Data Enabled Predictive Control (DeePC)

Dr. Keith Moffat

We can use data to find the valid I/O sequences

All and only the valid I/O sequences are of the form

$$\begin{bmatrix} \mathbf{u}_{L} \\ \mathbf{y}_{L} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{u}_{L}^{(1)} & \mathbf{u}_{L}^{(2)} & \mathbf{u}_{L}^{(3)} \\ \mathbf{y}_{L}^{(1)} & \mathbf{y}_{L}^{(2)} & \mathbf{y}_{L}^{(3)} \\ H \end{bmatrix}}_{H} \cdots \underbrace{\mathbf{u}_{L}^{(L+n)}}_{H} g$$

with $g \in \mathbb{R}^{L+n}$.

Hankel matrices

A very efficient use of your data points:



All system trajectories of length *L* are generated as

 $\mathcal{H}_L(\mathbf{u})$ $\mathcal{H}_L(\mathbf{y})$ g



For DeePC, we divide each Hankel matrix (input and output) into

- one Hankel matrix for the past measurements
- one Hankel matrix for the future measurements

$$egin{bmatrix} \mathcal{H}_L(\mathbf{u}_{ ext{data}})\ \mathcal{H}_L(\mathbf{y}_{ ext{data}}) \end{bmatrix} g = egin{bmatrix} \mathbf{u}\ \mathbf{u}\ \mathbf{y}_{ ext{past}}\ \mathbf{y} \end{bmatrix}$$

MPC DeePC K-1 $\min_{\mathbf{u},\mathbf{x},\mathbf{y}} \quad \sum_{k=0}^{K-1} y_k^2 + r u_k^2$ $\min_{\mathbf{u},\mathbf{x},\mathbf{y},\mathbf{g}} \quad \sum_{k=0}^{n} y_k^2 + r u_k^2$ subject to $\begin{bmatrix} \mathcal{H}_L(\mathbf{u}_{data}) \\ \mathcal{H}_L(\mathbf{y}_{data}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{past} \\ \mathbf{u} \\ \mathbf{y}_{past} \\ \mathbf{v} \end{bmatrix}$ subject to $x_{k+1} = Ax_k + Bu_k$ $y_k = C x_k$ $X_0 = X$ $u_k \in \mathcal{U}_k \ \forall k$ $u_k \in \mathcal{U}_k \ \forall k$ $\mathbf{y}_k \in \mathcal{Y}_k \ \forall \mathbf{k}$ $\mathbf{y}_{\mathbf{k}} \in \mathcal{Y}_{\mathbf{k}} \ \forall \mathbf{k}$

This week...

- MIMO DeePC
- A slide correction
- DeePC with noisy data



What is the difference between the two formulations?

MIMO DeePC Hankel Matrix

Block Hankel matrix:

$w_{\rm d}^2(1)$	$w_{\rm d}^2(2)$	$w_{\rm d}^2(3)$	$w_{\rm d}^2(4)$	$w_{\rm d}^2(5)$	$w_{\rm d}^2(6)$
$w_{\rm d}^2(2)$	$w_{\rm d}^2(3)$	$w_{\rm d}^2(4)$	$w_{\rm d}^2(5)$	$w_{\rm d}^2(6)$	$w_{\rm d}^2(7)$
$w_{\rm d}^2(3)$	$w_{\rm d}^2(4)$	$w_{\rm d}^2(5)$	$w_{\rm d}^2(6)$	$w_{\rm d}^2(7)$	$w_{\rm d}^2(8)$
$w_{\rm d}^2(4)$	$w_{\rm d}^2(5)$	$w_{\rm d}^2(6)$	$w_{\rm d}^2(7)$	$w_{\rm d}^2(8)$	$w_{\rm d}^2(9)$
$w_{\rm d}^2(5)$	$w_{\rm d}^2(6)$	$w_{\rm d}^2(7)$	$w_{\rm d}^2(8)$	$w_{\rm d}^2(9)$	$w_{\rm d}^2(10)$

Markovsky, Ivan, Linbin Huang, and Florian Dörfler. "Data-driven control based on the behavioral approach: From theory to applications in power systems." *IEEE Control Systems* (2023).

MIMO DeePC



What is the maximum rank for a Hankel matrix for a MIMO system with noiseless data?

• the rank of H stops increasing after mL + n

$$\underbrace{\begin{bmatrix} \mathbf{u}_{L}^{(1)} & \mathbf{u}_{L}^{(2)} & \mathbf{u}_{L}^{(3)} & \dots \\ \mathbf{y}_{L}^{(1)} & \mathbf{y}_{L}^{(2)} & \mathbf{y}_{L}^{(3)} & \dots \end{bmatrix}}_{H} g$$

This is true only in the **noiseless case**.

The effect of noise on the Hankel matrix

• the rank of *H* stops increasing after mL + n

$$\underbrace{\begin{bmatrix} \mathbf{u}_{L}^{(1)} & \mathbf{u}_{L}^{(2)} & \mathbf{u}_{L}^{(3)} & \cdots \\ \mathbf{y}_{L}^{(1)} & \mathbf{y}_{L}^{(2)} & \mathbf{y}_{L}^{(3)} & \cdots \\ H \end{bmatrix}}_{H} g$$

This is true only in the **noiseless case**.

In the presence of **noise in the data**, the rank of H will increase past mL + n.

MIMO Systems: it is helpful to know the "lag" of the system

Lag of a system The lag of a system is the smallest ℓ such that $\mathcal{O}_{\ell} = \begin{bmatrix} C \\ CA \\ \cdots \\ CA^{\ell-1} \end{bmatrix}$ has rank *n*.

• A past trajectory of length $T_{ini} \ge \ell$ is needed as "initial conditions".

For observable SISO systems, $\ell = n$.

The (misled) optimism of DeePC with noisy data

If there is noise in the data, the DeePC optimization can (and most likely) will choose trajectories that are not feasible for the true system.

Often, these trajectories use large g vector values. Why?

The (misled) optimism of DeePC with noisy data

Consider the (pathological) case in which one of the outputs is not controllable.

$$y_k = C x_k = \begin{bmatrix} \star & \star & \star \\ 0 & 0 & 0 \end{bmatrix} x_k$$

All output sequences should be of the form

$$y_k = \begin{bmatrix} \star \\ 0 \end{bmatrix}$$

However, if the data contains noise, some columns in the data matrix may present a non-zero second coordinate in the output

$$\underbrace{\begin{bmatrix} \mathbf{u}_{L}^{(1)} & \mathbf{u}_{L}^{(2)} & \mathbf{u}_{L}^{(3)} & \dots \\ \mathbf{y}_{L}^{(1)} & \mathbf{y}_{L}^{(2)} & \mathbf{y}_{L}^{(3)} + \epsilon & \dots \end{bmatrix}}_{H} g$$

The (misled) optimism of DeePC with noisy data

If there is noise in the data, the DeePC optimization can (and most likely) will choose trajectories that are not feasible for the true system.

Often, these "misled optimistic" trajectories use large g vector values. Why?

Hand-wavy intuition:

The over-optimism comes from the g vector selecting trajectories that are only possible after adding noise. The noise contributes smaller singular values than the true data. To activate the trajectories that are created by noise and correspond to small singular values, the g vector must be large.

How can you avoid selecting the "misled optimistic" trajectories?

$$\begin{split} \min_{\mathbf{u},\mathbf{y},g,\hat{\mathbf{u}}_{data},\hat{\mathbf{y}}_{data}} & \sum_{k=0}^{K-1} \|\mathbf{y}_k\|_Q^2 + \|\mathbf{u}_k\|_R^2 \\ \text{subject to} & (\hat{\mathbf{u}}_{data},\hat{\mathbf{y}}_{data}) = \operatorname{argmin} \|\hat{\mathbf{u}}_{data} - \mathbf{u}_{data}\|^2 + \|\hat{\mathbf{y}}_{data} - \mathbf{y}_{data}\|^2 \\ & \operatorname{subject to} & \operatorname{rank} \left(\begin{bmatrix} \mathcal{H}_L(\hat{\mathbf{u}}_{data}) \\ \mathcal{H}_L(\hat{\mathbf{y}}_{data}) \end{bmatrix} \right) = L + n \\ & \begin{bmatrix} \mathcal{H}_L(\hat{\mathbf{u}}_{data}) \\ \mathcal{H}_L(\hat{\mathbf{y}}_{data}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{past} \\ \mathbf{u} \\ \mathbf{y}_{past} \\ \mathbf{y} \end{bmatrix} \\ & u_k \in \mathcal{U}_k \ \forall k \\ & y_k \in \mathcal{Y}_k \ \forall k \end{split}$$

Step 2: Convex relaxation \rightarrow drop the rank constraint

$$\begin{split} \min_{\mathbf{u},\mathbf{y},g,\hat{\mathbf{u}}_{data},\hat{\mathbf{y}}_{data}} & \sum_{k=0}^{K-1} \|y_k\|_Q^2 + \|u_k\|_R^2 \\ \text{subject to} & (\hat{\mathbf{u}}_{data},\hat{\mathbf{y}}_{data}) = \operatorname{argmin} \|\hat{\mathbf{u}}_{data} - \mathbf{u}_{data}\|^2 + \|\hat{\mathbf{y}}_{data} - \mathbf{y}_{data}\|^2 \\ & \left[\frac{\mathcal{H}_L(\hat{\mathbf{u}}_{data})}{\mathcal{H}_L(\hat{\mathbf{y}}_{data})} \right] g = \begin{bmatrix} \mathbf{u}_{past} \\ \mathbf{u} \\ \mathbf{y}_{past} \\ \mathbf{y} \end{bmatrix} \\ & \|g\|_0 \leq L + n \\ & u_k \in \mathcal{U}_k \ \forall k \\ & y_k \in \mathcal{Y}_k \ \forall k \end{split}$$

Is the 0-norm constraint convex?

Step 3: from sparsity of *g* to a bound on $||g||_1$

$$\begin{array}{ll} \min_{\mathbf{u},\mathbf{y},g,\hat{\mathbf{u}}_{\text{data}},\hat{\mathbf{y}}_{\text{data}}} & \sum_{k=0}^{K-1} \|y_k\|_Q^2 + \|u_k\|_R^2 \\ \text{subject to} & \begin{bmatrix} \mathcal{H}_L(\mathbf{u}_{\text{data}}) \\ \mathcal{H}_L(\mathbf{y}_{\text{data}}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{\text{past}} \\ \mathbf{u} \\ \mathbf{y}_{\text{past}} \\ \mathbf{y} \end{bmatrix} \\ & \|g\|_1 \leq \alpha \\ & u_k \in \mathcal{U}_k \ \forall k \\ & y_k \in \mathcal{Y}_k \ \forall k \end{array}$$

Is the 1-norm constraint convex?

Will the 1-norm still induce sparsity on g?

The unit ball for some different norms (and pseudo-norms)

$$\|\mathbf{x}\|_p := \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}$$
 $(p = 2)$ $(p = 1)$ $(p = 1)$ $(p = 1)$ $(p = 1)$

The 1-norm and sparsity



Fif. 3.11 from Elements of Statistical Learning by Hastie, Tibshirani, and Friedman

Step 3: from sparsity of *g* to a bound on $||g||_1$

$$\begin{array}{ll} \min_{\mathbf{u},\mathbf{y},g,\hat{\mathbf{u}}_{\text{data}},\hat{\mathbf{y}}_{\text{data}}} & \sum_{k=0}^{K-1} \|y_k\|_Q^2 + \|u_k\|_R^2 \\ \text{subject to} & \begin{bmatrix} \mathcal{H}_L(\mathbf{u}_{\text{data}}) \\ \mathcal{H}_L(\mathbf{y}_{\text{data}}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{\text{past}} \\ \mathbf{u} \\ \mathbf{y}_{\text{past}} \\ \mathbf{y} \end{bmatrix} \\ & \|g\|_1 \leq \alpha \\ & u_k \in \mathcal{U}_k \ \forall k \\ & y_k \in \mathcal{Y}_k \ \forall k \end{array}$$

Is the 1-norm constraint convex? Yes.

Will the 1-norm still induce sparsity on g? Yes.

Do we know α ?

Step 4: Replace the bound on $||g||_1$ with a regularization term

$$\begin{split} \min_{\mathbf{u},\mathbf{x},g} \quad \sum_{k=0}^{K-1} y_k^2 + r u_k^2 + \lambda_g \|g\|_1 \\ \text{subject to} \quad \begin{bmatrix} \mathcal{H}_L(\mathbf{u}_{\text{data}}) \\ \mathcal{H}_L(\mathbf{y}_{\text{data}}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{\text{past}} \\ \mathbf{u} \\ \mathbf{y}_{\text{past}} \\ \mathbf{y} \end{bmatrix} \\ u_k \in \mathcal{U}_k \ \forall k \\ y_k \in \mathcal{Y}_k \ \forall k \end{split}$$

This is regularized DeePC

What is the effect of large λ_g ? Small λ_g ?

Choosing λ_g



Figure 4.1: Predicted and realized errors (relative to the ground-truth optimal performance and averaged over 100 data sets) with 1-norm regularizer $\lambda ||g||_1$.

Dörfler, Florian, Jeremy Coulson, and Ivan Markovsky. "Bridging direct and indirect data-driven control formulations via regularizations and relaxations." *IEEE Transactions on Automatic Control* 68.2 (2022): 883-897.

An (often) necessary slack variable for DeePC with noisy data

$$\min_{\mathbf{u},\mathbf{x},g} \quad \sum_{k=0}^{K-1} y_k^2 + r u_k^2 + \lambda_g \|g\|_1 + \lambda_\sigma \|\sigma\|_2$$

subject to
$$\begin{bmatrix} \mathcal{H}_L(\mathbf{u}_{data}) \\ \mathcal{H}_L(\mathbf{y}_{data}) \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{past} \\ \mathbf{u} \\ \mathbf{y}_{past} + \sigma \\ \mathbf{y} \end{bmatrix}$$
$$u_k \in \mathcal{U}_k \ \forall k$$
$$y_k \in \mathcal{Y}_k \ \forall k$$

Why is σ often necessary when there is noise in the data?