Accelerating reinforcement-learning-based flow control using transfer learning – 30 ECTS

Saeed Salehi

Division of Applied Thermodynamics and Fluid Mechanics, Department of Management and Engineering, Linköping University

About the job

Explore the cutting edge at the intersection of Artificial Intelligence (AI) and fluid flows by performing your thesis with us. We aim to develop and accelerate flow-control strategies using state-of-the-art machine learning coupled with Computational Fluid Dynamics (CFD). As a thesis worker at the Division of Applied Thermodynamics and Fluid Mechanics, you will benefit from active knowledge exchange and contribute to our research environment. All thesis projects include a local room and necessary resources and are supported by an experienced supervisor throughout the project.

Background

The control of unstable flows is a central challenge in fluid mechanics, with broad relevance to drag reduction, heat transfer control, and noise mitigation. A canonical benchmark is the two-dimensional laminar wake behind a circular cylinder, where periodic vortex shedding emerges at moderate Reynolds numbers.

Machine learning has emerged as a powerful tool for flow control and optimization. In particular, reinforcement learning (RL) provides a data-driven control methodology where an agent interacts with an environment by observing states, selecting actions, and receiving rewards (Fig. 1). The combination of RL with deep neural networks, known as Deep Reinforcement Learning (DRL), has enabled the discovery of nonlinear feedback control laws in complex systems. DRL has been successfully applied to various flow-related tasks, including the control of bluff-body wakes[1, 2] and drag reduction in turbulent channels [3]. As an example, Fig. 2 displays the baseline and DRL-controlled laminar vortex shedding behind a cylinder. While such studies highlight DRL's potential for controlling fluid flows, they also underline a key limitation, since training requires an immense number of flow simulations, with over 99% of the cost attributed to the CFD computations themselves [4, 5].

To mitigate this computational burden, multifidelity learning strategies have been introduced. In this approach, agents are pretrained in inexpensive, low-fidelity environments (e.g.,

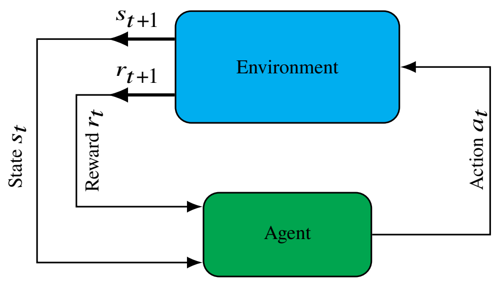

Figure 1: Schematic of the reinforcement learning (RL) framework.

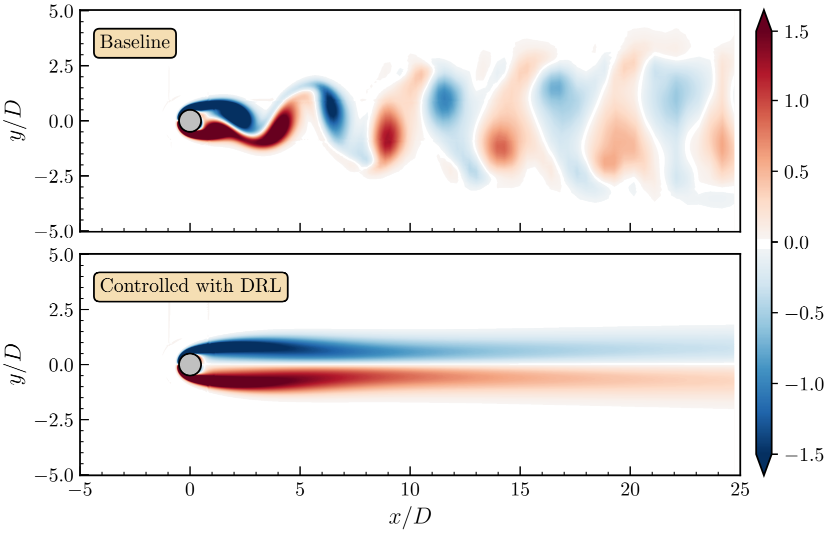

Figure 2: Baseline and DRL-controlled laminar vortex shedding behind a cylinder [6]

coarse meshes) and subsequently adapted to higher-fidelity settings [7, 8]. Fig. 2 provides a conceptual illustration of multifidelity DRL control of fluid flows. A central component is transfer learning, which enables reuse of knowledge from a source task in a related target task [9]. The most common technique, fine-tuning, often accelerates training but can suffer from catastrophic forgetting, where prior knowledge is overwritten during adaptation [10]. More structured frameworks such as Progressive Neural Networks (PNNs) [11] have been proposed to address this limitation by explicitly preserving earlier representations while extending the network with new trainable modules.

This thesis will build on these developments by investigating transfer-learning strategies for DRL-based control of canonical fluid flows.

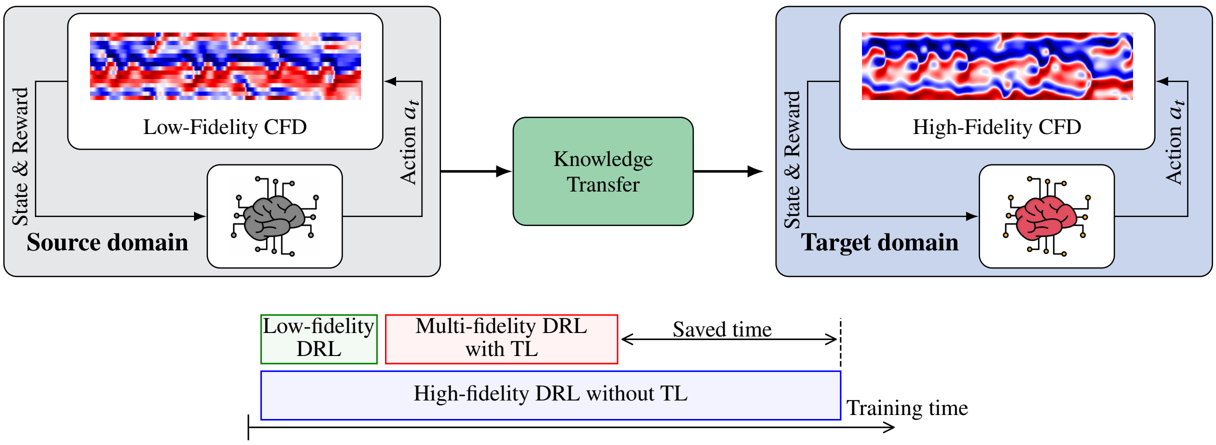

Figure 3: Conceptual illustration of the multifidelity DRL framework and its potential for computational acceleration

Scope

This thesis will investigate transfer learning for multifidelity DRL-based flow control in the canonical case of two-dimensional laminar vortex shedding behind a cylinder. The main fidelity parameter will be the mesh resolution, enabling systematic studies of how controllers trained in coarse, inexpensive environments can be adapted to finer simulations without losing previously acquired knowledge. If time allows, the work will also explore transfer across different physics, such as varying Reynolds numbers.

The CFD simulations will preferably be performed using OpenFOAM, although other open-source CFD codes may also be considered. These simulations will be coupled with a Python-based RL framework, such as Stable Baselines3 (PyTorch) or Tensorforce (TensorFlow). The open-source coupled CFD–DRL libraries developed by the supervisor (e.g., TensorforceFoam [6]) may serve as a starting point. In the same spirit, all codes and cases from this work will be released in a public open-source repository, ensuring reproducibility and facilitating future research.

The project will address the following research questions.

- Can a controller trained in low-fidelity be effectively transferred to higher-fidelity settings?
- How does fine-tuning compare with more structured transfer mechanisms such as PNNs?
- Can transfer learning across different physics accelerate learning?

The scope is deliberately limited to two-dimensional laminar flows, ensuring feasibility within a 30 ECTS master's thesis while still providing meaningful insights into the efficiency and robustness of transfer learning for DRL-based flow control. Extensions to higher Reynolds numbers may be considered if time permits.

General methodology

The project will follow a structured workflow, consisting of the following main stages.

1. **Project initiation**
   - 1.1 Review recent literature on DRL-based flow control and transfer learning.
   - 1.2 Establish a project plan with milestones and deliverables.
2. **Familiarization with the framework**
   - 2.1 Study the existing OpenFOAM-DRL framework developed by the supervisor. Alternatively, you can develop your own CFD-DRL framework with your preferred open-source packages.
   - 2.2 Read and understand the codebase to gain full control of the simulation-learning pipeline.
   - 2.3 Perform initial test runs to ensure correct setup and reproducibility.
3. **Running CFD-DRL benchmarks**
   - 3.1 Apply the coupled CFD and DRL framework to two-dimensional cylinder wake simulations at different mesh resolutions.
   - 3.2 Define and monitor performance metrics such as drag, lift fluctuations, and control energy.
   - 3.3 Analyze the results, visualize flow and learning behavior, and ensure reproducibility.
4. **Training and transfer learning**
   - 4.1 Train baseline DRL controllers on coarse-mesh simulations.
   - 4.2 Apply transfer learning strategies from coarse to fine meshes using conventional fine-tuning approaches.
   - 4.3 Implement the PNN transfer learning strategy for control of vortex shedding.
   - 4.4 If time allows, explore transfer across different Reynolds numbers or control objectives.

Saeed Salehi
saeed.salehi@liu.se
liu.se/en/employee/saesa86

5. **Analysis and documentation**
   5.1 Evaluate transfer efficiency in terms of learning speed, stability, and retention of knowledge.
   5.2 Compare the performance of different transfer strategies and identify strengths and limitations.
   5.3 Document the process, results, and conclusions in the thesis report.
   5.4 Clean up and prepare all codes and cases, and provide them openly in a public repository to ensure reproducibility and enable future research.

What you'll bring

You are pursuing a Master of Science degree in Mechanical Engineering, Aerospace Engineering (or related field), and have a solid background in fluid mechanics, numerical methods, programming, and CFD. You also have good analytical and communication skills and are a good team player.

Meriting experience

The following qualifications are considered meriting.

- Previous experience with open-source CFD, such as OpenFOAM.
- Courses or prior experience in machine learning.
- Familiarity with deep learning frameworks such as PyTorch or TensorFlow.
- Skills in Python programming and numerical libraries (e.g., NumPy, Matplotlib).
- Experience with visualization and data analysis in Python.
- Experience with Linux environments and high-performance computing.
- Experience with scientific writing in LaTeX.

Other details

- **Location:** It is necessary that you pursue your thesis at Linköping University.
- **Duration:** The duration of the thesis is 20 weeks (30 ECTS credits), starting in Jan. 2025.
- **Number of students:** Ideal for two students but can be adapted for one student.
- **Application period:** Reviewed on a rolling basis until the position is filled.

Be part of the innovative research culture – Apply now!

Be part of our innovative research culture and contribute to advancing the intersection of fluid mechanics and artificial intelligence. We welcome your application and look forward to supporting your thesis journey. Selection will be ongoing during the application period, so do not hesitate to apply. Please attach your CV, transcript of grades, and a personal letter describing your interest in the topic. Apply as soon as possible!

Contact

For more information about this thesis project, feel free to contact me.

Saeed Salehi saeed.salehi@liu.se
013 28 26 82 saeedsalehi.com
Visiting address: Room 3A:970, A-huset, Valla Campus, Linköping University
Further details on my research: saeedsalehi.com/research.

References

[1] J. Rabault, M. Kuchta, A. Jensen, U. Réglade, and N. Cerardi. "Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control". In: *Journal of Fluid Mechanics* 865 (2019), pp. 281–302. DOI: 10.1017/jfm.2019.62.

[2] F. Ren, J. Rabault, and H. Tang. "Applying deep reinforcement learning to active flow control in weakly turbulent conditions". In: *Physics of Fluids* 33.3 (2021), p. 37121. DOI: 10.1063/5.0037371.

[3] Z. Zhou, M. Zhang, and X. Zhu. "Reinforcement-Learning-Based Control of Turbulent Channel Flows at High Reynolds Numbers". In: *Journal of Fluid Mechanics* 1006 (Mar. 2025), A12. DOI: 10.1017/jfm.2025.27.

[4] J. Rabault and A. Kuhnle. "Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach". In: *Physics of Fluids* 31.9 (2019). DOI: 10.1063/1.5116415.

[5] P. Suárez, F. Alcántara-Ávila, A. Miró, J. Rabault, B. Font, O. Lehmkuhl, and R. Vinuesa. "Active Flow Control for Drag Reduction Through Multi-agent Reinforcement Learning on a Turbulent Cylinder at $\mathrm{Re}_D = 3900$". In: *Flow, Turbulence and Combustion* (Mar. 5, 2025). DOI: 10.1007/s10494-025-00642-x.

[6] S. Salehi. "An efficient intrusive deep reinforcement learning framework for OpenFOAM". In: *Meccanica* (2024). DOI: 10.1007/s11012-024-01830-1.

[7] M. Giselle Fernández-Godino. "Review of Multi-Fidelity Models". In: *Advances in Computational Science and Engineering* 1.4 (2023), pp. 351–400. DOI: 10.3934/acse.2023015.

[8] S. Bhola, S. Pawar, P. Balaprakash, and R. Maulik. "Multi-Fidelity Reinforcement Learning Framework for Shape Optimization". In: *Journal of Computational Physics* 482 (June 2023), p. 112018. DOI: 10.1016/j.jcp.2023.112018.

[9] Z. Wang, D. Fan, X. Jiang, M. S. Triantafyllou, and G. E. Karniadakis. "Deep Reinforcement Transfer Learning of Active Control for Bluff Body Flows at High Reynolds Number". In: *Journal of Fluid Mechanics* 973 (Oct. 2023), A32. DOI: 10.1017/jfm.2023.637.

[10] V. Campos, P. Sprechmann, S. Hansen, A. Barreto, S. Kapturowski, A. Vitvitskyi, A. P. Badia, and C. Blundell. *Beyond Fine-Tuning: Transferring Behavior in Reinforcement Learning*. 2021. eprint: 2102.13515.

[11] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. "Progressive Neural Networks". Preprint available on arXiv. 2016. arXiv: 1606.04671.

Linköping University
Department of Management and Engineering
Division of Applied Thermodynamics and Fluid Mechanics

Saeed Salehi
saeed.salehi@liu.se
liu.se/en/employee/saesa86