uc3m PHD

Model predictive control applied to turbulent flows

PhD candidate: Luigi Marra[1]
Thesis advisors: Stefano Discetti[1], Andrea Meilán-Vila[2]

[1] *Department of Aerospace Engineering, UC3M*
[2] *Department of Statistics, UC3M*

Doctoral meetings, June 3, 2025

PREDATOR-CM-UC3M

uc3m Universidad Carlos III de Madrid

Comunidad de Madrid

A brief introduction on fluid flow control

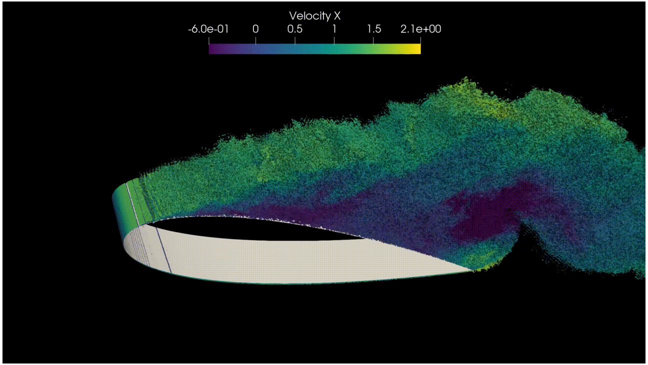

[Puri et al. (2018)]

[1] K. Puri, M. Laufer, H. Müller-Vahl, D. Greenblatt, & S. H. Frankel. (2017). Computations of Active Flow Control Via Steady Blowing Over a NACA-0018 Airfoil: Implicit LES and RANS Validated Against Experimental Data. *AIAA 2018-0792, 2018 AIAA Aerospace Sciences Meeting*, January 2018.

A brief introduction on fluid flow control

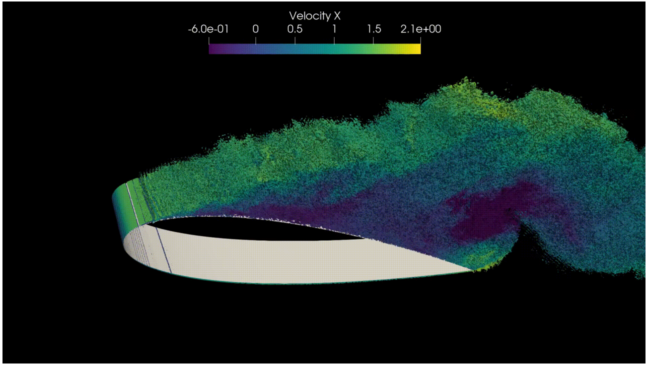

[**Puri et al. (2018)**]

Control goals:

- Drag reduction
- Lift increase
- Mixing layer control
- Noise reduction
- Mixing enhancement

[1] K. Puri, M. Laufer, H. Müller-Vahl, D. Greenblatt, & S. H. Frankel. (2017). Computations of Active Flow Control Via Steady Blowing Over a NACA-0018 Airfoil: Implicit LES and RANS Validated Against Experimental Data. *AIAA 2018-0792, 2018 AIAA Aerospace Sciences Meeting*, January 2018.

A brief introduction on fluid flow control

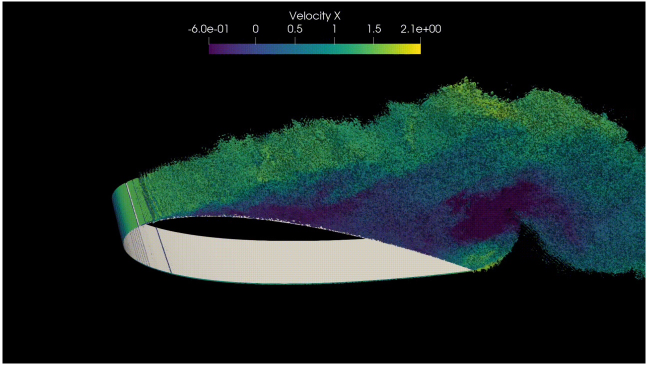

[Puri et al. (2018)]

Control goals:

- Drag reduction
- Lift increase
- Mixing layer control
- Noise reduction
- Mixing enhancement

Control strategies:

- Aerodynamic shape optimization
- Passive control
- Active control

[1] K. Puri, M. Laufer, H. Müller-Vahl, D. Greenblatt, & S. H. Frankel. (2017). Computations of Active Flow Control Via Steady Blowing Over a NACA-0018 Airfoil: Implicit LES and RANS Validated Against Experimental Data. *AIAA 2018-0792, 2018 AIAA Aerospace Sciences Meeting*, January 2018.

Model Predictive Control (MPC): working principle

Prediction
(Opponent's future moves)

Optimization
(The move I consider best)

⇓

Action
(my move)

Moving Horizon
(next opponent move)

Rakovic, S.V and Levine, W.S. (2017). Handbook of model predictive control. *Springer*.

Model Predictive Control (MPC): working principle

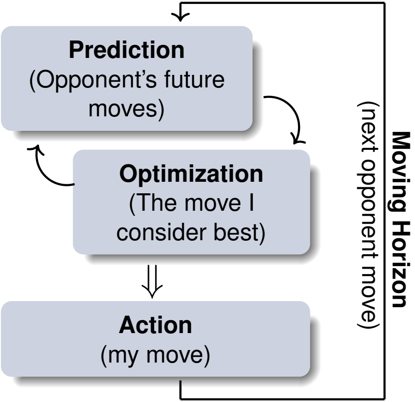

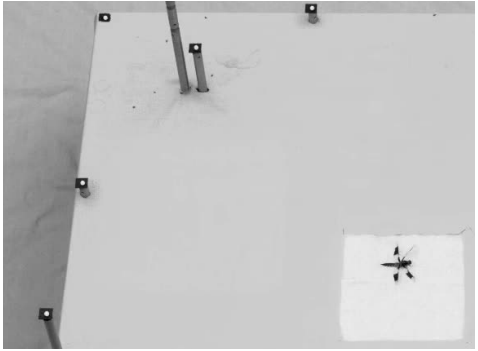

[Mischiati et al. (2017)]

Rakovic, S.V and Levine, W.S. (2017). Handbook of model predictive control. *Springer*.

Mischiati, M., Lin, HT., Herold, P. et al. Internal models direct dragonfly interception steering. Nature 517, 333–338 (2015).

Model Predictive Control (MPC): working principle

- **Working principle**: Optimal control problem over a receding horizon with constraints.

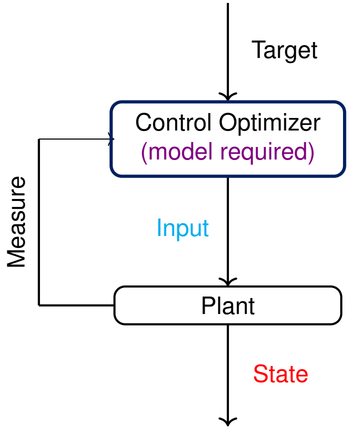

- **Working principle**: Optimal control problem over a receding horizon with constraints.

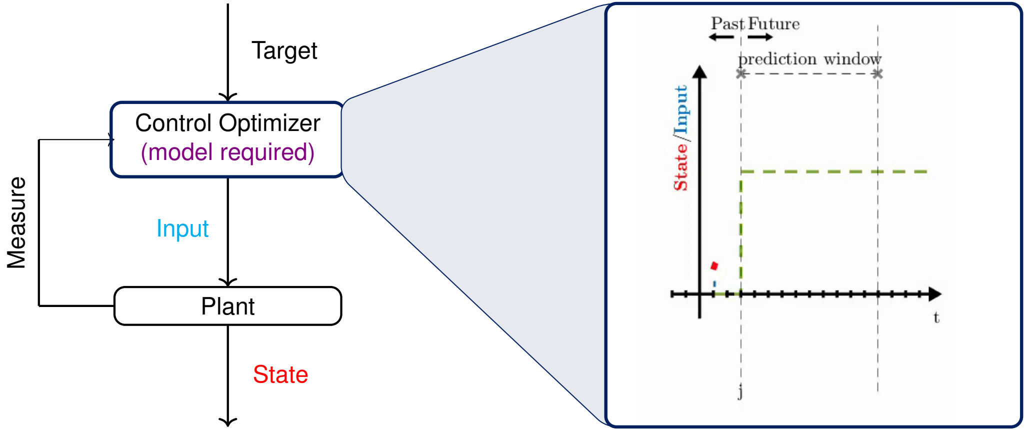

Model Predictive Control (MPC): working principle

- **Working principle**: Optimal control problem over a receding horizon with constraints.

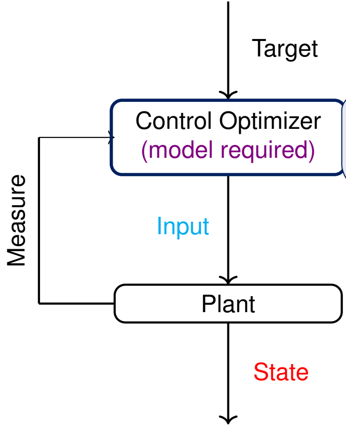

Challenges of plant modeling:

- Multi-dimensional
- Chaotic
- Nonlinear
- Multi-scale
- Time delays

- **Model order reduction** required
- **Trade-off** between **accuracy** and **efficiency**

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication

"Self-tuning model predictive control for wake flows", JFM

Journal publication

"Actuation manifold from snapshots data", JFM

MPC for learning

MPC + RL

In development...

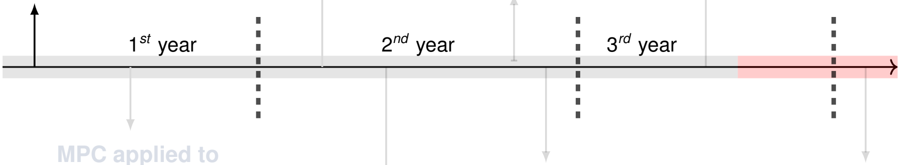

MPC applied to 2D laminar flow

(fluidic pinball)

Experimental validation

(fluidic pinball)

In development...

Research stay

3 months at CNRS - LISN

Experiment on lab-scale problem;

Timeline

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

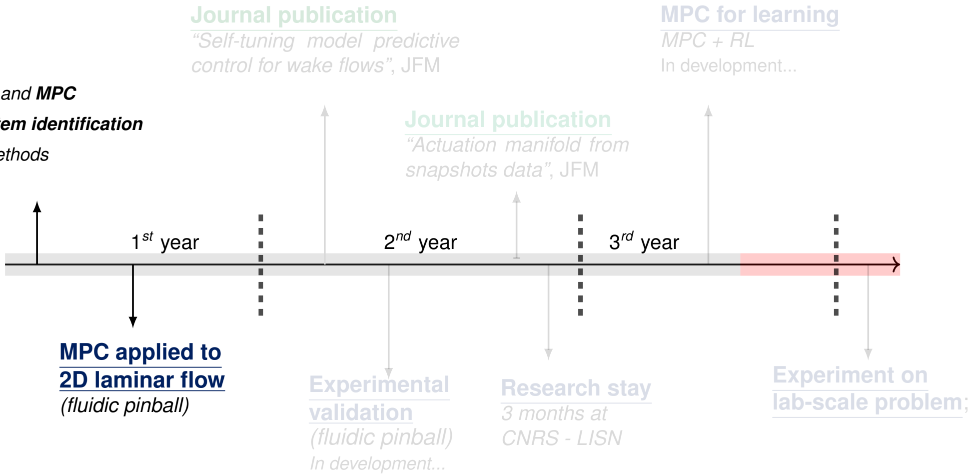

Timeline

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

MPC for learning
MPC + RL
In development...

Journal publication
"Actuation manifold from snapshots data", JFM

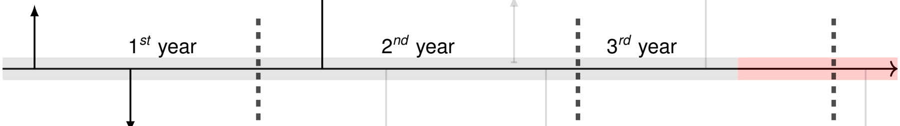

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem;

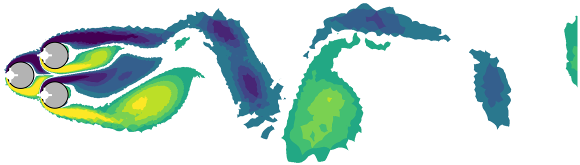

- Wide **variety** of **actuation** mechanism and chaotic dynamics

- **Control goal:** drag reduction / lift stabilization
- **Control actuation:** independent rotation of the three cylinders

Deng, N., Noack, B., Morzyński, M., and Pastur, L. (2022). Cluster-based hierarchical network model of the fluidic pinball – cartographing transient and post-transient, multi-frequency, multi-attractor behaviour. J. Fluid Mech., 934, A24.

Control approach

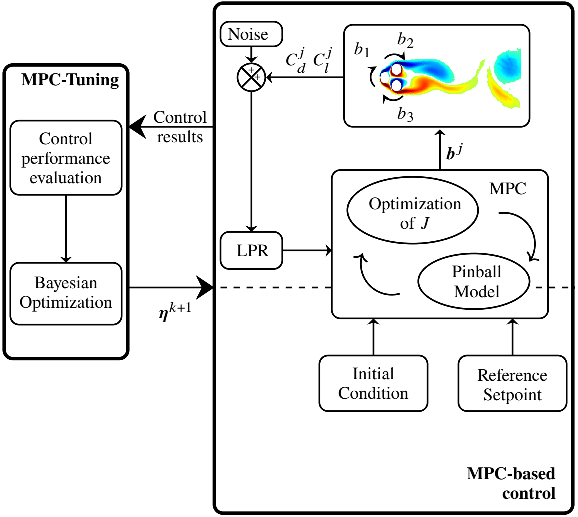

- Sparse Identification of Nonlinear DYnamics (**SINDY**) for **force modeling**
- Bayesian optimization (**BO**) for MPC **hyperparameter selection**
- Local polynomial regression (**LPR**) for **noise robustness**

Hewing, L., Wabersich, K. P., Menner, M., and Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. Annu. rev. control robot., 3, 269-296.

Nottingham, Q. J., and Cook, D. F. (2001). Local linear regression for estimating time series data. Comput. Stat. Data Anal., 37(2), 209-217.

Results

Hand-selected parameters

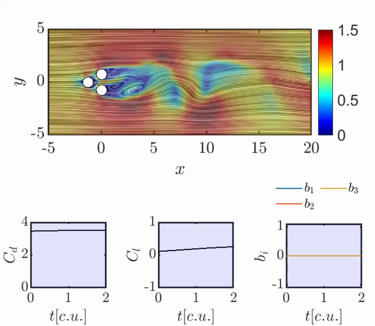

Drag reduction
Lift to 0

Automatic parameter selection

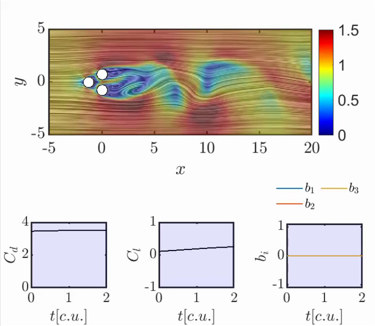

Drag reduction
Lift to 0

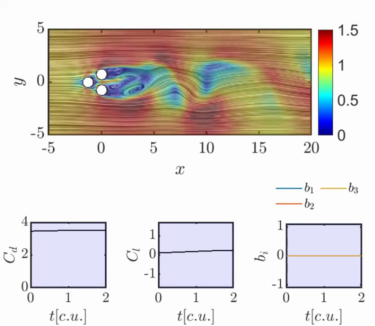

Drag reduction
Lift to sinusoidal signal

Timeline

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

MPC for learning
MPC + RL
In development...

Journal publication
"Actuation manifold from snapshots data", JFM

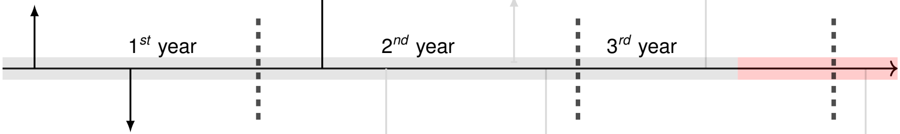

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem;

Timeline

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

Journal publication
“Self-tuning model predictive control for wake flows”, JFM

MPC for learning
MPC + RL
In development...

Journal publication
“Actuation manifold from snapshots data”, JFM

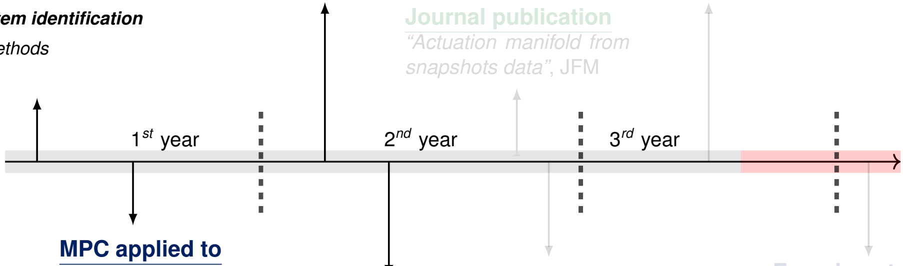

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem;

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

Journal publication
“Self-tuning model predictive control for wake flows”, JFM

Journal publication
“Actuation manifold from snapshots data”, JFM

MPC for learning
MPC + RL
In development...

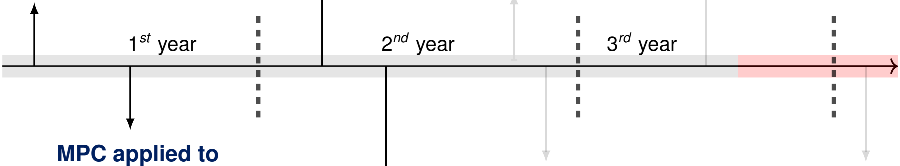

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem;

Timeline

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

Journal publication
"Actuation manifold from snapshots data", JFM

MPC for learning
MPC + RL
In development...

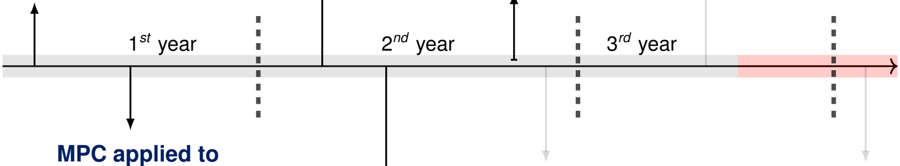

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem;

Model order reduction with ISOMAP

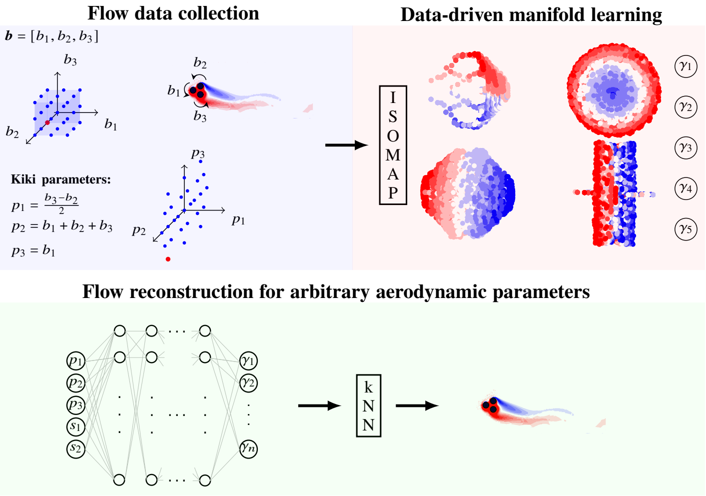

- **Interpretable** low-dimensional **manifold** of **controlled** flows
- Accurate flow **estimation** using a **limited** number of **sensors** for control

Timeline

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

MPC applied to 2D laminar flow
(fluidic pinball)

1st year

Journal publication
"Actuation manifold from snapshots data", JFM

Experimental validation
(fluidic pinball)
In development...

2nd year

MPC for learning
MPC + RL
In development...

Research stay
3 months at CNRS - LISN

3rd year

Experiment on lab-scale problem;

Timeline

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

Journal publication
"Actuation manifold from snapshots data", JFM

MPC for learning
MPC + RL
In development...

1^{st} year — 2^{nd} year — 3^{rd} year

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Supervision by **Onofrio Semeraro** and **Lionel Mathelin**

Timeline

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

Journal publication
"Actuation manifold from snapshots data", JFM

MPC for learning
MPC + RL
In development...

1^{st} year | 2^{nd} year | 3^{rd} year

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Supervision by **Onofrio Semeraro** and **Lionel Mathelin**

Learning from MPC

Training a policy from **MPC** via **Imitation** and **Reinforcement** Learning (RL) strategies.

Strategic **goals**:

- **Safe** and **efficient learning**
- **Real-time** (fast) control in experiments
- Generalization **beyond MPC horizon**
- **Scalability** to complex systems

Source: Dettmers T. (2016) *Deep Learning in a Nutshell: Reinforcement Learning*, NVIDIA.

Timeline

Literature review

- *Control theory and **MPC***
- *Nonlinear **system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

Journal publication
"Actuation manifold from snapshots data", JFM

MPC for learning
MPC + RL
In development...

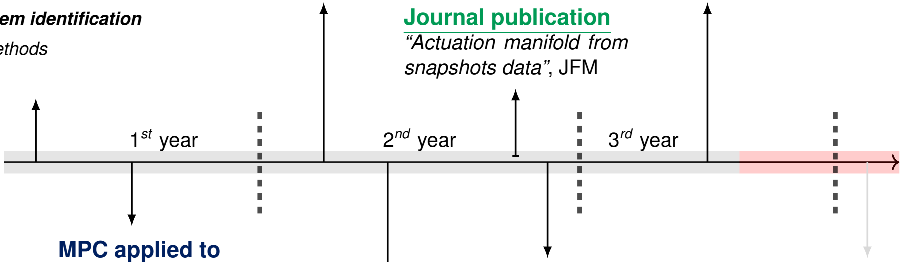

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem

Timeline

Literature review

- *Control theory and* ***MPC***
- *Nonlinear* ***system identification***
- *Data-driven methods*

Journal publication
"Self-tuning model predictive control for wake flows", JFM

Journal publication
"Actuation manifold from snapshots data", JFM

MPC for learning
MPC + RL
In development...

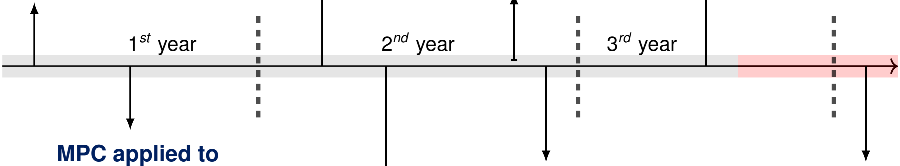

MPC applied to 2D laminar flow
(fluidic pinball)

Experimental validation
(fluidic pinball)
In development...

Research stay
3 months at CNRS - LISN

Experiment on lab-scale problem

Contributions

- **Journal articles and code/datasets**
  - ✓ Marra L., Meilán-Vila A., Discetti S. **Self-tuning model predictive control for wake flows**. Journal of Fluid Mechanics. 2024; 983:A26. 10.1017/jfm.2024.47 (Dataset available in Zenodo and code available in GitHub).
  - ✓ Marra L., Cornejo-Maceda G. Y., Meilán-Vila A., Guerrero V., Rashwan S., Noack B. R., Discetti S., Ianiro A. **Actuation manifold from snapshot data**. Journal of Fluid Mechanics. 2024; 996:A26. 10.1017/jfm.2024.593 (Dataset available in Zenodo and code available in GitHub).

  - ✓ Chang H., Marra L., Cornejo Maceda G. Y., Jiang P., Chen J., Liu Y., Hu G., Chen J., Ianiro A., Discetti S., Meilán-Vila A. and Noack B. R. **Machine-learned flow estimation with sparse data—Exemplified for the rooftop of an unmanned aerial vehicle vertiport**. Physics of Fluids. 2024; 36:125198. 10.1063/5.0242007

- **Conference contributions**
  - ✓ **1st International Conference on Mathematical Modelling in Mechanics and Engineering** (ICME), Sep 8–10, 2022, Belgrade, Serbia
  - ✓ **Math 2 Product: Emerging Technologies in Computational Science for Industry, Sustainability and Innovation** (M2P), May 30 – Jun 1, 2023, Taormina, Italy
  - ✓ **1st Joint Workshop on Functional Data Analysis and Nonparametric Statistics** (JW-FDA-NP), Jun 6–9, 2023, Miraflores de la Sierra, Spain
  - ✓ **2nd Spanish Fluid Mechanics Conference** (SFMC), Jul 2–5, 2023, Barcelona, Spain
  - ✓ **APS Annual Meeting 2024**, Nov 24–26, 2024, Salt Lake City, UT, USA
  - ✓ **2nd ERCOFTAC SIG54 Workshop “Machine Learning for Fluid Dynamics”**, Apr 2–4, 2025, London, UK

Contributions

- **Courses Attended**
  - **Machine Learning for Fluid Mechanics: Analysis, Modeling, Control and Closures**, Von Karman Institute and ULB Lecture Series, Jan 29 – Feb 2, 2024.

- **Dissemination Activities**
  - ✓ Marra L., Rodríguez-Asensio A., Meilán-Vila A., Discetti S. *Descubre el encanto de controlar el agua*, **Viernes STEM**, Universidad Carlos III de Madrid.
  - ✓ Marra L., Rodríguez-Asensio A., Meilán-Vila A., Discetti S. *Descubre la magia de controlar el aire o el agua*, **Madrid Science and Innovation Week**, Nov 15, 2023.

  - ✓ Marra L. Participation in **Thesis Talk 2023**: *Controlling fluids as in the game of chess*. Video
  - ✓ Marra L. Participation in **Thesis Talk 2024**: *Aprendiendo de la naturaleza a controlar la turbulencia*. Video

Acknowledgements

The authors acknowledge the support from the funding under ***"Orden 5067/2023, del consejero de educación, ciencia y universidades por la que se convocan ayudas para la contratación de personal investigador predoctoral en formación para el año 2022"***

Comunidad de Madrid

uc3m | Universidad Carlos III de Madrid

Any questions ?

Some additional slides...

Self-tuning model predictive control for wake flows

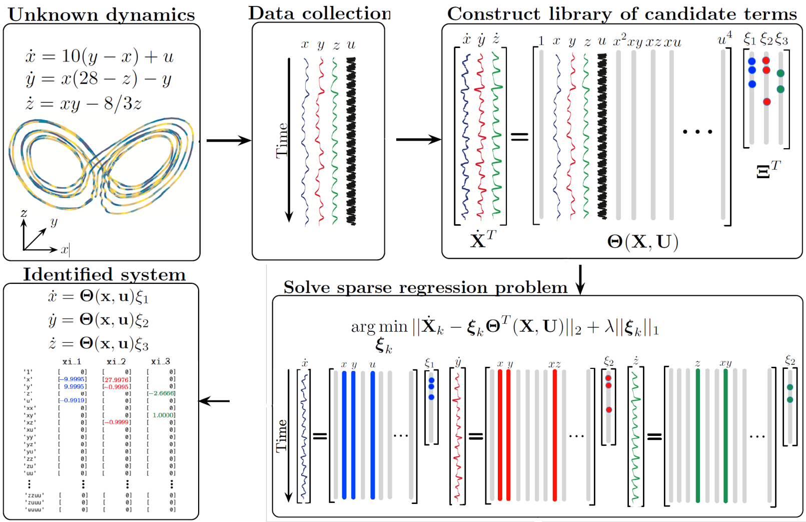

	xi_1	xi_2	xi_3
'1'	0	0	0
'x'	−9.9995	27.9976	0
'y'	9.9995	−0.9995	0
'z'	0	0	−2.6666
'u'	−0.9919	0	0
'xx'	0	0	0
'xy'	0	0	1.0000
'xz'	0	−0.9999	0
'xu'	0	0	0
'yy'	0	0	0
'yz'	0	0	0
'yu'	0	0	0
'zz'	0	0	0
'zu'	0	0	0
'uu'	0	0	0
⋮	⋮	⋮	⋮
'zzuu'	0	0	0
'zuuu'	0	0	0
'uuuu'	0	0	0

MPC results: ideal measurements

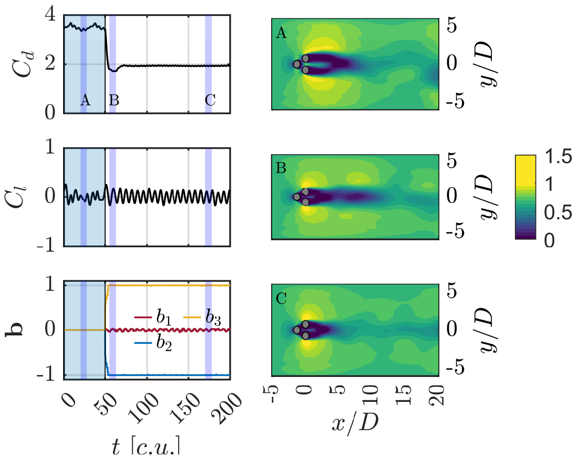

Output control mechanism

- **Boat-tailing** (drag reduction)
- **Phasor control** (lift stabilization)

Results

- $E(C_d)$ reduced by 43.5%
- $\sigma(C_d)$ reduced by 81.3%
- $\sigma(C_l)$ reduced by 3.89%
- $E(C_l) \approx 0$

Li, Y., Cui, W., Jia, Q., Li, Q., Yang, Z., Morzyński, M., and Noack, B. R. (2022). Explorative gradient method for active drag reduction of the fluidic pinball and slanted Ahmed body. J. Fluid Mech., 932, A7.

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|_{\boldsymbol{Q}}^2 + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_b}^2 + \|\Delta \boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_{\Delta b}}^2 \right)$$

- w_p prediction horizon length
  - $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
  - $\|\boldsymbol{x}\|_{\boldsymbol{H}}^2 = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
  - $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
  - $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}' \boldsymbol{H} \boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|_{\boldsymbol{Q}}^2 + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_b}^2 + \|\Delta \boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_{\Delta b}}^2 \right)$$

- w_p prediction horizon length
- $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
- $\|\boldsymbol{x}\|_{\boldsymbol{H}}^2 = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
- $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
- $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
  - $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
  - $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
- $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \dots, w_p$ conditioned to measure in t^j
- $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

- **Errors state predictions**
- **Actuation** cost
- **Input variability**

$$\|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} = \left(\hat{\boldsymbol{c}}^{j+k|j}\right)' \begin{bmatrix} Q_{C_d} & 0 \\ 0 & Q_{C_l} \end{bmatrix} \hat{\boldsymbol{c}}^{j+k|j}$$

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
- $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
- $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

- **Errors state predictions**
- **Actuation cost**
- **Input variability**

$$\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} = \left(\boldsymbol{b}^{j+k|j}\right)' \begin{bmatrix} R_{b_1} & 0 & 0 \\ 0 & R_{b_2} & 0 \\ 0 & 0 & R_{b_3} \end{bmatrix} \boldsymbol{b}^{j+k|j}$$

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|_{\boldsymbol{Q}}^2 + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_b}^2 + \|\Delta \boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_{\Delta b}}^2 \right)$$

- w_p prediction horizon length
- $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \dots, w_p$ conditioned to measure in t^j
- $\|\boldsymbol{x}\|_{\boldsymbol{H}}^2 = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

- **Errors state** predictions
- **Actuation** cost
- **Input variability**

$$\|\Delta \boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_{\Delta b}}^2 = \left(\Delta \boldsymbol{b}^{j+k|j}\right)' \begin{bmatrix} R_{\Delta b_1} & 0 & 0 \\ 0 & R_{\Delta b_2} & 0 \\ 0 & 0 & R_{\Delta b_3} \end{bmatrix} \Delta \boldsymbol{b}^{j+k|j}$$

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|^2_{\boldsymbol{Q}} + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_b} + \|\Delta \boldsymbol{b}^{j+k|j}\|^2_{\boldsymbol{R}_{\Delta b}} \right)$$

- w_p prediction horizon length
  - $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \dots, w_p$ conditioned to measure in t^j
  - $\|\boldsymbol{x}\|^2_{\boldsymbol{H}} = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

All **parameters** included in a single **vector** $\eta \in \mathbb{R}^{N_\eta}$:
$\eta = \left[w_p \,, Q_{C_d} \,, Q_{C_l} \,, R_{b_1} \,, R_{b_2} \,, R_{b_3} \,, R_{\Delta b_1} \,, R_{\Delta b_2} \,, R_{\Delta b_3}\right]$.

MPC cost function

MPC cost function:

$$\mathcal{J}_{MPC}(\boldsymbol{b}) = \sum_{k=0}^{w_p} \|\hat{\boldsymbol{c}}^{j+k|j}\|_{\boldsymbol{Q}}^2 + \sum_{k=1}^{w_p} \left(\|\boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_b}^2 + \|\Delta \boldsymbol{b}^{j+k|j}\|_{\boldsymbol{R}_{\Delta b}}^2 \right)$$

- w_p prediction horizon length
  - $\hat{\boldsymbol{c}}^{j+k|j}$ predictions of $\boldsymbol{c}$ in timesteps t^{j+k}, $k = 1, \ldots, w_p$ conditioned to measure in t^j
  - $\|\boldsymbol{x}\|_{\boldsymbol{H}}^2 = \boldsymbol{x}'\boldsymbol{H}\boldsymbol{x}$
- $\boldsymbol{Q}$, $\boldsymbol{R}_b$, $\boldsymbol{R}_{\Delta b}$ positive and semi-positive definite diagonal weight matrices

All **parameters** included in a single **vector** $\boldsymbol{\eta} \in \mathbb{R}^{N_\eta}$:
$\boldsymbol{\eta} = \left[w_p\,, Q_{C_d}\,, Q_{C_l}\,, R_{b_1}\,, R_{b_2}\,, R_{b_3}\,, R_{\Delta b_1}\,, R_{\Delta b_2}\,, R_{\Delta b_3}\right]$.

Control **results** $\tilde{\boldsymbol{c}} = \tilde{\boldsymbol{c}}(\boldsymbol{\eta})$ are **dependent** on the choice of the **hyperparameter vector**

Tuning parameters cost function

Running the control algorithm for N_t timesteps, global control **performance** can be **assessed** by the following cost function:

$$\mathcal{J}_{BO}(\eta) = \frac{1}{N_t} \sum_{k=1}^{N_c} \sum_{j=1}^{N_t} \left(\tilde{c}_k^j(\eta)\right)^2$$

Tuning parameters cost function

Running the control algorithm for N_t timesteps, global control **performance** can be **assessed** by the following cost function:

$$\mathcal{J}_{BO}(\eta) = \frac{1}{N_t} \sum_{k=1}^{N_c} \sum_{j=1}^{N_t} \left(\tilde{c}_k^j(\eta)\right)^2$$

Parameters in η are **optimized** by **maximizing** control **performance** the minimizing $\mathcal{J}_{BO}$.

Bayesian optimization

The **tuning problem** is explained by the following **optimization** problem:

$$\boldsymbol{\eta}_{opt} = \underset{\boldsymbol{\eta} \in H}{\arg\min} \quad \mathcal{J}_{BO}(\boldsymbol{\eta})$$

- $H \subset \mathbb{R}^{N_\eta}$ is a **hyper-rectangle** of the type $\boldsymbol{\eta} \in \left[\boldsymbol{\eta}^{min}, \boldsymbol{\eta}^{max}\right]$
- $\mathcal{J}_{BO}$ behaves as a "**black box**" function

Bayesian optimization

The **tuning problem** is explained by the following **optimization** problem:

$$\boldsymbol{\eta}_{opt} = \arg\min_{\boldsymbol{\eta}\in H} \ \mathcal{J}_{BO}(\boldsymbol{\eta})$$

- $H \subset \mathbb{R}^{N_\eta}$ is a **hyper-rectangle** of the type $\boldsymbol{\eta} \in \left[\boldsymbol{\eta}^{min}, \boldsymbol{\eta}^{max}\right]$
- $\mathcal{J}_{BO}$ behaves as a "**black box**" function

Bayesian optimization builds a **probabilistic model** of $\mathcal{J}_{BO}$

- **Gaussian process** (GP) is used

Given the **data** the posterior **distribution** is **updated**

An **acquisition process** iteratively **proposes** a new **sampling point** in the domain in order to find the minimum.

Balance between **exploration** and **exploitation**.

Results: tuning and noise effects

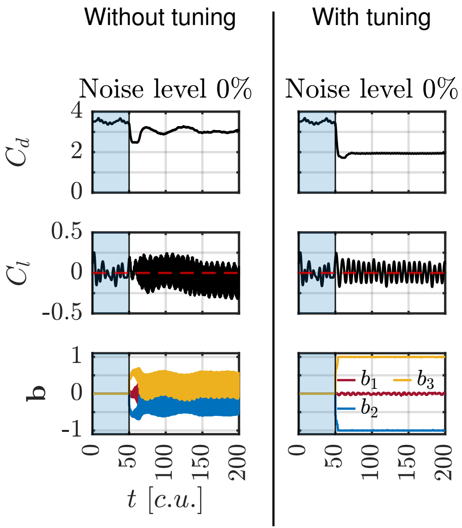

- **Poor performance without** paramaters **tuning**
- **Results** remains **unchanged** with increasing **noise**

Results: tuning and noise effects

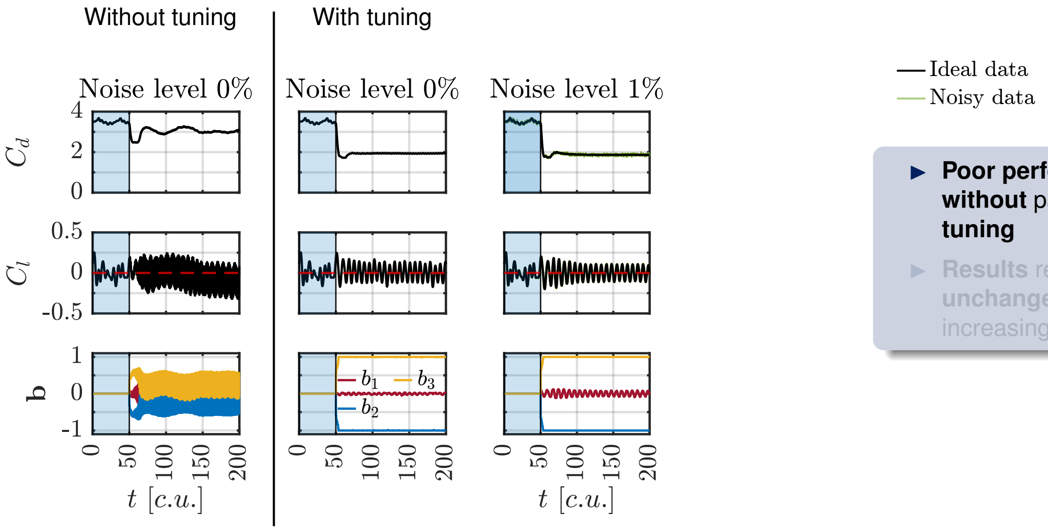

- **Poor performance without** paramaters **tuning**
- **Results** remains **unchanged** with increasing **noise**

Results: tuning and noise effects

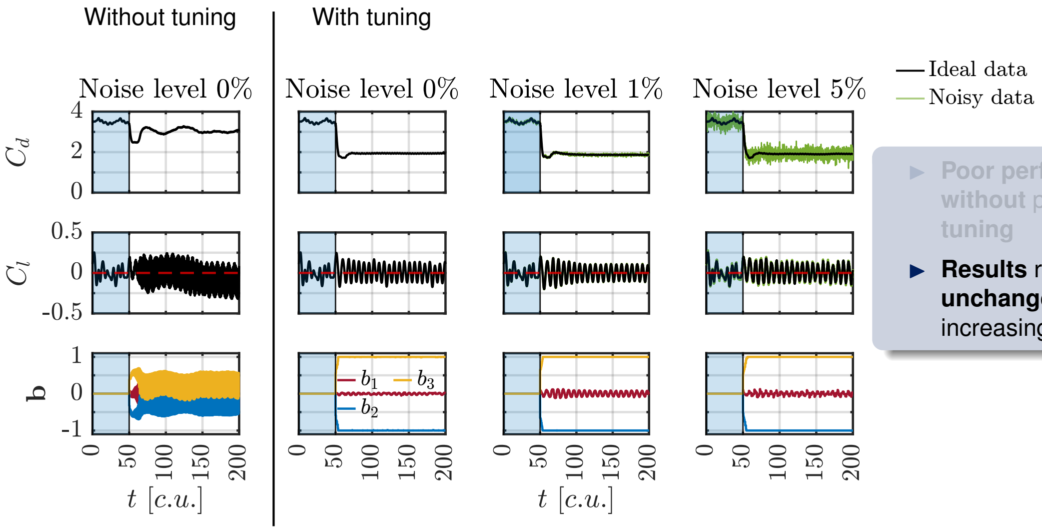

- **Poor performance without** paramaters **tuning**
- **Results** remains **unchanged** with increasing **noise**

Actuation manifold from snapshots data

Methodology flowchart

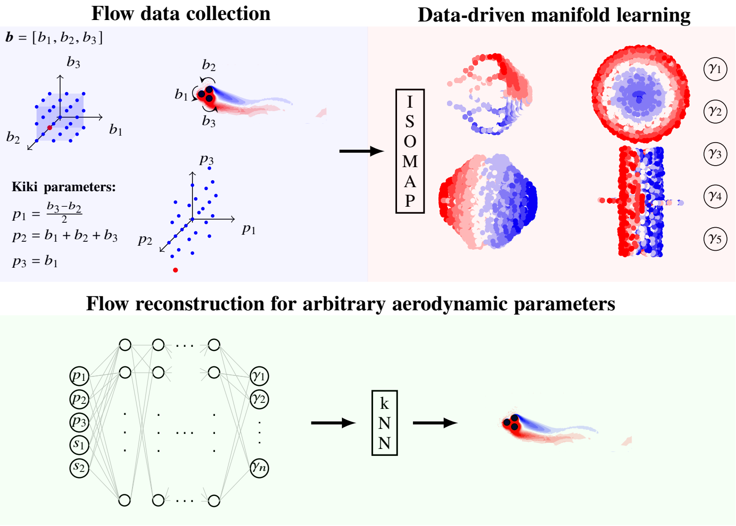

ISOMAP

- Snapshots data $\boldsymbol{u}_i$, $i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

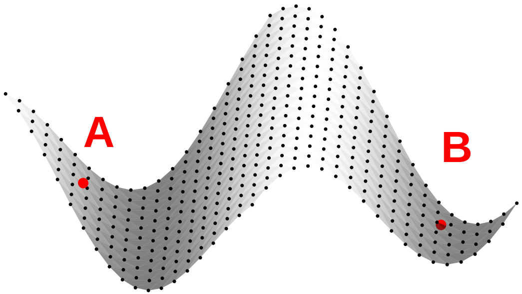

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

ISOMAP

- Snapshots data $\boldsymbol{u}_i, i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

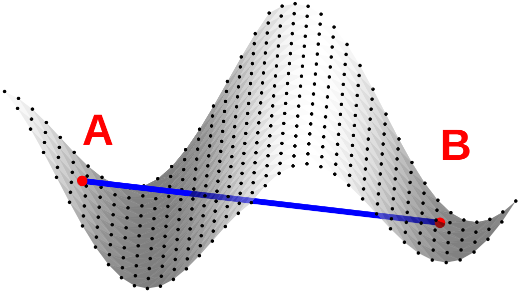

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

ISOMAP

- Snapshots data $\boldsymbol{u}_i$, $i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

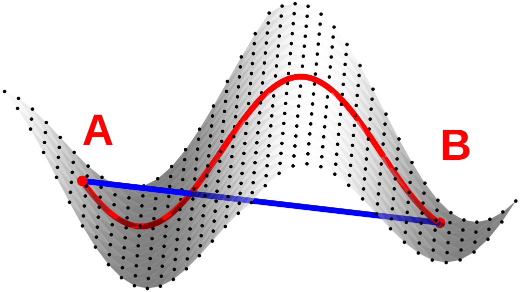

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

ISOMAP

- Snapshots data $\boldsymbol{u}_i$, $i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

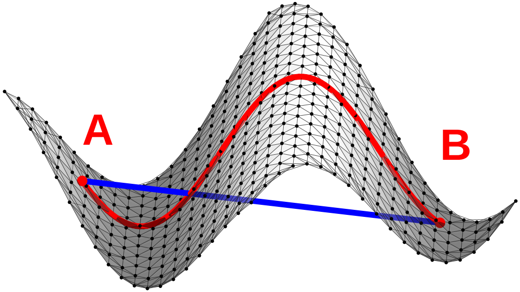

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

ISOMAP

- Snapshots data $\boldsymbol{u}_i, i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

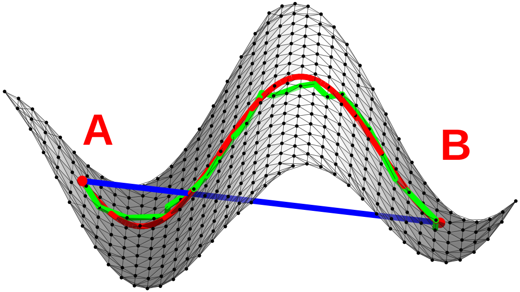

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

- Snapshots data $\boldsymbol{u}_i$, $i = 1, \ldots, M$
- Build the **euclidean distance** matrix $\mathbf{D}_E$ among snapshots
- **Approximate geodesic distance** matrix $\mathbf{D}_G$:
  - Construct **neighbourhood graph**
  - Compute **shortest path** across graph (e.g., Floyd-Warshall method)
- **Project** data into low-dimensional space (MDS) retaining n coordinates

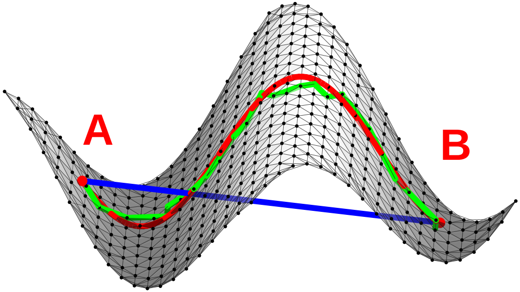

Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. science, 290(5500), 2319-2323.

Actuation manifold sections

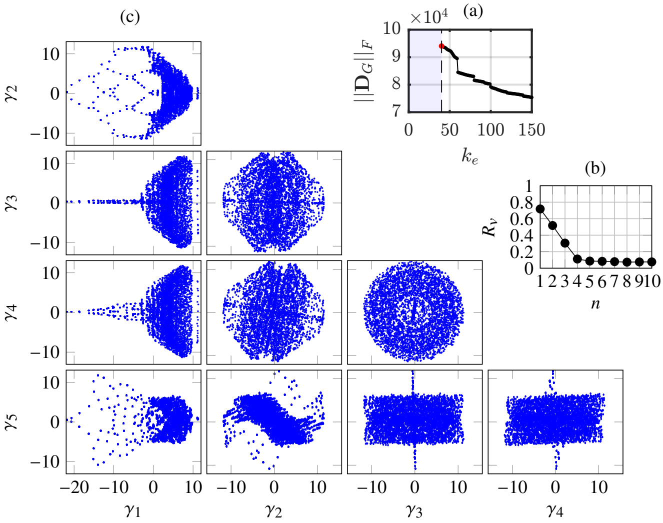

Actuation manifold interpretation

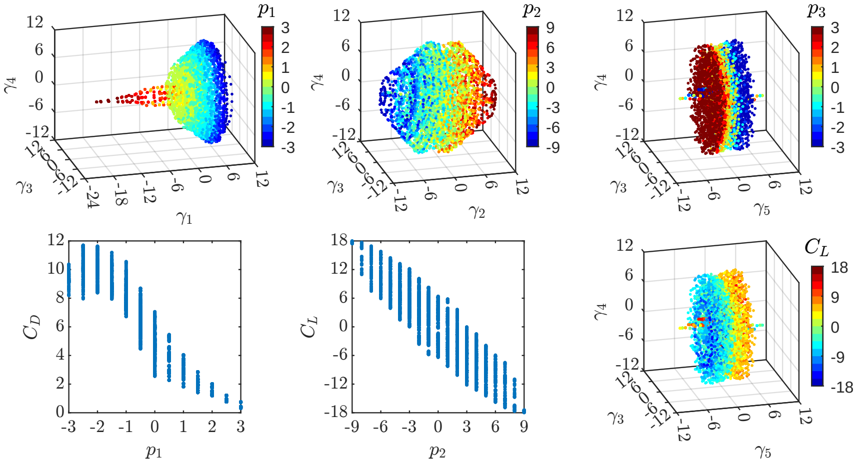

ISOMAP pseudomodes

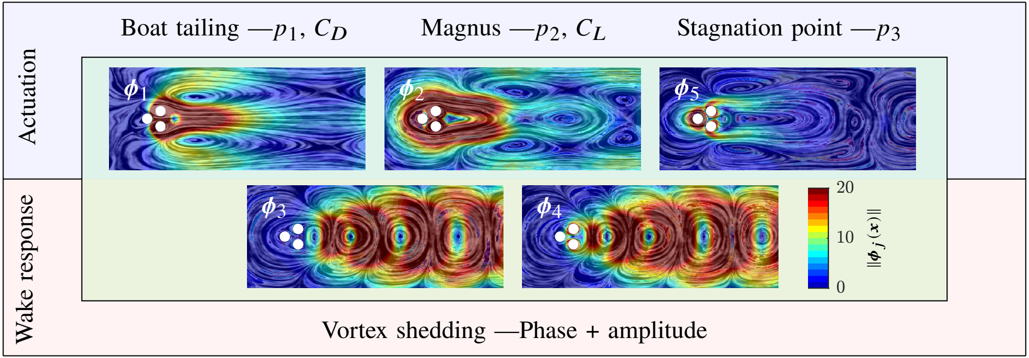

$p_1 = \frac{b_3 - b_2}{2}$ (Base-bleeding/boat-tailing)

$p_2 = b_1 + b_2 + b_3$ (Magnus)

$p_3 = b_1$ (Stagnation point control)

Decoding performance

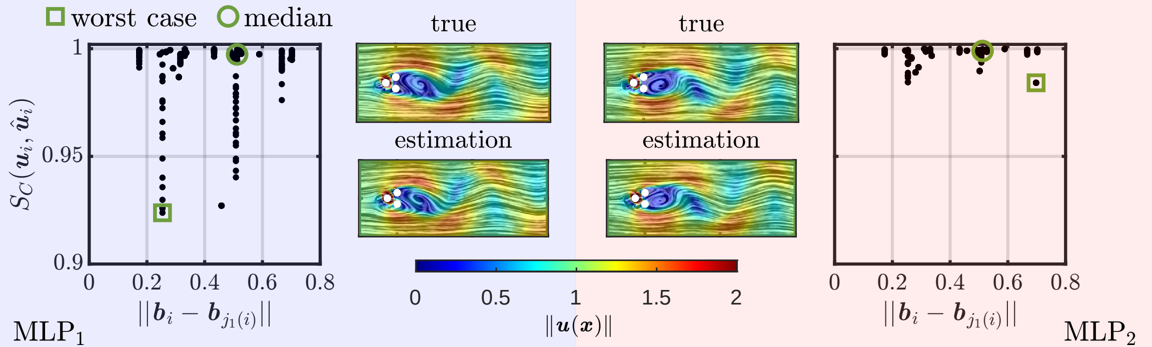

$$S_C(\boldsymbol{u}_i, \boldsymbol{u}_j) = \frac{\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle}{||\boldsymbol{u}_i||||\boldsymbol{u}_j||}$$

MPC, imitation and reinforcement learning

The idea

The ‘expert’
(MPC)

→

The ‘student’
(NN policy)

- ☺ Safe
- ☺ Greedy optimal
- ☹ Computational cost

- ☹ Unsafe training with reinforcement learning (RL)
- ☺ Faster (direct mapping)
- ☺ Less data hungry after imitation

The method

Control authority on expert (MPC)

Experience replay buffer

Parallel policy learning

Control authority on trained policy

Policy fine-tuning (RL)

Train the policy: Initially BC (supervised learning)

Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state

Aggregate data: Add new state-action pairs to the experience buffer

Retrain: Train the model on the updated experience buffer

The method

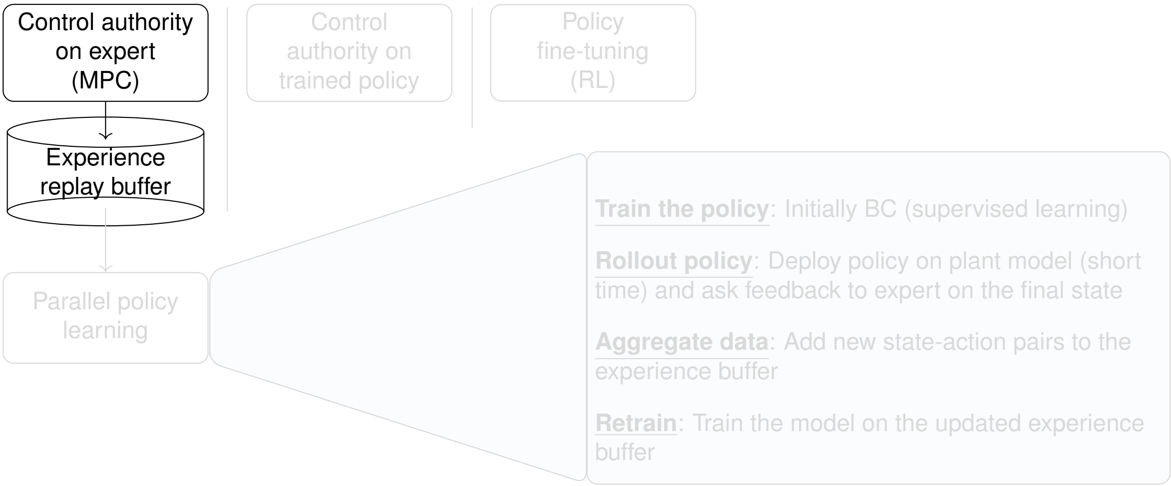
Control authority on expert (MPC)
Experience replay buffer
Parallel policy learning
Control authority on trained policy
Policy fine-tuning (RL)
Train the policy: Initially BC (supervised learning)
Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state
Aggregate data: Add new state-action pairs to the experience buffer
Retrain: Train the model on the updated experience buffer

The method

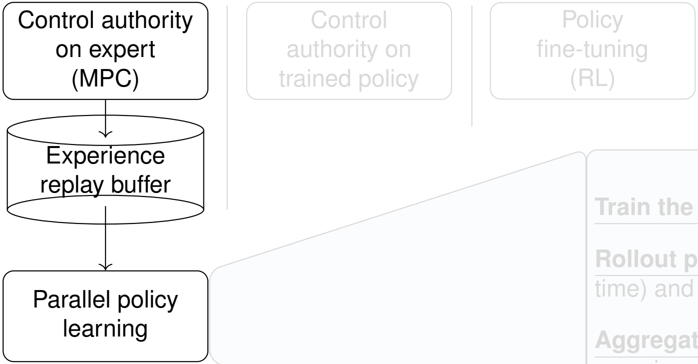

Train the policy: Initially BC (supervised learning)

Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state

Aggregate data: Add new state-action pairs to the experience buffer

Retrain: Train the model on the updated experience buffer

The method

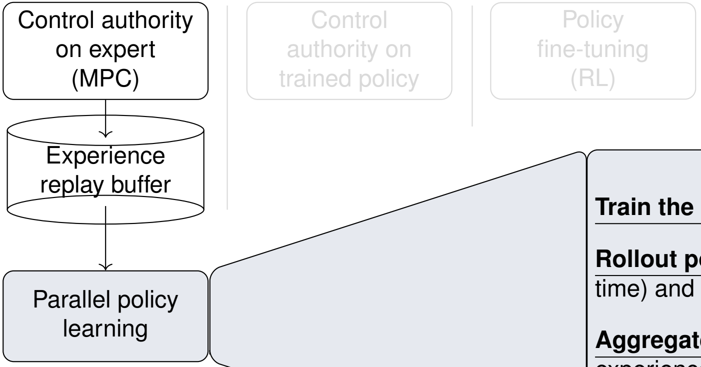

Train the policy: Initially BC (supervised learning)

Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state

Aggregate data: Add new state-action pairs to the experience buffer

Retrain: Train the model on the updated experience buffer

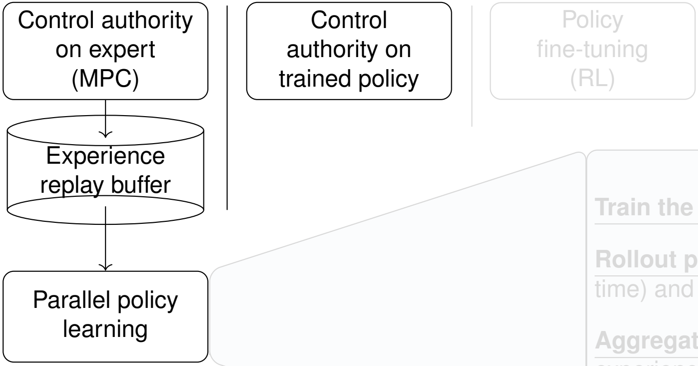

Train the policy: Initially BC (supervised learning)

Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state

Aggregate data: Add new state-action pairs to the experience buffer

Retrain: Train the model on the updated experience buffer

The method

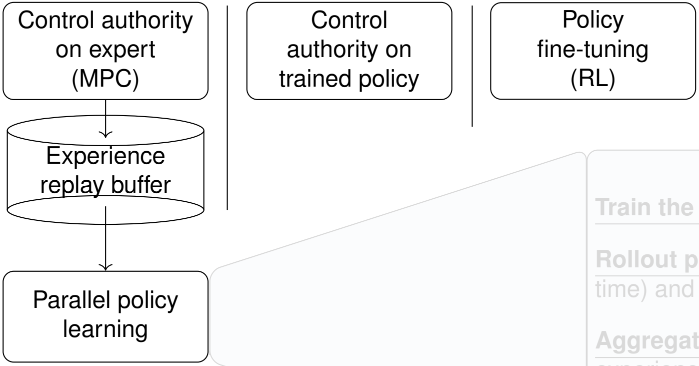

Train the policy: Initially BC (supervised learning)

Rollout policy: Deploy policy on plant model (short time) and ask feedback to expert on the final state

Aggregate data: Add new state-action pairs to the experience buffer

Retrain: Train the model on the updated experience buffer

Test case

Kuramoto–Sivashinsky (KS) equation:

$$\frac{\partial x}{\partial t} + x\frac{\partial x}{\partial \xi} = -\frac{\partial^2 x}{\partial \xi^2} - \frac{\partial^4 x}{\partial \xi^4} + \phi$$

- State x, time t, spatial coordinate ξ
- Domain in $[0, L]$, with $L = 22$
- $x(\xi, t) = x(\xi + L, t)$
- Sampling on 64 collocation points
- ϕ is a 4-dimensional Gaussian supported actuation

Learning task:

Guide the KS **solution** from a random initial condition to the **unstable equilibrium** point E_3.
Failure $\rightarrow$ not reaching the target within a threshold in an episode.

Bucci, M. A., *et al.* (2019). Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A, 475(2231), 20190351.*

Kuramoto–Sivashinsky (KS) equation:

$$\frac{\partial x}{\partial t} + x\frac{\partial x}{\partial \xi} = -\frac{\partial^2 x}{\partial \xi^2} - \frac{\partial^4 x}{\partial \xi^4} + \phi$$

- State x, time t, spatial coordinate ξ
- Domain in $[0, L]$, with $L = 22$
- $x(\xi, t) = x(\xi + L, t)$
- Sampling on 64 collocation points
- ϕ is a 4-dimensional Gaussian supported actuation

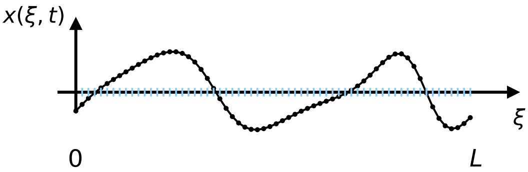

Learning task:

Guide the KS **solution** from a random initial condition to the **unstable equilibrium** point E_3.
Failure $\rightarrow$ not reaching the target within a threshold in an episode.

Bucci, M. A., *et al.* (2019). Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A, 475(2231), 20190351.*

Kuramoto–Sivashinsky (KS) equation:

$$\frac{\partial x}{\partial t} + x\frac{\partial x}{\partial \xi} = -\frac{\partial^2 x}{\partial \xi^2} - \frac{\partial^4 x}{\partial \xi^4} + \phi$$

- State x, time t, spatial coordinate ξ
- Domain in $[0, L]$, with $L = 22$
- $x(\xi, t) = x(\xi + L, t)$
- Sampling on 64 collocation points
- ϕ is a 4-dimensional Gaussian supported actuation

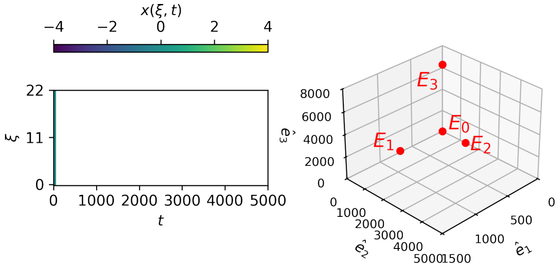

$\hat{e}_1, \hat{e}_2, \hat{e}_3 \rightarrow$ squared dominant Fourier coefficients

Learning task:

Guide the KS **solution** from a random initial condition to the **unstable equilibrium** point E_3.
Failure $\rightarrow$ not reaching the target within a threshold in an episode.

Bucci, M. A., *et al.* (2019). Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A, 475(2231), 20190351.*

Test case

Kuramoto–Sivashinsky (KS) equation:

$$\frac{\partial x}{\partial t} + x\frac{\partial x}{\partial \xi} = -\frac{\partial^2 x}{\partial \xi^2} - \frac{\partial^4 x}{\partial \xi^4} + \phi$$

- State x, time t, spatial coordinate ξ
- Domain in $[0, L]$, with $L = 22$
- $x(\xi, t) = x(\xi + L, t)$
- Sampling on 64 collocation points
- ϕ is a 4-dimensional Gaussian supported actuation

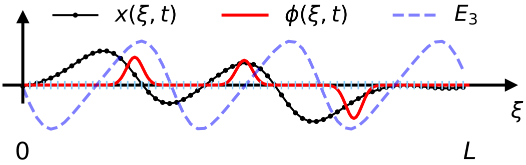

Learning task:

Guide the KS **solution** from a random initial condition to the **unstable equilibrium** point E_3.
Failure $\rightarrow$ not reaching the target within a threshold in an episode.

Bucci, M. A., *et al.* (2019). Control of chaotic systems by deep reinforcement learning. *Proceedings of the Royal Society A, 475(2231), 20190351.*

The expert

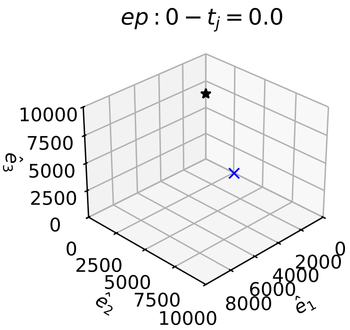

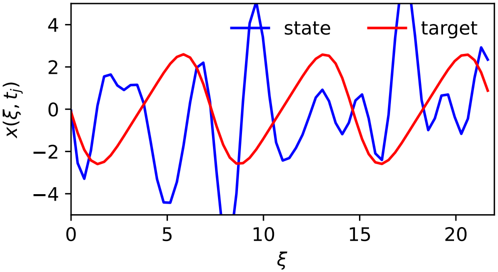

* Plant model from data-driven Operator Inference with POD on 17 modes and model dependencies up to 3^{rd} polynomial order

** MPC with full-state feedback

Kramer, B., *et al.* (2024). Learning nonlinear reduced models from data with operator inference. *Annual Review of Fluid Mechanics, 56(1), 521-548.*

The student

W/o policy rollouts on plant model...

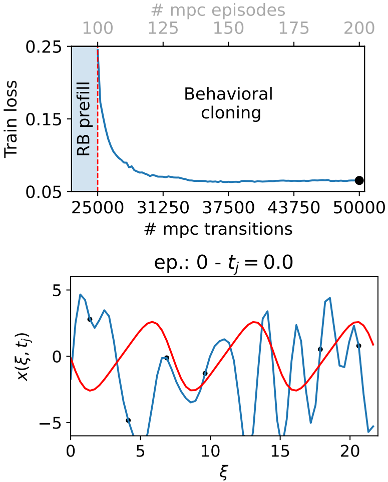

* Policy with sparse sensors feedback (8 equispaced sensors)

The student

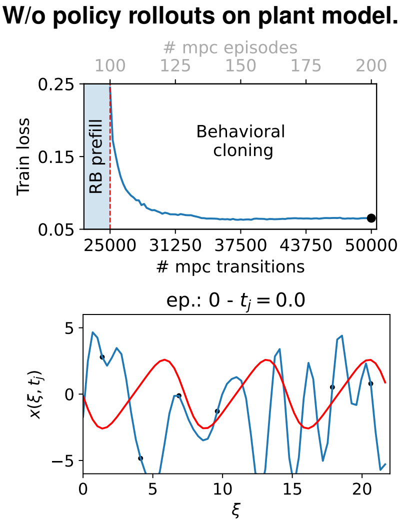

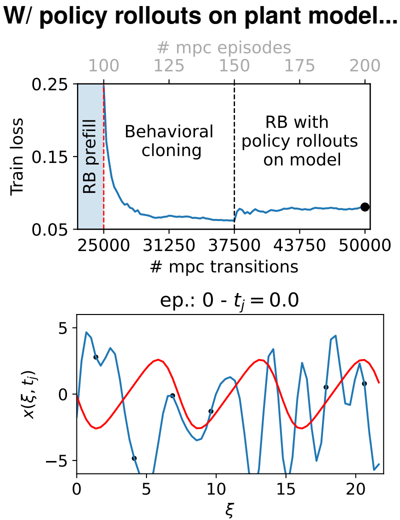

* Policy with sparse sensors feedback (8 equispaced sensors)