


**EUROPEAN  
MECHANICS  
SOCIETY**

# Data-driven active control in flows: from model-based to reinforcement learning methods



## Chairpersons

**Onofrio Semeraro  
Stefano Discetti  
Lionel Mathelin  
Miguel Alfonso Mendez**

**June 08-10, 2026  
Paris, France**

The open-source L<sup>A</sup>T<sub>E</sub>X template, `AMCOS_booklet`, used to generate this booklet is available at [https://github.com/maximelucas/AMCOS\\_booklet](https://github.com/maximelucas/AMCOS_booklet)

## Data-driven active control in flows: from model-based to reinforcement learning methods

Data-driven methods are transforming active flow control, enabled by advances in experimental hardware, computational power, and modern sensing and actuation strategies. Approaches such as reduced-order modelling, enhanced with data-driven closures and neural networks, offer accurate predictions while retaining physical insight. At the same time, reinforcement learning has emerged as a powerful framework for control design, learning effective policies directly through interaction with complex fluid systems without relying on prior models. The workshop will emphasise the theoretical aspects and practical applications of model-based and model-free control, following to main topics: Modeling and Prediction of Fluid Flows; Model-based and Model-free Strategies for Flow Control. The aim of the colloquium is to highlight recent breakthroughs in data-driven flow control and stimulate discussion on the most promising methodologies.

To recognize outstanding contributions, a special issue associated with the colloquium is planned on European Journal of Mechanics / B Fluids (Eur. J. Mech. B Fluids).

### Organizers

Onofrio Semeraro, CNRS-LISN, France    Stefano Discetti, UC3M, Spain  
Lionel Mathelin, CNRS-LISN, France    Miguel A. Mendez, VKI-ULB, Belgium

### Scientific and local committee

Michele Alessandro Bucci, SafranTech, France    Georgios Rigas, Imperial College, UK  
Jean-Christophe Loiseau, ENSAM, France    Taraneh Sayadi, CNAM, France  
Iraj Mortazavi, CNAM, France

The European Mechanics Society is an international non-governmental non-profit scientific organization. The objective of the Society is to engage in all activities intended to promote in Europe the development of mechanics as a branch of science and engineering. The present colloquium is part of the European Mechanics Society Colloquia series. It has been partially funded by CNRS. We thank for the support the CNAM and the Laboratoire Interdisciplinaire des Sciences du Numérique (LISN).



# Timetable

We plan 20 minutes slots for contributed presentations, consisting of 15 minutes for the talk followed by 4-5 minutes for discussion and transition. Speakers are encouraged to use their own laptops. For keynote lectures, the 50-minute slot will be divided into 40 – 45 minutes for the presentation and 5-10 minutes for discussion.

CT: Contributed Talk, IS: Invited Speaker.

Day 1 – 8 June 2026			
8:30–9:00	Registration		
9:00–9:10	Welcome remarks		
Session 1 – Data-driven modelling 1			
Chairman: Stefano Discetti			
9:10–10:00	IS	<b>George Haller</b>	Data-driven modeling and control via spectral submanifolds
10:00–10:20	CT	<b>James King</b>	Parametric spectral submanifolds across Hopf bifurcations with applications to fluid dynamics
10:20–10:40	CT	<b>Luc Pastur</b>	Bistable flow dynamics of airfoil stall under varying angle of attack: A stochastic model with multiplicative noise
10:40–11:10	Coffee Break		

Session 2 – Data-driven modelling 2			
Chairman: Stefano Discetti			
11:10–11:30	CT	<b>Alessandro Franchini</b>	Mean resolvent analysis of a stochastic axisymmetric jet
11:30–11:50	CT	<b>Francesco Fossella</b>	Multiscale data assimilation in turbulence
11:50–12:10	CT	<b>Aymane Lahgazi</b>	Uncertainty-aware latent-dynamics ROM for state estimation in unsteady flows
12:10–12:30	CT	<b>Valerio Di Domenico</b>	POD-Galerkin reduced order method for unsteady, turbulent, buoyant flows with parametric boundary conditions
12:30–14:00	Lunch		

<b>Session 3 – Data-driven modelling 3</b>			
<b>Chairman: Miguel A. Mendez</b>			
14:00–14:50	IS	<b>Olga Fink</b>	Momentum-conserving physics-informed graph neural networks for dynamical systems
14:50–15:10	CT	<b>Eloi Lebrech</b>	Sparse data conditioned diffusion model for PDE solving
15:10–15:30	CT	<b>Iacopo Tirelli</b>	From spikes to dynamics: SNN-based latent space identification for wake flow prediction

15:30–16:00	<b>Coffee Break</b>		
<b>Session 4 – Data-driven modelling 4</b>			
<b>Chairman: Miguel A. Mendez</b>			
16:00–16:20	CT	<b>Niccoló Tonioni</b>	Navigating the intermittency: A generative surrogate for long-horizon forecasting of minimal flow unit from sparse measurements
16:20–16:40	CT	<b>Carlos Sanmiguel-Vila</b>	Sparse-sensing-state estimation of vortex-gust airfoil interactions
16:40–17:00	CT	<b>Lukasz Klotz</b>	Dominant recurrent carrier of turbulence within spatially-localized turbulent structures
17:00–17:20	CT	<b>Marco Castelletti</b>	Estimating turbulent channel flow from wall measurements

<b>Day 2 – 9 June 2026</b>			
<b>Session 5 – Model-based control 1</b>			
<b>Chairman: George Rigas</b>			
9:00–09:50	IS	<b>Denis Sipp</b>	Koopman-based model-order reduction: coherent structures, frequencies and damping rates in stochastic cavity flow
09:50–10:10	CT	<b>Luigi Marra</b>	Model predictive control in latent coordinates for partially observable systems
10:10–10:30	CT	<b>Colin Leclercq</b>	Increasing non-linear stability of flows using optimal control

10:30–11:00	<b>Coffee Break</b>		
	<b>Session 6 – Model-based control 2</b>		
	<b>Chairman: Colin Leclercq</b>		
11:00–11:20	CT	<b>Philippe Gilotte</b>	Drag and lift optimization with a surrogate model
11:20–11:40	CT	<b>Luigi Marra</b>	Safe learning by combination of MPC and reinforcement learning
11:40–12:00	CT	<b>Defne E. Ozan</b>	A sliding mode observer perspective on chaos synchronisation of turbulent flows
12:00–12:20	CT	<b>Guy Y. Cornejo Maceda</b>	Actuation manifold for control of a flapping wing under dynamic flow conditions
12:20–12:40	CT	<b>Alicia Rodríguez-Asensio</b>	Can we predict the transient trajectories of the actuated fluidic pinball?

12:40–14:00	<b>Lunch</b>		
	<b>Session 7 – Reinforcement Learning 1</b>		
	<b>Chairman: Taraneh Sayadi</b>		
14:00–14:20	CT	<b>Vaibhav Chaturvedi</b>	Reinforcement learning for olfactory navigation in a turbulent flow
14:20–14:40	CT	<b>Jingran Qiu</b>	Adaptive shape control for microswimmer navigation in turbulence
14:40–15:00	CT	<b>Zisong Zhou</b>	Two-way regulation of turbulent heat transfer via interpretable bang-bang control discovered by reinforcement learning
15:00–15:20	CT	<b>Babak Mohammadikalakoo</b>	Data-driven suppression of naturally developing Tollmien–Schlichting waves using plasma actuation
15:20–15:40	CT	<b>Matteo Tomasetto</b>	Physics-enhanced reinforcement learning for real-time optimal control of dynamical systems
15:40–16:00	<b>Coffee Break</b>		
16:00–18:00	<b>Free visit at CNAM museum</b>		
19:30–22:30	<b>Social dinner</b>		

Day 3 – 10 June 2026			
Session 8 – Control theory			
Chairman: Onofrio Semeraro			
09:20–09:40	CT	<b>Andrés Marcos</b>	Compressed sensing for launchers and satellites' parameter and fault estimation
09:40–10:00	CT	<b>J. Simon Kern</b>	Control-then-reduce approach to optimal control using dynamical low-rank approximation
10:00–10:20	CT	<b>Lorenzo Schena</b>	Reinforcement twinning for wind farm flow control using differentiable dynamic wake models and learned value corrections

10:20–10:50 Coffee Break			
Session 9 – Reinforcement Learning 2			
Chairman: Lionel Mathelin			
10:50–11:10	CT	<b>Felice Manganelli</b>	Multi-agent reinforcement learning for wind farm optimization in atmospheric boundary layer
11:10–11:30	CT	<b>Junjie Zhang</b>	Learning observers for partially observable flow control
11:30–11:50	CT	<b>Defne E. Ozan</b>	Data-assimilated model-informed reinforcement learning
11:50–12:10	CT	<b>Yannick Lecomte</b>	Bayesian reinforcement twinning: a multi-fidelity framework for reciprocal learning between digital twins and control
12:10–12:30	CT	<b>Guillermo Suarez</b>	Active flow control via model-based reinforcement learning
12:30–12:50	CT	<b>Gerardo Paolillo</b>	Deep reinforcement learning for combined shape and tangential blowing optimization with wind tunnel validation

12:50–14:00 Lunch			
Session 10 – Concluding session			
Chairman: Iraj Mortazavi			
14:00–14:50	IS	<b>Elie Hachem</b>	Coupling reinforcement learning and CFD to support decision making
14:50–15:00	Final remarks		

# List of Abstracts – Talks

## Monday 8th – Day 1 - Data-driven modelling and reduced-order models

### Data-driven modeling and control via spectral submanifolds

George Haller

ETH Zürich, Zürich, Switzerland

The recent concept of spectral submanifolds (SSMs) uncovers very low-dimensional attractors in virtually all dynamics problems of physical importance. A data-driven identification of the reduced dynamics on these SSMs gives a mathematically justified way to construct accurate and predictive reduced-order models for solids, fluids, and controls without the use of governing equations. I review these concepts and recent progress in applying them to the reduced modeling of fluid flows from numerical and experimental data. I also show very recent comparisons of SSM-based model-predictive control with reinforcement learning (RL)-based control in soft robotics. Additionally, I show results comparing the use of SSM-reduced models with the use of recurrent neural networks and transformers in training RL policies.

# Parametric Spectral Submanifolds across Hopf Bifurcations with Applications to Fluid Dynamics

James King<sup>1</sup>, Bálint Kaszás<sup>1,2</sup>, William Jussiau<sup>3</sup>, George Haller<sup>1</sup>

<sup>1</sup>Institute for Mechanical Systems, ETH Zurich, Zurich, Switzerland

<sup>2</sup>Department of Mechanical Engineering, Stanford University, Stanford, USA

<sup>3</sup>Department of Engineering Cybernetics, NTNU, Trondheim, Norway

We investigate the persistence and regularity of spectral submanifolds (SSMs [1]) in high-dimensional parametric dynamical systems undergoing a Hopf bifurcation. By analyzing how resonances in the linearized spectrum near bifurcation points limit the existence and smoothness of SSMs, a phenomenon that has been mostly overlooked, we show that low-order Taylor coefficients of the SSM expansion and the associated reduced dynamics persist smoothly through the bifurcation. This analysis generalizes to any local bifurcation and provides a clear estimate of the parameter ranges over which a parametric SSM model can be justified, thus illustrating how globally the model can be extended despite the presence of resonances near criticality. We demonstrate these findings on multiple examples, including a data-driven SSM approach to the lid-driven cavity flow (See Figure 1). For that problem, we construct a parametric SSM-reduced model that accurately captures the full transition to periodic dynamics and the critical Reynolds number. These results provide a mathematical foundation for robust data- and equation-driven model reduction of fluid flows across bifurcations, enabling an accurate prediction of nonlinear dynamics across critical parameter regimes.

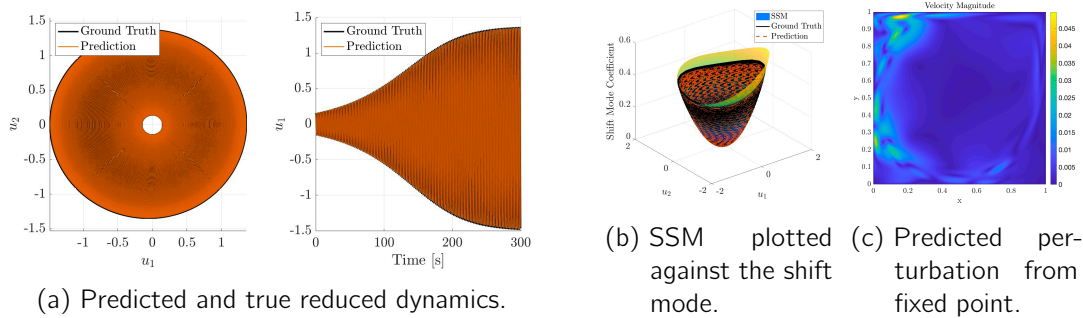


Figure 1: Parametric SSM predictions for the lid-driven cavity flow at an unseen Reynolds number.

1. George Haller, Modeling Nonlinear Dynamics from Equations and Data – with Applications to Solids, Fluids, and Controls. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2025.

## **Bistable flow dynamics of airfoil stall under varying angle of attack: A stochastic model with multiplicative noise**

Edouard Boujo<sup>1</sup>, Ivan Kharsansky Atallah<sup>2</sup>, Luc Pastur<sup>2</sup>

<sup>1</sup>École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland,

<sup>2</sup>IMSIA, ENSTA Paris, Institut Polytechnique de Paris, F-91120 Palaiseau, France

We model the intermittent, bistable stall dynamics of an airfoil under varying angle of attack  $\alpha$ . Experiments at  $Re \approx 10^5$  show random transitions between two flow states, high-lift attached state A and low-lift detached state D, with very long residence times [1]. We propose a one-dimensional Langevin equation using the lift coefficient  $C_l(t)$  as a scalar observable. The deterministic potential depends continuously on  $\alpha$  and, crucially, the stochastic forcing is multiplicative, i.e. depends on the flow state (A or D). The model, identified based on the flow statistics and dynamics [2], reproduces the S-shaped lift curve  $C_l(\alpha)$  as well as the flow dynamics. It also predicts that the bistable region is delimited by two saddle-node bifurcations and yields their normal form.

1. I. Kharsansky, L.R. Pastur, R. Monchaux, and L. Zimmer. Phys. Rev. F, 9(6):063902, 2024.
2. E. Boujo and N. Noiray. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 473(2200), 2017.

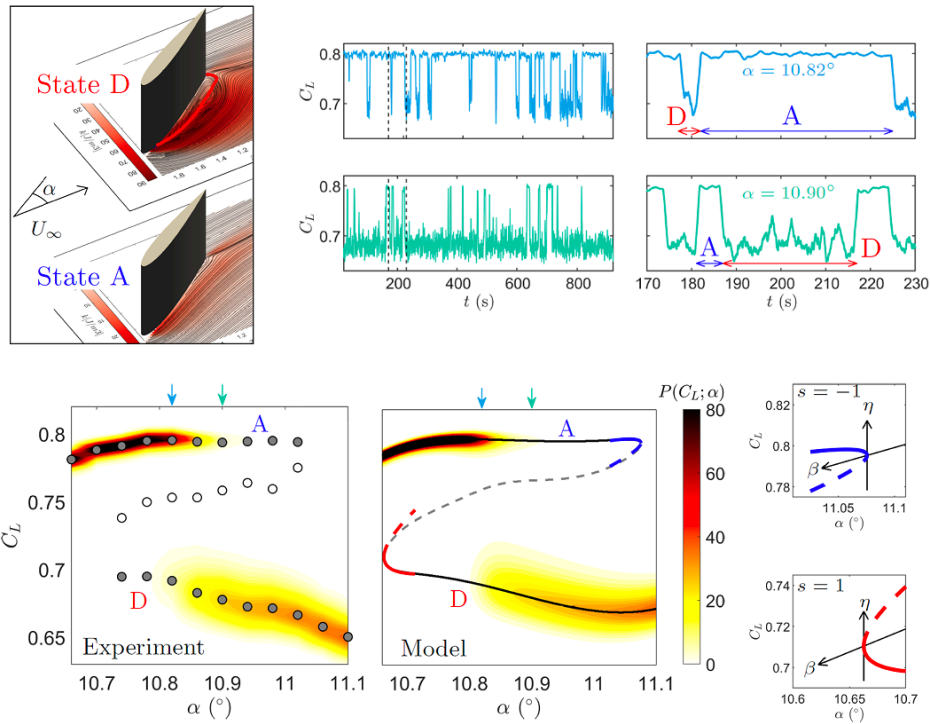


Figure 2: **Top left:** bistable states A (high-lift attached flow) and D (low-lift detached flow; red: backflow region boundary) from conditional PIV. **Top right:** transitions between states A and D in the experimental lift coefficient signals  $C_L(t)$ , for two different angles of attack  $\alpha$ . Left: 15 min window; right: 1 min window. **Bottom left and middle:** experimental and identified lift coefficient as a function of  $\alpha$  in the bistable range. Contours: PDF  $P(C_L; \alpha)$ . Gray circles: maxima of the PDF in states A and D; white circles: minimum of the PDF. Solid black lines: stable states A and D; dashed grey line: unstable branch of the double saddle-node bifurcation. **Bottom right:** zoom on the two local saddle-node bifurcations described by the identified normal forms.

## Mean resolvent analysis of a stochastic axisymmetric jet

Alessandro Franchini<sup>1</sup>, Nicolas Alferex<sup>2</sup>, Alessandro Bongarzone<sup>1</sup>, Cedric Content<sup>1</sup>, Laurent Cordier<sup>3</sup>, Denis Sipp<sup>1</sup>, Colin Leclercq<sup>1</sup>

<sup>1</sup>ONERA, DAAA, Meudon, France

<sup>2</sup>CNAM, Dynfluid, Paris, France

<sup>3</sup>Institut Pprime, Poitiers, France

Classical resolvent analysis of turbulent flows is based on the singular value decomposition of the operator  $\mathbf{R}_{\bar{Q}} = (i\omega\mathbf{I} - \mathbf{J}_{\bar{Q}})^{-1}$ , where  $\mathbf{J}_{\bar{Q}}$  is the Jacobian of the Navier–Stokes equations linearized about the mean flow. For each forcing frequency  $\omega$ , this framework provides optimal forcing and response modes, as well as associated energy gains. However, the approach is known to be formally ill-posed [1]: the poles of  $\mathbf{R}_{\bar{Q}}$  depend on the chosen formulation of the compressible Navier–Stokes equations. An alternative and rigorous framework has recently been proposed in the form of the mean resolvent operator  $\mathbf{R}_0$  [2]. Instead of linearizing about the mean flow, the dynamics are linearized along unsteady trajectories, leading to an operator that predicts the mean linear response of the unsteady flow (Figure 3(a-b)). The formalism is rooted in the well-established linear response theory of statistical physics [3]. Recent developments have demonstrated the feasibility of this approach for periodic flows and highlighted, in some cases, significant differences between classical and mean resolvent analyses in terms of optimal gains and mode shapes [4]. The present work aims at extending mean resolvent analysis to a stochastically forced axisymmetric jet at  $M = 0.1$  and  $Re_D = 1000$  (Figure 1). A projection strategy onto a low-dimensional subspace spanned by optimal modes of the classical resolvent  $\mathbf{R}_{\bar{Q}}$  is used to solve the associated eigenvalue problem at an affordable computational cost, as in [4].

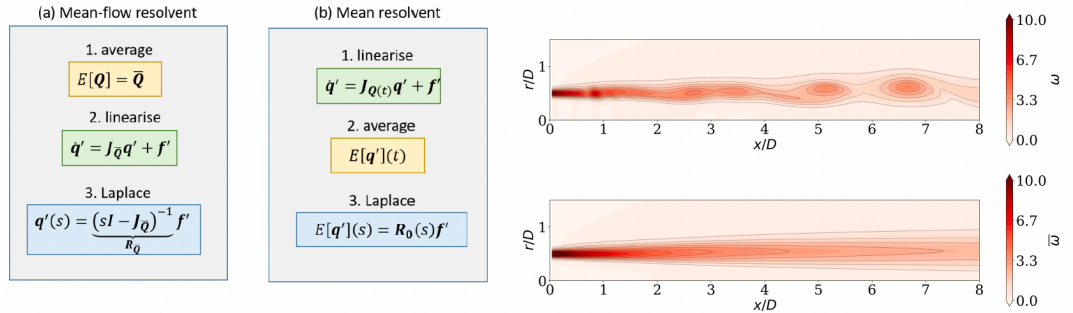


Figure 3: Left: (a) Resolvent about mean flow (b) Mean resolvent [4]. Top right: instantaneous snapshot of the vorticity  $\omega$  of the stochastically forced jet. Bottom right: mean vorticity  $\bar{\omega}$ .

1. U. Karban, B. Bugeat, E. Martini, A. Towne, A. V. G. Cavalieri, L. Lesshafft, A. Agarwal, P. Jordan, and T. Colonius. Ambiguity in mean-flow-based linear analysis. *J. Fluid Mech.*, 900:R5, 2020.
2. C. Leclercq and D. Sipp. Mean resolvent operator of a statistically steady flow. *J. Fluid Mech.*, 968:A13, 2023.

3. D. Ruelle. A review of linear response theory for general differentiable dynamical systems. *Nonlinearity*, 22(4):855, 2009.
4. A. Bongarzone, C. Content, D. Sipp, and C. Leclercq. Adjoint-free method for mean resolvent analysis of periodic flows. *arXiv preprint arXiv:2503.08401*, 2025.

## Multiscale Data Assimilation in Turbulence

Francesco Fossella<sup>1,2,3</sup>, Luca Biferale<sup>1,2</sup>, Alberto Carrassi<sup>4</sup>, Massimo Cencini<sup>5</sup>, Vikrant Gupta<sup>6</sup>

<sup>1</sup>Dept. of Physics, University of Rome Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

<sup>2</sup>INFN Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

<sup>3</sup>LTCI, Telecom Paris, 19 Place Marguerite Perey, 91120 Palaiseau, France

<sup>4</sup>Dept. of Physics and Astronomy, University of Bologna, Viale Carlo Berti Pichat 6/2, 40127 Bologna, Italy

<sup>5</sup>CNR-ISC, Istituto dei Sistemi Complessi, Via dei Taurini 19, 00185 Rome, Italy

<sup>6</sup>Guangdong Technion – Israel Institute of Technology, 241 Daxue Road, Shantou, Guangdong, 515063, China

Data Assimilation (DA) is an essential tool for integrating sparse measurements to reconstruct and improve predictions of physical model states. This study applies DA techniques to systems ranging from Shell Models [1] to the more realistic Rayleigh-Bénard convection, which provide ideal examples for studying fully developed turbulence and, more in general, nontrivial multi-scale systems. We apply DA methods, including the Ensemble Kalman Filter (EnKF) [2], Ensemble 4D-Var, and Nudging, to shell models to explore how measuring mesoscales (inertial range scales) can enhance predictions of large- and small-scale intermittent variables. Our study systematically varies observation frequencies and identifies the key scales that must be measured for statistical synchronization. We demonstrate [3] the superiority of the EnKF, showing that measurements taken at frequencies exceeding the characteristic frequency of the observed scales enable full synchronization of all larger scales, provided at least two adjacent mesoscale levels are measured. We then move to the 2D Rayleigh-Bénard flow, where, by using temperature-only measurements, we assess the performance of each method across various turbulence regimes and data sparsity levels, from both Eulerian and Lagrangian perspectives. Reconstruction accuracy is evaluated in terms of dynamical and statistical fidelity, focusing on spectral multiscale reconstruction.

1. A. Pomyalov I. Procaccia V. S. L'vov, E. Podivilov and D. Vandembroucq. Improved shell model of turbulence. *Phys. Rev. E*, 58(2):1811–1822, 1998.
2. F. C. Vossepoel G. Evensen and P. Jan van Leeuwen. *Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter Estimation Problem*. Springer, 2022.
3. A. Carrassi M. Cencini F. Fossella, L. Biferale and Vikrant Gupta. Multiscale data assimilation in turbulent models. *Physical Review E*, 113(2):024208, 2025.

# Uncertainty-Aware Latent-Dynamics ROM for State Estimation in Unsteady Flows

Aymane Lahgazi<sup>1</sup>, Denis Sipp<sup>2</sup>, Iraj Mortazavi<sup>1</sup>, Taraneh Sayadi<sup>1</sup>

<sup>1</sup>CNAM, M2N Lab, Paris, France

<sup>2</sup>ONERA, DAAA, Institut Polytechnique de Paris, Meudon, France

Active flow control requires fast state estimation and robust prediction, typically based on sparse measurements. We propose an uncertainty-aware, parameter-conditioned reduced-order model built using limited observational data and applied to unsteady incompressible flows. The model’s goal is to enable real-time forecasting and provide reliable surrogates that can be integrated into a model-based control framework. The architecture comprises the following three components: (i) a probabilistic encoder that infers a low-dimensional latent state distribution from sparse sensor observations and a conditioning variable  $\mu$  (representing operating conditions and/or control inputs); (ii) a latent dynamics model that enables accurate multi-step rollouts in latent space; and (iii) a decoder that reconstructs full-field quantities of interest from the latent trajectory. Feature-wise linear modulation (FiLM) [1] is used to inject conditioning on  $\mu$  throughout the architecture, applying parameter-dependent per-feature scaling and shifting to intermediate representations. The probabilistic formulation explicitly captures uncertainties in model prediction, producing calibrated confidence estimates [2] that are critical for robust decision making under partial observability and distribution shift. The approach is illustrated on the canonical laminar vortex-shedding flow past a circular cylinder, using sparse velocity measurements to reconstruct and forecast the unsteady wake. This framework provides an efficient, uncertainty-aware state estimator and predictor that can be embedded in closed-loop workflows for flow control and online inference.

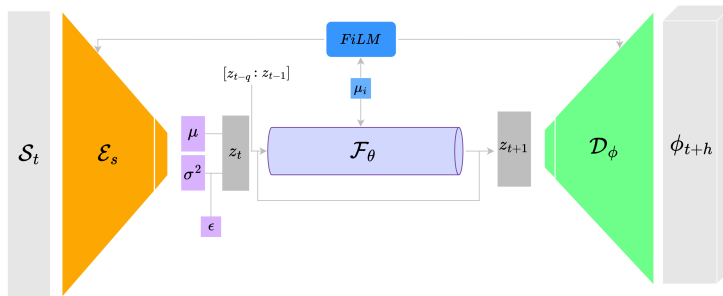


Figure 4: Parameter-conditioned probabilistic ROM.

1. Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In AAAI Conference on Artificial Intelligence, 2018.
2. Ismaël Zighed, Nicolas Thome, Patrick Gallinari, and Taraneh Sayadi. Up-drom: Uncertainty-aware and parametrised dynamic reduced-order model – application to unsteady flows, 2025.

## POD-Galerkin Reduced Order method for unsteady, turbulent, buoyant flows with parametric boundary conditions

Valerio Di Domenico<sup>1,2</sup>, Miguel A. Mendez<sup>1</sup>, Lilla Koloszar<sup>1</sup>, Yann Bartosiewicz<sup>2</sup>, Matilde Fiore<sup>1</sup>

<sup>1</sup>von Karman Institute for Fluid Dynamics, 1640 Sint-Genesius-Rode, Belgium,

<sup>2</sup>Institute of Mechanics, Materials, and Civil Engineering (iMMC), Université catholique de Louvain (UCLouvain), 1348 Louvain-la-Neuve, Belgium

We present a Galerkin Projection-Proper Orthogonal Decomposition (POD) reduced-order method for unsteady, turbulent, buoyant flows with parametric boundary conditions. The modeling approach is intended for application to the 1 : 5 scale model of the MYRRHA nuclear reactor pools, with the goal of enabling real-time modeling and control strategies for nuclear thermal hydraulics. The method builds upon the work of [1] and [2], implementing the predictive turbulence treatment proposed by [3], in which the coefficients of turbulent eddy-viscosity modes are computed via RBF interpolation. In contrast to similar approaches, parametric boundary conditions (BCs) are imposed using the hyperparameter-free method proposed in [4]. This approach constructs linear combinations of POD basis functions that vanish on the parametrized BCs during the projection phase. The boundary conditions are then enforced by adding  $N_{BC}$  equations to the reduced system, where  $N_{BC}$  is the number of parametrized BCs. Snapshots are generated for a cavity-flow test case using an incompressible URANS formulation with the Boussinesq approximation. The ROM is evaluated over a series of transient simulations at different Reynolds and Richardson numbers, triggering different stratification and mixing regimes. Results show that the ROM captures the interplay between thermal and inertial effects induced by the boundary conditions, reproducing the thermal profile along the cavity height with a relative  $L_2$  error of  $O(10^{-3})$  in the temperature field. Depending on the number of retained modes, the model achieves a significant ( $\approx O(10^2)$ ) speed-up compared to the CFD simulations, reaching faster-than-real-time performance, therefore paving the way for control-oriented applications.

1. Sabrina Kelbij Star, Giovanni Stabile, Sokratia Georgaka, Francesco Belloni, Gianluigi Rozza, and Joris Degroote. POD-Galerkin reduced order model of the Boussinesq approximation for buoyancy driven enclosed flows. In Int. Conf. Math. Comput. Methods Appl. Nucl. Sci. Eng., M C, pages 2452–2461. American Nuclear Society, 2019.
2. Sokratia Georgaka, Giovanni Stabile, Kelbij Star, Gianluigi Rozza, and Michael Bluck. A hybrid reduced order method for modelling turbulent heat transfer problems. *Computers & Fluids*, 208:104615, 2020.
3. Saddam Hijazi, Giovanni Stabile, Andrea Mola, and Gianluigi Rozza. Data-driven POD-Galerkin reduced order model for turbulent flows. *Journal of Computational Physics*, 416:109513, 2020.
4. Max D. Gunzburger, Janet S. Peterson, and John N. Shadid. Reduced-order modeling of time dependent PDEs with multiple parameters in the boundary data. *Computer Methods in Applied Mechanics and Engineering*, 196(4):1030–1047, 2007.

## Momentum-Conserving Physics-Informed Graph Neural Networks for Dynamical Systems

Olga Fink

EPFL Lausanne, Lausanne, Switzerland

Accurate and interpretable modeling of multi-body dynamical systems is a fundamental challenge in domains ranging from robotics and aerospace to biophysics and materials science. Traditional physics-based approaches are often computationally expensive and difficult to scale, while purely data-driven methods like graph neural networks (GNNs) may lack physical consistency and generalization. This talk presents Dynami-CAL GraphNet, a new physics-informed GNN framework that explicitly integrates conservation laws, specifically, the pairwise conservation of linear and angular momentum, into its architecture. By leveraging edge-local reference frames that are equivariant to rotations and translations, our model produces physically consistent predictions and offers interpretable insights into the forces and moments governing each interaction. We demonstrate the effectiveness of Dynami-CAL GraphNet across a wide spectrum of tasks. Beyond standard 3D granular systems with inelastic collisions, we systematically evaluated the model on complex, real-world datasets, including human body motion prediction and protein molecular dynamics simulations. In all cases, Dynami-CAL GraphNet was benchmarked against several established baseline methods. Our results show not only stable error accumulation over extended prediction horizons and superior maintenance of physical constraints, but also a strong ability to extrapolate to previously unseen system configurations and interaction regimes, a key capability for robust deployment in real-world scenarios. This talk will highlight how embedding physical principles within machine learning architectures enables not only accuracy and interpretability, but also robust extrapolation to previously unseen scenarios, opening new avenues for real-time, scalable, and generalizable modeling of complex systems in science and engineering.

## Sparse data conditioned diffusion model for PDE Solving

Eloi Lebre<sup>1</sup>, Marcello Meldi<sup>2</sup>, Taraneh Sayadi<sup>1</sup>

<sup>1</sup>Conservatoire National des Arts et Métiers, Paris, France

<sup>2</sup>Univ. Lille, CNRS, ONERA, Arts et Métiers ParisTech, Centrale Lille, UMR 9014- LMFL- Laboratoire de Mécanique des fluides de Lille - Kampé de Fériet, F-59000, Lille, France

In many fluid mechanics applications, direct access to full-state flow measurements is not feasible due to experimental limitations, sensor sparsity, or computational constraints. Instead, only partial observations are typically available, obtained from a limited number of measurement locations and only for some physical quantities. These sparse measurements can be incorporated within a data assimilation framework to infer the underlying full-state dynamics of the system. Accurate reconstruction of the full flow field from limited observational data remains a significant challenge, particularly for strongly nonlinear and chaotic systems. In this work, we investigate the potential of conditional diffusion models for reconstructing full-state flow fields from sparse observational data. Diffusion models [1] belong to a class of generative models that learn the data distribution by reversing a gradual stochastic noising process. During training, structured data are progressively perturbed by noise according to a predefined stochastic process, and the model learns to iteratively denoise the samples to recover the original data distribution. Owing to their stochastic formulation and ability to represent complex, high-dimensional probability distributions, diffusion models are particularly well suited for dynamic systems exhibiting multiscale interactions and strong unsteady behavior, such as the ones relevant for realistic applications in fluid mechanics. More recently, they have also demonstrated promising capabilities in data assimilation tasks by incorporating observational constraints directly into the sampling process ([2], [3]). In this work, we follow the conditional diffusion framework proposed by [4] and modify the architecture to enable reconstruction of full-state fields using only sparse observational data. We evaluate the performance of this approach in reconstructing high-dimensional flow fields from limited observational inputs and discuss its potential as a probabilistic data-assimilation tool for complex fluid systems. Experiments are realised on the 1D Kuramoto-Sivashinsky equation:  $\partial_t u + \partial_x^2 u + \nu \partial_x^4 u + \frac{1}{2} \partial_x(u^2) = 0$ , with  $u$  representing the velocity. Diffusion-based generative models have recently been used as surrogate models to autoregressively generate full-state flow realizations from past observations (see Figure 5).




Figure 5: Kuramoto–Sivashinsky dynamics compared with predictions, with a snapshot of the probability density function.

1. Jonathan Ho et al. Denoising diffusion probabilistic models, 2020.
2. Aliaksandra Shysheya et al. On conditional diffusion models for PDE simulations, 2024.
3. Caspar Meijer and Lydia Y. Chen. The rise of diffusion models in time-series forecasting, 2024.
4. G. Kohl et al. Benchmarking autoregressive conditional diffusion models for turbulent flow simulation. *Neural Networks*, 199, 2026.

# From spikes to dynamics: SNN-based latent space identification for wake flow prediction

Iacopo Tirelli, Stefano Discetti, Andrea Ianiro

Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, 28911, Spain

Deep learning has greatly advanced fluid mechanics, but the computational cost of state-of-the-art architectures remains high. Spiking Neural Networks (SNNs, [1]) offer an energy-efficient alternative by representing information as binary, event-driven spike trains, enabling ultra-fast, low-power computation on neuromorphic hardware. However, the binary encoding inherently introduces information loss and training challenges, and despite recent progress in classification tasks, SNN applications in fluid mechanics remain limited. We present a physics-informed SNN-based pipeline for flow reconstruction and prediction. A convolutional SNN encoder maps rate-coded spike representations of velocity fields to a compact vector bottleneck latent space, from which a convolutional decoder reconstructs the full flow field. Velocity fluctuations from DNS of the fluidic pinball at  $Re = 150$  are encoded via rate coding into four polarity channels  $[u^+, u^-, v^+, v^-]$ , mimicking the ON/OFF response of neuromorphic event cameras. A convolutional Leaky-Integrate-and-Fire network compresses the resulting spike trains into a 32-dimensional latent vector, and a convolutional decoder with a learned depthwise spatial filter reconstructs the full velocity field. Divergence-free and vorticity penalties are applied during training to enforce physical consistency. Preliminary results show that the architecture achieves a reconstruction error of 10%, while zero-shot generalization to  $Re = 120$  yields a 16% relative error. The learned latent representation also enables efficient dynamics prediction: a Long Short-Term Memory trained with a window of 2 shedding periods  $T_{shed}$  achieves relative errors of 15.3%, 36.2%, 54%, and 80.1% at prediction horizons  $h = 1$ ,  $h = 0.5T_{shed}$ ,  $h = T_{shed}$  and  $h = 2T_{shed}$ , respectively (reconstruction floor is  $\approx 10\%$ ). Future work will extend the pipeline to more complex turbulent flows, with architectural refinements for both the SNN encoder and the dynamics predictor, targeting full neuromorphic hardware deployment.

**Acknowledgments** This work was supported by the project SPANDRELS (SParse AND paRsimonious Event-based fLOW Sensing), which has received funding from the European Union's Horizon Europe research and innovation program under grant agreement No 101171280 (ERC-2024-COG).

1. Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.

# Navigating the Intermittency: A Generative Surrogate for Long-Horizon Forecasting of Minimal Flow Unit from Sparse Measurements

Niccolo Tonioni<sup>1</sup>, Marcial Sanchis-Agudo<sup>2</sup>, Mohammad Umair<sup>3,2</sup>, Lionel Agostini<sup>1</sup>, Laurent Cordier<sup>1</sup>, Ricardo Vinuesa<sup>4,2</sup>

<sup>1</sup>Curiosity Team, Pprime Institute, CNRS, Université de Poitiers, ISAE-ENSMA, Bât. H2 - Bd. Marie & Pierre Curie TSA 51124, 86073, Poitiers

<sup>2</sup>FLOW, Mekanik, KTH Royal Institute of Technology, SE-100 44, Stockholm, Sweden

<sup>3</sup>Lab. des Écoulements Géophysiques et Industriels, BP 53, Grenoble 38041, France

<sup>4</sup>Dept. of Aerospace Eng., University of Michigan, Ann Arbor, Michigan, United States

This work introduces an equation-free, data-driven reduced-order model for wall-bounded turbulent flows, designed for potential integration into active control loops. The framework integrates a  $\beta$ -VAE-GAN for non-linear spatial dimensionality reduction with a sensor-conditioned Easy Attention Transformer for temporal evolution. Evaluated on the Minimal Flow Unit at  $Re_\tau = 200$  at  $y^+ = 14$ , the architecture compresses the flow into a four-variable latent space while preserving the turbulent kinetic energy and integral length scales. This learned manifold autonomously isolates the low-frequency signatures of the near-wall intermittent regeneration cycle. By conditioning the Transformer on three sensors, the model sustains accurate latent trajectories over extended horizons from a minimal initialization window. End-to-end inference reconstructs flow fields near the upper limit of the compression stage, with quadrant analysis confirming the accurate capture of dominant ejection and sweep events. While extreme bursts are slightly attenuated, the framework's ability to track alternating active and quiescent phases demonstrates its potential as a robust state-estimator for future model-based control applications.




Figure 6: Temporal evolution of the spatial maximum of the instantaneous Reynolds shear stress,  $(-u'v')_{max}^+$ , comparing the DNS reference data (gray) and the end-to-end framework prediction (sea-green).

## Sparse-sensing-state estimation of vortex-gust airfoil interactions

Alberto Solera-Rico<sup>1,2</sup>, Arnau Miró<sup>3,4</sup>, Oriol Lehmkühl<sup>4</sup>, Carlos Sanmiguel Vila<sup>1,2</sup>

<sup>1</sup>Aerial Platforms Department, Spanish National Institute for Aerospace Technology (INTA), San Martín de la Vega, 28330, Madrid, Spain,

<sup>2</sup>Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, 28911, Madrid, Spain

<sup>3</sup>TUAREG Research Group, Universitat Politècnica de Catalunya (UPC), Carrer Colom 11, 08222 Terrassa, Spain

<sup>4</sup>Barcelona Supercomputing Centre (BSC), CASE Large-Scale Computational Fluid Dynamics, Placa d'Eusebi Güell, 1-3, 08034 Barcelona, Spain

Active flow control in realistic aerodynamic environments requires more than actuation alone: it also requires access to the instantaneous flow state from a limited number of practical sensors. This challenge becomes particularly acute under gust disturbances and clear-air turbulence, where the flow experienced by the lifting surface is highly unsteady, spatially complex, and only partially observable. In such conditions, full-field information is unavailable in operation, while dense sensing is difficult to justify in terms of hardware integration, robustness, and data-rate constraints. This contribution explores sparse-sensing state estimation as a necessary step toward data-driven aerodynamic control. The underlying premise is that, although gust-driven aerodynamic flows remain high dimensional, their dominant response may still be organized around low-dimensional coherent structures [1]. If such latent flow states can be inferred from surface-mounted measurements, they may provide compact and physically meaningful representations for downstream estimation, prediction, and future control design [2]. The study is based on DNS data at  $Re = 5000$  for a NACA0012 airfoil subjected to transient gust disturbances generated using the high-order code SOD2D. Taylor-type vortical structures are introduced as an idealized inflow perturbation to emulate localized unsteady loading. At this stage, the objective is not to assess a specific closed-loop strategy, but rather to investigate the observability of the gust-response dynamics from sparse wall information and to examine whether reduced state representations can be identified in a form suitable for later model-based or reinforcement-learning frameworks. By focusing on the relation between sparse sensing, flow-state inference, and reduced descriptions of unsteady aerodynamic response, this work aims to contribute to ongoing efforts in data-driven active flow control and to stimulate discussion on how simulation-based learning can support future estimation and control under gust disturbances.

**Acknowledgments** Work produced with the support of a 2024 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation, project PREVENT grant n.LEO24-2-15988-ING-ING-203. The Foundation takes no responsibility for the opinions, statements and contents of this project, which are entirely the responsibility of its authors.

1. Kai Fukami and Kunihiko Taira. Grasping extreme aerodynamics on a low-dimensional manifold. *Nature Communications*, 14(1):6480, 2023.
2. Kai Fukami, Yuta Iwatani, Soju Maejima, Hiroyuki Asada, and Soshi Kawai. Compact representation of transonic airfoil buffet flows with observable-augmented machine learning. *Journal of Fluid Mechanics*, 1021:A39, 2025.

## Dominant recurrent carrier of turbulence within spatially-localized turbulent structures

Lukasz Klotz

Warsaw University of Technology, IAAM, Poland

Despite being an almost century-old concept, only the recent advances in numerical computations have enabled the practical application of periodic orbit theory to chaotic dynamics embedded within a fluctuating turbulent phase. Similarly, recent experiments provide support for the relevance of periodic orbits in turbulent dynamics across weak and moderate regimes. However, no experimental evidence has been reported to date for the dominant recurrent orbit as a precursor of the turbulent dynamics in the early transitional range of Reynolds numbers. In this paper, we present the first experimental validation for such a dominant recurrent carrier of turbulence in Couette-Poiseuille flow - a simple wall-bounded shear flow introduced recently to study transition to turbulence [1]. The advantage of zero-mean flow allows us to generate fluctuating patches of turbulence that are nearly stationary in the laboratory frame of reference, and to quantify the dominant period of the turbulence regeneration cycle. This sheds light on the underlying dynamical processes that sustain the turbulent phase on the short time scale, providing the first experimental signature of recurrent dynamics underlying the regeneration cycle of turbulence proposed by F. Waleffe nearly three decades ago.




Figure 7: Flow visualizations of a doubly-localised turbulent spot (panel a) and an oblique turbulent band (panel b). Green rectangle conceptually marks the minimal flow unit supporting Waleffe's [2] self-sustaining cycle within the localized turbulent structures following; panel c) illustrates spatial structure at three different instants and periodic time evolution of G. Kawahara & S. Kida's recurrent orbit representing regeneration cycle of turbulence (data inferred from [3]).

1. L. Klotz, A. M. Pavlenko, and J. E. Wesfreid. Experimental measurements in plane Couette-Poiseuille flow: dynamics of the large- and small-scale flow. *J. Fluid Mech.*, 912, 2021.
2. F. Waleffe. On a self-sustaining process in shear flows. *Phys. Fluids*, 9(4):883-900, 1997.
3. G. Kawahara and S. Kida. Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.*, 449:291-300, 2001.

## Estimating turbulent channel flow from wall measurements

Marco Castelletti<sup>1</sup>, Federica Gattere<sup>2</sup>, Paolo Luchini<sup>3</sup>, Maurizio Quadrio<sup>1</sup>

<sup>1</sup>DAER, Politecnico di Milano, via La Masa 34, 20156 (MI), Italy

<sup>2</sup>Gran Sasso Science Institute, viale Crispi 7, 67100 (AQ), Italy

<sup>3</sup>DIMEC, Università di Salerno, via Ponte don Melillo 1, 84084 (SA), Italy

This work considers the state estimation in a turbulent channel flow, with the aim of predicting the velocity inside the channel by filtering wall-based measurements through a designed statistical estimator. Accurate estimation from wall measurements is fundamental as for any practical implementation of control strategy which have to be based solely on quantities measurable at the wall. The present estimators extend those proposed by [1] and [2], and introduces optimal, causal, linear time-invariant (LTI) filters for the problem at hand, obtained without introducing assumptions beyond the linearity of the estimator. For the first time, a causal Wiener filter is successfully demonstrated. The estimator is constructed directly from the co-spectra between wall measurements and the internal velocity field, obtained from direct numerical simulations. In this way, the governing equations are effectively replaced by a linear model that reproduces the same second-order (space-time) statistics of the turbulent flow. This approach differs from previous studies, where the predictor is derived from linearized representations of the Navier–Stokes equations, such as the Orr–Sommerfeld–Squire equations or the resolvent operator. Results indicate that the presence of an equation-based linear model is not required to achieve accurate predictions, as confirmed by comparison with the approaches of [1] and [2]. Prediction accuracy decreases with distance from the wall and depends on the Reynolds number, reflecting the increasing influence of large-scale attached motions at higher Reynolds numbers (Figure 8). The proposed method shows a significant improvement over [1]. Compared with [2], it provides better near-wall predictions but worse performance in the bulk, owing to the limited statistical sample used in [2].




Figure 8: Correlation between the true  $u$  and estimated  $\tilde{u}$  streamwise velocity obtained with a causal multi-input filter using the wall measurements  $\tau_x$ ,  $\tau_z$ , and  $p$  at one wall.

1. M. Chevalier, J. Höpfner, T. R. Bewley, and D.S. Henningson. State estimation in wall-bounded flow systems. Part 2. Turbulent flows. *J. Fluid Mech.*, 552:167–187, 2006.
2. F. R. Amaral, A. V. G. Cavalieri, E. Martini, P. Jordan, and A. Towne. Resolvent-based estimation of turbulent channel flow using wall measurements. *J. Fluid Mech.*, 927:A17, November 2021.

## Tuesday 9th – Day 2 - Model-based and model free control

### **Koopman-based model-order reduction: coherent structures, frequencies and damping rates in stochastic cavity flow**

Denis Sipp

MONHADE, Inria-ONERA team, ONERA, Institut Polytechnique de Paris, Palaiseau, France

This presentation explores how Koopman-based model-order reduction can extract coherent structures, frequencies, and damping rates from stochastic flow data. Starting from cross-correlation and cross-spectral density matrices, we show how realization theory reveals finite-dimensional dynamics hidden inside turbulent signals. The approach is connected to Mori-Zwanzig theory, Ruelle-Pollicott resonances, and spectral POD. A stochastic two-dimensional cavity flow is used as a testbed. Particular attention is paid to robustness, bootstrap sensitivity, and the separation between discrete resonances and continuous-spectrum effects. Finally, similarity transforms and invariant graphs are used to assess how close the identified dynamics are to a Markovian reduced model.

## Model Predictive Control in latent coordinates for partially observable systems

Luigi Marra<sup>1</sup>, Onofrio Semeraro<sup>2</sup>, Lionel Mathelin<sup>2</sup>, Andrea Meilán-Vila<sup>3</sup>, Stefano Discetti<sup>1</sup>

<sup>1</sup>Department of Aerospace Engineering, Universidad Carlos III de Madrid, Spain

<sup>2</sup>CNRS, Laboratoire interdisciplinaire des sciences du numérique (LISN), Université Paris-Saclay, France

<sup>3</sup>Department of Statistics, Universidad Carlos III de Madrid, Spain

The application of Model Predictive Control (MPC) to high-dimensional and partially observable flows remains challenging due to the computational cost of predictive models and the limited availability of measurements. In this work, we propose a scalable data-driven control strategy that integrates reduced-order modelling, state estimation, and nonlinear MPC within a unified latent-space formulation. A compact predictive model of the system dynamics is constructed using Operator Inference applied to Proper Orthogonal Decomposition modes, enabling a non-intrusive identification of low-dimensional latent dynamics directly from data. The resulting reduced-order model is embedded in a nonlinear MPC loop, where control inputs are optimised over a receding horizon in the latent coordinates, significantly reducing the computational burden compared with full-order formulations. To address partial observability, an Unscented Kalman Filter is employed to reconstruct the latent state from sparse and noisy measurements, thereby enabling output-feedback MPC without requiring access to the full system state. The proposed framework is assessed on canonical chaotic systems described by the one- and two-dimensional Kuramoto-Sivashinsky equations. The results demonstrate robust closed-loop performance despite severe measurement limitations (see Fig. 9).

**Acknowledgments** This work is supported by the funding 'Orden 3789/2022, del Vicepresidente, Consejero de Educación y Universidades, por la que se convocan ayudas para la contratación de personal investigador predoctoral en formación para el año 2022'

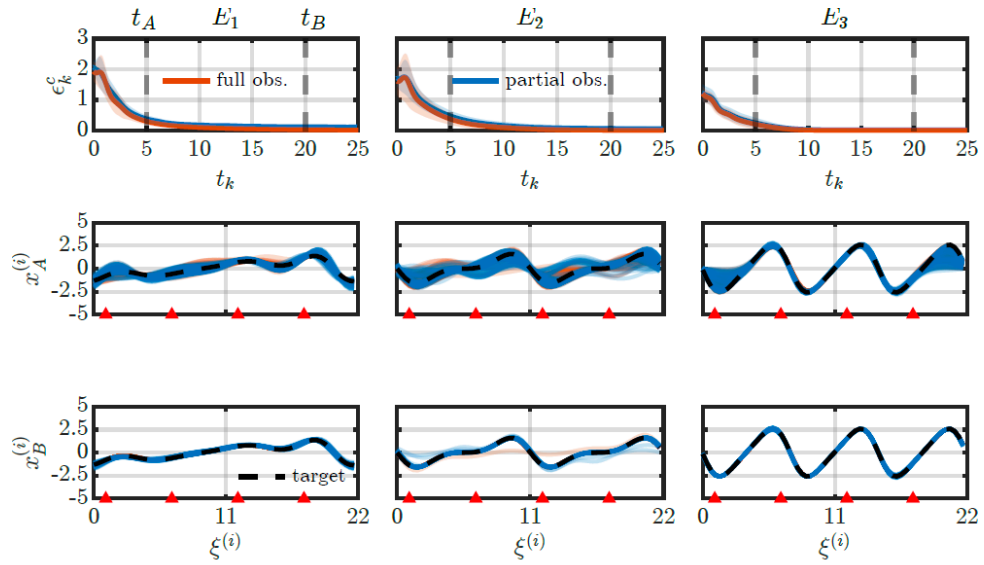


Figure 9: Application of the latent-space MPC around the three unstable equilibrium points of the 1D Kuramoto-Sivashinsky system, indicated with  $E_1$  (left column),  $E_2$  (middle column), and  $E_3$  (right column) in a configuration with 4 sensors, measurement noise  $\sigma_{\nu\nu} = 0.1$ , and sampling interval  $T_s = 0.1$ . Top row: control error of the true state with respect to the target. Center and bottom rows: state evolution with MPC in full and partial observability at two different time instants  $t_A$  and  $t_B$ .

## Increasing non-linear stability of flows using optimal control

Mathieu Salmon<sup>1</sup>, Colin Leclercq<sup>1</sup>, Jean-Christophe Robinet<sup>2</sup>

<sup>1</sup>ONERA DAAA, Institut Polytechnique de Paris, 8 rue des Vertugadins, 92190 Meudon

<sup>2</sup>DynFluid Laboratory, Arts et Métiers Institute of Technology, 151 boulevard de l'Hôpital, 75013, Paris.

We investigate the possibility of increasing the volume of the basin of attraction of a stable fixed point using full-state linear feedback control. For a fixed controller gain, two approaches are compared: one where we maximize the energy of the minimal seed (point of minimal energy on the basin boundary), and one where we maximize the energy of the edge state (attractor relative to the edge manifold but saddle in the full phase space); see figure 10*left*. The methods are applied to a four-dimensional reduced-order model of the self-sustaining process in wall-bounded shear flows [1, 2]. In that context, increasing the size of the basin of attraction of the laminar base flow amounts to making it more robust to subcritical to transition to turbulence caused by finite-amplitude perturbations. The two approaches provide a significant increase in the volume of the basin of attraction, the edge state approach being superior (see figure 10*right*). Even though both optimal controllers predominantly combat the lift-up mechanism, which is purely linear, they also exploit nonlinear mechanisms such as damping secondary sinuous perturbations to the streaks. This allows the basin boundary to be expanded further than if only linear transient growth was minimized.




Figure 10: Left: Two *local* approaches for *globally* expanding the basin of attraction of a laminar base flow at the origin: increase edge state ( $\mathbf{q}_{ES}$ ) energy or minimal seed ( $\mathbf{q}_{MS}$ ) energy. Right: Probability of relaminarization as a function of initial perturbation energy  $E_0$ ; ES (edge state optimization), MS (minimal seed optimization), LTG (linear transient growth minimization). Vertical lines indicate minimal seed energies for the ES and MS cases.

1. F. Waleffe. On a self-sustaining process in shear flows. *Phys. Fluids*, 9(4):883–900, 1997.
2. C. Cossu, L. Brandt, S. Bagheri, and D.S. Henningson. Secondary threshold amplitudes for sinuous streak breakdown. *Phys. Fluids*, 23(7), 2011.

## Drag and lift optimization with a surrogate model

Philippe Gilotte<sup>1</sup>, Ceyhan Erdem<sup>2</sup>, Iraj Mortazavi<sup>3</sup>

<sup>1</sup>TourbillOnde

<sup>2</sup>Friendship Systems

<sup>3</sup>M2N CNAM Paris

Aerodynamic improvements in the transport industry can now involve shape morphing technics involving radius basis functions, included inside a full optimization process. This methodology, already tested for automotive industry [1] is now presented for an optimization example of lift and drag performances, according to the relative positions between the wing, the flap and the slat profiles at landing conditions. After validation of the grid convergence thanks to comparison with available results on the baseline configuration [2], the CAESES software is used to drive a Sobol sequence which enables to create a Kriging model, derived from CFD results obtained with ultraFluidX, a LES lattice Boltzmann solver. A gradient descent is then performed with the surrogate model to select the set of geometrical parameters leading to the optimum drag and lift coefficients. Streamlines show importance of the vortex strength and location behind the flap depending on the geometrical parameter values (see figure 11). A DMD analysis is then performed with a more refined grid of the baseline case, to identify acoustic source locations on which applying this shape optimization procedure at specific frequencies.

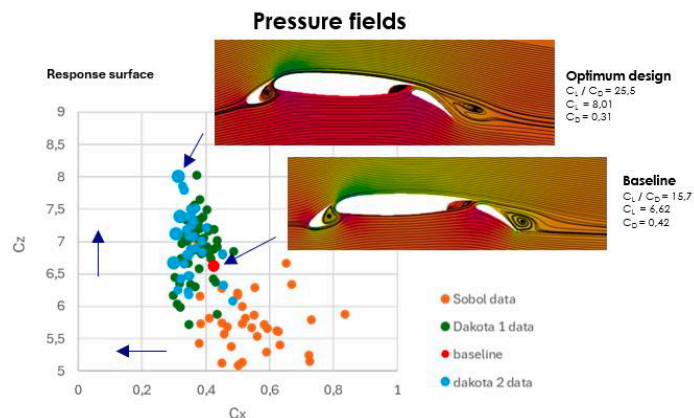


Figure 11: drag and lift optimization results with a Kriging model

1. Erdem, C.; Eulalie, Y.; Gilotte, P.; Harries, S.; Nayeri, C.N.; Aerodynamic Optimization of a Reduced Scale Model of a Ground Vehicle with a Shape Morphing Technique. *Fluids* 2022, 7, 166. <https://doi.org/10.3390/fluids7050166>.
2. Manoha, E.; Pott-Pollenske, M.; LEISA2: an experimental database for the validation of numerical predictions of slat unsteady flow and noise. AIAA 2015-3137. 21st AIAA/CEAS Aeroacoustics Conference. June 2015.

## Safe learning by combination of MPC and reinforcement learning

Luigi Marra<sup>1</sup>, Onofrio Semeraro<sup>2</sup>, Lionel Mathelin<sup>2</sup>, Andrea Meilán-Vila<sup>3</sup>, Stefano Discetti<sup>1</sup>

<sup>1</sup>Department of Aerospace Engineering, Universidad Carlos III de Madrid, Spain

<sup>2</sup>CNRS, Lab. interdisciplinaire des sciences du numérique (LISN), Univ. Paris-Saclay, France

<sup>3</sup>Department of Statistics, Universidad Carlos III de Madrid, Spain

Safe and data-efficient control design is a key requirement for the practical deployment of machine learning strategies in active flow control, where unstable dynamics, actuation constraints and limited sensing may lead to unsafe exploration. In this contribution, we propose a hybrid control framework that combines model predictive control (MPC) and reinforcement learning (RL) through imitation-based learning. MPC is employed as a stabilising expert controller based on a reduced-order model of the flow, providing constraint-aware optimal trajectories that are used to initialise a feedback policy via behavioural cloning. The policy is progressively improved by aggregating expert-corrected rollouts generated on the plant model, and subsequently refined through RL fine-tuning once sufficient performance and robustness are achieved. The methodology is demonstrated on the control of the chaotic 1D Kuramoto-Sivashinsky equation. The control objective is to drive the system from random initial conditions to one of its unstable equilibrium points. The plant model of the MPC is a data-driven reduced-order model obtained by operator inference. The results show that the proposed MPC-guided learning approach substantially improves training safety and convergence speed compared with standalone RL. In addition, the learned policy retains good performance under sparse sensing conditions, highlighting its potential relevance for realistic flow-control configurations. Figure 12 shows the learning performance and policy assessment for the proposed MPC-guided learning strategy. The running average of failing episodes shows a rapid decrease after imitation-based initialisation, indicating improved training safety and faster convergence. The violin plots of the average episode reward show that the warm-started policy achieves improved performance and reduced variability, highlighting the benefit of combining MPC-driven demonstrations with reinforcement learning for flow-control policy design.




Figure 12: Left: Evolution of the average reward per episode during RL training as a function of the number of environment transitions. Right: Distribution of the average episode reward for policies obtained after imitation learning, after standalone RL, and after RL fine-tuning of the imitation-initialised policy.

**Acknowledgments** This work is supported by the funding 'Orden 3789/2022, del Vicepresidente, Consejero de Educación y Universidades, por la que se convocan ayudas para la contratación de personal investigador predoctoral en formación para el año 2022'

## A sliding mode observer perspective on chaos synchronisation of turbulent flows

Defne E. Ozan, Andrew Wynn

Dept. of Aeronautics, Imperial College London, SW7 2AZ, London, UK

The state estimation of turbulent flows is challenging due to their chaotic dynamics, as infinitesimal prediction errors grow exponentially. However, it has been remarkably observed that when a signal from one chaotic system, i.e., the master, is fed to another identical system (or subsystem), i.e., the slave, the dynamics of the slave system synchronises with that of the master system [1]. This phenomenon has been similarly observed in turbulent flows [2, 3]. From a control theory perspective, the problem of chaos synchronisation can be viewed as the design of an observer [4, 3]. In this work, we first show that synchronisation achieved by direct substitution of the observed state variables into the slave system, e.g., [3], is equivalent to a sliding mode observer. In a sliding mode observer, the observation error is kept on a zero surface with the observer dynamics being effectively constrained to a reduced-order manifold [5]. Chaos synchronisation is then guaranteed to occur when the equivalent error dynamics of the reduced system are stable. Beyond direct substitution, the observer perspective enables the design of feedback mechanisms with desirable properties such as faster convergence or more robust disturbance rejection. State space transformations to companion forms via a nonlinear observability matrix can simplify the design of observer gains [6]. We explore neural network strategies for efficient data-driven mappings from the observation history to these companion forms, where such transformations may be analytically intractable for high-dimensional systems. We demonstrate our results on the Lorenz system (Fig. 1) and a 9-dimensional low-order model of a turbulent shear flow.

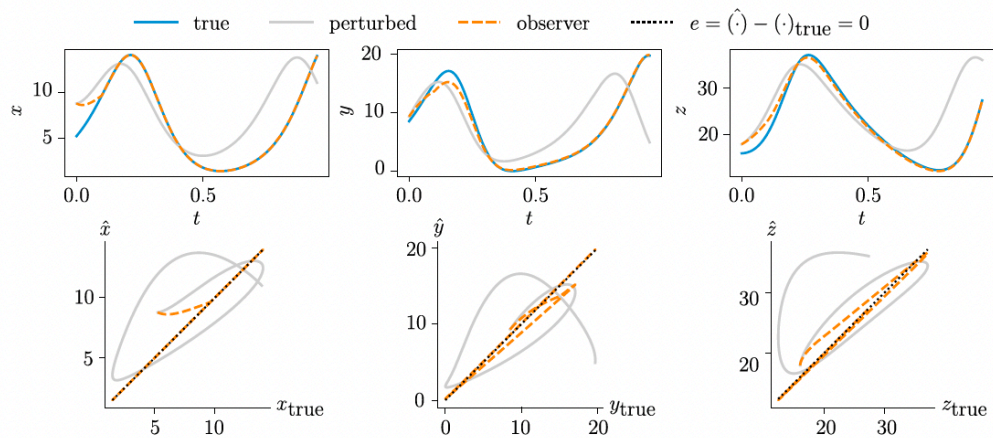


Figure 13: Sliding mode observer for the Lorenz system, where the state variable  $x$  is observed

1. L. M. Pecora and T. L. Carroll. Synchronization in chaotic systems. *Phys. Rev. Lett.*, 64(8):821–824, 1990
2. C. C. Lalescu, C. Meneveau, and G. L. Eyink. Synchronization of Chaos in Fully Developed Turbulence. *Phys. Rev. Lett.*, 110(8):084102, 2013.
3. T. A. Zaki. Turbulence from an Observer Perspective. *Annu. Rev. Fluid Mech.*, 57(1):311–334, 2025.
4. H. Nijmeijer and I. M. Y. Mareels. An observer looks at synchronization. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 44(10):882–890, 1997.
5. J.-J. E. Slotine, J. K. Hedrick, and E. A. Misawa. On Sliding Observers for Nonlinear Systems. In 1986 American Control Conference, pages 1794–1800, 1986.
6. S. Drakunov and V. Utkin. Sliding mode observers. Tutorial. In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, volume 4, pages 3376–3378 vol.4, 1995.

## Actuation manifold for control of a flapping wing under dynamic flow conditions

Guy Y. Cornejo Maceda, Lucas Schneeberger, Zhiyuan Wang, Iacopo Tirelli, Andrea Ianiro, Stefano Discetti

Department of Aerospace Engineering, Universidad Carlos III de Madrid, Av. de la Universidad, 30, Leganés, 28911, Madrid, Spain

Effective model predictive control (MPC) of unsteady aerodynamic flows requires compact, accurate models that capture the essential dynamics under actuation. In this study, we identify a low-dimensional actuation manifold for flapping wing dynamics, providing a principled foundation for the development of control-oriented reduced-order models. Building on the actuation manifold framework introduced by Marra et al. [1] for actuated flows under steady conditions, we extend this concept to unsteady flows, a regime of direct relevance to real-world flow control applications. The experimental setup consists of a two-dimensional NACA 0018 airfoil mounted on a six-degrees-of-freedom hexapod robotic platform in a fan-array wind generator. Unsteady flow conditions are imposed through prescribed global wing motions, while flow control is achieved by superimposing low-amplitude angle-of-attack variations. Particle image velocimetry (PIV), synchronized with unsteady pressure measurements, provides simultaneous access to the flow field and aerodynamic loads. The actuation manifold is extracted through a systematic three-step approach: (i) acquisition of a rich dataset spanning a broad range of dynamical regimes and control inputs; (ii) motion-compensated post-processing of PIV data to ensure a consistent reference frame; and (iii) nonlinear dimensionality reduction to reveal the intrinsic low-dimensional structure of the controlled flow. The resulting manifold offers a compact, data-driven representation of the input-output dynamics of the system, directly amenable to model identification for MPC. This work, therefore, addresses a key bottleneck in model-based flow control: the construction of low-dimensional models that remain valid across a wide range of operating conditions and actuation inputs.

**Acknowledgments** This activity is part of the project EXCALIBUR (Grant No PID2022-138314NBI00), funded by MCIU/AEI/ 10.13039/501100011033 and by "ERDF A way of making Europe".

1. L. Marra, G. Y. Cornejo Maceda, A. Meilán-Vila, V. Guerrero, S. Rashwan, B. R. Noack, S. Discetti, and A. Ianiro. Actuation manifold from snapshot data. *J. Fluid Mech.*, 996:A26, 2024.

## Can we predict the transient trajectories of the actuated fluidic pinball?

Alicia Rodríguez-Asensio<sup>1</sup>, Guy Y. Cornejo Maceda<sup>1</sup>, Bernd R. Noack<sup>2,3,1</sup>, Stefano Discetti<sup>1</sup>, Andrea Ianiro<sup>1</sup>

<sup>1</sup>Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, 28911, Spain

<sup>2</sup> College of Mechatronics and Control Engineering, Shenzhen University, Canghai campus, 518060, Shenzhen, PR China

<sup>3</sup> Guangdong Province VTOL Aircraft Manufacturing Innovation Center, 518060, Shenzhen, PR China

We develop a feature-based manifold model to describe the transient aerodynamic response of the actuated fluidic pinball under time-varying cylinder rotation. For steady actuations, the flow evolves toward limit cycles characterized by a steady mean drag coefficient and an oscillatory lift coefficient. These variables define a steady actuation manifold that describes the evolution of the global wake state with respect to the rotation speed. However, under unsteady actuation, the wake cannot follow this manifold instantaneously because of the intrinsic relaxation time scales associated with vortex-shedding dynamics and mean-flow deformation. Based on the steady actuation manifold and the time-delay representation, we aim to predict the evolution of the flow features of the fluidic pinball under time-varying actuation, following the procedure illustrated in Figure 14.

**Acknowledgments** This work is part of the project EXCALIBUR (Grant No PID2022-138314NB100), funded by MCIU/AEI/ 10.13039/501100011033 and by "ERDF A way of making Europe".

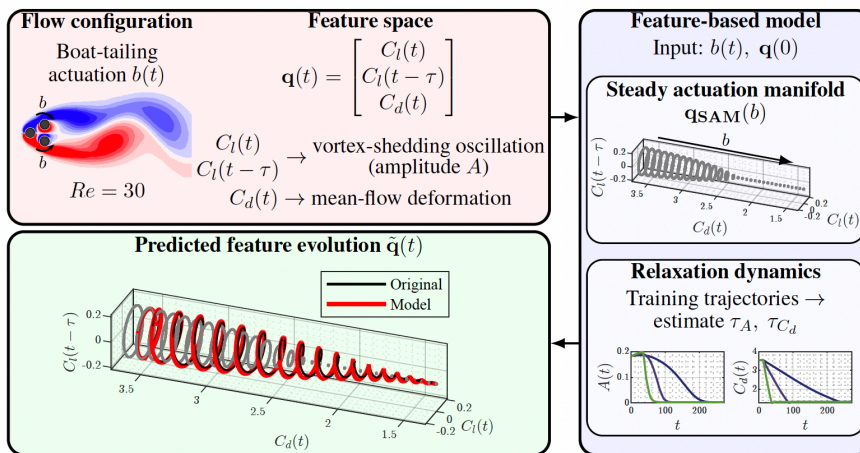


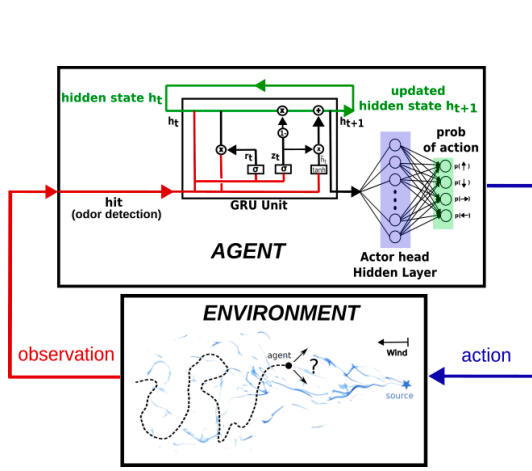
Figure 14: Schematic of the feature-based manifold model for predicting the transient wake response of the actuated fluidic pinball. The state is represented by a low-dimensional feature vector composed of lift and drag coefficients. Controlled dynamics are modeled using the steady actuation manifold together with relaxation time scales identified from transient trajectories.

# Reinforcement Learning for Olfactory Navigation in a Turbulent Flow


Vaibhav Chaturvedi, Christophe Eloy, Aurore Loisy

Aix-Marseille Univ, CNRS, Centrale Méd, IRPHE, Marseille, France

Olfactory navigation is the problem of locating an odor source using only odor detections as observations. In turbulent flows, this problem is very challenging because odor patches are quickly dispersed by the flow. Despite having only access to rare odor detections, numerous insects can locate odor sources from hundreds of meters away – a feat that is not well understood yet and remains an open challenge in robotics. To discover strategies for olfactory navigation, we propose a novel approach based on Reinforcement Learning (RL). We model the finite memory of our agent, encoding information about odor detections, as the hidden state of a GRU cell (Fig. 15a). Actions are predicted from this hidden state by the actor head within our Proximal Policy Optimization (PPO) framework. We use the TURB-Smoke[1] database, generated from high-fidelity DNS of scalar transport in a turbulent flow, to generate realistic odor detections for training our agent at locating the source. To analyze the policy, we project the learned hidden state (32-dimensions) on a 2D latent space using an autoencoder. In this latent space (Fig. 15b) we observe a clear geometrically-separated classification of the actions. Although not as good as the original learned policy, the latent space projection of the policy does retain comparable performance across various metrics.



(a) RL framework.



(b) Policy in 2D latent space where the colors indicate the action taken at each state; the overlaid curve traces the trajectory of a representative search episode.

1. L. Biferale, F. Bonaccorso, N. Cocciaglia, R. A. Heinonen, and L. Piro. Turb-smoke. a database of lagrangian pollutants emitted from point-sources and dispersed in turbulent flows. arXiv:2507.22749, 2025.

## Data-driven suppression of naturally developing Tollmien–Schlichting waves using plasma actuation

Jingran Qiu<sup>1</sup>, Lorenzo Piro<sup>2</sup>, Luca Biferale<sup>2</sup>, Massimo Cencini<sup>3,4</sup>, Bernhard Mehlig<sup>1</sup>, Kristian Gustavsson<sup>1</sup>

<sup>1</sup>Dept. of Physics, Gothenburg University, 41296 Gothenburg, Sweden

<sup>2</sup>Dept. of Physics & INFN, University of Rome "Tor Vergata", Via della Ricerca Scientifica 1, 00133 Rome, Italy

<sup>3</sup>Istituto dei Sistemi Complessi, CNR, Via dei Taurini 19, 00185 Rome, Italy

<sup>4</sup>INFN "Tor Vergata", Via della Ricerca Scientifica 1, 00133 Rome, Italy

Navigation in turbulent environments poses a central challenge for both biological and artificial microswimmers. Whereas prior work has predominantly emphasized modulation of propulsion or steering, the navigational potential of active morphological adaptation remains largely unexplored. We consider a deformable spheroidal microswimmer that dynamically adjusts its aspect ratio while seeking to maximise net displacement from its release point in two-dimensional stochastic and turbulent flows. A reinforcement learning (RL) framework is employed to infer feedback policies conditioned on swimmer orientation and local velocity-gradient information. The resulting control strategies consistently outperform fixed-shape and short-time-optimal benchmarks across flow regimes. Notably, policies trained in stochastic flows retain their effectiveness when deployed in fully resolved turbulent fields, indicating a degree of cross-regime generality. Motivated by the structure of the learned feedback, we derive a minimal analytical model that isolates the governing navigation mechanisms and quantitatively reproduces performance trends in both stochastic and turbulent environments. These findings establish active morphological modulation as a robust and physically interpretable control strategy for microswimmer navigation in complex flows, extending the design space beyond conventional motility- and steering-based paradigms [1].

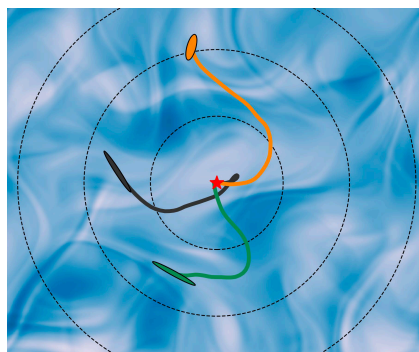


Figure 16: Illustration of shape-changing microswimmers tasked with escaping as far as possible from their initial position while navigating in a turbulent flow.

1. J. Qiu, L. Piro, L. Biferale, M. Cencini, B. Mehlig, and K. Gustavsson. Adaptive shape control for microswimmer navigation in turbulence. arXiv:2603.08201 (2026).

## Two-way regulation of turbulent heat transfer via interpretable bang-bang control discovered by reinforcement learning

Zisong Zhou, Xiaojue Zhu

Max Planck Institute for Solar System Research, 37077 Göttingen, Germany

Robust bidirectional control of turbulent heat transfer remains a fundamental challenge in flow regulation [1]. We investigate this problem in three-dimensional turbulent Rayleigh-Bénard convection at Rayleigh numbers up to 108, using deep reinforcement learning (DRL) coupled with direct numerical simulations. The agent senses near-wall temperature fluctuations at the thermal boundary-layer height and prescribes spatially distributed bottom-wall temperature perturbations in closed loop. The learned policy enables genuine two-way regulation: the Nusselt number can be increased by more than 30% [2] or reduced by over 10% in fully turbulent regimes, outperforming classical sinusoidal heating strategies. Analysis of the trained controller reveals that its behavior effectively reduces to a simple sign-based switching rule. Wall forcing aligned with the sign of the local near-wall temperature fluctuation enhances plume emission and transport, while opposite-sign forcing suppresses convective flux. Based on this insight, we formulate a compact bang-bang-type analytic controller that reproduces the DRL performance with minimal complexity. Our results demonstrate how data-driven reinforcement learning can uncover a universal and interpretable feedback principle for two-way active control of turbulent heat transfer.

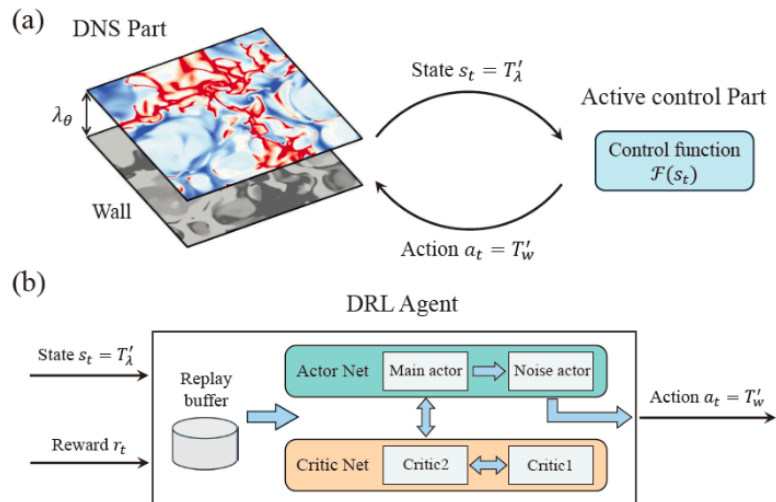


Figure 17: Schematic of the DRL-based active control framework for turbulent convection.

1. Lohse, Detlef, and Ke-Qing Xia. Small-scale properties of turbulent Rayleigh-Bénard convection. *Annual review of fluid mechanics* 42.1 (2010): 335-364.
2. Zhou, Zisong, and Xiaojue Zhu. Deep reinforcement learning control unlocks enhanced heat transfer in turbulent convection. *Proceedings of the National Academy of Sciences* 122.37 (2025):e2506351122.

# Data-driven suppression of naturally developing Tollmien–Schlichting waves using plasma actuation

Babak Mohammadikalakoo<sup>1</sup>, Nguyen Anh Khoa Doan<sup>2</sup>, Mario Kotsonis<sup>1</sup>

<sup>1</sup>Dept. of Flow Physics and Technology, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands

<sup>2</sup>Dept. of Aeronautics, Imperial College London, London, United Kingdom

This work investigates real-time suppression of naturally developing Tollmien-Schlichting (TS) waves in a flat-plate boundary layer using plasma actuation and deep reinforcement learning (DRL) to delay laminar-turbulent transition. The flow is simulated by solving the perturbation form of the incompressible Navier-Stokes equations with a finite-volume method. Naturally developing TS wave packets are generated by upstream broadband forcing and convect downstream toward a plasma actuator modeled through experimentally derived body-force distributions [1]. The controller is formulated as a multi-step DRL policy trained with proximal policy optimization (PPO), and updates the coefficients of a causal finite-impulse-response (FIR) filter in real time to map upstream wall-pressure measurements to actuator voltage, as illustrated in Fig. 18(a). Different sensor layouts are considered to quantify sensing effects. The framework is assessed under reduced observability, actuator saturation, reward-delay mismatch, and frozen-policy deployment to unseen forcing conditions. Results show stable convergence and broadband attenuation, with the best sparse-sensing configuration (DRL4) achieving a 64.5% reduction in downstream RMS wall-pressure fluctuations, corresponding to  $-9.0\text{dB}$  attenuation. As shown in Fig. 18(b), the reduction is localized mainly downstream of the actuator, consistent with convective wave cancellation. Using more than four sensors does not improve performance, indicating that compact sensing is sufficient. A policy trained only under broadband excitation also retains strong zero-shot performance on unseen single- and multi-tone forcings, yielding 53 – 64% RMS reduction in the main test cases. These findings demonstrate that multi-step DRL provides a robust and experimentally relevant approach for suppressing convective instabilities in boundary-layer flows.

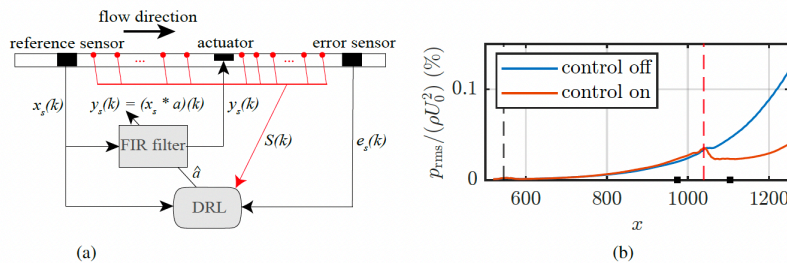


Figure 18: (a) Controller schematic, red dots: state observation sensors, (b) Streamwise distribution of the normalized RMS wall-pressure fluctuations. Reference and error sensor locations (black squares).

1. Kotsonis et al., AIAA Journal 51 (2013)

## Physics-enhanced reinforcement learning for real-time optimal control of dynamical systems

Matteo Tomasetto<sup>1</sup>, Nicolò Botteghi<sup>2</sup>, Gabriele Bruni<sup>2</sup>, Andrea Manzoni<sup>2</sup>

<sup>1</sup>Department of Mechanical Engineering, Politecnico di Milano, Italy

<sup>2</sup>MOX – Department of Mathematics, Politecnico di Milano, Italy

Deep Reinforcement Learning (DRL) [1,2] has recently emerged as a promising feedback control strategy for complex physical systems governed by differential equations. However, DRL is typically sample inefficient as many environment interactions are required to synthesize optimal control strategies, and limited to low-dimensional state and action variables, due to the curse of dimensionality entailed by the exploration and exploitation in high-dimensional spaces. In this work, we present a novel Physics-Enhanced Reinforcement Learning (PEARL) paradigm tailored to the control of (possibly high-dimensional and parametric) dynamical systems. PEARL leverages automatic differentiation-based policy learning algorithms, such as, e.g., Short-Horizon Actor–Critic (SHAC) [3], to guide policy learning through adjoint-based sensitivities, enabling optimal control strategies after a few environment interactions. Through applications on challenging optimal control problems, we show that PEARL (i) can effectively steer differentiable environments, outperforming state-of-the-art DRL algorithms, (ii) is extremely sample efficient, thanks to the physics-guided optimization, (iii) generalizes across multiple scenarios when dealing with parametric systems, and (iv) can cope with high-dimensional problems characterized by large state and action spaces, without requiring low-dimensional state representations or multi-agent strategies.

1. Deep Reinforcement Learning: An Overview, arXiv, arxiv:1701.07274, 2018
2. A survey on physics informed reinforcement learning: Review and open problems, Expert Systems with Applications 287, 128166, 2025
3. Accelerated Policy Learning with Parallel Differentiable Simulation, Proceedings of International Conference on Learning Representations, 2022

# Wednesday 10th – Day 3 - Control theory and model free control

## Compressed Sensing for Launchers and Satellites' Parameter and Fault Estimation

Andrés Marcos, Diego Navarro-Tapia

Dept. of Aerospace Eng., Univ. Carlos III de Madrid (UC3M), Leganés (Madrid), Spain

Compressed Sensing (CS) [1] is a machine learning method that in the past two decades has had a large impact and widespread use in quite a few fields from biomedical to sound/video processing, see Figure 1 and [2]. In the aerospace domain, the lead in the development and application of this method has been made by the fluid dynamics & mechanics community with almost no cross-fertilization with the classical system identification & control communities. In this talk, we present two applications of CS to the estimation of parameters and/or faults in aerospace vehicles, specifically a launcher [3] and a CubeSat [4] (a small satellite with a form factor multiple of  $1U = a$  cube of  $10 \times 10 \times 10$  cms). The results indicate that CS has also the potential to address some of the shortcomings of current estimation and fault identification methods for these systems and that this is a potentially viable venue for collaboration across the aforementioned communities.

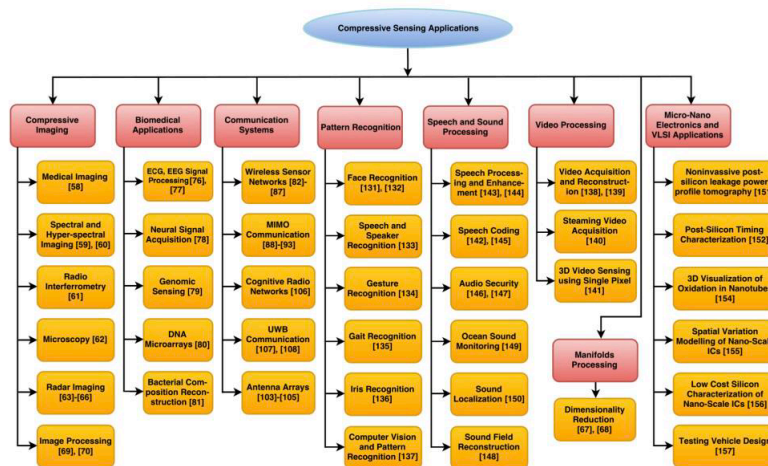


Figure 19: Major applications of CS, source: [2].

1. S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proc. Natl. Acad. Sci. USA 113, 3932 (2016).
2. Rani, M. et al, A Systematic Review of Compressive Sensing: Concepts, Implementations and Applications. IEEE Access, Vol. 6, 2018.
3. D. Navarro-Tapia, A. Marcos., J. Belhadj. Identification of ascent phase launcher dynamics using machine learning. ESA GNC Conference 2023, Sopot, Poland.
4. J. Ilarraz-Zuazo, A. Marcos. Machine Learning for Parameter Estimation in CubeSats: A preliminary study for ST3LLARsat1. CEAS GNC Conference 2026, Madrid, Spain

## Control-then-Reduce approach to optimal control using Dynamical Low-Rank Approximation.

J. Simon Kern, and Jean-Christophe Loiseau

DynFluid, Arts et Métiers Institute of Technology, Paris, France

The Linear Quadratic Regulator (LQR) is a classic benchmark for model-based state-feedback control for continuous-time dynamical systems represented by Linear Time-Invariant (LTI) models including the system dynamics  $A$ , actuators  $B$  to apply the control law and sensors  $C$  to measure the partial system state. Central to both LQR and Kalman filtering is the solution of matrix-valued Algebraic Riccati equations of the general form

$$A^T B + PA - PB^T R^{-1} B P + Q = 0 \quad (0.1)$$

where  $(\cdot)^T$  denotes transposition,  $R$  is the control cost and  $Q$  is the state cost. While the theory of optimal control is well established, the application of optimal control to very high-dimensional problems, such as those encountered in fluid mechanics, is a persistent problem, as the full Riccati-type matrix equations that need to be solved for the computation of the controller gains are intractable. However, in many practical cases the solutions to the Riccati equations are of low rank and there are essentially two alternatives to take advantage of this fact. The first, *Reduce-then-Design*, consists of first constructing a reduced-order model for the system dynamics, either directly using the matrix representation or using low-rank data-driven approaches, and then constructing a controller using standard tools for the small-scale Riccati equations of the reduced problem. The second approach, *Design-then-Reduce*, consists in taking advantage of the solution structure to directly approximate the low-rank algebraic Riccati equations to obtain the low-rank controller. In this work, staying within the *Design-then-Reduce* framework, we take a different route towards a reduced controller considering the differential Riccati equation

$$\dot{P} = A^T B + PA - PB^T R^{-1} B P + Q \quad (0.2)$$

which, for a stabilizable and detectable system, converges to a unique steady state  $P_\infty$  which is the solution of (1). We do not time-integrate  $P$  directly, but instead use a version of Dynamical Low-Rank Approximation (DLRA) [1], which allows for the efficient time-integration of a low-rank SVD-like representation of the time-dependent matrix solution  $P(t)$ . We employ an operator-splitting technique [2] in order to combine this matrix approximation scheme with existing stand-alone integrators for the system dynamics  $A$ . The toolchain based on the in-house package for the numerical linear algebra `LightKrylov` [3] is first developed and tested on the Ginzburg-Landau model equation. Subsequently, the tools are linked to `Nek5000` [4], a high-order spectral-element solver for the incompressible Navier-Stokes equations, to compare the efficacy, accuracy and scalability of the computation of the low-rank controller for high-dimensional flow problems using the standard benchmark case of two-dimensional cylinder flow.

1. O. Koch and C. Lubich. Dynamical Low-Rank Approximation. *SIAM Journal on Matrix Analysis and Applications*, 29(2):434–454, January 2007.
2. C. Lubich and I. V. Oseledets. A projector-splitting integrator for dynamical low-rank approximation. *BIT Numerical Mathematics*, 54(1):171–188, March 2014.
3. J. S. Kern, R. S. Frantz, and J.-Ch. Loiseau. `LightKrylov`: Lightweight implementation of Krylov subspace techniques in modern Fortran. *Journal of Open Source Software*, 11(117):9623, January 2026.
4. P. F. Fischer, J. W. Lottes, and S. G. Kerkemeier. `NEK` – fast high-order scalable CFD, 2008.

## Reinforcement Twinning for Wind Farm Flow Control using Differentiable Dynamic Wake Models and Learned Value Corrections

Lorenzo Schena<sup>1,2</sup>, Benito Dello Iacono<sup>1</sup>, Wim Munters<sup>1</sup>, Timothy Vaerstraten<sup>2</sup>, Jan Helsen<sup>2</sup>, Miguel Alfonso Mendez<sup>1,2,4</sup>

<sup>1</sup>the von Karman Institute (VKI), Waterloosesteenweg 72, 1640 Sint-Genesius-Rode, Belgium

<sup>2</sup>Vrije Universiteit Brussel (VUB), Pleinlaan 2, 1050 Elsene, Brussels

<sup>3</sup>Aero-Thermo-Mechanics Laboratory, Université Libre de Bruxelles, 1050 Elsene, Brussels

<sup>4</sup>Aerospace Engineering Research Group, Universidad Carlos III de Madrid, Leganés, Spain

The optimization of wind farm power production is fundamentally a large-scale flow control problem, governed by complex wake interactions and turbulent atmospheric boundary layer dynamics. Traditional model-based control strategies rely on simplified engineering wake models that suffer from steady-state assumptions or calibration mismatch. Conversely, purely model-free reinforcement learning approaches require prohibitive amounts of plant interaction to learn causal flow physics, even for the simplest tasks [1]. Building upon the concept of Reinforcement Twinning [2], this work presents a hybrid, closed loop control architecture that interleaves model-based physics and data-driven learning across strictly separated time-scales. At its foundation, the framework utilizes sparse, turbine-local measurements to continuously update a dynamic, differentiable partial differential equation (PDE) wake model, tracking the evolving flow state. The action-value landscape is then formalized as a decomposition of a model-based contribution and a learned residual. The physics-based value estimate is obtained by model-based rollouts, while a data-driven network acts as a value corrector, learning the discounted residual gap between the PDE-predicted returns and the actual sampled plant rewards to capture epistemic structural mismatch. The control policy is then optimized via a deterministic policy gradient that explicitly sums these two distinct causal streams: a model-based adjoint sensitivity–backpropagated directly through the PDE’s wake transport characteristics – and the stochastic, model-free residual correction. We validate this Reinforcement Twinning approach online in a high-fidelity environment comprising four IEA-15 MW wind turbines. By coupling large-eddy simulations (LES) with OpenFAST, the setup explicitly resolves turbine aeroelasticity and actuator dynamics. This ensures that the discovered control strategies effectively mitigate wake deficits while remaining physically realizable under realistic operational constraints.

1. Tuhfe Göocmen, Jaime Liew, Elie Kadoche, Nikolay Dimitrov, Riccardo Riva, Søren Juhl Andersen, Alan WH Lio, Julian Quick, Pierre-Elouan Réthoré, and Katherine Dykes. Data-driven wind farm flow control and challenges towards field implementation: A review. *Renewable and Sustainable Energy Reviews*, 216:115605, 2025.
2. Lorenzo Schena, Pedro A Marques, Romain Poletti, Samuel Ahizi, Jan Van den Berghe, and Miguel A Mendez. Reinforcement twinning: From digital twins to model-based reinforcement learning. *Journal of Computational Science*, 82:102421, 2024.

## Multi-Agent Reinforcement Learning for Wind Farm Optimization in Atmospheric Boundary Layer

Felice Manganelli<sup>1</sup>, Onofrio Semeraro<sup>2</sup>, Stefano Leonardi<sup>3</sup>,  
Stefania Cherubini<sup>1</sup>, Pietro De Palma<sup>1</sup>


<sup>1</sup>Dept. of Mechanics, Mathematics and Management, Polytechnic University of Bari, 70126, Bari, Italy

<sup>2</sup>LISN-CNRS, Université Paris-Saclay, Orsay, France


<sup>3</sup> Dept. of Mechanical Engineering, University of Texas at Dallas, Richardson, TX, 75080, USA

This study investigates the application of Multi-Agent Reinforcement Learning (MARL) to wind farm yaw control by benchmarking nine algorithms across two wind farm configurations of increasing complexity: a canonical three-turbine in-line array and the irregular Ablaincourt Energies onshore wind farm ( $M = 7$  turbines). Three reinforcement learning algorithms – Proximal Policy Optimisation (PPO) [1], Soft Actor-Critic (SAC) [2], and Twin Delayed Deep Deterministic Policy Gradient (TD3) [3] – are evaluated under three multi-agent coordination paradigms: independent learning, centralised training with decentralised execution (CTDE), and fully centralised control. Training is performed using the FLORIS engineering wake model, and the resulting control policies are subsequently transferred to a high-fidelity Large-Eddy Simulation (LES) environment based on the UTD-WF solver [4], with reinforcement learning integration enabled through the SmartRedis framework [5]. This setup allows a systematic assessment of policy robustness and sim-to-real transfer fidelity.

In the three-turbine case, all algorithms converge consistently, with the centralised variants of PPO and SAC achieving power gains of approximately 25% relative to the baseline operation. In contrast, the more complex Ablaincourt layout reveals substantial differences in algorithmic robustness: PPO and centralised SAC achieve successful convergence in all 20 training runs, whereas all TD3 variants exhibit recurrent failures associated with ineffective credit assignment in the multi-agent setting. Transfer to LES demonstrates strong agreement with FLORIS predictions for the simpler configuration, with discrepancies below 1.2% for PPO- and SAC-based controllers. For the Ablaincourt wind farm, LES results exceed FLORIS predictions by approximately 10%, highlighting a systematic modelling gap that is attributed to turbulence-enhanced wake recovery effects not captured by the Gaussian wake formulation.



(a) Ablaincourt – LES



(b) Ablaincourt – FLORIS

Figure 20: Hub-height streamwise velocity for the Ablaincourt farm: LES time-averaged (left) vs FLORIS steady-state (right), for the CPPO policy at steady-state yaw.

1. John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
2. Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In International conference on machine learning, pages 1861–1870. PMLR, 2018.
3. S. Fujimoto, H. van Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. In Proceedings of the 35th International Conference on Machine Learning. PMLR, 2018.
4. F. Bernardoni, U. Ciri, M. A. Rotea, and S. Leonardi. Identification of wind turbine clusters for effective real time yaw control optimization. Journal of Renewable and Sustainable Energy, 2021.
5. Sam Partee, Matthew Ellis, Alessandro Rigazzi, Andrew E. Shao, Scott Bachman, Gustavo Marques, and Benjamin Robbins. Using machine learning at scale in numerical simulations with SmartSim: An application to ocean climate modeling. Journal of Computational Science, 62:101707, 2022.

## Learning Observers for Partially Observable Flow Control

Junjie Zhang<sup>1</sup>, Anton Plaksin<sup>1</sup>, Jacob Page<sup>2</sup>, Georgios Rigas<sup>1</sup>

<sup>1</sup>Department of Aeronautics, Imperial College London

<sup>2</sup>Department of Mathematics, University of Edinburgh

Partially observable control problems are common in fluid mechanics, where the measurements available during deployment are often much more limited than those accessible in simulation or controlled experiments. This is especially relevant in active flow control for bluff-body drag reduction, where pressure sensors can be mounted on the body, but informative wake measurements are typically unavailable in practice. In this work, we propose an observer-based reinforcement learning framework that exploits privileged training-time information to improve control under partial observability. Our approach augments an off-policy reinforcement learning controller with an auxiliary observer trained to predict privileged observations from a finite history of deployable sensor measurements. During training, the observer learns to reconstruct the richer measurement set available in simulation, while the controller uses this information to improve policy learning. During deployment, the policy operates using only the history of body-mounted sensor measurements, without access to wake information. Applied to active drag reduction in bluff-body flows, the proposed method improves both learning and final control performance relative to reinforcement learning based only on partial measurements. The framework also enables offline analysis of sensor importance, providing a practical route toward reduced sensing requirements in real-world flow-control applications.

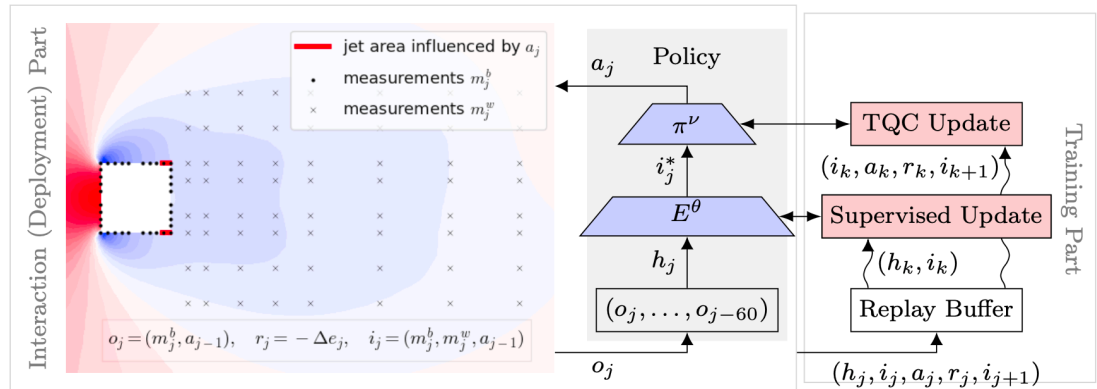


Figure 21: Overview of the proposed observer-based reinforcement learning framework for partially observable flow control. An observer network learns to predict privileged training-time observations from a history of deployable sensor measurements. The predicted information is used during training to improve the reinforcement learning policy, while deployment relies only on the partial sensor history.

## Learning Observers for Partially Observable Flow Control

Defne E. Ozan<sup>1</sup>, Andrea Nóvoa<sup>1</sup>, Georgios Rigas<sup>1</sup> and Luca Magri<sup>1,2</sup>

<sup>1</sup>Department of Aeronautics, Imperial College London, London SW7 2AZ, UK

<sup>2</sup>DIMEAS, Politecnico di Torino, Torino 24 10129, Italy

In this work based on [1], we propose a framework that tackles partial observability in the real-time active feedback control of spatio-temporally chaotic systems, which is challenging because of high dimensionality and unpredictability. The data-assimilated model-informed reinforcement learning (DA-MIRL) transforms the partially observable Markov decision process [2] into a fully observed process by integrating: (i) a predictive, approximate model of the system's high-dimensional dynamics; (ii) sequential data assimilation to correct the model's state in real-time using the available observations (here, the ensemble Kalman Filter [3]); and (iii) an off-policy actor-critic reinforcement learning (RL) algorithm to learn an optimal control strategy based on the corrected state estimates, (here, the Deep Deterministic Policy Gradient). We test the DA-MIRL on the Kuramoto-Sivashinsky (KS) equation with the goal of stabilizing the system's chaotic dynamics, similarly to [4]. We employ two different models for the environment: (i) a physics-based coarse-grained model; and (ii) a fully data-driven model, the control-aware echo state network, which is designed here for computational efficiency. The DA-MIRL learns control policies with 60% fewer sensors than required by model-free RL under noisy and infrequent observation conditions (Figure 22) and it is robust across different choices of models and chaotic regimes. This work opens opportunities for the RL-based control of partially observable systems.

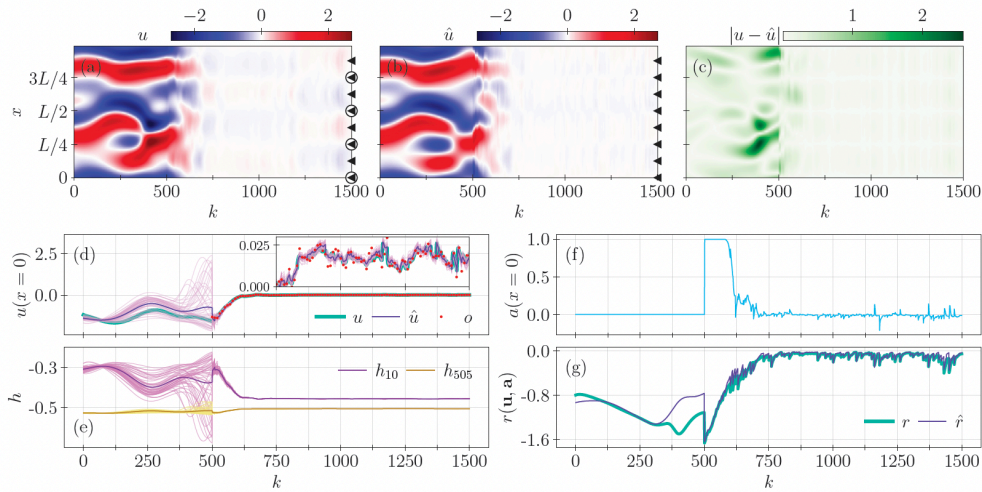


Figure 22: DA-MIRL using a CA-ESN as the predictive model successfully stabilizes the chaotic dynamics of the KS equation when only four sensor measurements are available. The ensemble prediction,  $\hat{u}$ , synchronizes with the true system,  $u$ , once the observations become available ( $k = 500$ ).

1. D. E. Ozan, A. Nóvoa, G. Rigas, and L. Magri. Data-assimilated model-informed reinforcement learning. Proc. A, 481(2327):20250476, 2025.

2. Vikram Krishnamurthy. Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing. Cambridge University Press, 1 edition, 2016.
3. Geir Evensen. Data Assimilation: The Ensemble Kalman Filter. Springer, Berlin, Heidelberg, 2009.
4. M. A. Bucci, O. Semeraro, A. Allauzen, G. Wisniewski, L. Cordier, and L. Mathelin. Control of chaotic systems by deep reinforcement learning. Proc. A, 475(2231):20190351, 2019.

## Bayesian Reinforcement Twinning: A Multi-Fidelity Framework for Reciprocal Learning Between Digital Twins and Control

Yannick Lecomte<sup>1,2</sup>, Alessandro Grava<sup>1</sup>, Emanuele Garone<sup>2</sup>, Miguel A. Mendez<sup>1,3,4</sup>

<sup>1</sup>Environmental and Applied Fluid Dynamics, von Karman Institute for Fluid Dynamics, Belgium

<sup>2</sup>Dept. of Control Engineering and System Analysis, Université Libre de Bruxelles, Belgium

<sup>3</sup>Aero-Thermo-Mechanics Laboratory, Université Libre de Bruxelles, Belgium

<sup>4</sup>Experimental Aerodynamics and Propulsion Lab, Universidad Carlos III de Madrid, Spain

Unmanned Aerial Vehicles (UAVs) are increasingly employed in applications requiring high levels of autonomy and reliability. However, their operation often involves strongly nonlinear dynamics, actuator couplings, and environmental disturbances that challenge the design of robust control policies. Purely model-based strategies can suffer from modeling inaccuracies, while model-free learning approaches typically require large amounts of experimental data and may explore unsafe or inefficient regions of the control space. In this work, we propose a framework for Bayesian Reinforcement Twinning, a hybrid methodology that combines bidirectional links between digital twin modeling and model-free control. The approach leverages the digital twin, consider as low-fidelity, to guide the exploration of the model-free policy search, consider as high-fidelity, within a multi-fidelity Gaussian framework. This effectively constrains the learning process toward promising regions improving data efficiency while avoiding unsafe exploration. In contrast to standard multi-fidelity approaches, the proposed framework introduces a reciprocal interaction between the physical system and its digital counterpart. Specifically, observations collected from the real system are used to update the digital twin online. This bidirectional exchange enables continuous refinement of the digital twin across the policy search space, leading to progressively improved predictions and more informed exploration–exploitation trade-offs. A demonstration of the proposed approach is carried out on a dedicated test bench replicating the attitude dynamics of a UAV (see Fig. 23). This setup features two propellers mounted on the ends of a beam pivoting on a central fulcrum, targeting a set point in terms of attitude, using the motor torque as control actuation and relying only on measurements of the attitude angle.




Figure 23: Picture of the balance test bench at the von Karman Institute.

## Active Flow Control via Model-Based Reinforcement Learning

Guillermo Suarez, Emre Özkaya, Nico R. Gauger

RPTU University Kaiserslautern-Landau

In this work we investigate reinforcement learning (RL) for active flow control (AFC) on the Schäfer and Turek benchmark [1]: a 2D cylinder in a channel flow equipped with two synthetic jets operating under zero-net-mass-flux actuation. This configuration provides a well established test case for drag and vortex-shedding mitigation while retaining moderate computational cost for computational fluid dynamics (CFD). Reinforcement learning is particularly attractive for AFC because it formulates the problem as a global optimization task and does not require gradient information from the Navier-Stokes solver. However, RL relies on extensive interaction with the environment. When the environment is a CFD solver, this results in an excessive computational time, making direct training impractical for realistic flow configurations. To address this limitation, we explore replacing the CFD environment with surrogate models that approximate the flow dynamics. Several data-driven approaches are investigated, including Dynamic Mode Decomposition, and Long Short-Term Memory networks. While these methods can capture certain flow features, they are not robust enough to represent the nonlinear, transient dynamics, or they require a considerable amount of data for training. We therefore investigate the model-based meta-policy optimization (MB-MPO) framework [2]. Instead of relying on a single surrogate, it trains an ensemble of surrogate models to represent uncertainties in the learned dynamics. A meta-policy is then optimized across this ensemble, improving robustness and generalization before being deployed in the real CFD environment. The proposed approach aims to significantly reduce training cost while preserving a good performance in the real CFD environment. We assess the performance of MB-MPO and compare it against model-free algorithms such as Proximal Policy Optimization and Soft Actor-Critic. Through these comparisons, we analyze the effectiveness of model-based approaches for active flow control and discuss under which conditions surrogate-based reinforcement learning can provide a practical advantage over model-free methods.

## Deep Reinforcement Learning for Combined Shape and Tangential Blowing Optimization with Wind Tunnel Validation

Piergiorgio Scavella, Fabiana Ruggiano, Gerardo Paolillo, Tommaso Astarita, Gennaro Cardone, Carlo Salvatore Greco

University of Naples Federico II, Italy

Recent advances in data-driven control have demonstrated the effectiveness of Deep Reinforcement Learning (DRL) for aerodynamic optimization [1]. In previous work [2], DRL was applied to airfoil shape design, achieving performance comparable to state-of-the-art gradient-free methods and converging toward classical high-lift configurations. Here, the framework is extended to a coupled optimization of airfoil geometry and tangential jet actuation. The agent simultaneously adjusts B-spline control points and jet parameters to maximize lift at low angle of attack. The key contribution is the experimental validation of the DRL-optimized configuration. Wind tunnel tests combine wall-pressure measurements, quantitative infrared thermography (IRT) and Particle Image Velocimetry (PIV). Pressure data enable the reconstruction of the  $C_p$  distribution and the estimation of lift, while IRT allows the identification of transition and separation, revealing an extended laminar region upstream of the jet in the controlled case. PIV measurements show suppression of the separated shear layer and flow reattachment when tangential blowing is activated, confirming effective separation control (Figure 24). Results coming from pressure matching with numerical results and PIV flow field analysis confirm that the DRL-driven integrated optimization framework yields measurable aerodynamic improvements while preserving physical consistency between numerical and experimental flow features.

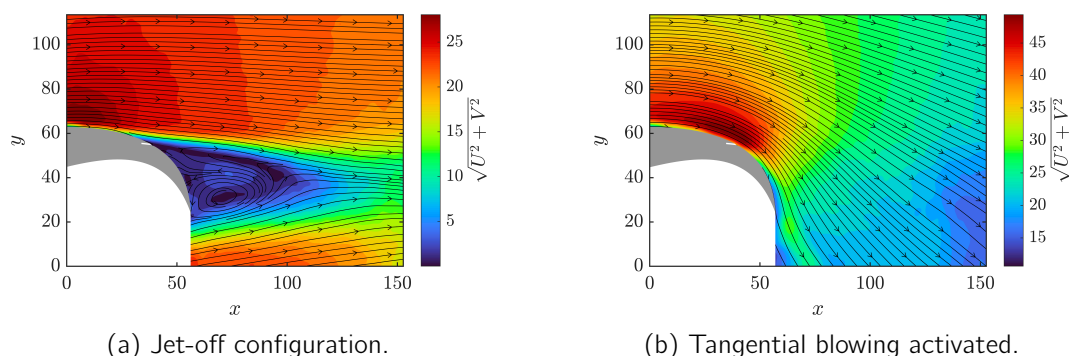


Figure 24: Parametric SSM predictions for the lid-driven cavity flow at an unseen Reynolds number.

1. Jonathan Viquerat, Jean Rabault, Alexander Kuhnle, Hassan Ghraieb, Aurélien Larcher, and Elie Hachem. Direct shape optimization through deep reinforcement learning. *Journal of Computational Physics*, 428:110080, 2021.
2. P Scavella, G Paolillo, and CS Greco. Deep reinforcement learning-based airfoil design and optimization: An aerodynamic analysis. *Aerospace Science and Technology*, page 110638, 2025.

## **Coupling Reinforcement Learning and CFD to support decision making**

Elie Hachem

Mines Paris – PSL, France

This talk presents recent work combining computational fluid dynamics (CFD), advanced numerical methods, and reinforcement learning for flow and energy related problems. The focus is on data-driven optimization strategies based on high-fidelity simulations, stabilized finite element methods, and adaptive anisotropic meshing. Several applications will be discussed, including reinforcement-learning approaches for flow and thermal systems. In particular, recent work will be presented on the optimisation of industrial furnaces using reinforcement learning coupled with CFD-based simulations, as well as a framework for the optimization of photovoltaic panel configurations under strong wind conditions. In the latter case, the control problem is formulated as a reinforcement learning task in which the agent interacts with a CFD environment to identify panel arrangements that reduce aerodynamic loads and flow-induced fluctuations. The presentation will discuss the numerical framework, optimization methodology, and perspectives for data-driven control of complex fluid and thermal systems.

## Useful info

The workshop will be hosted at the Conservatoire national des arts et métiers (CNAM), located in one of the last medieval architectural area of Paris, Le Marais, in the 3rd arrondissement of Paris, in the buildings of the former Benedictine priory of Saint-Martin-des-Champs.

**Access.** The nearest metro stations are Réaumur–Sébastopol (lines 3 and 4) and Arts et Métiers (lines 3 and 11). Bus services 20, 38, and 47 also stop nearby.

**Amphitheatre.** The colloquium will take place in the Amphi Gaston Planté, located at 2 Rue Conté (Gate 35, 1st floor).

**Registration.** Registration will open at 08.30 on Monday, June 8, outside the amphitheatre.

**Coffee breaks.** Coffee and refreshments will be served outside the amphitheatre.

**Lunches.** Lunch will be provided at the CNAM restaurant (Gates 31–33 on the map).

**CNAM Museum.** A free visit is offered at the CNAM museum starting at 16.00 on June 9. The museum closes at 18.00.

**Social dinner.** The social dinner will take place at Café Léonard, 57 Rue de Turbigo, on June 9, starting at 19:30.

## Plan du centre Cnam Paris





