Session 7

Reinforcement Learning and Flow Control

S7-A4	Isabel Scherl — Towards Developing Reduced-Order Models for Active Flow Control from Experimental Data
S7-A7	Miguel Alfonso Mendez — Reinforcement Twinning for Control of Transport-Dominated
S7-P2	Valerio Di Domenico — Development of Reduced Order Models for the Thermal Hydraulics of Nuclear Reactors
S7-P5	Priyam Chakraborty — Learning to Swim Upstream: Koopman-Physics AI for Smart Microswimmer Navigation

Towards Developing Reduced-Order Models for Active Flow Control from Experimental Data

[1]Isabel Scherl*; [1]James Bercaw; [2]Alec J. Linot

[1]Department of Mechanical and Industrial Engineering, University of Massachusetts, Amherst, MA USA

[2] Department of Chemical and Biomolecular Engineering, University of Massachusetts, Amherst, MA USA

Abstract

In nearly all practical aerodynamic applications, unsteadiness causes deviations from nominal behavior which requires actuation to correct. These turbulent, unsteady flows are complex and high-dimensional, though reduced order models (ROMs) can still represent these systems because the dissipative nature of the Navier-Stokes Equations causes turbulent trajectories to lie in a low-dimensional subspace. Machine learning methods have been used to build these ROMs. These models can successfully capture the dominant patterns and provide a good approximation of the flow physics while still preserving the unsteady nature of the problem. The primary goal of developing these ROMs is to integrate them into active flow control environments because algorithms for identifying effective active, closed-loop control policies require fast, accurate models. Much of the work to date has developed these models using simulation data. In this work, we will primarily use experimentally acquired data. Specifically, flow measurements collected from a pitching foil experiment will be used to build the reduced-order model using manifold learning methods. These methods use autoencoders to identify the low-dimensional space that describes the system dynamics. The measured force coefficients will be mapped to a model built from flow field measurements. This will result in fast, accurate ROMs that can be used to rapidly forecast the system or to discover control strategies.

Keywords: reduced-order modeling, flow control

Reinforcement Twinning for Control of Transport-Dominated Systems with Partial Observations

Miguel Alfonso Mendez[1,2,3] Lorenzo Schena[1,4] Romain Poletti[1]

[1]*von Karman Institute for Fluid Dynamics, Waterloosesteenweg 72, Sint-Genesius-Rode, Belgium*

[2]*Université Libre de Bruxelles (ULB), Avenue Franklin Roosevelt 50, 1050 Brussels, Belgium*

[3]*Vrije Universiteit Brussel (VUB), Pleinlaan 2,1050 Brussels, Belgium*

[4]*Experimental Aerodynamics and Propulsion Lab, Universidad Carlos III de Madrid, Leganés, 28911, Spain*

Abstract

This work introduces a reinforcement twinning framework [1,2] for the control of systems governed by nonlinear partial differential equations under partial observability and delayed actuation. The approach combines a model-based optimal control loop, relying on a differentiable surrogate of the governing dynamics, with a model-free reinforcement learning component trained from interaction data. The two components are coupled through shared trajectories and consistency measures, enabling reciprocal improvement of both the policy and the digital twin.

The methodology is assessed on a distributed control problem based on a convective Burgers equation with localized actuators, capturing key features of transport-dominated systems such as advection-induced delays, non-local coupling, actuator dynamics, and sparse sensing. The resulting control problem is intrinsically non-Markovian from the perspective of the available observations, requiring the policy to exploit temporal structure and anticipate downstream effects.

Results show that the learned policies outperform steady and quasi-steady strategies by coordinating actuation in time, reducing spatial imbalance, and improving overall power extraction. These findings highlight the potential of reinforcement twinning as a general framework for real-time control of complex physical systems with limited sensing and modeling uncertainty.

Keywords:

Reinforcement Twinning, PDE-constrained control, Partial observability, Hybrid control, Digital twin, Model-based Reinforcement Learning

References:

[1] Schena, L., Marques, P. A., Poletti, R., Ahizi, S., Van den Berghe, J., & Mendez, M. A. (2024). Reinforcement twinning: From digital twins to model-based reinforcement learning. Journal of Computational Science, 82, 102421.

[2] Poletti, R. Schena, L., Koloszar, L., Degroote, J., & Mendez, M. A. (2025). Reinforcement twinning for hybrid control of flapping-wing drones., arXiv:2505.18201.

Development of Reduced Order Models for the Thermal Hydraulics of Nuclear Reactors

[1,2]V. Di Domenico*; [1]M. A. Mendez; [1]L. Koloszar; [2]Y.Bartosiewicz; [1]M. Fiore

[1]*Environmental and Applied Fluid Dynamics Department, von Karman Institute for Fluid Dynamics, Sint-Genesius-Rode, Belgium;*

[2]*Institute of Mechanics, Materials and Civil Engineering (iMMC), University of Louvain, Louvain-la-Neuve, Belgium*

Keywords: Reduced Order Model, Proper Orthogonal Decomposition, Galerkin Projection, Nuclear Reactors

Abstract. Real-time modelling and forecasting of new-generation nuclear reactors are a key enabler for the reduction of operational costs and increased safety. However, current thermal-hydraulic codes either lack the necessary physics to model flow and thermal stratification in coolant pools or are too computationally expensive. In this work, we propose a computationally efficient framework based on Proper Orthogonal Decomposition and Galerkin projection to model parametric, turbulent, unsteady flows with buoyancy effects. Numerical results show promising behavior and decreased computational cost compared to traditional methods, as the model is able to qualitatively predict typical natural and forced convection patterns.

1. Introduction

Pool-type reactors are a cornerstone of the upcoming Generation IV of nuclear reactors and are characterized by large volumes of coolant within a single vessel, where the primary components are immersed. From a thermal-hydraulic perspective, the coolant can undergo more complex phenomena compared to the traditional loop-type ones. Large, three-dimensional volumes of coolant can in fact experience natural, mixed, and forced convection regimes under nominal and accidental conditions.

Experimental facilities study both fundamental coolant physics and design-specific behaviour. A key facility is MYRRHABELLE [1], a water-based 1:5 scale model of the MYRRHA reactor, which allows Particle Image Velocimetry measurements as well as temperature measurements at different locations.

In addition to experimental facilities, numerical models are crucial for the design and safety assessment of nuclear reactors. The main numerical tools for nuclear thermal-hydraulic analysis are system codes, which rely on 0D or 1D nodalization of the fluid domain. This approach is accurate for loop-type reactors; however, both the 0D/1D approaches and the multidimensional methods such as RELAP5-3D have proven to be less accurate in the prediction of stratified and strongly three-dimensional flows [2].

3D Computational Fluid Dynamics (CFD) models are more accurate for complex geometries and flow fields. However, their computational cost is not yet compatible with the requirements for safety assessment, uncertainty quantification, and real-time applications. System codes and CFD can be coupled together, with the system code providing the boundary conditions for the CFD model [3].

Reduced Order Modelling (ROM) offers an alternative approach: using Reduced Basis (RB) projection to build low-cost, low-dimensional representations of complex physical systems [4]. One technique for generating such reduced bases is the Proper Orthogonal Decomposition (POD) [5], [6]. In the context of reduced order models, the optimality and orthogonality properties of the POD have made it one of the most used techniques for the generation of the basis. An in-depth analysis of the creation of low-dimensional Galerkin models based on POD can be found in [7].

In nuclear applications, ROM and POD methods are gaining attention both for the modelling of thermal hydraulics [8] and neutronics [9] problems. For thermal hydraulics, ROMs must capture unsteadiness, turbulence, and buoyancy effects while maintaining numerical stability. Fundamental works in this direction are [10], [11], [12], [13], [14], where the authors tackled increasingly complex flows with one or more of the aforementioned requirements; however, to the authors' knowledge, no single, unified model has been produced to treat all of them simultaneously.

In this work, we develop a POD-Galerkin framework for the simulation of parametric, unsteady, turbulent flows with buoyancy effects, with the ultimate goal of developing a Reduced Order Model for the hot and cold pools of the MYRRHABELLE experimental facility.

2. Full Order Model

The full order model (FOM) for unsteady, turbulent, incompressible flows with energy transport is:

$$\begin{cases} \nabla \cdot \mathbf{u} = 0\,, \\ \dfrac{\partial \boldsymbol{u}}{\partial t} + \nabla \cdot (\boldsymbol{u} \otimes \boldsymbol{u}) = -\nabla p_{rgh} - \boldsymbol{g}\beta(\theta - \theta_{ref}) + \nabla \cdot ((\nu + \nu_t)(\nabla \mathbf{u} + (\nabla \mathbf{u}^{\boldsymbol{T}}))), \\ \dfrac{\partial \theta}{\partial t} + \nabla \cdot (\boldsymbol{u}\theta) = \nabla \cdot ((\alpha + \alpha_t)\nabla\theta). \end{cases} \tag{2.1}$$

Here, $\mathbf{u}, p_{rgh}$, and θ are the averaged fields of velocity, shifted kinematic pressure ($p_{rgh} = p - \mathbf{g} \cdot \mathbf{r}$), and temperature. ν_t and α_t are the turbulent viscosity and turbulent thermal diffusivity, respectively. Equations are solved in a domain Ω using finite volume discretization with the unsteady Reynolds-averaged Navier-Stokes (URANS) approach.

The turbulence model chosen to represent the Reynolds stresses is the $k - \omega\ SST$ model, while the turbulent heat fluxes are modelled with a Simple Gradient Hypothesis and a constant turbulent Prandtl number Pr_t:

$$-\overline{\boldsymbol{u}'\boldsymbol{u}'} = 2\nu_t\big(\boldsymbol{\nabla u} + (\boldsymbol{\nabla u})^{\boldsymbol{T}}\big)/2 - \frac{2}{3}k\boldsymbol{I}, \quad \overline{\boldsymbol{u}'\theta'} = -\alpha_t\nabla\theta = -\frac{\nu_t}{Pr_t}\nabla\theta. \tag{2.2}$$

The $k - \omega\ SST$ model augments the full order model with two additional transport equations governing k (turbulent kinetic energy) and ω (specific dissipation rate) respectively, which are here omitted for brevity.

3. Reduced Order Model

We propose a reduced order model that extends and integrates the approaches in [11],[12], [14].

3.1 POD-Projection based Reduced Order Modelling

Using typical reduced basis formalism, an approximation of rank N_i of each flow quantity can be expressed as a linear combination of time-dependent coefficients and spatially varying bases:

$$\begin{aligned} &\boldsymbol{u}(\boldsymbol{x}, t) \approx \boldsymbol{u}_r(\boldsymbol{x}, t) = \sum_{i=i}^{N_u} a_i(\mathrm{t})\phi_{\mathrm{i}}(\boldsymbol{x}), && p_{rgh}(\boldsymbol{x}, t) \approx \mathrm{p}_{\mathrm{rgh_r}}(\mathbf{x}, \mathrm{t}) = \sum_{\mathrm{i}=1}^{\mathrm{N_p}} b_i(t)\chi_i(\boldsymbol{x}), \\ &\theta(\boldsymbol{x}, t) \approx \theta_r(\boldsymbol{x}, t) = \sum_{i=1}^{N_\theta} c_i(t)\psi_i(\boldsymbol{x}), && \nu_t(\boldsymbol{x}, t) \approx \nu_{t_r}(\boldsymbol{x}, t) = \sum_{i=1}^{N_{\nu_t}} d_i(t)\xi_i(\boldsymbol{x})\,, \end{aligned} \tag{3.1}$$

Where the Greek letter terms represent the spatial bases and the Latin letter terms represent the temporal coefficients. In this work, the spatial bases (or modes) and their temporal evolution are obtained from precomputed full order model solutions using the POD via the snapshot method [6].

To obtain a reduced order model, the linear combinations in equation (3.1) are substituted in the system of equations (2.1), which is projected onto the POD subspace spanned by its bases. In this work, the projection operator is taken to be the L^2 orthogonal projection, denoted by $\langle \cdot,\cdot \rangle_{L^2(\Omega)}$. The projection operation, coupled with the expression of variables as their POD expansion (equation (3.1)) transforms the original set of PDEs into ODEs whose unknowns are the temporal coefficients of the POD modes:

$$\begin{cases} \boldsymbol{M\dot{a}} = \nu(\boldsymbol{B} + \boldsymbol{B_T})\boldsymbol{a} - \boldsymbol{a}^{\boldsymbol{T}}\boldsymbol{Ca} + \boldsymbol{d}^{\mathrm{T}}\boldsymbol{C}_T\boldsymbol{a} - \boldsymbol{Kb} - \boldsymbol{Hc}, \\ \boldsymbol{W\dot{c}} = \alpha\boldsymbol{Yc} - \boldsymbol{a}^T\boldsymbol{Qc} + \dfrac{\boldsymbol{d}^T}{Pr_t}\boldsymbol{Y_T a}\,. \end{cases} \tag{3.2}$$

For more details on the general derivation and composition of the terms in equation (3.2), the reader is referred to [12],[14],[15],[16]. As an example, the buoyancy related term $\boldsymbol{H}$ in equation (3.2) is a 2D matrix defined as:

$$H_{i,j} = \langle \phi_i, (\boldsymbol{g} \cdot \boldsymbol{r})\, \boldsymbol{\nabla}(1 - \beta(\psi_i - \theta_{ref})) \rangle_{L^2(\Omega)}, \tag{3.3}$$

The total number of unknowns N is given by the sum of the number of modes retained for each quantity: $N = N_u + N_p + N_\theta + N_{\nu_t}$. The system of equations (3.2) only provides $N_u + N_\theta$, and thus $N_p + N_{\nu_t}$ additional equations are required to compute the coefficients $\boldsymbol{b}$ and $\boldsymbol{d}$.

The relevance of the pressure term in Galerkin-based POD models has been shown in [17]. Among the different strategies to include pressure contributions [13],[18],[19], the projection of a pressure Poisson equation (PPE) is adopted in this work as it directly replicates the full-order model. This leads to an additional set of N_p equations for the pressure coefficients $\boldsymbol{b}$ which are added to the system of equations (3.2).

To add N_{ν_t} equations for the $\boldsymbol{d}$ coefficients, a Radial Basis Function (RBF) treatment based on the one proposed in [14] is adopted. Given a vector of parameters $\boldsymbol{\mu}$, which in this work are the parametric boundary conditions, and M different solutions of the full order model, each with a different set of parametric boundary conditions, the turbulent viscosity field is decomposed as:

$$\nu_t(\boldsymbol{x}, t; \boldsymbol{\mu}) \approx \sum_{j=1}^{M} \overline{d_j}(\boldsymbol{\mu})\, \overline{\xi_j}(\boldsymbol{x}) + \sum_{i=1}^{N\nu_t} d_i(t) \xi_i(\boldsymbol{x}), \tag{3.4}$$

where the fields $\left[\overline{\xi_j}\right]_{j=1}^{M}$ are given by the time averaged ν_t fields of the M different full order solutions, and $[\xi_i]_{i=1}^{N_{\nu_t}}$ are the POD modes of the fluctuating component. The $\overline{d_j}$ coefficients are computed by linear interpolation during the solution of the ROM, as they only depend on the parameters $\boldsymbol{\mu}$. The d_i time dependent coefficients are computed using RBF interpolation, denoted as $f(\cdot)$, based on the assumption that the coefficients d_i can be approximated as $d_i \approx f_i(\boldsymbol{a}(t))$.
Finally, introducing both the decomposition of ν_t given by equation (3.4) and the PPE-based closure for the pressure field in the Galerkin projection, the final system of ODEs for the ROM is:

$$\begin{cases} \boldsymbol{M}\dot{\boldsymbol{a}} = \nu(\boldsymbol{B} + \boldsymbol{B_T})\boldsymbol{a} - \boldsymbol{a^T}\boldsymbol{C}\boldsymbol{a} + \boldsymbol{d}^{\mathrm{T}}\boldsymbol{C}_T\boldsymbol{a} + \overline{\boldsymbol{d}}^{\mathrm{T}}\overline{\boldsymbol{C}_T}\boldsymbol{a} - \boldsymbol{K}\boldsymbol{b} - \boldsymbol{H}\boldsymbol{c}, \\ \boldsymbol{W}\dot{\boldsymbol{c}} = \alpha\boldsymbol{Y}\boldsymbol{c} - \boldsymbol{a}^T\boldsymbol{Q}\boldsymbol{c} + \dfrac{\boldsymbol{d}^T}{Pr_t}\boldsymbol{Y_T}\boldsymbol{a} + \dfrac{\overline{\boldsymbol{d}}^T}{Pr_t}\overline{\boldsymbol{Y}}_{\boldsymbol{T}}\boldsymbol{a}\,, \\ \boldsymbol{D}\boldsymbol{b} - \nu\boldsymbol{N}\boldsymbol{a} + \boldsymbol{L}\dot{\boldsymbol{a}} + \boldsymbol{a^T}\boldsymbol{G}\boldsymbol{a} - \boldsymbol{d}^{\mathrm{T}}(\boldsymbol{C}_{T3} + \boldsymbol{C}_{T4})\boldsymbol{a} - \overline{\boldsymbol{d}}^{\mathrm{T}}(\overline{\boldsymbol{C}_{T3}} + \overline{\boldsymbol{C}_{T4}})\boldsymbol{a} + \boldsymbol{H}_{PPE}\boldsymbol{c} = \boldsymbol{0}, \end{cases} \tag{3.5}$$

Note that due to the splitting made in equation (3.4), new terms appear in the ROM equations, namely the overline terms. For a detailed derivation of the projected pressure equation, the reader is referred to [20].

To impose parametric boundary conditions, the two most common strategies are penalty methods [21] and the *lifting function* method [15]. However, in this work, a third approach is adopted, originally proposed in [22], which is based on the creation of a linear combination of the POD modes that vanishes on the boundary. Unlike the penalty method, which requires a penalty parameter, or the lifting function method, which requires a non-trivial choice of lifting function for the temperature field, this approach requires no input from the user, as the linear combination is automatically computed through the complete QR factorization of the original POD basis.

4. Numerical results

To test the proposed framework, the numerical test case illustrated in Figure 1 was created. The cavity flow is governed by the two parameters U_{in}, θ_{in}, the inlet velocity and inlet temperature, respectively, which directly determine the Richardson number $Ri = \frac{g\beta L(\theta_0 - \theta_{in})}{U^2}$. The initial conditions are represented by an established isothermal flow pattern with $U_0 = 0.6$ m/s, $\theta_0 = 300$K. Six full order cases for different combinations of (U_{in}, θ_{in}) were run from this initial condition ($M = 6$) until either thermal stratification or mixing is established, exploring the range of $Ri = \{0.21, 0.43, 0.49, 0.87, 0.98, 1.74\}$. A total of 2720 snapshots were collected and used for the POD and the Galerkin projection. The ROM was implemented in C++ using the ITHACA-FV library [15],[16] and run in serial on an AMD EPYC 7551 32-Core Processor; the online solution of the ROM was performed with a fixed time step of 0.01 seconds.

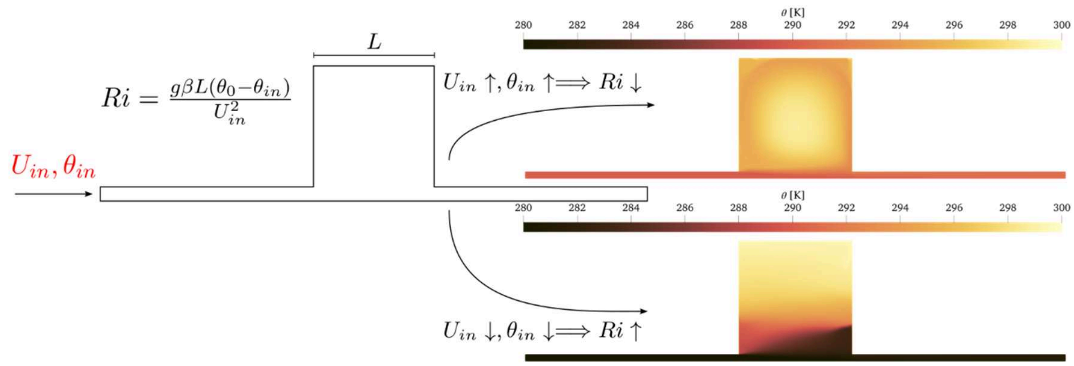

Figure 1. Numerical test case. The red quantities are controllable parameters that determine whether the flow reaches thermal mixing (top, low Ri numbers) or thermal stratification (bottom, higher Ri numbers).

Temperature profiles along the vertical midline of the cavity are reported in Figure 2, at the latest time step of each of the numerical simulations. Results shown were obtained using $N_u = 15, N_{p_{rgh}} = N_{\nu_t} = 5, N_\theta = 20$. The computational time saving in the online phase with this number of modes was $\sim \mathcal{O}(10)$, with each simulation requiring ≈ 0.25 seconds to simulate 1 second of physical time.

In all cases, the POD model is able to qualitatively predict the temperature profile and the onset of stratification or mixing, depending on the parametric boundary conditions, while remaining numerically stable during the run. Depending on the number of retained modes, the accuracy may vary, but in general the relative error for the velocity field is higher during the initial transient phase, and then decreases as the flow approaches a stable state. This phenomenon is likely due to the complexity of the velocity field, compounded by the influence of θ and ν_t in the momentum equation, and by imperfect RBF interpolation procedure. Figure 3 shows the L^2 relative error for velocity and temperature for two cases with $Ri = 0.87$ and $Ri = 1.74$, and in both cases the velocity error decreases as the flow approaches the final state.

In the current implementation, the RBF shape and smoothness parameters are set equal for all ν_t modes and are selected by trial and error. A further improvement would be to assign different parameters to each mode and tune them using an optimization procedure such as leave-one-out cross-validation.

Increasing the number of velocity modes N_u reduces the error but also causes an increase in computational cost and memory usage, particularly due to the nonlinear terms in equations (3.5), and must be balanced with the requirements of the specific application.

5. Conclusion

The proposed framework is able to qualitatively predict the evolution of the system towards stratification dominated or mixing dominated conditions depending on the parametric boundary conditions; moreover, depending on the number of modes retained, there is considerable ($\mathcal{O}(10)$ to $\mathcal{O}(10^3)$) computational time reduction, enabling faster-than-real-time usage.

To move towards real-life applications in nuclear reactors, the model must incorporate energy sinks and sources such as heat exchangers and heating elements. In addition, the application of this kind of ROM to real-life scenarios hinges on the ability to couple them with different codes that simulate other components of the reactor; therefore, this is a path that must be explored further [23].

In a Digital Twinning context, methods to combine the ROM prediction with real-time sensor data must be developed to guarantee continuous cooperation between the real and digital environments.

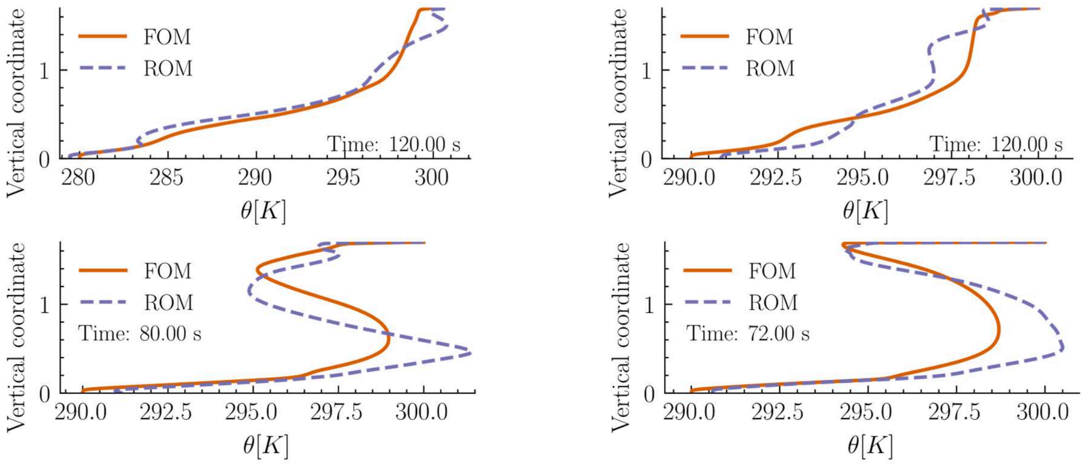

Figure 2. Temperature profiles measured along the vertical centreline of the cavity at the end of each simulation. From the top left, in clockwise order: Ri=1.74, Ri=0.87, Ri=0.49, Ri=0.21

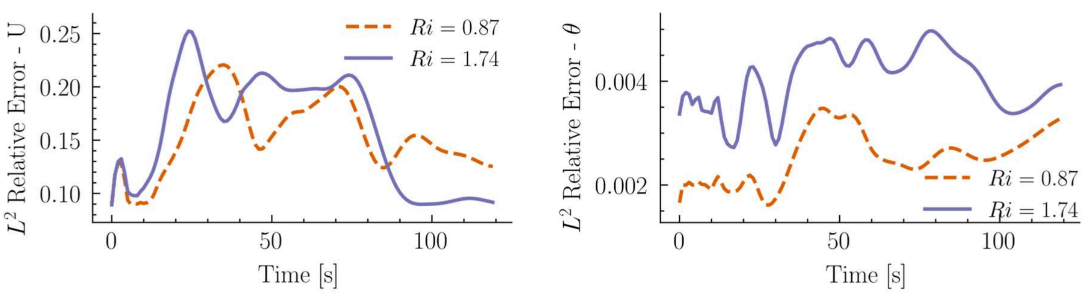

Figure 3. L^2 relative errors for velocity (left) and temperature (right)

6. References

[1] P. Planquart and K. Van Tichelen, “Experimental investigation of accidental scenarios using a scale water model of a HLM reactor,” Nuclear Engineering and Design, vol. 346, pp. 10–16, May 2019. doi: 10.1016/j.nucengdes.2019.02.016

[2] V. Narcisi, F. Giannetti, and G. Caruso, “Investigation on RELAP5-3D© capability to predict thermal stratification in liquid metal pool-type system and comparison with experimental data,” Nuclear Engineering and Design, vol. 352, p. 110 152, Oct. 2019. doi: 10.1016/j.nucengdes.2019.110152

[3] P. Andrea, G. Francesco, M. Marigrazia, M. Daniele, and F. Nicola, “STH/CFD Coupled Simulation of the Protected Loss of Flow Accident in the CIRCE-HERO Facility,” en, Applied Sciences, vol. 10, no. 20, p. 7032, Jan. 2020. doi: 10.3390/app10207032

[4] A. Quarteroni, A. Manzoni, and F. Negri, Reduced Basis Methods for Partial Differential Equations (UNITEXT). Cham: Springer International Publishing, 2016, vol. 92. doi: 10.1007/978-3-319-15431-2

[5] J. L. Lumley, “Computational Modeling of Turbulent Flows,” en, in Advances in Applied Mechanics, vol. 18, Elsevier, 1979, pp. 123–176. doi: 10.1016/S0065-2156(08)70266-7

[6] L. Sirovich, “Turbulence and the dynamics of coherent structures. I. Coherent structures,” en, Quarterly of Applied Mathematics, vol. 45, no. 3, pp. 561–571, 1987. doi: 10.1090/qam/910462

[7] D. Rempfer, “On Low-Dimensional Galerkin Models for Fluid Flow,” en, Theoretical and Computational Fluid Dynam-

ics, vol. 14, no. 2, pp. 75–88, Jun. 2000. doi: 10.1007/s001620050131

[8] S. Lorenzi, A. Cammi, L. Luzzi, and G. Rozza, “A reduced order model for investigating the dynamics of the Gen-IV LFR coolant pool,” Applied Mathematical Modelling, vol. 46, pp. 263–284, Jun. 2017. doi: 10.1016/j.apm.2017.01.066

[9] A. Sartori, A. Cammi, L. Luzzi, and G. Rozza, “Reduced Basis Approaches in Time-Dependent Non-Coercive Settings for Modelling the Movement of Nuclear Reactor Control Rods,” en, Communications in Computational Physics, vol. 20, no. 1, pp. 23–59, Jul. 2016. doi: 10.4208/cicp.120914.021115a

[10] S. K. Star, G. Stabile, G. Rozza, and J. Degroote, “A POD-Galerkin reduced order model of a turbulent convective buoyant flow of sodium over a backward-facing step,” Applied Mathematical Modelling, vol. 89, pp. 486–503, Jan. 2021. doi: 10.1016/j.apm.2020.07.029

[11] S. K. Star, G. Stabile, S. Georgaka, F. Belloni, G. Rozza, and J. Degroote, “Pod-Galerkin reduced order model of the Boussinesq approximation for buoyancy-driven enclosed flows,” English, in Int. Conf. Math. Comput. Methods Appl. Nucl. Sci. Eng., M C, American Nuclear Society, 2019, pp. 2452–2461.

[12] S. Georgaka, G. Stabile, K. Star, G. Rozza, and M. J. Bluck, “A hybrid reduced order method for modelling turbulent heat transfer problems,” Computers & Fluids, vol. 208, p. 104 615, Aug. 2020. doi: 10.1016/j.compfluid.2020.104615

[13] L. Vergari, A. Cammi, and S. Lorenzi, “Reduced order modeling approach for parametrized thermal-hydraulics problems: Inclusion of the energy equation in the POD-FV-ROM method,” Progress in Nuclear Energy, vol. 118, p. 103 071, Jan. 2020. doi: 10.1016/j.pnucene.2019.103071

[14] S. Hijazi, G. Stabile, A. Mola, and G. Rozza, “Data-driven POD-Galerkin reduced order model for turbulent flows,” Journal of Computational Physics, vol. 416, p. 109 513, Sep. 2020. doi: 10.1016/j.jcp.2020.109513

[15] G. Stabile and G. Rozza, “Finite volume POD-Galerkin stabilised reduced order methods for the parametrised incompressible Navier–Stokes equations,” Computers & Fluids, vol. 173, pp. 273–284, Sep. 2018. doi: 10.1016/j.compfluid.2018.01.035

[16] G. Stabile, S. Hijazi, A. Mola, S. Lorenzi, and G. Rozza, “Pod-Galerkin Reduced Order Methods for CFD Using Finite Volume Discretisation: Vortex Shedding Around a Circular Cylinder,” Communications in Applied and Industrial Mathematics, vol. 8, no. 1, pp. 210–236, Dec. 2017, arXiv:1701.03424 [math]. doi: 10.1515/caim-2017-0011

[17] B. R. Noack, P. Papas, and P. A. Monkewitz, “The need for a pressure-term representation in empirical Galerkin models of incompressible shear flows,” en, Journal of Fluid Mechanics, vol. 523, pp. 339–365, Jan. 2005. doi: 10.1017/S0022112004002149

[18] F. Ballarin, A. Manzoni, A. Quarteroni, and G. Rozza, “Supremizer stabilization of POD–Galerkin approximation of parametrized steady incompressible Navier–Stokes equations,” en, International Journal for Numerical Methods in Engineering, vol. 102, no. 5, pp. 1136–1161, 2015, eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.4772. doi: 10.1002/nme.4772

[19] I. Akhtar, A. H. Nayfeh, and C. J. Ribbens, “On the stability and extension of reduced-order Galerkin models in incompressible flows: A numerical study of vortex shedding,” en, Theoretical and Computational Fluid Dynamics, vol. 23, no. 3, pp. 213–237, Jul. 2009. doi: 10.1007/s00162-009-0112-y

[20] A. Ivagnes, G. Stabile, and G. Rozza, “Data-driven closure strategies for parametrized reduced-order models via deep operator networks,” Computer Methods in Applied Mechanics and Engineering, vol. 455, p. 118 900, Jun. 2026. doi: 10.1016/j.cma.2026.118900

[21] S. Sirisup and G. E. Karniadakis, “Stability and accuracy of periodic flow solutions obtained by a POD-penalty method,” Physica D: Nonlinear Phenomena, vol. 202, no. 3, pp. 218–237, Mar. 2005. doi: 10.1016/j.physd.2005.02.006

[22] M. D. Gunzburger, J. S. Peterson, and J. N. Shadid, “Reduced-order modeling of time-dependent PDEs with multiple parameters in the boundary data,” Computer Methods in Applied Mechanics and Engineering, vol. 196, no. 4, pp. 1030–1047, Jan. 2007. doi: 10.1016/j.cma.2006.08.004

[23] S. K. Star, G. Spina, F. Belloni, and J. Degroote, “Development of a coupling between a system thermal–hydraulic code and a reduced order CFD model,” Annals of Nuclear Energy, vol. 153, p. 108 056, Apr. 2021. doi: 10.1016/j.anucene.2020.108056

Learning to Swim Upstream: Koopman-Physics AI for Smart Microswimmer Navigation

[1]Priyam Chakraborty*; [2]Prabhat Kumar Mishra; [3]Suman Chakraborty

[1]*Department of Aerospace Engineering, Indian Institute of Technology Kharagpur*

[2]*Department of Artificial Intelligence, Indian Institute of Technology Kharagpur*

[3]*Department of Mechanical Engineering, Indian Institute of Technology Kharagpur*

Abstract

Microswimmers hold promise for targeted drug delivery, yet steering them through physiological flows remains a fundamental challenge. In realistic Poiseuille flow, propulsion, intrinsic chirality, shear, wall interactions, and Brownian noise interact nonlinearly, yielding stochastic trajectories that hinder navigation. In our earlier work [1], the swimmer was modeled as an active Brownian particle with state evolving under advection , self-propulsion , and rotation , while a reinforcement learning (RL) agent (in a Markov Decision Process) applied discrete control pulses that reset its orientation to maximize a navigation reward. However, RL exploration suffers from the curse of dimensionality and limited sample efficiency: repeated noisy simulations across broad flow and chirality ranges make training and rapid strategy discovery costly.

To overcome this bottleneck, we introduce a Koopman-enabled Physics-Informed Neural Network (KooPINN) architecture for deep learning. It provides a systematic way to validate training accuracy while incorporating physics-informed and Koopman losses to improve sample efficiency. Governing physics is embedded directly into training by constraining residuals of the translational/rotational dynamics using the Poiseuille velocity field, the + rotation law, and wall-impenetrability conditions.

Trained on ABP trajectories, the KooPINN provides orders-of-magnitude faster inference across flow and chirality regimes, reproducing upstream migration and the swinging-to-tumbling transition. It enables rapid estimation of time-to-target and path spread, accelerating RL policy/value learning and sensitivity studies, and its latent modes reveal interpretable links between learned control and fluid-mechanical dynamics. This framework bridges RL, Koopman operator theory, and physics-informed learning toward safe, real-time, physics-aware microrobot navigation in noisy in vivo flows for next-generation biomedical therapies.

Keywords: Microswimmer, Koopman operator, physics-informed neural network, reinforcement learning

Reference

1. **Priyam Chakraborty**, Rahul Roy & Shubhadeep Mandal (2025). Smart navigation of microswimmers in Poiseuille flow via reinforcement learning. *European Physical Journal E*, 48(6), 30. DOI: https://doi.org/10.1140/epje/s10189-025-00496-1