ParCFD2024
35th International Conference on Parallel Computational Fluid Dynamics
Sep. 02-04, 2024, Bonn, Germany

HydroGym-GPU: From 2D to 3D Benchmark Environments for Reinforcement Learning in Fluid Flows

Christian Lagemann[a,*], Mario Rüttgers[b], Miro Gondrum[c], Matthias Meinke[c], Wolfgang Schröder[c], Andreas Lintermann[b] and Steven L. Brunton[a]

[a]University of Washington, Department of Mechanical Engineering, Stevens Way NE, Box 352600, 98195, Seattle, WA, USA
[b]Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Wilhelm-Johnen-Straße, 52425 Jülich, Germany
[c]Institute of Aerodynamics and Chair of Fluid Mechanics, RWTH Aachen University, Wüllnerstr. 5a, 52062 Aachen, Germany

ARTICLE INFO[†]

Keywords:
Reinforcement Learning;
Flow Control;
Graphics Processing Units;
Lattice Boltzmann Method

ABSTRACT

Fluid flow modeling and control is a significant modern challenge with potential impacts across science, technology, and industry. Improved flow control could enhance drag reduction, mixing, and noise reduction in areas like transportation, energy, and medicine. However, progress in flow control is currently hindered by the lack of systematically standardized benchmarks and the high computational cost of fluid simulations. While two-dimensional problems have been extensively studied, three-dimensional simulations with larger meshes are rarely considered due to the need for highly parallelized and specialized solvers. As a result, the engineering burden of encapsulating these simulations in benchmark environments has proven to be a significant barrier. In this paper, a GPU-based extension of the HydroGym platform coupling the multiphysics solver framework m-AIA with a state-of-the-art reinforcement learning platform is presented for fluid flow control problems. Based on the highly-parallelized lattice Boltzmann solver, which is part of m-AIA, a new set of three-dimensional, non-differentiable fluid flow environments is added that extend existing flow control challenges to a new level of physical and computational complexity.

1. Introduction

The effective control of fluid dynamics is a critical challenge in many scientific, technological, and industrial systems and improved flow control has the potential to dramatically enhance performance in domains as diverse as energy, transportation, security, and medicine. For example, turbulent wall-bounded fluid flows are of significant importance for numerous engineering [1, 2, 3, 4, 5] and biomedical applications [6, 7, 8, 9, 10], e.g., in the context of reducing the CO_2 emissions in the transportation sector or enhancing disease prevention and monitoring in human medicine. Moreover, understanding and controlling the dynamics of mixing processes and multi-phase flows, such as microscopic fibers or gas bubbles in turbulent flows, is crucial for various environmental and industrial applications, including pollution control [11, 12], marine biology [13, 14], and chemical engineering [15, 16, 17]. However, controlling fluid flow is notoriously difficult due to the nonlinear and multiscale nature of fluid dynamics, which leads to high-dimensional and non-convex control problems, but rapid advancements in machine learning have significantly improved our ability to tackle these complex optimization challenges [18]. For example, reinforcement learning has recently achieved notable successes in various modern tasks, including decision-making in planning [19], robotics [20], and protein design [21]. A major factor driving these advances is the development of scalable reinforcement learning frameworks and standardized environments, which facilitate direct comparisons between learned policies.

In contrast, progress in reinforcement learning for flow control has been limited by the scarcity of such platforms. To overcome this fundamental limitation, a new scalable and extensible platform called HydroGym was recently developed that closes the loop between efficient flow solvers, flow control benchmark problems, and state-of-the-art reinforcement learning. However, rooting in a finite-element based PDE solver called Firedrake, the existing HydroGym environments are limited to a variety of two-dimensional benchmark flows since running and validating advanced three-dimensional simulations is not feasible due to scalability issues of the Firedrake solver.

[†]This paper is part of the ParCFD 2024 Proceedings. A recording of the presentation is available on YouTube. The DOI of this document is 10.34734/FZJ-2025-02455 and of the Proceedings 10.34734/FZJ-2025-02175.

*Corresponding author
clage@uw.edu (C. Lagemann); m.ruettgers@fz-juelich.de (M. Rüttgers); m.gondrum@aia.rwth-aachen.de (M. Gondrum); m.meinke@aia.rwth-aachen.de (M. Meinke); office@aia.rwth-aachen.de (W. Schröder); a.lintermann@fz-juelich.de (A. Lintermann); sbrunton@uw.edu (S.L. Brunton)
ORCID(s): 0000-0003-1150-4987 (C. Lagemann); 0000-0003-3917-8407 (M. Rüttgers); 0000-0002-2796-6644 (M. Gondrum); 0000-0003-4812-8495 (M. Meinke); 0000-0002-3472-1813 (W. Schröder); 0000-0003-3321-6599 (A. Lintermann); 0000-0002-6565-5118 (S.L. Brunton)

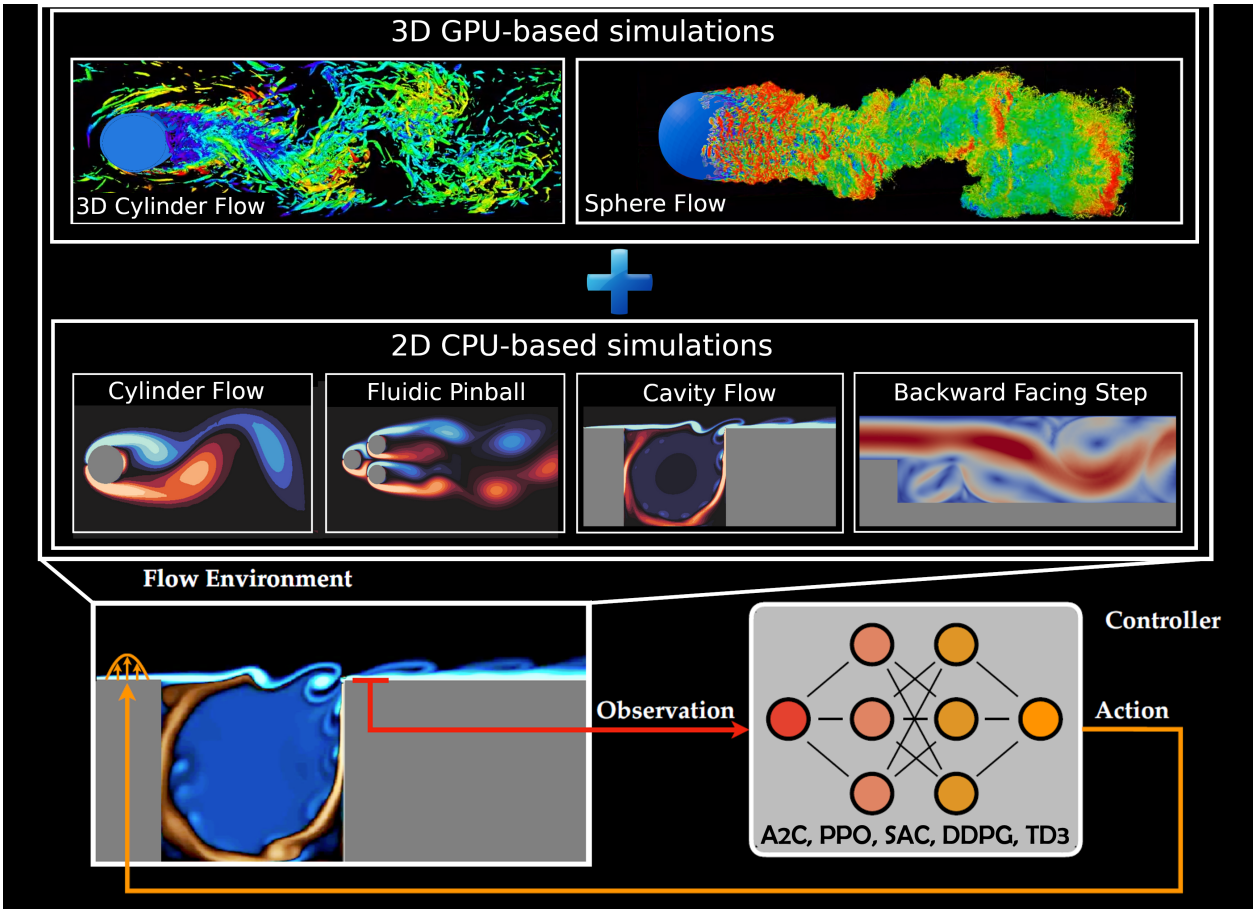

Figure 1: Flow configurations and benchmark algorithms included in the initial and extended HydroGym-GPU platform.

To tackle this limitation, a new GPU-accelerated extension bridging the lattice Boltzmann (LB) solver implemented in the multiphysics solver framework m-AIA[1] (formerly known as Zonal Flow Solver - ZFS [22, 23]) with the extensible HydroGym platform is introduced. Based on previous, successful RL applications [24] using this highly parallelized LB solver, a novel collection of three-dimensional, non-differentiable fluid flow environments has been incorporated. These additions elevate the existing flow control challenges to an unprecedented level of physical and computational complexity.

2. HydroGym-GPU

Recent advances have combined reinforcement learning with flow control, focusing on three main areas: controlling individual environments, developing strategies to navigate through flow environments, and using multi-agent reinforcement learning to learn components of numerical solvers. Despite these impressive results, the studies have been limited to specific flow control environments and lack the diversity found in modern machine learning frameworks. Hence, to effectively apply modern reinforcement learning to a broader class of fluid flow control problems, it appears to be beneficial to train policies across multiple environments. As a result, this approach allows for fine-tuning existing agents for future applications and significantly reduces computational efforts, especially for three-dimensional simulations. Therefore, integrating various flow control environments with different computational complexities requires scaling both the environments and the reinforcement learning agents to optimally use available computational resources.

One solution to tackle these problems is the usage of highly efficient and parallelized CFD solvers combined with modern data exchange protocols suitable for large-scale HPC application to enable efficient communication between RL agents and environments during runtime. Therefore,

[1] *multiphysics - Aerodynamisches Institut Aachen*

the present work leverages the existing LB solver of the m-AIA framework offering compute efficient and largely scalable simulations. It benefits from a hybrid parallelization approach based on *MPI* and a shared memory model either based on *OpenMP* or on the parallel algorithms defined in the C++ standard 17 (PSTL) implemented in the *NVIDIA HPC SDK* allowing for i) a hardware-agnostic implementation on CPU and GPU and ii) favorable strong and weak scaling on modern HPC architectures. The LB solver operates on hierarchical unstructured Cartesian grids, which are generated using a massively parallel grid generator [25] being part of m-AIA. The discretized form of the *Boltzmann* equation is solved with the *Bhatnagar-Gross-Krook* (BGK) approximation of the right-hand side collision process [26], i.e.,

$$f_i(\boldsymbol{x}+\boldsymbol{\xi_i}\delta t, t+\delta t) - f_i(\boldsymbol{x},t) = -\omega(f_i(\boldsymbol{x},t) - f_i^{eq}(\boldsymbol{x},t)), \quad (1)$$

is solved for the particle probability distribution functions f_i (PPDFs) at neighboring fluid cells at locations $\boldsymbol{x}+\boldsymbol{\xi_i}\delta t$. They are functions of the location vector $\boldsymbol{x} = (x_1, x_2, x_3)^T$, the discrete molecular velocity vector $\boldsymbol{\xi_i} = (\xi_{i1}, \xi_{i2}, \xi_{i3})^T$, and the time and time increment t and δt. The collision frequency is expressed by ω.

To enable highly efficient communication between our LB solver and the RL agents, we extended m-AIA's MPI interface leveraging the Multiple Program Multiple Data (MPMD) mode. The MPMD interface has the great advantage of executing different programs across multiple processors and heterogeneous compute hardware, e.g., CPU-CPU, CPU-GPU, GPU-GPU setups while facilitating complex, distributed computations. Moreover, MPMD supports the use of different programming languages and tools within the same application. This interoperability is extremely beneficial for the present HydroGym benchmark platform as it facilitates efficient communication between different m-AIA flow environments and standard, Python-based RL libraries. Moreover, it allows for frequent data exchange and coordination between different computational tasks with barely any computational overhead enabling very efficient training runs. Finally, it also enables multi-agent/multi-environments training protocols in a straightforward fashion without requiring changes in the code. To this end, all communications between environments and agents as wells as inter-environment communication (if necessary) is handled by our MPMD interface, while relevant inter-agent communication, e.g., gradient and weight sharing, is realized using existing deep learning libraries (JAX, PyTorch, TensorFlow, etc.).

Leveraging this computational efficiency of the GPU-accelerated LB solver and the MPMD communication, three-dimensional fluid flow environments with grid sizes on the order of $100 \cdot 10^6$ cells are added to the HydroGym platform, paving the way for novel transfer learning and control applications in the largely unexplored chapter of three-dimensional flow control. Precisely, two distinct drag reduction scenarios which are typically considered as benchmark problems in CFD code development are implemented:

3D Cylinder Flow: The first test case of our HydroGym-GPU extension targets the flow around a smooth circular cylinder which is characterized by a large range of interesting fluid mechanics phenomena as the REYNOLDS number (Re) is increased from a low to high Re-number regime, e.g., $100 < Re < 10^5$. In more detail, the flow develops from a two-dimensional steady wake to three-dimensional unsteady vortex shedding, followed by wake transition, shear layer instability and boundary layer transition. The present flow simulation setup is validated to be accurate for $Re <$ 4,000. To facilitate interactions between the RL agent and the flow environment, multiple mass sources are equally distributed across the circumference of the cylinder, each being independently controllable. Note that the mass sources extend over the entire spanwise direction. The state space, which the RL agent observes, can include a mixture of an arbitrary number of probes, including velocity, vorticity, pressure, and density sensors, that can be distributed in the entire computational domain. The reward metric is calculated based on the weighted total force value in all three dimensions similar to two-dimensional counterparts.

3D Sphere Flow: The second test case consists of an unsteady flow past a sphere. In the subcritical regime, e.g. $800 < Re < 3.7 \cdot 10^5$, the dynamics show a variety of complex flow patterns including a thin laminar boundary layer, flow separation at a location that is not known a priori, transition to turbulence in thin shear layers, and an unsteady recirculation zone followed by a turbulent wake. Without faithful representation of these flow features, it is not possible to accurately predict aerodynamic or hydrodynamic loading on complex-geometry objects. This is important since a drag reduction setup similar to the cylinder flow case is considered here. In the current implementation, the flow simulations are validated for $Re < 10^4$. To enable interactions between the RL agent and the flow environment, up to 32 individually controllable point mass sources are distributed across the sphere surface. Similar to the cylinder test case, the observational state space can comprise an arbitrary number of velocity, vorticity, pressure, and density sensors. Again, the reward is based on the integral force value in all three dimensions.

3. Results

In the following, we briefly present validation results for the introduced 3D flow environments and compare them to

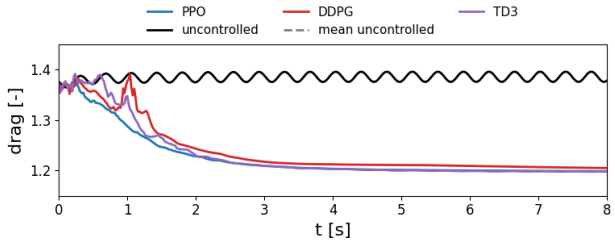

(a) Drag distribution over time for a single test episode.

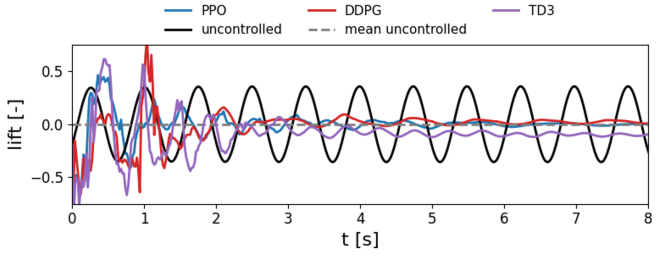

(b) Lift distribution over time for a single test episode.

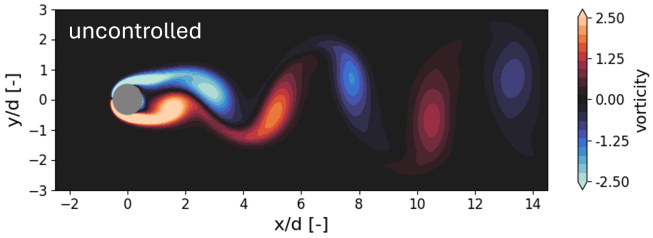

(c) Instantaneous vorticity snapshot of an uncontrolled flow environment.

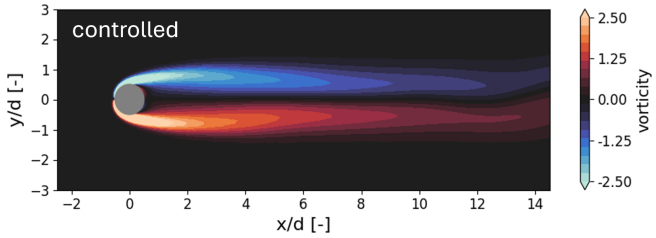

(d) Instantaneous vorticity snapshot of a flow environment controlled by a PPO agent.

Figure 2: Exemplary test results of on- (PPO) and off-policy (DDPG, TD3) agents for the 2D cylinder flow test case at $Re = 100$.

literature. Afterwards, we shortly discuss training results for different agents interacting with a 2D cylinder flow at $Re = 100$ and outline how successfully learned control policies can be leveraged in transfer learning experiments to reduce computational costs in complementing 3D environments.

Validation: A grid refinement study has been conducted for 3D simulations with $Re = 200$ by comparing the drag coefficient $C_d = (2 \cdot F_x)/(\rho \cdot U \cdot A)$ to the results in [27], with the temporally averaged force in the streamwise direction acting on the cylinder F_x, the density ρ, the inlet velocity U, and the cross-sectional area of the cylinder A. The grids are locally refined with 2 cubic refinement patches that capture the wake region and 3 cylindrical refinement patches for the near-wall regions. The results are shown in Tab. 1. Since an additional gain in accuracy of only 0.1% for the simulation with the fine grid compared to the medium grid an increased grid size from $40 \cdot 10^6$ to $80 \cdot 10^6$ cells cannot be justified and the medium grid is used for training the RL agents.

Training results: To investigate the effectiveness of our GPU-enhanced benchmark platform, we performed multiple tests for a variety of RL agents proposed in literature and conducted extensive hyperparameter optimizations and transfer learning experiments across different flow environments. However, for sake of brevity, here we only elaborate on the selected test case of a flow around cylinder at a $Re = 100$. Extensive results for all other test cases will be discussed in follow-up work.

Grid resolution	No. cells	C_d [dev. to $C_d = 1.338$ [27]]
Coarse	$20 \cdot 10^6$	1.180 [−11.8%]
Medium	$40 \cdot 10^6$	1.317 [−1.6%]
Fine	$80 \cdot 10^6$	1.318 [−1.5%]

Table 1
Grid refinement study for 3D simulations at $Re = 200$.

2D flow environment: Figure 2 exhibits exemplary test results for different RL agents including one on-policy (PPO) and two off-policy methods (DDPG and TD3). As illustrated in Fig. 2d, all agents learn a competitive control policy that stabilizes the wake and mostly reduce the lift fluctuations caused by the vortex shedding (see Fig. 2b). Overall, a total drag reduction of approximately 12% can be achieved across all agents (see Fig. 2a) which is in line with previous studies reported in literature. More importantly, by leveraging the highly efficient m-AIA LB solver in combination with the developed MPMD interface, our HydroGym-GPU extension requires only approximately 40 minutes on a single NVIDIA A100 GPU for a full training cycle (400 episodes, each containing 20 vortex shedding periods), marking a new milestone for comparable flow environments in terms of efficiency.

Transfer learning to 3D flow environments: In the following, we highlight another interesting aspect about the

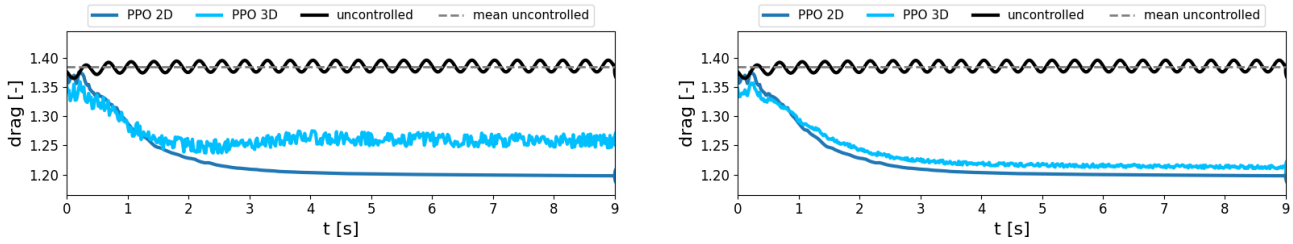

(a) zero-shot transfer learning (no training)

(b) fine-tuned control policy (20 episodes)

Figure 3: Transfer learning experiments from 2D to 3D environments leveraging (a) zero-shot applications without further training in the target environment and (b) fine-tuning experiments with limited adaptation (20 training episodes) in the target environment.

HygroGym-GPU platform enabling transfer learning experiments in a straightforward fashion. Here, we demonstrate this feature in two experiments - first, we directly transfer a control policy learned in the 2D flow case to the corresponding 3D environment, e.g. a zero-shot transfer with no subsequent training/adaptation, and second, we fine-tune the 2D flow control policy for 20 episodes in the 3D environment. Results are shown in Fig. 3 and exhibit promising trends for future compute intensive 3D flow environments. That is, even in the zero-shot transfer, the agent can mostly suppress the vortex shedding in the wake and achieves a drag reduction of approximately 8 % (see Fig. 3a). Considering that the agent has never explored the 3D environment before and consequently could never adapted to it, these results indicate a good performance for future, more general control agents leveraging foundation models across multiple flows simultaneously.

Furthermore, we can improve the performance of the control policy in the 3D environment using a fine-tuning training step in the new test environment. Precisely, we now allow the agent to adapt its network parameters to the new environment in a limited training sequence (max. 20 episodes for the 3D test case). Results are shown in Fig. 3b. The fine-tuned control agent can achieve a similar drag reduction performance compared to the original 2D test case. As a result, this hybrid transfer learning strategy requires only a fraction of the computational costs of training procedures exploring exclusively 3D environments, but still learns competitive control policies. Hence, pre-training in a simpler 2D environment with low or moderate compute requirements followed by a fine-tuning sequence in more challenging 3D environments massively reduces the training time and costs. This finding is particularly promising for future HydroGym-GPU extensions which will investigate various complex and compute intensive 3D flow cases exhibiting turbulence dominated flow features.

4. Conclusion

In this contribution, a novel GPU-accelerated extension of the recently introduced HydroGym platform is presented, connecting the multi-physics solver framework m-AIA with this adaptable RL benchmark. By utilizing a highly parallelized LB solver, a representative set of three-dimensional, non-differentiable fluid flow scenarios has been integrated and validated. Furthermore, we outlined that hybrid transfer learning strategies, which pre-train an agent in a simple flow environment of moderate computational complexity and fine-tune the policy in the target environment, require only a fraction of the computational costs compared to training procedures exploring costly 3D environment, yet learn competitive control policies. In future, we will extend our highly efficient HydroGym-GPU platform and introduce more challenging and real-world oriented test cases such as aviation applications, noise reduction in aeroacoustics settings, and mixing enhancement in multi-phase flows. Additionally, future benchmarks will investigate model-based RL agents that use latent dynamical models to uncover physical mechanisms [28], as part of our quest for general foundation models in fluid flow control. These advancements will significantly elevate the complexity of existing flow control challenges, introducing new levels of physical and computational intricacy.

Acknowledgements

The research leading to these results partially has been conducted in the HANAMI project, which receives funding from the European Union Horizon Europe Programme - Grant Agreement Number 101136269 under the call HORIZON-EUROHPC-JU-2022-INCO-04. The authors gratefully acknowledge the computing time granted by the JARA Vergabegremium and provided on the JARA Partition part of the supercomputer JUWELS [29] at *Forschungszentrum Jülich*. Moreover, this research was partially funded by the German Research Foundation within the Walter Benjamin fellowship LA 5508/1-1 (CL) and CL and SLB

acknowledge support from the National Science Foundation AI Institute in Dynamic Systems (grant number 2112085) and the Boeing Company. Furthermore, the authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. for funding this project by providing computing time on the GCS Supercomputers (CL, MM, MG, WS).

References

[1] I. Marusic, R. Mathis, N. Hutchins, Predictive Model for Wall-Bounded Turbulent Flow, Science 329 (5988) (2010) 193–196. doi:10.1126/science.1188765.

[2] P. Ricco, M. Skote, M. A. Leschziner, A review of turbulent skin-friction drag reduction by near-wall transverse forcing, Progress in Aerospace Sciences 123 (2021) 100713. doi:10.1016/j.paerosci.2021.100713.

[3] E. Lagemann, M. Albers, C. Lagemann, W. Schröder, Impact of Reynolds Number on the Drag Reduction Mechanism of Spanwise Travelling Surface Waves, Flow, Turbulence and Combustion 113 (1) (2024) 27–40. doi:10.1007/s10494-023-00507-1.

[4] E. Mäteling, M. Albers, W. Schröder, How spanwise travelling transversal surface waves change the near-wall flow, Journal of Fluid Mechanics 957 (2023) A30. doi:10.1017/jfm.2023.54.

[5] E. Lagemann, S. L. Brunton, W. Schröder, C. Lagemann, Towards extending the aircraft flight envelope by mitigating transonic airfoil buffet, Nature Communications 15 (1) (2024) 5020. doi:10.1038/s41467-024-49361-3.

[6] J. Bellien, M. Iacob, V. Richard, J. Wils, V. Le Cam-Duchez, R. Joannidès, Evidence for wall shear stress-dependent t-PA release in human conduit arteries: role of endothelial factors and impact of high blood pressure, Hypertension Research 44 (3) (2021) 310–317. doi:10.1038/s41440-020-00554-5.

[7] G. Zhou, Y. Zhu, Y. Yin, M. Su, M. Li, Association of wall shear stress with intracranial aneurysm rupture: systematic review and meta-analysis, Scientific Reports 7 (1) (2017) 5331. doi:10.1038/s41598-017-05886-w.

[8] L. Adamo, O. Naveiras, P. L. Wenzel, S. McKinney-Freeman, P. J. Mack, J. Gracia-Sancho, A. Suchy-Dicey, M. Yoshimoto, M. W. Lensch, M. C. Yoder, G. García-Cardeña, G. Q. Daley, Biomechanical forces promote embryonic haematopoiesis, Nature 459 (7250) (2009) 1131–1135. doi:10.1038/nature08073.

[9] E. Tzima, M. Irani-Tehrani, W. B. Kiosses, E. Dejana, D. A. Schultz, B. Engelhardt, G. Cao, H. DeLisser, M. A. Schwartz, A mechanosensory complex that mediates the endothelial cell response to fluid shear stress, Nature 437 (7057) (2005) 426–431. doi:10.1038/nature03952.

[10] E. Lagemann, S. L. Brunton, C. Lagemann, Uncovering wall-shear stress dynamics from neural-network enhanced fluid flow measurements, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 480 (2292) (2024). doi:10.1098/rspa.2023.0798.

[11] J. González-Martín, N. J. R. Kraakman, C. Pérez, R. Lebrero, R. Muñoz, A state–of–the-art review on indoor air pollution and strategies for indoor air pollution control, Chemosphere 262 (2021) 128376. doi:10.1016/j.chemosphere.2020.128376.

[12] M. He, Z. Xu, D. Hou, B. Gao, X. Cao, Y. S. Ok, J. Rinklebe, N. S. Bolan, D. C. W. Tsang, Waste-derived biochar for water pollution control and sustainable development, Nature Reviews Earth & Environment 3 (7) (2022) 444–460. doi:10.1038/s43017-022-00306-8.

[13] C. D. O'Dowd, M. C. Facchini, F. Cavalli, D. Ceburnis, M. Mircea, S. Decesari, S. Fuzzi, Y. J. Yoon, J.-P. Putaud, Biogenically driven organic contribution to marine aerosol, Nature 431 (7009) (2004) 676–680. doi:10.1038/nature02959.

[14] A. Paytan, K. R. M. Mackey, Y. Chen, I. D. Lima, S. C. Doney, N. Mahowald, R. Labiosa, A. F. Post, Toxicity of atmospheric aerosols on marine phytoplankton, Proceedings of the National Academy of Sciences 106 (12) (2009) 4601–4605. doi:10.1073/pnas.0811486106.

[15] V. Hessel, H. Löwe, F. Schönfeld, Micromixers—a review on passive and active mixing principles, Chemical Engineering Science 60 (8-9) (2005) 2479–2501. doi:10.1016/j.ces.2004.11.033.

[16] V. Giurgiu, L. Beckedorff, G. C. Caridi, C. Lagemann, A. Soldati, Machine learning-enhanced PIV for analyzing microfiber-wall turbulence interactions, International Journal of Multiphase Flow 181 (2024) 105021. doi:10.1016/j.ijmultiphaseflow.2024.105021.

[17] E. Lagemann, J. Roeb, S. L. Brunton, C. Lagemann, A deep learning approach to wall-shear stress quantification: From numerical training to zero-shot experimental application (2024). arXiv:2409.03933.

[18] S. L. Brunton, J. N. Kutz, Data-Driven Science and Engineering, 2nd Edition, Cambridge University Press, 2022. doi:10.1017/9781009089517.

[19] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, Mastering the game of Go with deep neural networks and tree search, Nature 529 (7587) (2016) 484–489. doi:10.1038/nature16961.

[20] Y. Chebotar, K. Hausman, Y. Lu, T. Xiao, D. Kalashnikov, J. Varley, A. Irpan, B. Eysenbach, R. Julian, C. Finn, S. Levine, Actionable Models: Unsupervised Offline Reinforcement Learning of Robotic Skills (2021). arXiv:2104.07749.

[21] I. D. Lutz, S. Wang, C. Norn, A. Courbet, A. J. Borst, Y. T. Zhao, A. Dosey, L. Cao, J. Xu, E. M. Leaf, C. Treichel, P. Litvicov, Z. Li, A. D. Goodson, P. Rivera-Sánchez, A.-M. Bratovianu, M. Baek, N. P. King, H. Ruohola-Baker, D. Baker, Top-down design of protein architectures with reinforcement learning, Science 380 (6642) (2023) 266–273. doi:10.1126/science.adf6591.

[22] A. Lintermann, M. Meinke, W. Schröder, Zonal Flow Solver (ZFS): a highly efficient multi-physics simulation framework, International Journal of Computational Fluid Dynamics 34 (7-8) (2020) 458–485. doi:10.1080/10618562.2020.1742328.

[23] Insitute of Aerodynamics and Chair of Fluid Mechanics, RWTH Aachen University, multiphysics - Aerodynamisches Institut Aachen (2024). doi:10.5281/zenodo.13350585.

[24] M. Rüttgers, M. Waldmann, K. Vogt, J. Ilgner, W. Schröder, A. Lintermann, Automated surgery planning for an obstructed nose by combining computational fluid dynamics with reinforcement learning, Computers in Biology and Medicine 173 (2024) 108383. doi:10.1016/j.compbiomed.2024.108383.

[25] A. Lintermann, S. Schlimpert, J. Grimmen, C. Günther, M. Meinke, W. Schröder, Massively parallel grid generation on HPC systems, Computer Methods in Applied Mechanics and Engineering 277 (2014) 131–153. doi:10.1016/j.cma.2014.04.009.

[26] X. He, L.-S. Luo, Theory of the lattice Boltzmann method: From the Boltzmann equation to the lattice Boltzmann equation, Physical Review E 56 (6) (1997) 6811–6817. doi:10.1103/PhysRevE.56.6811.

[27] B. Rajani, A. Kandasamy, S. Majumdar, Numerical simulation of laminar flow past a circular cylinder, Applied Mathematical Modelling 33 (3) (2009) 1228–1247. doi:10.1016/j.apm.2008.01.017.

[28] K. Lagemann, C. Lagemann, S. Mukherjee, Invariance-based learning of latent dynamics, in: The Twelfth International Conference on Learning Representations, 2024.
URL https://openreview.net/forum?id=EWTFMkTdkT

[29] D. Alvarez, JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Juelich Supercomputing Centre, Journal of large-scale research facilities JLSRF 7 (2021) A183. doi:10.17815/jlsrf-7-183.