \begin{pmatrix} u_{2,3}^d \\ y_{2,3}^d \end{pmatrix} & \cdots \\ \vdots & \vdots & \vdots & \vdots \\ \begin{pmatrix} u_{T,1}^d \\ y_{T,1}^d \end{pmatrix} & \begin{pmatrix} u_{T,2}^d \\ y_{T,2}^d \end{pmatrix} & \begin{pmatrix} u_{T,3}^d \\ y_{T,3}^d \end{pmatrix} & \cdots \end{bmatrix}$$

raw data (every column is an experiment)

if and only if the trajectory matrix has rank  $m \cdot T + n$  for all  $T > \ell$



$$\begin{aligned} \text{set of all } T\text{-length trajectories} = \\ \left\{ (u, y) \in \mathbb{R}^{(m+p)T} : \exists x \in \mathbb{R}^{nT} \text{ s.t. } \right. \\ \left. x^+ = Ax + Bu, y = Cx + Du \right\} \end{aligned} = \text{colspan} \begin{bmatrix} \begin{pmatrix} u_{1,1}^d \\ y_{1,1}^d \end{pmatrix} & \begin{pmatrix} u_{1,2}^d \\ y_{1,2}^d \end{pmatrix} & \begin{pmatrix} u_{1,3}^d \\ y_{1,3}^d \end{pmatrix} & \dots \\ \begin{pmatrix} u_{2,1}^d \\ y_{2,1}^d \end{pmatrix} & \begin{pmatrix} u_{2,2}^d \\ y_{2,2}^d \end{pmatrix} & \begin{pmatrix} u_{2,3}^d \\ y_{2,3}^d \end{pmatrix} & \dots \\ \vdots & \vdots & \vdots & \vdots \\ \begin{pmatrix} u_{T,1}^d \\ y_{T,1}^d \end{pmatrix} & \begin{pmatrix} u_{T,2}^d \\ y_{T,2}^d \end{pmatrix} & \begin{pmatrix} u_{T,3}^d \\ y_{T,3}^d \end{pmatrix} & \dots \end{bmatrix}$$

all trajectories constructible from finitely many previous trajectories

- **standing on the shoulders of giants:**

classic Willems' result was only "if" & required further assumptions: Hankel, persistency of excitation, controllability

A note on persistency of excitation

Jan C. Willems<sup>a</sup>, Paolo Rapisarda<sup>b</sup>, Ivan Markovsky<sup>a,\*</sup>, Bart L.M. De Moor<sup>a</sup>

<sup>a</sup>ESAT, SCD/SISTA, K.U. Leuven, Kasteelpark Arenberg 10, B 3001 Leuven, Heverlee, Belgium

<sup>b</sup>Department of Mathematics, University of Maastricht, 6200 MD Maastricht, The Netherlands

Received 3 June 2004; accepted 7 September 2004

Available online 30 November 2004

- terminology **fundamental** is justified: motion primitives, subspace SysID, dictionary learning, (E)DMD, ... all implicitly rely on this equivalence
- many recent **extensions** to other **system classes** (bi-linear, descriptor, LPV, delay, Volterra series, Wiener-Hammerstein, ...), other **matrix data structures** (mosaic Hankel, Page, ...), & other **proof methods**

# Input design for Fundamental Lemma



**Definition:** The data signal  $u^d \in \mathbb{R}^{mT_d}$  of length  $T_d$  is **persistently**

**exciting of order  $T$**  if the Hankel matrix  $\begin{bmatrix} u_1 & \cdots & u_{T_d-T+1} \\ \vdots & \ddots & \vdots \\ u_T & \cdots & u_{T_d} \end{bmatrix}$  is of full rank.   
  *$mT$  rows*

*rank =  $mT$  only if  $T_d - T + 1 \geq mT \Rightarrow$  sufficiently rich & long data*

**Input design** [Willems et al, '05]: Controllable LTI system & persistently exciting input  $u^d$  of order  $T + n \implies \text{rank} \left( \mathcal{H} \begin{pmatrix} u^d \\ y^d \end{pmatrix} \right) = mT + n$ .

# Data matrix structures & preprocessing

Matrix structures

trajectory matrix:  $H(w)$   $w = \begin{pmatrix} u^d \\ y^d \end{pmatrix}$

requires independent experiments

$$= \begin{bmatrix} \text{trajectory 1 of length } T & \text{trajectory 2 of length } T & \dots \end{bmatrix}$$

page matrix:  $H(w) = \begin{bmatrix} w_1 & w_{T+1} & w_{2T+1} & \dots \\ w_2 & \vdots & \vdots & \\ \vdots & & & \\ w_T & w_{2T} & w_{3T} & \end{bmatrix}$

requires one long trajectory

Hankel matrix:  $H(w) = \begin{bmatrix} w_1 & w_2 & w_3 & \dots \\ w_2 & w_3 & w_4 & \\ \vdots & \vdots & \vdots & \\ w_T & w_{T+1} & w_{T+2} & \end{bmatrix}$  uses shift invariance

requires one short trajectory

→ Hankel structure (constant anti-diagonals) due to shift invariance

→ all these are valid matrices whose columns span the space of length- $T$  trajectories provided that their rank is  $mT+n$ . One can also combine these different data structures into block matrices...

Preprocessing: if the data  $w = \begin{pmatrix} u^d \\ y^d \end{pmatrix}$  is noisy, then likely all of the above matrices have full rank.

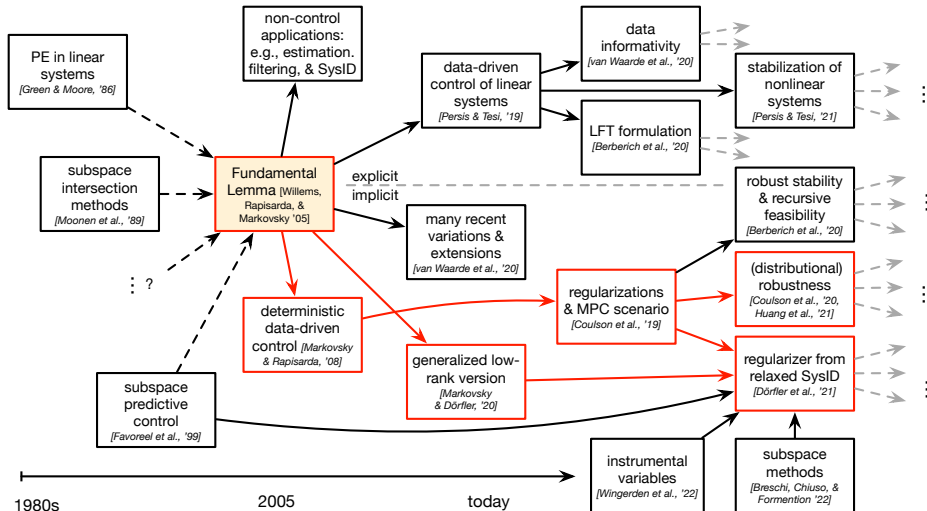
→ low-rank approximation: to denoise the data, find

for theory: the "closest" matrix so that its rank is  $mT+n$

look up the Eckhard → practical solution is singular value thresholding: remove Young theorem all but the  $mT+n$  dominant singular values of the SVD

→ hard problem for Hankel matrices if it is desired to keep the Hankel structure (and induce shift invariance)

# Bird's view & today's sample path through the accelerating literature



# Output Model Predictive Control (MPC)

$$\begin{aligned} & \underset{u, x, y}{\text{minimize}} && \sum_{k=1}^{T_{\text{future}}} \|y_k - r_k\|_Q^2 + \|u_k\|_R^2 \\ & \text{subject to} && \left. \begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k \end{aligned} \right\} \quad \forall k \in \{1, \dots, T_{\text{future}}\} \\ & && \left. \begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k \end{aligned} \right\} \quad \forall k \in \{-T_{\text{ini}} - 1, \dots, 0\} \\ & && \left. \begin{aligned} u_k &\in \mathcal{U} \\ y_k &\in \mathcal{Y} \end{aligned} \right\} \quad \forall k \in \{1, \dots, T_{\text{future}}\} \end{aligned}$$

**quadratic cost** with  
 $R \succ 0, Q \succeq 0$  & ref.  $r$

**model for prediction**  
with  $k \in [1, T_{\text{future}}]$

**model for estimation**  
with  $k \in [-T_{\text{ini}} - 1, 0]$  &  
 $T_{\text{ini}} \geq \text{lag}$  (many flavors)

**hard operational or  
safety constraints**

*“[MPC] has perhaps too little system theory and too much **brute force** [...], but MPC is an area where all aspects of the field [...] are in synergy.” – Willems '07*

Elegance aside, for a LTI plant, deterministic, & with known model, MPC is the **gold standard of control**.



# Data-enabled Predictive Control (**DeePC**)

$$\underset{g, u, y}{\text{minimize}} \quad \sum_{k=1}^{T_{\text{future}}} \|y_k - r_k\|_Q^2 + \|u_k\|_R^2$$

$$\text{subject to} \quad \mathcal{H} \begin{pmatrix} u^d \\ y^d \end{pmatrix} \cdot g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}$$

$$\left. \begin{array}{l} u_k \in \mathcal{U} \\ y_k \in \mathcal{Y} \end{array} \right\} \quad \forall k \in \{1, \dots, T_{\text{future}}\}$$

**quadratic cost** with  
 $R \succ 0, Q \succeq 0$  & ref.  $r$

**non-parametric  
model** for **prediction**  
and **estimation**

hard operational or  
safety **constraints**

- real-time measurements  $(u_{\text{ini}}, y_{\text{ini}})$  for estimation

updated **online**

- trajectory matrix  $\mathcal{H} \begin{pmatrix} u^d \\ y^d \end{pmatrix}$  from past  
experimental data

collected **offline**  
(could be adapted online)

→ **equivalent to MPC** in deterministic LTI case ...

**but needs to be robustified** in case of noise / nonlinearity !

# Regularizations to make it work

$$\underset{g, u, y, \sigma}{\text{minimize}} \quad \sum_{k=1}^{T_{\text{future}}} \|y_k - r_k\|_Q^2 + \|u_k\|_R^2 + \lambda_y \|\sigma\|_p + \lambda_g h(g)$$

$$\text{subject to} \quad \mathcal{H} \begin{pmatrix} u^d \\ y^d \end{pmatrix} \cdot g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma \\ 0 \\ 0 \end{bmatrix}$$

$$\left. \begin{array}{l} u_k \in \mathcal{U} \\ y_k \in \mathcal{Y} \end{array} \right\} \quad \forall k \in \{1, \dots, T_{\text{future}}\}$$

## measurement noise

→ infeasible  $y_{\text{ini}}$  estimate

→ estimation slack  $\sigma$

→ moving-horizon  
least-square filter

## noisy or nonlinear

(offline) **data matrix**

→ any  $\begin{pmatrix} u \\ y \end{pmatrix}$  feasible

→ add regularizer  $h(g)$

**Bayesian intuition:** regularization  $\Leftrightarrow$  prior, e.g.,  $h(g) = \|g\|_1$  sparsely selects {trajectory matrix columns} = {motion primitives}  $\sim$  low-order basis

**Robustness intuition:** regularization  $\Leftrightarrow$  robustifies, e.g., in a simple case *triangle inequality is tight since  $\Delta x$  can be offline with  $Ax=b$*

$$\min_x \max_{\Delta: \|\Delta\| \leq g} \|(A+\Delta)x - b\| \leq \min_x \max_{\|A\| \leq g} \|Ax - b\| + \|\Delta x\| = \min_x \|Ax - b\| + g\|x\|$$



# Regularization = relaxation of bi-level ID

minimize <sub>$u, y, g$</sub>  control cost( $u, y$ )

subject to  $\begin{bmatrix} u \\ y \end{bmatrix} = \mathcal{H} \left( \begin{bmatrix} \hat{u} \\ \hat{y} \end{bmatrix} \right) g$

where  $\begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix} \in \operatorname{argmin} \left\| \begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix} - \begin{pmatrix} u^d \\ y^d \end{pmatrix} \right\|$   
subject to  $\operatorname{rank}(\mathcal{H}(\begin{bmatrix} \hat{u} \\ \hat{y} \end{bmatrix})) = mL + n$

} optimal control




} system identification

↓ sequence of convex relaxations ↓

minimize <sub>$u, y, g$</sub>  control cost( $u, y$ ) +  $\lambda_g \cdot \|g\|_1$

subject to  $\begin{bmatrix} u \\ y \end{bmatrix} = \mathcal{H} \left( \begin{bmatrix} u^d \\ y^d \end{bmatrix} \right) g$

$\ell_1$ -regularization  
= relaxation of id  
smooth order selection



# Certainty-Equivalence Regularizer

ARX representation of predictor:

• recall  $y = \Phi_T x_{ini} + \mathcal{E}_T u$

where  $x_{ini}$  satisfies  $y_{ini} = \Phi_{T_{ini}} x_{ini} + \mathcal{E}_{T_{ini}} u_{ini}$

$\Rightarrow y = K \cdot \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \end{bmatrix}$ , where  $K$  is

learned from data (certainty equiv.):

$$K = \underset{\tilde{K}}{\operatorname{argmin}} \left\| \gamma_f - \tilde{K} \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix} \right\| = \gamma_f \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix}^+$$

(note:  $K$  should also satisfy structural constraints, which we dropped)

$\Rightarrow y = \gamma_f \cdot \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix}^+ \cdot \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \end{bmatrix}$

DeepPC representation of predictor:

$$\begin{bmatrix} u_{ini} \\ y_{ini} \\ u \\ y \end{bmatrix} = \begin{bmatrix} u_p \\ y_p \\ u_f \\ y_f \end{bmatrix} g \quad H(w) \text{ suitably partitioned}$$

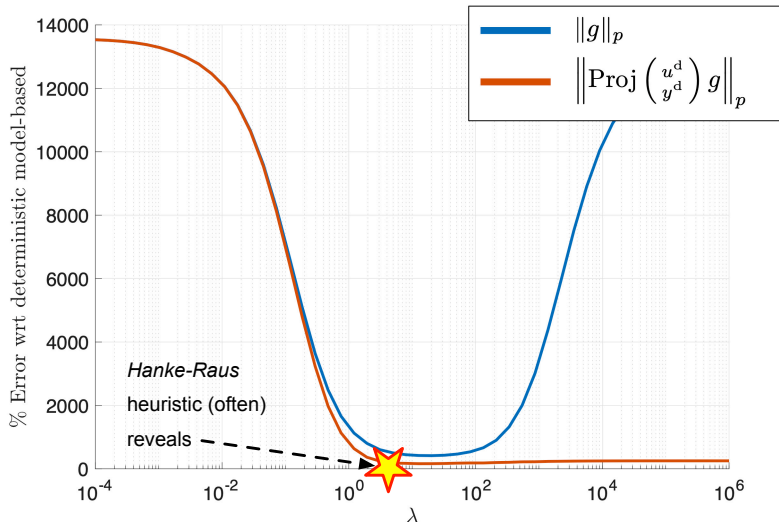
or  $y = \gamma_f g$  with  $g$  from  $\begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \end{bmatrix}$

$$\Rightarrow y = \underbrace{\gamma_f \cdot \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix}^+}_{\text{particular solution}} \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \end{bmatrix} + \underbrace{\gamma_f \cdot \ker}_{\text{homogeneous}} \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix}$$


these two coincide if we force that  $0 = \left\| \left( I - \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix}^+ \begin{bmatrix} u_p \\ y_p \\ u_f \end{bmatrix} \right) g \right\|$

add this regularizer to DeepPC  $\rightarrow \operatorname{Proj}(u^q)$  on nullspace



# Performance of regularizers applied to a stochastic LTI system




# Case study: wind turbine



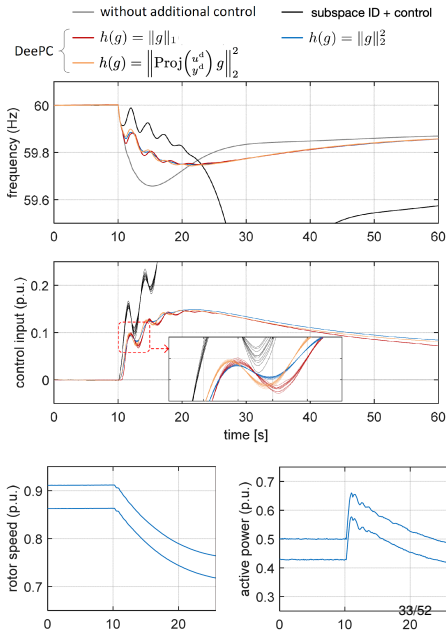
- detailed **industrial model**: 37 states & highly nonlinear (abc  $\leftrightarrow$  dq, MPTT, PLL, power specs, dynamics, etc.)
- turbine & grid model **unknown** to commissioning engineer & operator
- weak grid + PLL + fault  $\rightarrow$  **loss of sync**
- disturbance to be rejected by **DeePC**



# Case study ++ : wind farm

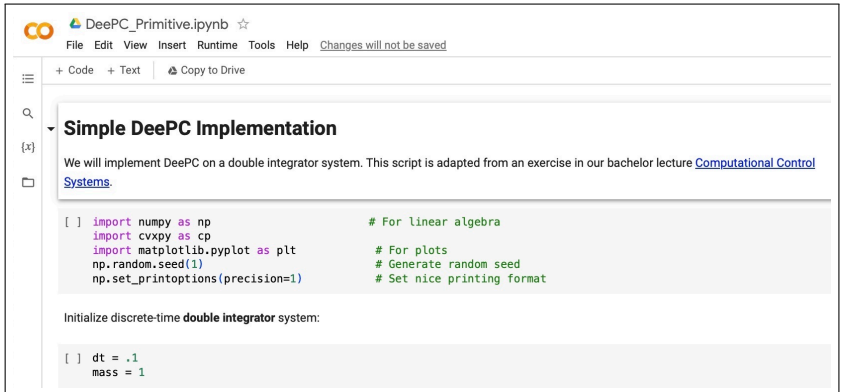


- **high-fidelity models** for turbines, machines, & IEEE-9-bus system
- **fast frequency response** via **decentralized DeePC** at turbines



# DeePC is easy to implement → try it !

→ simple script adapted from our ETH Zürich bachelor course on *Computational control*: <https://colab.research.google.com/drive/1URdRqr-Up0A6uDMjlU6gwms0AAP1lGId?usp=sharing>



The screenshot shows a Google Colab notebook interface. At the top, the title bar reads "DeePC\_Primitive.ipynb" with a star icon. Below the title bar is a menu bar with "File", "Edit", "View", "Insert", "Runtime", "Tools", and "Help". A status message "Changes will not be saved" is visible. The left sidebar contains icons for file management and search. The main content area has a tab labeled "+ Code" and a button "Copy to Drive". The notebook content is titled "Simple DeePC Implementation" and includes a paragraph explaining the implementation on a double integrator system, adapted from a lecture on "Computational Control Systems". Below the text is a code cell with Python imports for numpy, cvxpy, and matplotlib, along with comments for linear algebra, plots, random seed, and printing format. The code also initializes a discrete-time double integrator system with dt = 0.1 and mass = 1.

DeePC\_Primitive.ipynb ☆

File Edit View Insert Runtime Tools Help Changes will not be saved

+ Code + Text Copy to Drive

### Simple DeePC Implementation

We will implement DeePC on a double integrator system. This script is adapted from an exercise in our bachelor lecture [Computational Control Systems](#).

```
[ ] import numpy as np                # For linear algebra
import cvxpy as cp
import matplotlib.pyplot as plt      # For plots
np.random.seed(1)                   # Generate random seed
np.set_printoptions(precision=1)     # Set nice printing format
```


Initialize discrete-time **double integrator** system:

```
[ ] dt = .1
    mass = 1
```

# Towards a theory for nonlinear systems

**idea**: lift nonlinear system to large/ $\infty$ -dimensional bi-/linear system  
→ Carleman, Volterra, Fliess, Koopman, Sturm-Liouville methods  
→ nonlinear dynamics can be approximated by LTI on finite horizon

**regularization** singles out relevant features / basis functions in data



<https://www.research-collection.ethz.ch/handle/20.500.11850/493419>

# Works very well across case studies




quad copter fig-8 tracking



quadruped (by Fawcett, Afsari Amers, & Hamed)




greenhouse automation (by Automatoes)



combined cycle power plant (by P Mahdavi pour et. al)




robotic excavator




pendulum swing up




traffic coordination (by J. Wang et al.)




battery charging (by K. Chen et al.)




wind turbine control




grid-connected converter



synchronous motor drive



energy hub & building automation



power system oscillation damping



# Distributional robustification beyond LTI

- **problem abstraction**:  $\min_{x \in \mathcal{X}} c(\hat{\xi}, x) = \min_{x \in \mathcal{X}} \mathbb{E}_{\xi \sim \hat{\mathbb{P}}} [c(\xi, x)]$

where  $\hat{\xi}$  denotes *measured data with empirical distribution*  $\hat{\mathbb{P}} = \delta_{\hat{\xi}}$


⇒ **poor out-of-sample performance** of above sample-average solution  $x^*$  for real problem:  $\mathbb{E}_{\xi \sim \mathbb{P}} [c(\xi, x^*)]$  where  $\mathbb{P}$  is the *unknown distribution* of  $\xi$

- **distributionally robust** formulation accounting for all (possibly nonlinear) stochastic processes that could have generated the data

$$\inf_{x \in \mathcal{X}} \sup_{Q \in \mathbb{B}_\epsilon(\hat{\mathbb{P}})} \mathbb{E}_{\xi \sim Q} [c(\xi, x)]$$

where  $\mathbb{B}_\epsilon(\hat{\mathbb{P}})$  is an  **$\epsilon$ -Wasserstein ball** centered at empirical sample distribution  $\hat{\mathbb{P}}$ :

$$\mathbb{B}_\epsilon(\hat{\mathbb{P}}) = \left\{ \mathbb{P} : \inf_{\Pi} \int \|\xi - \hat{\xi}\|_p d\Pi \leq \epsilon \right\}$$




- **distributionally robustness**  $\equiv$  **regularization**: under minor conditions

$$\text{Theorem: } \underbrace{\inf_{x \in \mathcal{X}} \sup_{\mathbb{Q} \in \mathbb{B}_\epsilon(\hat{\mathbb{P}})} \mathbb{E}_{\xi \sim \mathbb{Q}} [c(\xi, x)]}_{\text{distributional robust formulation}} \equiv \underbrace{\min_{x \in \mathcal{X}} c(\hat{\xi}, x) + \epsilon \text{Lip}(c) \cdot \|x\|_p^*}_{\text{previous regularized DeePC formulation}}$$

distributional robust formulation

previous regularized DeePC formulation


**Cor:**  $\ell_\infty$ -robustness in trajectory space  
 $\iff \ell_1$ -regularization of DeePC



- **measure concentration**: average matrix


$\frac{1}{N} \sum_{i=1}^N \mathcal{H}_i(y^d)$  from i.i.d. experiments

$\implies$  ambiguity set  $\mathbb{B}_\epsilon(\hat{\mathbb{P}})$  includes true  $\mathbb{P}$   
 with high confidence if  $\epsilon \sim 1/N^{1/\dim(\xi)}$




# Further ingredients

- more **structured uncertainty sets**:  
tractable reformulations (relaxations)  
& performance guarantees



- distributionally robust probabilistic constraints**

$$\sup_{Q \in \mathbb{B}_\epsilon(\hat{\mathbb{P}})} \text{CVaR}_{1-\alpha}^Q \iff \text{averaging} + \text{regularization} + \text{tightening}$$



- replace (finite) moving horizon estimation via  $(u_{ini}, y_{ini})$  by **recursive Kalman filtering** based on optimization solution  $g^*$  as hidden state ...

how does DeePC relate to  
sequential SysID + control ?

surprise: **DeePC consistently  
beats models** across all our  
real-world case studies !

**why ?!?**

# Comparison: direct vs. indirect control

## *indirect ID-based data-driven control*

minimize control cost  $(u, y)$   
subject to  $(u, y)$  satisfy parametric model  
where  $\text{model} \in \arg\min \text{id cost } (u^d, y^d)$  } ID  
subject to  $\text{model} \in \text{LTI}(n, \ell)$  class

*ID projects data* on the set of LTI models

- with parameters  $(n, \ell)$
- removes noise & thus lowers variance error
- suffers bias error if plant is not  $\text{LTI}(n, \ell)$

## *direct regularized data-driven control*

minimize control cost  $(u, y) + \lambda \cdot \text{regularizer}$   
subject to  $(u, y)$  consistent with  $(u^d, y^d)$  data


- *regularization robustifies*  
→ choosing  $\lambda$  makes it work
- *no projection* on  $\text{LTI}(n, \ell)$   
→ no de-noising & no bias

*hypothesis*: ID wins in stochastic (variance) & DeePC in nonlinear (bias) case

# Case study: direct vs. indirect control


## *stochastic LTI case* → indirect ID wins

- LQR control of 5th order LTI system
- Gaussian noise with varying noise to signal ratio (100 rollouts each case)
- $\ell_1$ -regularized DeePC, SysID via N4SID, & judicious hyper-parameters

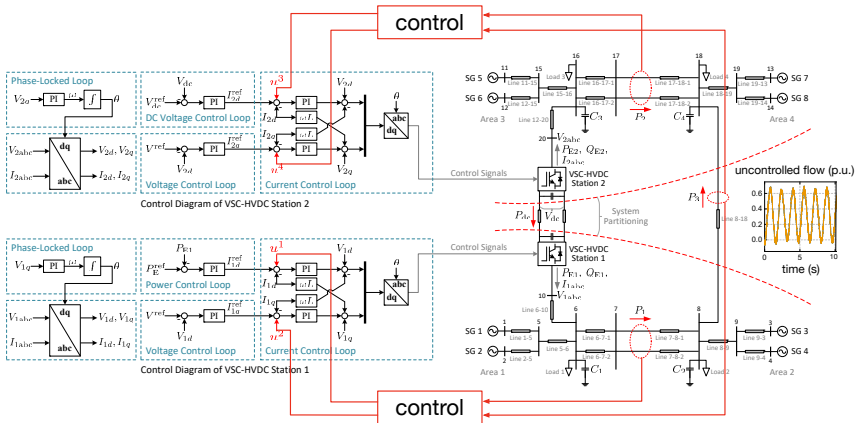


## *nonlinear case* → direct DeePC wins

- Lotka-Volterra + control:  $x^+ = f(x, u)$
- interpolated system  
 $x^+ = \epsilon \cdot f_{\text{linearized}}(x, u) + (1 - \epsilon) \cdot f(x, u)$
- same ID & DeePC as on the left & 100 initial  $x_0$  rollouts for each  $\epsilon$




# Power system case study revisited



- **complex** 4-area power **system**: large ( $n = 208$ ), few measurements (8), nonlinear, noisy, stiff, input constraints, & decentralized control
- **control objective**: damping of inter-area oscillations via HVDC link
- **real-time** MPC & DeePC prohibitive  $\rightarrow$  choose  $T$ ,  $T_{ini}$ , &  $T_{future}$  wisely

# Centralized control



DeePC  
PEM-MPC


= Prediction Error  
Method (PEM)  
System ID + MPC

$t < 10$  s : open loop  
data collection with  
white noise excitat.


$t > 10$  s : control



# Performance: DeePC wins (clearly!)




# DeePC hyper-parameter tuning




## *regularizer* $\lambda_g$

- for distributional robustness  $\approx$  radius of Wasserstein ball
- wide range of sweet spots  
 $\rightarrow$  choose  $\lambda_g = 20$




## *estimation horizon* $T_{\text{ini}}$

- for model complexity  $\approx$  lag
- $T_{\text{ini}} \geq 50$  is sufficient & low computational complexity  
 $\rightarrow$  choose  $T_{\text{ini}} = 60$



### *prediction horizon* $T_{\text{future}}$


- nominal MPC is stable if horizon  $T_{\text{future}}$  long enough  
 $\rightarrow$  choose  $T_{\text{future}} = 120$  and apply first 60 input steps



### *data length* $T$


- long enough for low-rank condition but  $\text{card}(g)$  grows  
 $\rightarrow$  choose  $T = 1500$   
 (data matrix  $\approx$  square)

# Computational cost



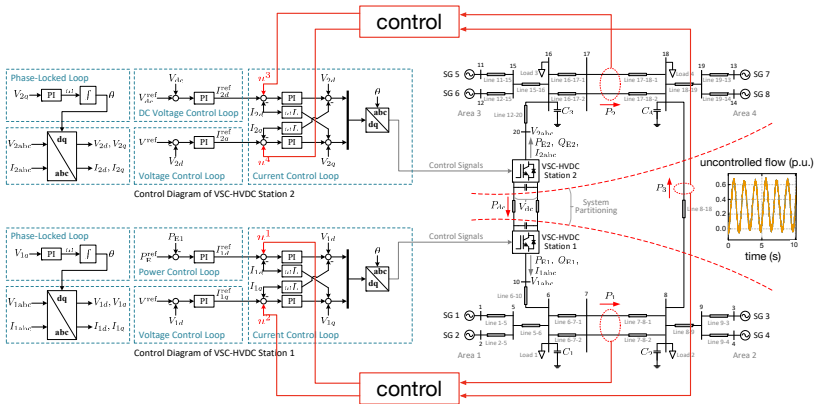
- $T = 1500$
  - $\lambda_g = 20$
  - $T_{\text{ini}} = 60$
  - $T_{\text{future}} = 120$  & apply first 60 input steps
  - sampling time = 0.02 s
  - solver (OSQP) time = 1 s (on Intel Core i5 7200U)
- ⇒ **implementable**

# Comparison: Hankel & Page matrix




- comparison baseline: Hankel and Page matrices of **same size**
- **performance**: Page consistency beats Hankel matrix predictors
- offline **denoising via SVD thresholding** works wonderfully for Page though obviously not for Hankel (entries are constrained)
- effects very pronounced for **longer horizon** (= open-loop time)
- **price-to-be-paid**: Page matrix predictor requires more data

# Decentralized implementation



- **plug'n'play MPC:** treat interconnection  $P_3$  as disturbance variable  $w$  with past disturbance  $w_{ini}$  measurable & future  $w_{future} \in \mathcal{W}$  uncertain
- for each controller **augment trajectory matrix** with disturbance data  $w$
- decentralized **robust min-max DeePC:**  $\min_{g,u,y} \max_{w \in \mathcal{W}}$

# Decentralized control performance



- colors correspond to different hyper-parameter settings (not discernible)
- ambiguity set  $\mathcal{W}$  is  $\infty$ -ball (box)
- for computational efficiency  $\mathcal{W}$  is downsampled (piece-wise linear)
- solver time  $\approx 2.6$  s

$\Rightarrow$  **implementable**


# Conclusions

## main take-aways

- matrix time series as predictive model
- robustness & side-info by regularization
- method that works in theory & practice

## ongoing work


- certificates for adaptive & nonlinear cases
- propagate optimal transport uncertainties
- applications with a true “business case”



**only catch** (no-free-lunch) : optimization problems become large  
→ models are compressed, de-noised, & tidied-up representations



# Florian's version of



# Thanks !

**Florian Dörfler**

mail: [dorfler@ethz.ch](mailto:dorfler@ethz.ch)

[\[link\]](#) to homepage

[\[link\]](#) to related publications