
Supplementary information

The HydroGym reinforcement learning platform for fluid dynamics

In the format provided by the
authors and unedited

Supplementary Information

The HydroGym Reinforcement Learning Platform for Fluid Dynamics

Christian Lagemann^{*1,2}, Sajeda Mokbel^{1,2}, Miro Gondrum³, Mario Rüttgers⁴, Yuning Wang⁵, Pol Suárez⁶, Ludger Paehler⁷, Deniz A. Bezgin⁷, Aaron B. Buhendwa⁷, Jared L. Callahan^{1,2}, Samuel Ahnert^{1,2}, Nicholas Zolman^{1,2}, Xiao Shao³, Jean-Christophe Loiseau⁸, Nikolaus A. Adams^{7,9}, Matthias Meinke^{3,10}, Wolfgang Schröder^{3,10}, Kai Lagemann^{11, 12}, Esther Lagemann^{1,2}, Ricardo Vinuesa^{*5,6}, and Steven L. Brunton^{*1,2}

¹*Department of Mechanical Engineering, University of Washington, Seattle, WA 98195, United States*

²*AI Institute in Dynamic Systems, University of Washington, Seattle, WA 98195, United States*

³*Chair of Fluid Mechanics and Institute of Aerodynamics, RWTH Aachen University, Aachen, Germany*

⁴*Data-Driven Fluid Engineering Laboratory, Inha University, Incheon, South Korea*

⁵*Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, United States*

⁶*FLOW, Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, Sweden*

⁷*School of Engineering and Design, Technical University of Munich, Munich, Germany*

⁸*Arts et Métiers Institute of Technology, CNAM, DynFluid, HESAM Université, Paris, France*

⁹*Munich Institute of Integrated Materials, Energy and Process Engineering, Technical University of Munich, Munich, Germany*

¹⁰*JARA Center for Simulation and Data Science, RWTH Aachen University, Aachen, Germany*

¹¹*MediaTek Research, London, United Kingdom*

¹²*Statistics and Machine Learning, German Center for Neurodegenerative Diseases, Bonn, Germany*

**Corresponding authors: christian.lagemann@rwth-aachen.de, rvinuesa@umich.edu, sbrunton@uw.edu*

Contents

1	Extended Introduction	3
2	Extended literature review	6
3	HydroGym - Numerical Backends and Implementation	9
3A	m-AIA - Lattice-Boltzmann solver implementation	10
3B	m-AIA – Finite-volume solver implementation	12
3C	Nek5000 implementation	15
3D	Incompressible JAX implementation	16
3E	JAX-Fluids implementation	17
3F	Firedrake implementation	18
4	Benchmarked RL Agents	19
4A	Proximal Policy Optimization (PPO)	20
4B	Deep Deterministic Policy Gradient (DDPG)	20
4C	Twin Delayed Deep Deterministic Policy Gradient (TD3)	21
4D	Gradient-Enhanced Proximal Policy Optimization (GPPO)	21
5	Environments	22
5A	Circular cylinder flow	22
5B	Fluidic pinball flow	26
5C	Open cavity flow	30
5D	Square cylinder flow	33
5E	Flow around a NACA 0012 airfoil	35
5F	Flow around a cube	41
5G	Flow around a sphere	43
5H	Zero-pressure gradient turbulent boundary layer	45
5I	Supercritical DRA2303 airfoil	48
5J	Flow through a stenotic pipe	49
5K	Turbulent channel flow	52
5L	Differentiable turbulent channel flow	54
5M	Kolmogorov flow	56
5N	Shock vector control in a single divergent nozzle	57

1 Extended Introduction

The control of fluid flows is a critical challenge across virtually every major industry, from aerospace and automotive to energy production and manufacturing. Improved flow control has the potential to transform trillion-dollar sectors by dramatically increasing energy efficiency, transportation systems, and industrial processes. However, fluid flow control remains notoriously challenging due to the nonlinear, multiscale nature of turbulent flows, resulting in control formulations that are high-dimensional, non-convex, and computationally intractable using traditional approaches [1, 2]. The convergence of recent breakthroughs in reinforcement learning (RL) with advances in differentiable physics simulation [3, 4] and specialized hardware acceleration has created an unprecedented opportunity to revolutionize fluid dynamics control through artificial intelligence (AI).

This revolution encompasses both the practical implementation of flow control and the fundamental scientific understanding that underpins it. AI-driven differentiable physics environments represent a transformative scientific tool for advancing our fundamental understanding of fluid dynamics. The key advantage of differentiable simulations lies in their ability to compute exact gradients of objective functions (such as drag, lift, or mixing efficiency) with respect to control inputs, geometric parameters, or initial conditions through automatic differentiation. This enables gradient-based optimization and sensitivity analysis that would require prohibitively many simulations with traditional adjoint methods or finite-difference approximations, particularly for high-dimensional control spaces and long time horizons.

Moreover, the combination of differentiable physics with RL creates a unique scientific tool: agents can autonomously discover non-intuitive control strategies and flow manipulation techniques that human intuition might overlook, effectively serving as computational experiments that probe the limits of what is physically achievable. These discovered strategies often reveal unexpected connections between control actuation and flow response, inspiring new hypotheses about underlying physical mechanisms. The gradient information also enables systematic inverse design—working backward from desired flow features to determine the forcing or geometry required to achieve them—which is invaluable for developing physics-based reduced-order models and improving turbulence closure schemes. By providing efficient pathways to explore the control-response landscape and extract interpretable patterns from successful strategies, these platforms establish a new methodology for scientific discovery that complements traditional CFD while offering unique capabilities for optimization, inverse problems, and automated discovery of control principles.

To understand the potential of this technological convergence, consider the vast economic significance of fluid flow control, which extends across industries representing over \$10 trillion in global market value. The industrial flow control sector alone represents a \$225 billion market*, with direct applications spanning aerospace systems valued at \$3.2 billion and projected to reach \$30 billion by 2032†, energy systems where industrial pumping accounts for nearly 20% of global electrical demand, and manufacturing processes across sectors responsible for 38% of industrial CO₂ emissions‡. Breakthrough successes in wind energy control have reduced costs from 55+ cents/kWh in 1980 to under 3 cents/kWh today§, creating thousands of jobs and generating \$10 billion in annual investments.

This economic imperative coincides with a remarkable period of breakthrough achievements in reinforcement learning, particularly in complex scientific applications that mirror the challenges of fluid dynamics control. Reinforcement learning has demonstrated transformative capabilities across increasingly complex

*McKinsey & Company, "Flow control: Sector at a crossroads?" <https://www.mckinsey.com/industries/industrials-and-electronics/our-insights/flow-control-sector-at-a-crossroads>

†Global Market Insights, "Aerospace & Defense Fluid Conveyance Systems Market Size, 2032." <https://www.gminsights.com/industry-analysis/aerospace-defense-fluid-conveyance-systems-market>

‡U.S. Department of Energy, "Transformative Pathways for US Industry," January 2025. <https://www.energy.gov/sites/default/files/2025-01/transformative-pathways-for-us-industry.pdf>

§American Clean Power Association, "Wind Power Facts and Statistics." <https://cleanpower.org/facts/wind-power/>

scientific domains, recently resulting in the 2024 Nobel Prizes in Physics for foundational neural network research that enables modern RL systems[¶], and in Chemistry for AlphaFold’s revolutionary protein structure prediction [5]. AlphaFold solved the 50-year protein folding challenge by achieving experimental-level accuracy within minutes rather than months, predicting structures for over 200 million proteins and fundamentally transforming biological research^{||}. Concurrent breakthroughs in plasma control demonstrate RL’s capability to master complex physical systems: DeepMind’s deployment of deep RL controllers on the TCV tokamak achieved real-time control of 100+ million°C plasmas at 10 kHz frequencies, enabling previously impossible configurations including sustained “droplets” and high-elongation shapes critical for fusion energy [6]. These systems overcame fundamental challenges including 19+ coupled magnetic coils, millisecond timescales with instability growth rates exceeding 1.4 kHz, and safety-critical constraints where failures could damage multi-billion dollar facilities.

These applications share common success factors that provide crucial insights for advancing fluid flow control. The success of RL across diverse domains has been catalyzed by standardized benchmark environments that democratize research access and accelerate algorithmic innovation. OpenAI Gym has attracted over 2 million users across 190 countries^{**}, while specialized environments, such as MuJoCo for robotics and Atari for discrete control established evaluation frameworks that transformed entire research communities [7]. These platforms enabled fair algorithmic comparison, reproducible research, and systematic progress tracking that proved essential for breakthrough developments including AlphaGo, robotics deployments, and large language model alignment through reinforcement learning from human feedback.

However, despite these remarkable successes in analogous complex domains, fluid dynamics has remained largely resistant to similar breakthroughs due to a constellation of fundamental technical barriers. The primary obstacle is computational intractability: training effective RL agents typically requires thousands of environment interactions, but each evaluation demands expensive computational fluid dynamics (CFD) simulations that can require hours on high-performance computing clusters. This computational bottleneck is magnified by the inherent sample inefficiency of model-free RL methods, making it practically impossible for complex turbulent flow problems. Moreover, traditional CFD approaches face severe limitations in real-time control applications, with typical simulations requiring wall-clock times orders of magnitude longer than the physical processes they model.

Beyond computational constraints, fundamental challenges in generalization and adaptability have further impeded progress. Current RL approaches trained at specific Reynolds numbers fail to transfer to different flow regimes, and agents optimized for particular geometries cannot adapt to modified configurations [8]. The transition from laminar to turbulent conditions introduces chaotic, irregular dynamics that destabilize training algorithms, while the multiscale nature of turbulence—spanning from energy-containing large scales to dissipative Kolmogorov microscales—necessitates extensive data collection across multiple scales, which amplifies computational costs. Traditional control approaches face complementary limitations: PID controllers with fixed parameters cannot adapt to highly nonlinear turbulent flows, while optimal control theory encounters severe time horizon limitations due to instability in adjoint equations for chaotic systems [9].

Compounding these technical challenges, the fluid dynamics community has lacked the comprehensive benchmark platforms that have enabled other domains. While computer vision benefits from standardized datasets, such as ImageNet, that enable systematic algorithmic comparison, fluid dynamics researchers work with limited, problem-specific configurations that prevent fair evaluation and reproducible research. Most existing fluid dynamics RL efforts have concentrated on computationally inexpensive 2D configurations at moderate Reynolds numbers, with the extension to highly turbulent flows, complex 3D geometries, and multiphysics applications representing the current frontier of research.

[¶]The Nobel Prize, “Press release: The Nobel Prize in Physics 2024.” <https://www.nobelprize.org/prizes/physics/2024/press-release/>

^{||}Google DeepMind, “AlphaFold.” <https://deepmind.google/science/alphafold/>

^{**}OpenAI, “OpenAI Gym Beta.” <https://openai.com/index/openai-gym-beta/>

Remarkably, the technological landscape has undergone a dramatic transformation that now makes it possible to overcome these historical barriers. GPU-accelerated differentiable fluid simulators now achieve orders-of-magnitude speedups over traditional CFD methods, with modern platforms enabling parallel environment training at scales previously requiring supercomputers [9]. The ML4CFD competition at NeurIPS 2024, featuring over 240 teams, demonstrated 300-600x speedups over traditional solvers while maintaining accuracy^{††}, indicating the maturity of these approaches for practical applications. Differentiable reinforcement learning methods specifically designed for continuous control problems address sample efficiency challenges, with approaches using analytic gradients from differentiable simulation to dramatically reduce sample complexity in multiphysics environments [10]. Physics-informed neural networks provide principled ways to incorporate fluid dynamics knowledge into RL systems, achieving superior performance in strongly nonlinear problems with sparse data [11].

The commercial sector has validated this technological readiness through concrete deployments and performance improvements. Major CFD vendors including Ansys and Siemens have launched GPU-accelerated solvers showing 20-33x performance improvements, while enterprise platforms provide production-grade tools for physics-informed machine learning. This convergence of mature differentiable simulation technology, breakthrough RL algorithms, and accessible hardware represents the first time that large-scale fluid dynamics RL training has become computationally feasible.

The urgency of capitalizing on this technological convergence is underscored by the massive economic stakes and immediate industrial needs for improved fluid control solutions. Advanced aerodynamic control in automotive and aerospace applications could generate billions in fuel savings across global fleets. For instance, removing side mirrors alone in vehicles (accounting for around 6% of form drag [12]) improves fuel economy by 1.5-2 miles per gallon, while comparable drag reductions in aircraft yield proportionally larger savings due to higher fuel consumption rate [13, 14]. Coordinated wind turbine controls can increase farm output by 4-5%^{‡‡}, representing substantial returns across the \$10 billion annual wind investment market. Industrial applications present even larger opportunities: pumping systems account for nearly 20% of global electrical energy demand, with potential energy savings up to 75% through advanced control. Manufacturing industries face direct competitive pressures where fluid control improvements translate to immediate economic advantages, from billion-dollar steel production operations to chemical processing facilities requiring millisecond response times for safety-critical applications.

Responding to this critical convergence of technological capability and economic opportunity, we introduce HydroGym, a benchmark flow control platform designed to bridge the gap between cutting-edge RL method development and fluid dynamics applications. HydroGym provides a flexible, solver-independent interface that seamlessly integrates with industry-standard RL frameworks while supporting both differentiable and non-differentiable fluid dynamics environments. Our platform abstracts computational fluid dynamics complexity from control algorithm development, enabling researchers from both communities to advance the field through standardized benchmarks and reproducible evaluation protocols.

The specific contributions of this work directly address each of the critical barriers that have prevented breakthrough progress in fluid dynamics RL:

- **Comprehensive benchmark platform:** Introduction of the first scalable RL platform specifically designed for fluid dynamics control, featuring a solver-independent architecture that supports diverse CFD backends and computational configurations.
- **Diverse environment suite:** 61+ non-differentiable flow control environments with increasing complexity, from canonical flows past cylinders to advanced turbulent control scenarios on 3D aircraft wings, plus an initial set of differentiable flow control environments enabling gradient-based RL methods.

^{††}NeurIPS 2024 ML4CFD Competition: Results and Retrospective Analysis." <https://arxiv.org/abs/2506.08516>

^{‡‡}U.S. Department of Energy, "Next-Generation Wind Technology." <https://www.energy.gov/eere/wind/next-generation-wind-technology>

- **Systematic algorithmic evaluation:** Comprehensive performance assessment using industry-standard algorithms including Proximal Policy Optimization (PPO), Deep Deterministic Policy Gradient (DDPG), and Twin Delayed DDPG (TD3), establishing baseline performance metrics for future research.
- **Democratized research access:** Platform enables RL researchers and fluids experts to collaborate via standardized interfaces, lowering the barrier to entry and accelerating interdisciplinary progress.
- **Extensible foundation:** Readily extensible framework supports integration of new CFD solvers, environments, and RL algorithms, to evolve with advancing technology and research needs.

Through this comprehensive approach, HydroGym addresses the specific technical barriers that have prevented breakthrough progress – computational intractability through state-of-the-art GPU-powered CFD solvers and efficient simulation interfaces, sample inefficiency through differentiable environments, generalization challenges through diverse benchmarks, and accessibility barriers through democratized platforms – thereby enabling the transition from traditional engineering approaches to AI-driven fluid control systems. The platform arrives at a critical juncture where enabling technologies have matured and economic pressures intensify, positioned to catalyze transformative progress comparable to breakthroughs in protein folding, plasma control, and other complex scientific domains. Just as AlphaFold demonstrated that the convergence of deep learning, domain expertise, and computational resources could solve seemingly intractable scientific challenges, HydroGym establishes the foundation for similar breakthroughs in fluid dynamics control. Success in this domain promises a fundamental transformation of how we approach complex flow systems across trillion-dollar industries, establishing artificial intelligence as an essential tool for addressing the most challenging problems in fluid mechanics and beyond.

2 Extended literature review

The intersection of reinforcement learning and flow control represents a paradigm shift in fluid system optimization. Traditional methods rely on predetermined strategies based on domain knowledge, while reinforcement learning enables autonomous discovery of optimal control policies through the direct interaction with flow environments. This capability is particularly valuable in turbulent flows, where traditional linear control theory fails due to high-dimensional, chaotic dynamics. Since the first successful demonstration of deep reinforcement learning for cylinder wake control [15], the field has experienced rapid growth. The progression from proof-of-concept studies to practical implementations reveals the maturation of core methodologies and sophisticated frameworks capable of handling realistic engineering problems. This evolution has been driven by advances in deep learning architectures, increased computational resources, and growing recognition of traditional control limitations for complex fluid systems.

Canonical Flow Problems and Foundational Methods. The journey toward practical reinforcement learning in flow control began with carefully chosen problems that balanced complexity with computational feasibility. The application of reinforcement learning to circular cylinder wake control was a first step, establishing foundational principles for the field [15]. The canonical von Kármán vortex street suppression provided a computationally tractable testbed, allowing researchers to develop and validate core methodologies before tackling more challenging applications.

Building on these initial successes, numerous studies consistently achieved 8 – 15% drag reduction with minimal actuation energy (0.5-1% of momentum deficit), demonstrating that neural agents could autonomously discover effective control strategies without prior knowledge of flow physics [16–20]. These results were encouraging because they showed that reinforcement learning could match or exceed the performance of carefully tuned traditional controllers without prior understanding of the underlying flow

physics. Crucially, experimental validation eliminated concerns about simulation-to-reality transfer that had plagued earlier computational studies. Laboratory studies using rotating cylinders achieved similar performance improvements [21], with drag reductions of 25-30% at Reynolds numbers up to 140,000. This success proved that reinforcement learning can handle noise, measurement uncertainties, and modeling errors inherent in real fluid systems.

Attention has since turned to more complex geometries that better represent practical engineering challenges. Square cylinder flows present additional complexity due to fixed separation points and pronounced pressure fluctuations [22–26]. Reinforcement learning policies often learn to employ multi-frequency patterns differing significantly from natural shedding frequencies, highlighting the discovery potential of learning-based methods.

The field has rapidly matured towards industrially relevant flows. Studies at Reynolds numbers exceeding 200,000 have achieved substantial drag reduction through multi-frequency control strategies [27]. Further, multi-agent reinforcement learning has shown superior performance compared to single-actuator approaches [28, 29], setting the stage for more complex applications requiring distributed control.

Shape Optimization and Advanced Applications. A natural next step in RL for fluids was to explore design optimization where the geometry itself becomes the control parameter. Reinforcement learning for shape optimization represents a major step toward autonomous design generation [30, 31]. This transition from controlling existing geometries to generating new ones opened new possibilities for fluid engineering. Modern frameworks parameterize complex shapes via Bézier curves or B-splines with hundreds of design variables [32, 33]; optimization in these spaces was previously intractable. The key advantage over conventional optimization lies in reinforcement learning’s ability to navigate high-dimensional design spaces without getting trapped in local optima, while simultaneously handling multiple competing objectives.

This multi-objective capability has proven particularly valuable for airfoil design, where competing requirements must be balanced across flight conditions. Reinforcement learning frameworks generate continuous Pareto fronts representing optimal trade-offs between lift and drag across ranges of Reynolds numbers and angles of attack [34]. Performance comparisons consistently demonstrate superiority over traditional optimization methods, with superior resistance to local minima and better generalization to untested operating conditions.

The success in two-dimensional applications naturally led to more ambitious and industrially relevant three-dimensional flows. Automotive and marine applications have leveraged reinforcement learning for complex three-dimensional geometries [35, 36], where traditional optimization approaches become computationally prohibitive. Marine applications have achieved significant reductions in acoustic signatures while maintaining hydrodynamic performance [37], while hull optimization frameworks generate designs meeting multiple constraints simultaneously [38], demonstrating the technology’s readiness for real-world engineering problems.

Separation Control and Turbulent Flow Applications. Flow separation control represents one of the most demanding applications, requiring precise manipulation of boundary layer dynamics in highly unsteady environments. Reinforcement learning has revealed sophisticated strategies surpassing traditional periodic actuation [39, 40], often discovering control mechanisms that human experts had not previously considered. Plasma actuators enable precise boundary layer manipulation [41], with learned policies discovering multi-frequency control strategies that simultaneously modify large-scale structures while fine-tuning smaller-scale phenomena. This multi-scale control represents a significant advance over traditional approaches that typically target single frequency ranges or specific flow scales.

The complexity of separation control is perhaps best illustrated in dynamic stall applications, where reinforcement learning has achieved complete stall prevention rather than mere mitigation [42]. These

applications require real-time adaptation to rapidly changing flow conditions, demonstrating the technology's potential for autonomous flight control. Multi-actuator systems coordinated by multi-agent frameworks enable spatially distributed manipulation with precision unattainable through centralized approaches.

The application of reinforcement learning to fully turbulent channel flow control has emerged as a canonical problem for high Reynolds number evaluation [43–46], representing conditions typical of many engineering applications. Successful agents achieve drag reduction levels of 20-40% across Reynolds numbers spanning two orders of magnitude, with performance scaling favorably with system size. The RL control mechanisms often differ fundamentally from traditional opposition control, employing sophisticated strategies that redistribute turbulent kinetic energy across different components and locations. Multi-agent reinforcement learning has addressed scalability for spatially distributed control, enabling parallelization across multiple locations while avoiding exponential complexity growth, pointing toward practical implementations in large-scale systems.

Bio-Inspired and Bio-medical Applications. Researchers have also explored how RL can provide insights into biological systems, while simultaneously investigating bio-inspired approaches for artificial applications. These studies yield insights into natural propulsion while revealing opportunities for artificial systems to potentially exceed biological performance. Self-propelled swimming studies reproduce complex behaviors including rheotaxis and Kármán gating [47–50], demonstrating that reinforcement learning can discover sophisticated control strategies evolved by natural swimmers. Multi-agent applications to collective swimming reveal coordination emerging from individual optimization without explicit communication [51, 52], suggesting new approaches for autonomous swarms.

Reinforcement learning for bio-medical fluids has mainly been used for respiration and hemodynamics. In the respiratory system, patient-specific respiratory flow simulations have been coupled with reinforcement learning to propose surgical interventions in the nasal cavity that optimize nasal resistance and improve the airway's capacity to heat incoming air [53, 54]. In the cardiovascular system, reinforcement learning agents have been trained to regulate nerve stimulation to maintain target heart rate and mean arterial pressure [55], to control the motor speed of cardiovascular pumps in order to restore pulsatility under varying heart failure conditions [56], or to design optimal stent placements in patient-specific aneurysm geometries [57]. Together, these studies highlight the potential of reinforcement learning and physics-based simulations to enable personalized medicine and support clinical decision-making.

Computational Enhancement and Algorithmic Advances. Parallel to advances in direct flow control applications, reinforcement learning has begun transforming the computational tools used to study fluid mechanics itself. RL for CFD solver optimization represents a paradigm shift toward intelligent computational fluid dynamics, where the simulation process itself becomes adaptive and self-optimizing. Adaptive mesh refinement applications treat grid optimization as sequential decision problems [58–60], with learned strategies demonstrating superior performance compared to conventional error-based approaches. High-order method integration enables simultaneous optimization of mesh topology and polynomial order distributions [61], potentially revolutionizing how CFD simulations are constructed and executed.

Concurrently, reinforcement learning has emerged as a transformative approach for turbulence modeling, addressing long-standing challenges in Reynolds-Averaged Navier-Stokes (RANS) closure and Large-Eddy subgrid scale modeling. By treating model parameter optimization as an adaptive learning problem where coefficients adjust based on local flow conditions [62–65], these approaches overcome the fundamental limitation of traditional turbulence models that rely on universal constants derived from canonical flows.

Multi-agent reinforcement learning frameworks have been successfully applied to wall-model development for large eddy simulations, enabling distributed optimization of near-wall turbulence representations while maintaining computational efficiency [66, 67]. Physics-informed reinforcement learning approaches

have integrated governing differential equations directly into the learning framework, ensuring that learned turbulence models satisfy fundamental conservation laws while improving predictive accuracy across diverse flow configurations [68, 69]. Data-driven model development leveraging high-fidelity simulation databases has demonstrated superior performance compared to traditional fixed-parameter approaches, with learned models providing better uncertainty quantification and adaptability to new flow regimes. Model-consistent training strategies that incorporate CFD solvers directly into the learning process have addressed the distribution shift problems that previously limited machine learning applications to turbulence modeling, enabling effective deployment in practical engineering simulations.

These advances represent a paradigm shift toward intelligent turbulence modeling where algorithms continuously adapt their closure strategies based on evolving flow physics rather than relying on universal constants derived from canonical flows. This adaptability is particularly valuable for complex engineering flows that deviate significantly from the canonical configurations used to calibrate traditional models.

Another promising direction targets model-based reinforcement learning approaches in flow control learning surrogate models of fluid dynamics to enable more sample-efficient policy optimization [70–72]. These lightweight, physics-informed models provide significant computational advantages over high-fidelity CFD simulations during policy training while maintaining sufficient accuracy for effective control strategy development, enabling practical deployment of reinforcement learning in scenarios where extensive interaction with expensive fluid solvers would otherwise be prohibitive.

The convergence of reinforcement learning with fluid mechanics has evolved from isolated demonstrations to a comprehensive framework capable of addressing fundamental challenges across the discipline, from turbulence closure modeling to autonomous system design. This progression demonstrates not just the maturation of individual techniques, but the emergence of a new paradigm for approaching fluid mechanics problems that emphasizes adaptation, learning, and autonomous optimization. Looking forward, the integration of physics-informed learning algorithms with high-performance computing promises to unlock previously intractable optimization problems in fluid systems, potentially revolutionizing fields ranging from renewable energy harvesting to biomedical device design. As these methodologies mature beyond laboratory settings toward real-world implementation, they represent a paradigmatic shift toward intelligent, adaptive fluid systems that autonomously optimize their performance in response to changing environmental conditions, fundamentally changing how we design and operate fluid mechanical systems.

3 HydroGym - Numerical Backends and Implementation

HydroGym provides multiple computational backends to accommodate diverse research requirements and computational resources in flow control applications. The framework’s modular architecture supports six distinct solver implementations: a high-performance lattice Boltzmann method (LBM) solver built on the m-AIA framework for large-scale simulations (Section 3A), a compressible finite-volume (FV) solver also within m-AIA targeting wall-bounded turbulent flows at higher Mach numbers (Section 3B), a spectral-element solver based on Nek5000 for incompressible high-Reynolds-number configurations requiring spectral accuracy (Section 3C), a differentiable incompressible JAX-based implementation optimized for gradient-enhanced machine learning workflows (Section 3D), a differentiable compressible solver based on JAX-Fluids for high-order flow control in compressible and multiphase regimes (Section 3E), and a Firedrake finite element backend designed for maximum code transparency and rapid prototyping (Section 3F). Each implementation targets specific use cases while maintaining a unified interface for reinforcement learning integration, enabling researchers to select the most appropriate computational approach based on their performance requirements, hardware constraints, and research objectives.

3A m-AIA - Lattice-Boltzmann solver implementation

The direct numerical simulations are conducted using a lattice Boltzmann method (LBM) embedded within the m-AIA solver framework, which has been continuously developed at the Institute of Aerodynamics of RWTH Aachen University for over two decades [73]. This method solves a discretized form of the Boltzmann equation, which describes the evolution of the particle distribution function $f(\mathbf{x}, \mathbf{v}, t)$ in time t and space $x = (x, y, z)$:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f = \Omega(f). \quad (1)$$

Here, $\Omega(f)$ is the collision operator, which models the redistribution of particle populations due to intermolecular collisions. Through a first-order Chapman-Enskog expansion of this equation, the Navier-Stokes equations can be recovered, establishing the method's connection to macroscopic fluid dynamics.

The discretization of the continuous Boltzmann equation results in the lattice Boltzmann equation (LBE), which evolves discrete particle populations f_i on a computational grid. The LBE update rule follows a two-step process consisting of collision and streaming phases:

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \Omega_i(f). \quad (2)$$

For the present test cases, a three-dimensional lattice with 27 discrete velocity vectors ($\mathbf{c}_i$) is employed, which is known as the D3Q27 model. From these discrete particle distributions, macroscopic fluid variables such as density ρ and velocity $\mathbf{u}$ are computed from the moments:

$$\rho = \sum_i f_i \quad \text{and} \quad \rho \mathbf{u} = \sum_i \mathbf{c}_i f_i. \quad (3)$$

Overall, m-AIA's LBM solver offers two different collision operators. The first approach is the widely-used Bhatnagar-Gross-Krook (BGK) model [74], which relaxes the particle distributions towards a local Maxwellian equilibrium f_i^{eq} with a single relaxation frequency ω_{BGK} :

$$\Omega_i = \omega_{\text{BGK}}(f_i^{\text{eq}} - f_i). \quad (4)$$

The relaxation frequency is related to the effective kinematic viscosity ν_{eff} of the fluid by $\omega_{\text{BGK}} = (\Delta t c_s^2) / (\nu_{\text{eff}} + 0.5 \Delta t c_s^2)$, where c_s is the lattice speed of sound. For isothermal and low-Mach number flows, the equilibrium distribution is given by:

$$f_i^{\text{eq}} = w_i \rho \left[1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{c}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right], \quad (5)$$

with w_i being the lattice-dependent weighting factors.

For enhanced stability in simulations of higher Reynolds number flows, an alternative cumulant-based collision operator is employed [75]. In this approach, the collision is performed in the space of statistical cumulants c_α , which are relaxed towards their equilibrium values with a set of distinct relaxation rates ω_α . The post-collision state in cumulant space is expressed as $c_\alpha^* = c_\alpha + \omega_\alpha(c_\alpha^{\text{eq}} - c_\alpha)$. Following the methodology of Geier et al. [75], all relaxation rates are set to unity except for one, which is set equal to ω_{BGK} to control the fluid viscosity.

For thermal simulations, a second set of particle distributions g_i is solved. To compute the thermal equilibrium distributions, three approaches have been validated and employed in respiratory flow simulations with m-AIA: (1) treating the temperature (T) solely as a passive scalar [31, 76], (2) including internal energy effects [77], and (3) considering total energy [53, 54]. Similar to Eq. 3, T is computed from the moments of g_i .

Furthermore, local grid refinement is implemented based on the cell-centered interface layout proposed by Eitel-Amor et al. [78] to efficiently resolve flow features across different length scales. This refinement

strategy involves adjusting the relaxation frequency on different grid levels to maintain a constant kinematic viscosity across the entire domain. In order to keep the Mach number constant across the grid transition, the time step is scaled proportionally to the local grid spacing $\Delta t \sim \Delta$. The viscosity is kept constant by setting the local relaxation frequency according to:

$$v_{\text{eff}} = \Delta t c_s^2 \left(\frac{1}{\omega_{\text{BGK}}} - \frac{1}{2} \right). \quad (6)$$

While the macroscopic quantities and thus the equilibrium state are interpolated in space and time, the non-equilibrium part is additionally rescaled ensuring consistency of the viscous stress tensor, see [79].

Boundary Conditions and actuation control. Solid boundaries within the fluid domain are handled using an interpolated bounce-back scheme [80]. This method provides a second-order accurate representation of the no-slip condition for arbitrarily shaped geometries. For the far-field boundaries, inlet and outlet conditions are imposed by setting the particle distributions to an equilibrium state based on prescribed velocity, density, or temperature values, while using second-order extrapolation for any unknown variables from the domain's interior.

Flow control is implemented by changing the inflow temperature in the thermal benchmark cases, and by imposing time-varying Dirichlet boundary conditions on the surface of the immersed body in the remaining cases. These conditions modify the standard no-slip wall condition by specifying a non-zero velocity vector, $\mathbf{u}_w$, which is then enforced through the interpolated bounce-back scheme. Two distinct types of actuation are employed: synthetic jet actuation and surface rotation. Both methods utilize a temporal ramping mechanism to ensure smooth transitions between control actions, which is critical for numerical stability.

Temporal Smoothing of Control Actions. To prevent numerical instabilities arising from discontinuous changes in boundary conditions, the control action is never applied instantaneously. Instead, the actuation magnitude $A(t)$ (representing either jet velocity or angular velocity) is smoothly interpolated from its value at the previous control step, A^{old} , to the new target value, A^{new} , over a predefined number of sub-iterations N_r .

Hence, the instantaneous actuation magnitude is calculated as:

$$A(t) = A^{\text{old}} + (A^{\text{new}} - A^{\text{old}}) \cdot R(\tau), \quad (7)$$

where $\tau = (t - t_{\text{update}})/(N_r \Delta t)$ is the normalized time within the ramping interval, such that $\tau \in [0, 1]$. The function $R(\tau)$ is a normalized ramp function that maps from 0 to 1. Two different ramp functions are available to tune the actuation response:

- **Hyperbolic Tangent Ramp:** This function provides a smooth, S-shaped transition that is gentle at the beginning and end of the ramp period. It is defined as:

$$R_{\text{tanh}}(\tau) = \frac{1}{2} (\tanh(b(\tau - 0.5)) + 1). \quad (8)$$

The parameter b controls the steepness of the ramp.

- **Exponential Ramp:** This function provides a ramp that starts slowly and accelerates, allowing for a more aggressive actuation towards the end of the interval. It is defined as:

$$R_{\text{exp}}(\tau) = \frac{e^{k\tau} - 1}{e^k - 1}. \quad (9)$$

The parameter k controls the curvature of the exponential ramp.

Synthetic Jet Actuation. This boundary condition models synthetic jets by prescribing a velocity profile over specified regions of the wall. Each jet j is defined by its center position $\mathbf{x}_j$, blowing angle α_j , and width W_j . The velocity profile is parabolic, ensuring it is maximal at the centerline and smoothly decays to zero at the jet's edges. The velocity vector $\mathbf{u}_w(\mathbf{x}, t)$ for a point $\mathbf{x}$ within the jet's active region is determined by the product of a spatial profile function $\mathbf{S}_j(\mathbf{x})$ and the time-varying jet velocity magnitude $A_j(t)$:

$$\mathbf{u}_w(\mathbf{x}, t) = A_j(t) \cdot \mathbf{S}_j(\mathbf{x}). \quad (10)$$

The spatial profile vector $\mathbf{S}_j(\mathbf{x})$ combines the jet's shape and direction:

$$\mathbf{S}_j(\mathbf{x}) = \left(1 - \left(\frac{2d_\perp}{W_j} \right)^2 \right) \mathbf{t}_j, \quad (11)$$

where $\mathbf{t}_j = (\cos \alpha_j, \sin \alpha_j)$ is the unit vector in the direction of the jet, and $d_\perp$ is the perpendicular distance from the point $\mathbf{x}$ to the jet's centerline. This profile is applied for locations where $d_\perp \leq W_j/2$. The jet magnitude $A_j(t)$ is supplied by the control agent and interpolated using one of the temporal smoothing functions described by Eq. 7.

Surface Rotation Actuation. This boundary condition simulates the rotation of solid surfaces, such as a cylinder. The actuation is defined by an angular velocity $\omega(t)$ around a specified center of rotation $\mathbf{x}_c$ and axis of rotation $\hat{\mathbf{a}}$. The resulting tangential velocity at any point $\mathbf{x}$ on the boundary is given by the cross product:

$$\mathbf{u}_w(\mathbf{x}, t) = \omega(t) \hat{\mathbf{a}} \times (\mathbf{x} - \mathbf{x}_c). \quad (12)$$

For a 2D simulation in the $x - y$ plane, the rotation axis is fixed to $\hat{\mathbf{a}} = (0, 0, 1)$, and the control action is the scalar angular velocity $\omega_z(t)$, which serves as the actuation magnitude $A(t)$. The velocity vector at the wall simplifies to:

$$\mathbf{u}_w(x, y, t) = \omega_z(t) (-(y - y_c), (x - x_c)). \quad (13)$$

The angular velocity $\omega_z(t)$ provided by the control agent is temporally smoothed between control updates.

MPI communications and GPU acceleration. HydroGym's C++ code implementation benefits from hybrid parallelization based on MPI and shared memory models, i.e., OpenMP and higher-level parallelism features introduced by C++20 parallel algorithms. This allows for a hardware-agnostic implementation on both CPU and GPU-based architectures, relying on GPU backends provided in, e.g., NVIDIA HPC SDK or AMD ROCm HIPSTDPAR. The hierarchical unstructured Cartesian grids are generated using a massively parallel grid generator. Previous work has achieved favorable strong and weak scaling on modern HPC systems [81]. To facilitate reproducibility and ease of deployment, we provide (pre-built) containerized environments and pre-configured setup files that enable users to readily deploy the HydroGym framework with minimal configuration overhead.

3B m-AIA – Finite-volume solver implementation

Beyond the lattice Boltzmann method, m-AIA also includes a compressible finite-volume (FV) solver targeting direct numerical simulation and large-eddy simulation of wall-bounded turbulent flows. The Navier-Stokes equations are discretized on structured, body-conforming curvilinear grids with a cell-centered finite-volume scheme. This grid topology is advantageous for configurations such as flat plates and airfoils, where strong velocity gradients are concentrated near solid surfaces and can be efficiently captured by anisotropic mesh refinement in the wall-normal direction.

Governing equations. The FV solver advances the compressible Navier-Stokes equations cast in an arbitrary Lagrangian-Eulerian (ALE) formulation, accommodating domains whose boundaries may deform over time. In integral form, these conservation laws for mass, momentum, and energy read

$$\frac{d}{dt} \int_{V(t)} \mathbf{Q} dV + \oint_{A(t)} (\mathbf{H}_{\text{inv}} + \mathbf{H}_{\text{vis}}) \cdot \mathbf{n} dA = 0, \quad (14)$$

with the state vector $\mathbf{Q} = (\rho, \rho \mathbf{u}, \rho E)^\top$ collecting the conserved quantities. Here, ρ denotes the density, $\mathbf{u} = (u, v, w)^\top$ the velocity, and E the specific total energy, related to the specific internal energy e by $\rho E = \rho e + \frac{\rho}{2} |\mathbf{u}|^2$. The tensors $\mathbf{H}_{\text{inv}}$ and $\mathbf{H}_{\text{vis}}$ represent inviscid and viscous fluxes through the surface $A(t)$ whose outward unit normal is $\mathbf{n}$. Their combined contribution reads

$$\mathbf{H} = \mathbf{H}_{\text{inv}} - \mathbf{H}_{\text{vis}} = \begin{pmatrix} \rho(\mathbf{u} - \mathbf{u}_A) \\ \rho \mathbf{u}(\mathbf{u} - \mathbf{u}_A) + p \mathbf{I} \\ \rho E(\mathbf{u} - \mathbf{u}_A) + p \mathbf{u}_A \end{pmatrix} + \begin{pmatrix} 0 \\ \boldsymbol{\tau} \\ \boldsymbol{\tau} \mathbf{u} + \mathbf{q} \end{pmatrix}, \quad (15)$$

with $\mathbf{u}_A$ the velocity of the moving control-volume surface, $\boldsymbol{\tau}$ the viscous stress tensor, $\mathbf{q}$ the heat flux vector, and p the pressure. Assuming a Newtonian fluid with vanishing bulk viscosity, the stress tensor takes the form $\boldsymbol{\tau} = 2\mu \mathbf{S} - \frac{2}{3}\mu(\nabla \cdot \mathbf{u})\mathbf{I}$, where $\mathbf{S} = \frac{1}{2}[\nabla \mathbf{u} + (\nabla \mathbf{u})^\top]$ is the rate-of-strain tensor and μ is obtained from Sutherland's law. Heat conduction obeys Fourier's law at a fixed Prandtl number $Pr_0 = 0.72$, and the thermodynamic closure consists of the ideal gas relation $T = \gamma p / \rho$ together with $e = p / [(\gamma - 1)\rho]$, using $\gamma = 1.4$ for air.

Curvilinear coordinate transformation. To resolve the steep wall-normal gradients characteristic of boundary layer flows, the equations are formulated on structured body-fitted grids. A mapping $(x, y, z) \rightarrow (\xi, \eta, \zeta)$ with unit spacing in computational space aligns coordinate lines with curved solid boundaries. Transforming the differential Navier-Stokes equations into this coordinate system and making use of the metric-identity relations [82] recovers a fully conservative form,

$$\frac{\partial \hat{\mathbf{Q}}}{\partial t} + \frac{\partial \hat{\mathbf{E}}}{\partial \xi} + \frac{\partial \hat{\mathbf{F}}}{\partial \eta} + \frac{\partial \hat{\mathbf{G}}}{\partial \zeta} = 0, \quad (16)$$

with the scaled state vector $\hat{\mathbf{Q}} = J\mathbf{Q}$, where J is the reciprocal of the Jacobian determinant. The transformed fluxes combine contributions from all physical directions weighted by the metric coefficients, e.g.,

$$\hat{\mathbf{E}} = J(\xi_t \mathbf{Q} + \xi_x \mathbf{E} + \xi_y \mathbf{F} + \xi_z \mathbf{G}), \quad (17)$$

and $\hat{\mathbf{F}}, \hat{\mathbf{G}}$ follow analogously.

When the grid moves in time, as is the case for surface-actuation problems, the geometric conservation law (GCL) must be enforced to avoid non-physical mass, momentum, or energy production in the discrete equations [82]. This is achieved by evaluating the temporal metrics $(J\xi_t), (J\eta_t), (J\zeta_t)$ from the volumes swept by individual cell faces between successive time levels.

Inviscid flux discretization. The inviscid fluxes at cell interfaces are evaluated with the advection upstream splitting method (AUSM) [83]. AUSM separates each interface flux into two contributions: a convective component, which is upwind-biased according to the local flow direction, and a pressure component, which is centrally averaged. Specifically, the discretized inviscid flux in the ξ -direction reads

$$\hat{\mathbf{E}}_{\text{inv}} = \frac{1}{2} \left[M_{L/R} \left(\hat{\mathbf{E}}_L^c + \hat{\mathbf{E}}_R^c \right) + |M_{L/R}| \left(\hat{\mathbf{E}}_L^c - \hat{\mathbf{E}}_R^c \right) \right] + \hat{\mathbf{E}}_{L/R}^p, \quad (18)$$

where the interface Mach number $M_{L/R} = \frac{1}{2}(M_L + M_R)$ determines whether the left or right convective flux is used, and the pressure flux is computed as $\hat{\mathbf{E}}_{L/R}^p = \frac{1}{2}(p_L + p_R)J(\xi_x, \xi_y, \xi_z, \xi_t)^\top$. Second-order accuracy is achieved through a monotonic upstream-centered scheme for conservation laws (MUSCL) [84] extrapolation of the primitive variables to the cell interfaces. The left and right states at an interface $i + \frac{1}{2}$ are reconstructed from neighboring cell-center values using a distance-weighted interpolation that accounts for non-uniform grid spacing.

Viscous flux discretization. Because the viscous terms contain nested first-order derivatives, their discretization produces second-order operators that propagate information omnidirectionally. A central scheme is therefore appropriate, and a modified cell-vertex (MCV) approach is adopted. Each interface flux is obtained by averaging the fluxes evaluated at the four vertices of that face:

$$\left. \frac{\partial \hat{\mathbf{E}}_{\text{vis}}}{\partial \xi} \right|_{i,j,k} = \frac{1}{4} \sum_{m,n \in \{-\frac{1}{2}, +\frac{1}{2}\}} \hat{\mathbf{E}}_{\text{vis}, i+\frac{1}{2}, j+m, k+n} - \frac{1}{4} \sum_{m,n \in \{-\frac{1}{2}, +\frac{1}{2}\}} \hat{\mathbf{E}}_{\text{vis}, i-\frac{1}{2}, j+m, k+n}, \quad (19)$$

where the velocity and temperature gradients at each vertex are computed with standard central differences. Compared with a naive evaluation at cell centers, this vertex-based averaging lowers the number of metric-term evaluations while retaining second-order accuracy.

Temporal integration. Discretizing the spatial operators converts the transformed Navier-Stokes equations into a system of ordinary differential equations,

$$\frac{\partial \hat{\mathbf{Q}}}{\partial t} = \text{RHS} \left(t; \hat{\mathbf{Q}} \right) \quad (20)$$

which is marched forward in time with a five-stage, low-storage Runge-Kutta method whose stage coefficients are $\alpha_k = (1/4, 1/6, 3/8, 1/2, 1)$, providing second-order accuracy in time. When the grid deforms, all geometric quantities—metric terms, Jacobians, and interface volume fluxes—are updated at every Runge-Kutta stage so that temporal accuracy is preserved.

Implicit large-eddy simulation. When the mesh is too coarse to resolve all turbulent scales, the solver relies on a monotonically integrated large-eddy simulation (MILES) strategy [85, 86]. Rather than adding an explicit subgrid-scale model, MILES exploits the fact that the upwind-biased AUSM scheme already introduces controlled numerical diffusion that acts preferentially on the smallest resolved scales, effectively mimicking the energy drain of the unresolved turbulent cascade. Extensive validation across a range of turbulent configurations has confirmed that this implicit modelling yields predictions of comparable quality to those obtained with explicit subgrid closures [87].

MPI parallelization and GPU acceleration. For parallel execution, the structured grid is partitioned into blocks of roughly equal size via a balanced cut-tree algorithm that simultaneously minimizes the surface area shared between neighboring blocks, thereby reducing communication overhead. Ghost-cell layers whose depth matches the stencil width surround each block, so that all numerical operators can be applied uniformly across the interior without special near-boundary treatment. Ghost-cell values are synchronized across blocks through MPI after every Runge-Kutta stage. As described for the LBM solver, the FV code also supports hybrid parallelism via OpenMP and C++20 parallel algorithms, allowing it to run on GPU-accelerated architectures through the NVIDIA HPC SDK or AMD ROCm HIPSTDPAR backends.

Boundary conditions and actuation control. At solid walls, adiabatic no-slip conditions enforce the fluid velocity to match the wall velocity and prescribe vanishing wall-normal gradients of pressure and temperature. Far-field in- and outflow boundaries employ a characteristic condition based on the local contravariant Mach number [88], supplemented by sponge layers on the pressure field to absorb acoustic reflections. Where required, periodic conditions are applied in the spanwise direction, and synthetic turbulence generation following the reformulated method of Roidl et al. [89] provides fully turbulent inflow data from precursor RANS solutions.

For flow control, the FV solver supports the same synthetic jet actuation described in Section 3A, in which a parabolic velocity profile is imposed on designated wall segments and the jet magnitude is smoothly ramped between successive control updates. In addition, a streamwise traveling transversal surface wave boundary condition is available, following the formulation introduced by Albers et al. [14, 90–93] for compressible turbulent boundary layer drag reduction. In the original fixed-parameter formulation, the wall is deflected in the wall-normal direction according to

$$y|_{\text{wall}}(x, t) = g(x) A \cos\left(\frac{2\pi}{\lambda} x - \frac{2\pi}{T} t\right), \quad (21)$$

where A is the wave amplitude, λ the streamwise wavelength, T the period, and $g(x)$ is a piecewise cosine envelope that smoothly ramps the wave on and off over prescribed streamwise intervals. The corresponding wall velocity $v|_{\text{wall}} = dy/dt|_{\text{wall}}$ is imposed consistently, and the body-fitted grid deforms to track the moving surface while satisfying the geometric conservation law.

For integration with the HydroGym reinforcement learning framework, this boundary condition is generalized so that the wave amplitude A , wavelength λ , and phase speed $c = \lambda/T$ become time-varying control parameters set by the RL agent at each control step. Because abrupt changes in any of these quantities would introduce discontinuities in the grid velocity field and compromise numerical stability, a cross-fade transition strategy is employed. When the agent issues new target values $(A^{\text{new}}, \lambda^{\text{new}}, c^{\text{new}})$, the actuation is not switched instantaneously. Instead, a blending factor

$$\alpha(\tau) = \frac{1}{2}(1 - \cos(\pi\tau)), \quad \tau = n_{\text{step}}/N_{\text{trans}} \in [0, 1], \quad (22)$$

is advanced over N_{trans} sub-iterations, where n_{step} counts the elapsed transition steps. The instantaneous wave parameters are then obtained by interpolation between their previous and target values, e.g., $A(t) = A^{\text{old}} + \alpha(A^{\text{new}} - A^{\text{old}})$, and analogously for the phase speed. Rather than interpolating the wavelength directly, the wavenumbers $k^{\text{old}} = 2\pi/\lambda^{\text{old}}$ and $k^{\text{new}} = 2\pi/\lambda^{\text{new}}$ are blended, which avoids singularities for large wavelengths and yields a smoother spatial transition. The wall displacement during the cross-fade is computed as a weighted superposition of the old and new waveforms,

$$y|_{\text{wall}}(x, t) = g(x) A(t) [(1 - \alpha) \cos(k^{\text{old}} \xi) + \alpha \cos(k^{\text{new}} \xi)], \quad (23)$$

where $\xi(x, t) = x - x_0 - d(t)$ is the position relative to the wave origin shifted by the accumulated phase distance $d(t) = \int_0^t c(t') dt'$, which is integrated numerically to ensure phase continuity across parameter changes. The grid velocity and acceleration fields required by the ALE formulation are obtained by consistent analytical time differentiation of Eq. (23), accounting for the time dependence of α , A , and d . This cross-fade mechanism guarantees C^1 -continuous grid motion in time and permits the RL agent to explore the full three-dimensional parameter space (A, λ, c) without risking solver instabilities from sudden actuation changes.

3C Nek5000 implementation

Mathematical Formulation and Discretization Nek5000 [94] is an incompressible Navier–Stokes solver that employs the spectral-element method (SEM) to achieve high numerical accuracy while maintaining

computational efficiency [95]. The computational domain is partitioned into hexahedral elements, within which the velocity and pressure fields are approximated by high-order Lagrange interpolants following the $\mathbb{P}_N\mathbb{P}_{N-2}$ formulation [96]. For a polynomial order of $N - 1$, the velocity field is represented on N^3 grid points per element distributed according to the Gauss–Lobatto–Legendre (GLL) quadrature rule. The pressure field is defined on a staggered grid comprising $(N - 2)^3$ points per element using Gauss–Legendre (GL) quadrature.

Temporal integration is performed using a third-order explicit extrapolation scheme (EXT3) for the nonlinear convective terms and a third-order implicit backward differentiation scheme (BDF3) for the viscous contributions. To mitigate aliasing errors, overintegration is applied by oversampling the evaluation of nonlinear terms with a factor of $3/2$ relative to the adopted polynomial order N in each spatial direction.

HydroGym currently provides two Nek5000-based environments: (1) a three-dimensional turbulent channel flow configuration solved using direct numerical simulation (DNS), and (2) a three-dimensional NACA 0012 wing section simulated using high-resolution large-eddy simulation (LES). Detailed descriptions of the turbulent channel flow and the NACA 0012 wing configurations are presented in Secs. 5K and 5E, respectively.

Boundary Conditions and actuation control A time-dependent wall-normal velocity component v_n is prescribed at the wall as the control actuation through a Dirichlet boundary condition. The spatial mean of the actuation signal is removed to enforce the zero-net-mass-flux (ZNMF) constraint, thereby ensuring global mass conservation. This boundary treatment models localized jet actuation that induces blowing or suction at the solid surface.

MPI communications and parallelization All Nek5000 environments are executed on CPU architectures, with parallelization achieved through a hybrid distributed-memory and shared-memory strategy based on OpenMPI. By taking advantage of the structured meshes and the tensor-product formulation inherent to the SEM, the solver exhibits strong scalability and robust parallel performance, as demonstrated in prior studies. The coupling with reinforcement learning agents is implemented through dynamic MPI-based communication, enabling stable and efficient interaction with large-scale simulations, such as three-dimensional wing configurations at high Reynolds numbers described in Sec. 5E. To enhance reproducibility and facilitate deployment, we provide pre-built containerized environments together with pre-configured setup files, allowing users to deploy the HydroGym framework with minimal configuration overhead.

3D Incompressible JAX implementation

Physics simulators built with modern automatic differentiation libraries make it possible to compute gradients of output quantities with respect to desired input parameters, even with complex physical processes in the middle. These gradients can then be integrated into gradient-based control schemes, one of them being gradient-enhanced reinforcement learning [97]. The HydroGym framework explores this capability by offering flow solvers written in JAX—a differentiable programming language with GPU acceleration—to model flows described by the incompressible Navier–Stokes equations:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \quad \nabla \cdot \mathbf{u} = 0, \quad (24)$$

where $\mathbf{u}$ denotes the velocity, p is the pressure, ν is the kinematic viscosity and $\mathbf{f}$ is an external forcing term. HydroGym currently provides two differentiable environments which model Equation 24: a three-dimensional turbulent channel flow and a two-dimensional forced turbulence problem known as the Kolmogorov flow [98]. A more detailed implementation description of both environments is given in Secs. 5M & 5L.

Software Architecture. Each environment employs a modular architecture that separates physical problem definitions from the reinforcement learning interface, following the design principles of HydroGym’s Firedrake environments (Sec. 3F) to ensure code clarity, extensibility, and ease of use. The differentiable environments differ by implementing a Gymnax interface [99] and inheriting reinforcement learning frameworks from PureJaxRL [100], enabling a fully synchronous training pipeline compatible with JAX’s automatic differentiation. This architecture allows users to compute gradients of any flow output with respect to tunable input parameters at any point during environment execution or RL training via JAX’s `grad` function, facilitating gradient-based optimization and control strategies.

GPU acceleration and parallelization. The JAX environments leverage GPU acceleration and parallelization to enable rapid RL training. The end-to-end framework supports parallel execution across multiple random seeds using JAX’s automatic vectorization functions such as `vmap` for batch processing. This enables training hundreds of agents simultaneously in a fraction of the time required for single-agent training in PyTorch [100, 101]. Additionally, just-in-time compilation via `jit` automatically optimizes code execution on GPU hardware.

3E JAX-Fluids implementation

JAX-Fluids is a fully-differentiable, high-order computational fluid dynamics (CFD) solver designed for compressible single- and two-phase flows [102, 103]. Written entirely in Python using the JAX library, it leverages automatic differentiation (AD) and hardware acceleration to seamlessly bridge traditional CFD with machine learning workflows. Within the HydroGym framework, JAX-Fluids provides a powerful backend for optimizing active flow control strategies in compressible regimes, enabling exact gradient computation through complex, highly-resolved flow physics.

Mathematical Formulation and Discretization. Unlike the incompressible solvers described previously, JAX-Fluids solves the compressible Navier-Stokes equations. The fluid state is described by the vector of conservative variables $\mathbf{U} = [\rho, \rho\mathbf{u}, E]^T$, representing density, momentum, and total energy per unit volume, respectively. The governing equations are formulated as:

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}^c(\mathbf{U}) = \nabla \cdot \mathbf{F}^d(\mathbf{U}) + \mathbf{S}(\mathbf{U}), \quad (25)$$

where $\mathbf{F}^c$ and $\mathbf{F}^d$ denote the convective and diffusive fluxes, and $\mathbf{S}$ represents source terms such as external forcings or body forces. The system is closed using an appropriate equation of state, typically the ideal or stiffened gas law.

The spatial domain is discretized using a finite-volume method on structured Cartesian grids, which support arbitrary one-dimensional mesh stretching. Convective fluxes are evaluated using a high-order Godunov-type approach, pairing advanced shock-capturing spatial reconstruction schemes (e.g., WENO5-Z, TENO6-A) with an approximate HLLC Riemann solver. Diffusive fluxes are approximated via high-order central finite differences. Time integration of the semi-discrete system is performed using explicit total-variation-diminishing Runge-Kutta (TVD-RK) schemes. Furthermore, JAX-Fluids natively supports multiphase flows, offering both a sharp-interface level-set method and a five-equation diffuse-interface model. Positivity-preserving limiters are used to ensure robust integration in the presence of strong shocks or large density ratios.

The level-set implementation in JAX-Fluids additionally functions as a conservative immersed boundary method (IBM), enabling flow simulations around complex geometries. The IBM is particularly well suited for active flow control problems, as control actions (e.g., blowing or suction) can be naturally incorporated through interface exchange terms.

Software Architecture and Machine Learning Integration. JAX-Fluids departs from traditional loop-based programming in favor of the array programming paradigm, utilizing JAX’s NumPy-like API. By adhering strictly to functional programming constraints—such as pure functions and static control flow operations—the entire simulation pipeline is end-to-end differentiable.

This architecture allows researchers to compute exact gradients of arbitrary scalar quantities of interest (e.g., time-averaged drag, aerodynamic efficiency, or an RL reward function) with respect to control inputs, boundary conditions, or neural network weights. Gradients are backpropagated through the complete temporal simulation trajectory using JAX’s `grad` or `value_and_grad` transformations. To mitigate memory bottlenecks associated with unrolling long rollout trajectories during the backward pass, JAX-Fluids employs gradient checkpointing (`jax.checkpoint`). This makes it highly straightforward to couple the solver with deep learning libraries like Flax and Optax, facilitating the joint optimization of parameterized control policies directly within the CFD environment.

HPC and Parallelization. To accommodate the vast computational resources required for three-dimensional compressible turbulence and complex flow control scenarios, JAX-Fluids incorporates a hardware-agnostic parallelization strategy tailored for modern HPC clusters. The computational domain is partitioned using a homogeneous domain decomposition strategy, distributing grid blocks across multiple XLA devices (GPUs or TPUs) via the `jax.pmap` transformation.

Crucially, inter-block communications for face, edge, and vertex halo updates are executed exclusively through JAX’s collective permutation primitives (`jax.lax.ppermute`). This purely JAX-native communication ensures that the automatic differentiation graph remains intact across distributed, multi-node compute clusters. This enables highly efficient weak scaling on supercomputers—such as NVIDIA A100 GPU clusters and Google TPU pods—while simultaneously supporting the parallel evaluation of AD gradients for large-scale differentiable flow control tasks.

3F Firedrake implementation

The HydroGym framework provides an alternative implementation using Firedrake [104], a finite element framework built on PETSc that offers automatic code generation and optimization for variational problems. This implementation targets researchers requiring full access to solver internals and complete Python-based extensibility, prioritizing code transparency and flexibility over computational performance. The Firedrake backend implements two-dimensional flow control problems using mixed finite element methods for the incompressible Navier-Stokes equations.

Mathematical Formulation and Discretization. The solver employs Taylor-Hood elements with second-order continuous Galerkin (CG) elements for velocity and first-order CG elements for pressure, ensuring inf-sup stability. The discretized system takes the form:

$$\mathbf{M} \frac{\partial \mathbf{q}}{\partial t} + \mathbf{R}(\mathbf{q}) = 0, \quad (26)$$

where $\mathbf{q} = (\mathbf{u}, p)$ represents the mixed velocity-pressure state, $\mathbf{M}$ is the mass matrix, and $\mathbf{R}(\mathbf{q})$ is the nonlinear residual of the incompressible Navier-Stokes equations:

$$\mathbf{R}(\mathbf{q}) = \int_{\Omega} [-(\mathbf{u} \cdot \nabla) \mathbf{u} \cdot \mathbf{v} - \boldsymbol{\sigma}(\mathbf{u}, p) : \boldsymbol{\epsilon}(\mathbf{v}) + (\nabla \cdot \mathbf{u}) s] d\Omega, \quad (27)$$

where $(\mathbf{v}, s)$ are velocity and pressure test functions, $\boldsymbol{\sigma}(\mathbf{u}, p) = 2\nu\boldsymbol{\epsilon}(\mathbf{u}) - p\mathbf{I}$ is the Newtonian stress tensor, and $\boldsymbol{\epsilon}(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$ is the strain rate tensor. Time integration is performed using implicit schemes with automatic selection of solver parameters based on Reynolds number and grid resolution.

Software Architecture. The implementation follows a three-tier modular architecture that separates physical problem definition from numerical methods and reinforcement learning interfaces. This design enables extensibility while maintaining code clarity and facilitating rapid prototyping of new flow control scenarios.

The `PDEBase` class serves as an abstract template defining the core functionality required for PDE-based control problems. It encapsulates the physical configuration including domain discretization, boundary conditions, and objective function evaluation, while remaining solver-agnostic. Key abstract methods include mesh handling (`load_mesh`, `initialize_state`), state management (`set_state`, `copy_state`, `reset`), and problem-specific definitions (`get_observations`, `evaluate_objective`). The class provides a standardized interface through which control agents interact with the underlying fluid physics, handling the conversion between control inputs and boundary condition modifications. Additionally, it manages actuator dynamics and maintains checkpoint functionality for simulation restart capabilities.

The `TransientSolver` class implements the time-stepping algorithms for evolving the PDE state. It provides a `solve` method that integrates the flow equations over specified time intervals, supporting optional callback functions for data collection and visualization. The core `step` method advances the simulation by one time increment, accepting control inputs and updating the PDE state accordingly. This separation allows for different time integration schemes to be implemented without modifying the physical problem definition. The solver supports both open-loop time integration and closed-loop operation with feedback controllers of the form $u(t, y)$.

The `FlowEnv` class bridges the PDE framework with reinforcement learning by implementing the Farama Foundation Gymnasium interface. It wraps `PDEBase` and `TransientSolver` instances, translating between RL concepts (observations, actions, rewards) and CFD operations (state queries, boundary condition updates, objective evaluation). The environment handles episode management through reset functionality, computes rewards as the negative of the time-integrated objective function, and manages termination conditions. This design allows any combination of PDE problem and solver to be automatically instantiated as a Gym-compatible environment through configuration dictionaries.

Available Environments. The framework includes three benchmark flow configurations: flow around a circular cylinder with rotary control, flow around a cylinder with blowing/suction actuation, cavity flow with leading-edge control, and the three-cylinder pinball configuration with independent rotary actuation for each cylinder.

Computational Infrastructure The implementation supports MPI parallelization through Firedrake’s distributed mesh capabilities and leverages PETSc’s parallel linear algebra routines. While computational efficiency is secondary to code accessibility, the framework maintains reasonable performance for research applications through automatic optimization of generated finite element code and efficient sparse linear system solvers. The modular architecture enables researchers to extend the framework with new flow configurations by implementing problem-specific methods while inheriting the complete finite element infrastructure and RL integration capabilities. Similar to the m-AIA environments, we also provide pre-built container environments.

4 Benchmarked RL Agents

The HydroGym platform conducts comprehensive performance evaluation using four distinct reinforcement learning algorithms, representing both established model-free approaches and emerging gradient-enhanced

methodologies. For the non-differentiable environments, we employ TorchRL [105] as a stable baseline implementation for Proximal Policy Optimization (PPO), Deep Deterministic Policy Gradient (DDPG), and Twin Delayed Deep Deterministic Policy Gradient (TD3). Additionally, we utilize in-house implementations of PPO and Gradient-Enhanced Proximal Policy Optimization (GPPO) specifically designed for JAX-based differentiable environments [106]. This dual implementation strategy ensures optimal performance across both traditional CFD solvers and modern differentiable simulation frameworks while maintaining algorithmic consistency and reproducibility.

4A Proximal Policy Optimization (PPO)

PPO implements a trust region policy optimization approach that constrains policy updates through a clipped surrogate objective function to prevent destructive parameter changes during training [107]. The core innovation lies in its clipped probability ratio mechanism, which limits the extent of policy updates based on the ratio between new and old policy probabilities. The algorithm optimizes the following clipped objective:

$$L^{CLIP}(\theta) = \mathbb{E}_t [\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)] \quad (28)$$

where $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)}$ represents the probability ratio, A_t denotes the advantage estimate, and ϵ controls the clipping range (typically 0.2).

A Generalized Advantage Estimation (GAE) with $\lambda = 0.95$ for variance reduction in advantage calculations [108] is employed, combined with a value function loss term and entropy regularization to maintain exploration. Network architectures utilize fully connected layers with ReLU activations, featuring separate actor and critic networks. The algorithm demonstrates consistent convergence across most HydroGym flow configurations, successfully learning control policies for drag reduction tasks, cavity flow stabilization, and multi-body coordination in the fluidic pinball. Training hyperparameters are adapted per environment, with learning rates ranging from 5×10^{-5} to 1×10^{-3} , batch sizes of 16-48 experiences, and episode lengths matching the physical timescales of each flow configuration.

4B Deep Deterministic Policy Gradient (DDPG)

DDPG extends the deterministic policy gradient theorem to high-dimensional continuous action spaces through an actor-critic architecture combined with experience replay and target networks [109]. The algorithm involves four neural networks: an actor network $\mu(s|\theta^\mu)$ that deterministically maps states to actions, a critic network $Q(s, a|\theta^Q)$ that estimates action-value functions, and corresponding networks μ' and Q' to provide stable learning targets [110]. The critic is trained using temporal difference learning with the Bellman equation:

$$L(\theta^Q) = \mathbb{E}_{s,a,r,s' \sim \mathcal{D}} [(Q(s, a|\theta^Q) - y)^2] \quad (29)$$

where $y = r + \gamma Q'(s', \mu'(s'|\theta^{\mu'})|\theta^{Q'})$ represents the target value computed using target networks. The actor network is updated using the deterministic policy gradient [111]:

$$\nabla_{\theta^\mu} J \approx \mathbb{E}_{s \sim \mathcal{D}} [\nabla_{\theta^\mu} \mu(s|\theta^\mu) \nabla_a Q(s, a|\theta^Q)|_{a=\mu(s)}] \quad (30)$$

Exploration during training is achieved through additive Ornstein-Uhlenbeck noise $\mathcal{N}$ with parameters $\mu_{OU} = 0.15$, $\sigma = 0.2$, providing temporally correlated exploration in the action space [112]. Target networks are updated using soft updates with $\tau = 0.005$ to ensure stable learning dynamics. We employ replay buffers of size 10^4 to 10^5 experiences, with batch sampling of 256 transitions per update step. Network architectures feature 256-512 hidden units across two hidden layers, utilizing batch normalization for improved training stability [113].

4C Twin Delayed Deep Deterministic Policy Gradient (TD3)

TD3 addresses the overestimation bias inherent in DDPG through three key modifications: twin critic networks, delayed policy updates, and target policy smoothing [114]. The twin critic architecture maintains two independent Q-functions Q_{θ_1} and Q_{θ_2} , taking the minimum estimate for target value computation to reduce overestimation [115]:

$$y = r + \gamma \min_{i=1,2} Q_{\theta'_i}(s', \tilde{a}) \quad (31)$$

where $\tilde{a} = \mu'(s') + \text{clip}(\epsilon, -c, c)$ represents the smoothed target action with clipped noise $\epsilon \sim \mathcal{N}(0, \sigma)$. Policy updates occur less frequently than critic updates (typically every second iteration) to allow critics to converge before policy modification. This delayed update mechanism prevents policy updates from exploiting temporary critic errors. Target policy smoothing adds regularized noise to target actions during critic training, smoothing the value function over similar actions and reducing the impact of function approximation errors [114].

The complete TD3 update procedure alternates between critic training using both networks and periodic actor updates using the first critic network only. Hyperparameter selection follows $\sigma = 0.2$ for target noise, clipping bounds $c = 0.5$, and policy delay factor of 2.

4D Gradient-Enhanced Proximal Policy Optimization (GPPO)

GPPO represents a hybrid approach that integrates analytic gradient information from differentiable fluid simulations into the standard PPO framework [9]. This method exploits the differentiability of JAX-based flow solvers to compute exact policy gradients with respect to the flow dynamics, providing more informative gradient signals than traditional finite-difference estimates or likelihood ratio methods [116].

The algorithm combines analytic gradients $\nabla_{\theta} \mathbb{E}[R(\tau)]$ computed directly through the differentiable simulator with PPO's clipped surrogate objective. The enhanced gradient computation leverages automatic differentiation through the entire simulation trajectory [117]:

$$\nabla_{\theta} J = \nabla_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T R(s_t, a_t) \right] \quad (32)$$

where θ represents the tunable policy parameters, J is the objective, R is the reward, and the expectation gradient is computed analytically through the deterministic fluid dynamics, eliminating variance associated with policy gradient estimation. To provide a more clear example, consider the 3D channel flow, where the objective is to minimize the wall-shear stress, τ_w , and the reward signal was a function of τ_w . At a point in time t , the environment backward pass can compute:

$$\frac{\partial \tau_w}{\partial \theta} = \frac{\partial \tau_w}{\partial s_t} \frac{\partial s_t}{\partial \theta} + \frac{\partial \tau_w}{\partial a_t} \frac{\partial a_t}{\partial \theta} \quad (33)$$

which can be included in the optimization by adding τ_w to the loss term. An alternative approach to leveraging gradient information is to incorporate it directly into the reward function. Taking into account this specific objective, the wall shear stress gradient with respect to a previous action can be computed by a backward pass of the environmental step, giving $\nabla_{a_t} \tau_w$. This value can then augment the reward function as follows:

$$R(s_t, a_t) = -[\tau_w + (\nabla_{a_t} \tau_w \cdot a_t)], \quad (34)$$

which can guide the policy towards actions that minimize future wall shear stress. These methods were applied to both the 3D channel flow and the 2D Kolmogorov flow, and demonstrated both sample efficiency and improved performance in a variety of flow scenarios.

The implementation of GPPO maintains PPO’s trust region constraints while incorporating these enhanced gradients through a weighted combination scheme. The gradient weighting balances analytic precision with exploration requirements, using adaptive coefficients based on gradient magnitude and policy entropy. Network architectures are specifically designed for JAX compatibility, utilizing Flax layers with consistent initialization schemes across differentiable and non-differentiable components [118].

5 Environments

The HydroGym platform provides a comprehensive suite of fluid dynamics environments specifically designed for reinforcement learning benchmarking and active flow control research. These environments span a diverse range of canonical flow configurations that capture fundamental fluid mechanics phenomena while maintaining computational tractability for iterative RL training processes. The framework encompasses 61 distinct flow scenarios across 14 different geometric configurations, ranging from two-dimensional cylinder flows to complex three-dimensional turbulent channel cases. The environments are systematically organized to cover multiple Reynolds number regimes within each configuration, enabling the investigation of flow control strategies across different physical regimes—from laminar vortex shedding at lower Reynolds numbers to transitional and turbulent dynamics at higher values.

The benchmark suite includes both classical bluff body flows (cylinders, squares, cubes, spheres) that exhibit well-documented vortex shedding phenomena as well as more specialized configurations such as cavity flows with shear layer instabilities, NACA 0012 airfoils under gust interactions, shock-vector control for divergent nozzle flows, extreme-event-driven Kolmogorov flows and biomedically-relevant stenotic pipe flows. Additionally, the framework incorporates multi-agent reinforcement learning (MARL) scenarios for spatially-distributed control strategies and differentiable physics environments that enable gradient-based optimization approaches.

Each environment follows a standardized framework that includes configurable reward functions targeting specific control objectives (typically drag reduction and flow stabilization), flexible actuation mechanisms (jet-based control, surface rotation and displacement, or temperature modulation), and adaptable sensor configurations through point probe measurements of flow quantities. The temporal coupling between CFD solvers and reinforcement learning agents is carefully calibrated to match characteristic flow timescales while maintaining computational efficiency for training.

The systematic organization of these environments provides researchers with a progressive pathway from fundamental flow control problems to increasingly complex multi-physics scenarios, supporting both algorithm development and practical flow control applications. The following subsections provide detailed descriptions of each environment category, including the underlying flow physics, numerical implementation specifics, validation against literature benchmarks, and the specific RL environment setup including state spaces, action spaces, and reward formulations.

5A Circular cylinder flow

Characteristic physics. The flow around circular cylinders in the Reynolds number range 200-3,900 based on the cylinder diameter and the free-stream velocity exhibits rich dynamical behavior characterized by distinct regime transitions and complex three-dimensional instabilities that have profound implications for both fundamental fluid mechanics research and practical flow control applications. This transitional range encompasses the evolution from two-dimensional laminar vortex shedding to fully three-dimensional

turbulent wakes, presenting unique challenges for modeling, prediction, and control.

The flow physics within this range is governed by two critical transitions that fundamentally alter wake dynamics. The first occurs at Reynolds numbers around 190-200, where the initially two-dimensional von Kármán vortex street undergoes its first three-dimensional instability, known as Mode A. This elliptic instability manifests as large-scale spanwise undulations with characteristic wavelengths of approximately 3-4 cylinder diameters, originating in the primary vortex cores and representing the onset of three-dimensional wake behavior [119, 120]. As the Reynolds number continues to increase beyond $Re = 230 - 260$, a second transition emerges through Mode B instability. Unlike Mode A, this hyperbolic instability develops in the braid shear layers and exhibits much smaller spanwise wavelengths of approximately 0.8-1.0 cylinder diameters [121, 122]. Together, these successive transitions mark the beginning of a complex regime where two-dimensional approximations become progressively inadequate for accurate flow prediction. These three-dimensional instabilities create fundamental differences between 2D and 3D flow configurations that extend well beyond simple geometric considerations. While two-dimensional simulations remain computationally efficient and provide reasonable approximations for basic flow characteristics below $Re = 189$, they fundamentally cannot capture the governing physics once three-dimensional instabilities develop [123]. That is, the spanwise variations introduced by Mode A and Mode B instabilities significantly affect drag coefficients, vortex shedding frequencies, and heat transfer characteristics, with these effects becoming increasingly dominant as the Reynolds number increases. In the subcritical regime ($Re > 1,000$), the wake transitions to fully turbulent characteristics while maintaining coherent large-scale vortex shedding patterns. This creates a complex interplay between organized and chaotic motion that further challenges traditional modeling approaches [124, 125].

Overall, the rich dynamical behavior has established cylinder wakes as particularly valuable canonical test cases for developing and validating flow control strategies. The combination of the relatively simple geometry with complex physics provides an ideal benchmark for both traditional control methods and emerging machine learning approaches. Recent developments in deep reinforcement learning have demonstrated remarkable success, achieving drag reductions of 8 – 25% using various actuator configurations including synthetic jets, rotational oscillation, and plasma actuators [15, 17]. The Reynolds number range 200-3,900 proves especially well-suited for these applications because the flow exhibits sufficient complexity to challenge control algorithms while remaining computationally tractable for iterative training processes.

However, the practical implementation of flow control in this Reynolds number range faces several specific challenges stemming from the complex transitional flow physics. Three-dimensional instabilities introduce spanwise variations that complicate sensor placement and actuator design, while the coexistence of laminar and turbulent regions creates difficulties in developing robust control strategies. Energy efficiency considerations become particularly critical, as the power required for active control can exceed the benefits gained through drag reduction, especially at lower Reynolds numbers within this range [126]. Furthermore, the sensitivity of the wake behavior to boundary conditions and geometric details makes experimental validation of numerical control studies challenging, often requiring careful attention to computational domain size and boundary condition implementation. To address these challenges, multi-agent deep reinforcement learning frameworks have been designed to target the three-dimensional control problems, though most successful demonstrations remain limited to lower Reynolds numbers within the transitional range [29].

To summarize, the established nature of cylinder wake physics makes this configuration particularly valuable as emerging computational methods and experimental techniques continue to provide new insights into flow control possibilities. The comprehensive understanding of fundamental mechanisms enables researchers to confidently assess whether novel approaches capture the essential physics or merely exploit computational artifacts [127–129]. Therefore, the well-documented cylinder wake serves as a valuable benchmark scenario for validating new innovations before applying them to more complex geometries and flow physics.

(a) $Re = 200$				(b) $Re = 1,000$				(c) $Re = 3,900$			
Reference	St	$\overline{C_D}$	$\overline{C_L}$	Reference	St	$\overline{C_D}$	$\overline{C_L}$	Reference	St	$\overline{C_D}$	$\overline{C_L}$
Present study	0.196	1.341	± 0.659	Present study	0.234	1.533	± 1.449	Present study	0.21	1.68	± 1.61
Meneghini et al.[130]	0.196	1.30	–	Colagrossi et al.[135]	0.231	1.507	± 1.457	Chatzimanolakis et al.[16]	–	1.75	–
Ding et al.[131]	0.196	1.348	± 0.659	Durante et al.[136]	0.235	1.491	± 1.46	Durante et al.[136]	0.19	1.98	± 1.73
Braza et al.[132]	0.200	1.40	± 0.75	Scott & Durst[137]	0.2276	1.54	± 1.225	(Re = 1×10^4)			
Posdziech & Grundmann[133]	0.194	1.311	± 0.66	Chatzimanolakis et al.[16]	–	1.61	–				
Harichandan & Roy[134]	0.192	1.32	± 0.602	von Wahl et al.[138]	0.23	1.53	–				

Table 1: Validation of 2D cylinder flow results at different Reynolds numbers. St: Strouhal number; $\overline{C_D}$: Time-averaged drag coefficient; $\overline{C_L}$: Time-averaged lift coefficient; $\pm$: Oscillation amplitude; –: Data not reported or not applicable

Numerical setup and validation. The cylinder benchmark scenario comprises 2D and 3D low Mach number flows ($Ma < 0.2$) over circular cylinders at different Reynolds number, i.e. $Re \in [100, 200, 1,000, 3,900]$. For the 2D cases, the computational domain is defined by $(x/D, y/D) \in [0, 40] \times [-10, 10]$ with the cylinder being centered at $(x/D, y/D) = (10, 0)$ and x representing the streamwise coordinate direction. The grid spacing is determined by the grid refinement level n via $\Delta_n = \Delta_0/2^n$ with $\Delta_0^{(2D)} = 40D$. Three refinement patches decrease the grid spacing up to a level of $n = 14$, see Fig. 1, corresponding to more than 320 cells per cylinder diameter.

In 3D the domain is $(x/D, y/D, z/D) \in [0, 32] \times [0, 16] \times [-2, 2]$, being periodic in the z -direction while the center of the cylinder is located at $(x/D, y/D) = (8, 8)$ and its axis is aligned in z -direction. The grid spacing is defined by $\Delta_0^{(3D)} = 32D$ and refinement levels of $n \in [9, 14]$ resolving the cylinder by 500 grid points.

At the inlet, a Dirichlet boundary condition is implemented with a prescribed constant velocity profile. The velocity vector is read from the properties file and applied uniformly across the inlet cross-section. The density at the inlet boundary cells is extrapolated from the interior domain using second-order accurate extrapolation to maintain consistency with the low Mach number assumption. The equilibrium distribution functions are calculated based on the prescribed velocity and extrapolated density values.

At the outlet, non-reflecting boundary conditions are applied minimizing spurious wave reflections. The velocity components are extrapolated from the interior domain using von Neumann conditions (zero normal gradient). The density is computed using a non-reflecting formulation:

$$\rho = \frac{\rho_{\text{old}} + \rho_{\text{ref}}/c_s(|\vec{v}| - |\vec{v}_{\text{old}}|) + \alpha\Delta t\rho_{\text{ref}}}{1 + \alpha\Delta t} \quad (35)$$

where ρ_{old} and $\vec{v}_{\text{old}}$ are values from the previous time step, c_s is the speed of sound, ρ_{ref} is a reference density term that can be specified as either density fluctuations or absolute density depending on the simulation requirements, and $\alpha = 1.0$ is the relaxation parameter.

Periodic boundary conditions are applied to the lateral boundaries (side walls) of the computational domain. This treatment ensures mass conservation and eliminates the need for additional boundary condition specifications at these surfaces, effectively creating an infinite domain in the lateral directions.

Based on a comprehensive validation study, our numerical simulations demonstrate good agreement with established literature values across multiple Reynolds numbers and flow configurations. For 2D cylinder flows, our results closely match reference data, with Strouhal numbers within 1 – 3% of published values and drag coefficients typically within 5% of established benchmarks. The 3D simulations show similarly strong validation, with our Strouhal numbers and drag coefficients falling well within the scatter of literature values across all tested Reynolds numbers. The close correspondence between our computed flow parameters and the extensive reference database confirms the accuracy and reliability of our numerical approach for cylinder flow simulations.

(a) $Re = 200$				(b) $Re = 1,000$				(c) $Re = 3,900$			
Reference	St	$\overline{C_D}$	$C_{L'}$	Reference	St	$\overline{C_D}$	$C_{L'}$	Reference	L_z	St	$\overline{C_D}$
Present study	0.192	1.318	0.418	Present study	0.213	1.102	0.374	Present study	4.0	0.217	1.074
Qu et al.[139]	0.183	1.25	0.35	Tong et al.[145]	0.215	1.08	0.2	Suárez et al.[28]	π	0.220	1.08
Posdziech & Grundmann[140]	0.182	1.24	–	Lei et al.[146]	–	1.112	0.325	Lehmkuhl et al.[148]	π	0.215	1.015
Carmo & Meneghini[141]	0.184	1.28	–	Labbé & Wilson[142]	0.215	1.200	–	Parnaudeau et al.[124]	23	0.208	–
Labbé & Wilson[142]	0.195	1.318	–	Zhao et al.[147]	0.21	1.17	0.335	Tremblay et al.[149]	π	0.22	1.03
Behara & Mittal[143]	0.195	1.38	–					Kravchenko & Moin[125]	π	0.21	1.04
Rajani et al.[144]	0.1936	1.338	0.4276								

Table 2: Comparison of 3D cylinder flow results at different Reynolds numbers. St: Strouhal number; $\overline{C_D}$: Time-averaged drag coefficient; $C_{L'}$: RMS lift coefficient; L_z : Span length; L_r/D : Recirculation length to diameter ratio; –: Data not reported or not applicable

Environment setup. The default reward objective r for the cylinder flow environment targets drag minimization and the reduction of lift force oscillations and reads $r = -|C_D| - \omega|C_L|$, where C_D, C_L denote the normalized drag and lift coefficients, and ω is a Reynolds number-dependent scaling factor with values of 0.1, 0.2, 1.0, and 1.0 for $Re = 100, 200, 1,000, 3,900$, respectively. Since the unstable symmetric state is also the minimum-drag configuration, this objective corresponds to stabilizing this unstable fixed point. This reward formulation serves as the default configuration, though additional reward shaping terms can be easily implemented by users according to their specific research objectives.

Two different actuation strategies are implemented to provide the RL agent with control authority over the flow dynamics. First, a zero-net mass-flux blowing/suction strategy is provided, comprised of two equal and opposite jets - one at the top and one at the bottom of the cylinder. The RL agent controls the mass flux of the upper jet, which automatically specifies the lower jet through the zero-net-mass constraint. In the default configuration, the jets are placed at 90° and 270° positions and cover 10° of the cylinder circumference each, though the HydroGym framework is designed to allow for flexible user-specified adjustments of the jet number, placement angles, and coverage areas. Second, a rotating surface actuator is implemented for which the RL agent directly controls the angular velocity of the cylinder surface. Both action spaces are normalized to the range $[-1, 1]$ to ensure consistent RL algorithm performance across different actuation modes.

The environment supports both two-dimensional and three-dimensional cylinder flow configurations across multiple Reynolds numbers ($Re = 100, 200, 1,000, 3,900$), enabling the investigation of flow control across different flow regimes from laminar vortex shedding to fully turbulent conditions. For three-dimensional cases at $Re = 1,000$ and $3,900$, the framework additionally supports multi-agent reinforcement learning through spanwise domain decomposition, where the cylinder is divided into 4 and 10 subcylinder segments respectively, each individually controlled by a dedicated agent. This multi-agent approach enables the investigation of cooperative control strategies and spanwise flow coordination. The framework provides flexibility in multi-agent coordination, allowing agents to share gradient information and experience buffer data if desired, though independent agent operation is also supported based on user requirements.

The state space representation consists of user-configurable point probes distributed throughout the computational domain, providing measurements of velocity components, pressure, density, and body forces at specified locations. By default, the state observations consist of the raw solver output, but different normalization strategies are implemented to normalize the state space in the range $[-1, 1]$ for consistent RL training performance. This flexible sensor configuration allows the investigation of minimal sensing strategies and optimal sensor placement for effective flow control.

Each training episode spans 200 control actions, corresponding to approximately 10 vortex shedding periods, providing sufficient time for the agent to observe and respond to flow dynamics. The temporal coupling between the CFD solver and RL agent employs Reynolds number-specific action frequencies: 200, 350, and 600 CFD timesteps per control action for 2D cases at $Re = 200, 1,000, 3,900$ respectively, and 400,




Figure 1: Computational domain setup, sensor probe distribution and actuation strategy for the circular cylinder flows: Three-dimensional view showing the probe distribution around the cylinder setup with vortical structures visualized through Q-criterion isosurfaces. The yellow dots in the zoom-in highlight sensor probe locations in a single spanwise plane. Active flow control is either performed using jet actuators positioned at the top and bottom of the cylinder or via surface rotation.

500, and 1,200 timesteps for corresponding 3D cases. These ratios are designed to match the characteristic timescales of vortex shedding dynamics while maintaining computational efficiency and control authority.

The environment framework provides a comprehensive testbed for investigating reinforcement learning-based active flow control strategies, supporting both fundamental research into control mechanisms and practical development of sensing and actuation approaches for cylinder wake management. All described configurations represent default settings that can be easily modified by users to accommodate specific research requirements, including custom reward functions, alternative actuation strategies, modified sensor placements, and adjusted temporal coupling parameters.

5B Fluidic pinball flow

Characteristic physics. The fluidic pinball serves as an exemplary test case for investigating complex wake dynamics and flow control strategies in transitional flow regimes. This benchmark configuration consists of three equal-diameter circular cylinders positioned at the vertices of an equilateral triangular arrangement, with one vertex directed toward the incoming flow [150]. Although geometrically simple, this arrangement demonstrates a diverse spectrum of nonlinear flow behaviors, including steady-state solutions, time-periodic oscillations, quasi-periodic motion, and fully chaotic dynamics across different operating parameters.

The fluidic pinball exhibits distinct flow regime transitions as the Reynolds number varies from 30 to 150, with critical thresholds that differ from classical single-cylinder wake behavior. Initial investigations by [151]

identified specific transition points unique to this multi-cylinder configuration. For Reynolds numbers below $Re \lesssim 18$, the flow field maintains global stability with symmetric recirculation zones forming downstream of the cylinder cluster. The transition from this steady state occurs via a supercritical Hopf bifurcation at approximately $Re = 18$ [151]. This early onset of instability results from the enhanced wake interactions between the three cylinders. The Hopf bifurcation introduces time-periodic vortex shedding characterized by complex conjugate eigenvalues crossing the stability boundary in the linearized flow equations. The resulting flow exhibits von Kármán-type vortex shedding with dimensionless frequency (Strouhal number) of approximately 0.2. Importantly, this initial instability preserves the mirror-plane symmetry about the configuration centerline, maintaining symmetric shedding patterns until $Re \approx 68$.

At $Re \approx 68$, the fluidic pinball undergoes a second critical transition through a supercritical pitchfork bifurcation [152]. This transition fundamentally alters the wake dynamics by breaking the mirror symmetry preserved during the initial periodic regime. The physical mechanism involves the destabilization of the jet flow between the two downstream cylinders, which becomes susceptible to preferential deflection in either the upward or downward direction. This symmetry-breaking creates two mirror-symmetric stable solutions - one with upward jet deflection and another with downward deflection. The system exhibits bistable behavior, where initial conditions determine which asymmetric state the flow adopts [152]. This represents a qualitative change from the previous symmetric periodic shedding to asymmetric limit cycle dynamics.

Beyond $Re = 104$, the fluidic pinball system develops additional complexity through a secondary Neimark-Sacker bifurcation [153]. This mathematical bifurcation introduces quasi-periodic motion by creating torus-like attractors in the system's phase space. The primary vortex shedding frequency becomes modulated by secondary low-frequency oscillations, establishing multi-frequency dynamics characteristic of weakly chaotic flows. The transition to fully chaotic motion occurs above $Re \approx 130$, following what was identified by [153] as a Newhouse-Ruelle-Takens scenario. In this Re number regime, the quasi-periodic torus structure becomes unstable and breaks down, leading to aperiodic motion with sensitive dependence on initial conditions - the hallmark of deterministic chaos.

Moreover, three-dimensional effects can introduce additional complexity beyond the two-dimensional analysis. Following the foundational work of [119] on cylinder wake transitions, the fluidic pinball exhibits spanwise instabilities originating from hyperbolic regions within the braid shear layers. These instabilities manifest as short-wavelength structures with characteristic spanwise dimensions comparable to the cylinder diameter. The three-dimensional modes create streamwise vortical structures that interact with the primary cross-stream vortex shedding. Multiple instability modes can coexist and compete, leading to complex scenarios including cooperative destabilization and vortex reconnection phenomena that cannot be captured by two-dimensional computational models.

To summarize, the fluidic pinball's value as a benchmark derives from the combination of geometric simplicity and dynamical richness, making it suitable for testing advanced control strategies while remaining computationally tractable. The independent rotation capability of each cylinder provides multiple actuation inputs, enabling implementation of various flow control mechanisms established in the literature. Classical flow control approaches applicable to this configuration include boundary layer manipulation through surface blowing (analogous to base bleeding techniques), high-amplitude periodic forcing for vortex synchronization, and circulation control via cylinder rotation utilizing Magnus force effects. Additionally, modern approaches such as opposition control and model-based feedback strategies can be evaluated.

Hence, the fluidic pinball presents several fundamental challenges for flow control design. The succession of bifurcations means that control strategies must adapt to qualitatively different flow physics as operating conditions change. Systems designed for the symmetric periodic regime may prove ineffective or counterproductive when applied to the asymmetric bistable regime. The multi-frequency dynamics create coupling between different temporal scales, where actuation at one frequency can influence the entire spectral content of the flow response. This cross-frequency coupling necessitates control architectures that account for broadband nonlinear interactions rather than targeting individual frequency components. Furthermore, the

presence of multiple coexisting attractors (in the bistable regime) and hysteresis effects between flow states require control strategies capable of managing state transitions and maintaining desired operating points despite external disturbances.

Extension to three-dimensional control introduces requirements for spanwise-distributed actuation to address the various three-dimensional instability modes. The transition between different three-dimensional regimes creates scenarios where two-dimensional control approaches become inadequate, necessitating fundamentally different control frameworks. Control strategies optimized for one three-dimensional mode may lose effectiveness or become destabilizing when the flow transitions to alternative modal states. This highlights the importance of robust adaptive control approaches that maintain performance across regime boundaries.

Numerical setup and validation. The fluidic pinball configuration consists of three circular cylinders arranged in an equilateral triangle formation, creating a fundamentally different flow topology compared to the single cylinder case. Relative to the pinball center at $(x/D, y/D) = (10.0, 0.0)$ for the 2D simulations and $(8.0, 0.0)$ for the 3D counterpart, the cylinders are positioned with centers at $(x/D, y/D) = (-1.299, 0.0)$, $(0.0, 0.75)$, and $(0.0, -0.75)$, where the upstream cylinder interacts with two downstream cylinders to generate complex wake dynamics and vortex interactions.

The computational setup employs the same domain dimensions and boundary condition treatments as described for the single cylinder benchmark. However, the fluidic pinball operates at lower Reynolds numbers $Re \in [30, 100, 150]$, allowing for a uniform refinement strategy with $n = 12$ across all cases. This refinement level provides adequate resolution to capture the intricate flow features arising from the multi-cylinder interaction, including wake merging, vortex pairing, and the characteristic frequency modulation effects.

The triangular arrangement leads to inherently asymmetric flow patterns, where the downstream cylinders experience different flow conditions due to their positioning relative to the upstream cylinder's wake. This asymmetry manifests in the lift coefficient distributions and requires careful validation against established benchmarks to ensure accurate representation of the flow physics. Validation results are summarized in Table 3, demonstrating good agreement with the reference data of Wang et al. [154]. The simulations accurately predict the natural frequencies across all Reynolds numbers. The time-averaged drag coefficients show good agreement (within 2% deviation), while the asymmetric lift coefficients on the downstream cylinders are well-captured, confirming a proper resolution of the wake interaction dynamics. The 3D results maintain similar accuracy levels, validating the three-dimensional extension of the fluidic pinball configuration.

Re	Reference	Dim	f_0	$\overline{C_D}$	$\overline{C_{L,2}}$	$\overline{C_{L,3}}$
30	Present study	2D	0.064	2.651	-0.361	0.361
	Present study	3D	0.061	2.597	-0.358	0.359
	Wang et al.[154]	2D	0.065	2.656	-0.365	0.365
100	Present study	2D	0.088	2.904	-0.079	0.110
	Present study	3D	0.087	2.795	-0.0107	0.119
	Wang et al.[154]	2D	0.090	2.874	-0.0745	0.108
150	Present study	2D	0.120	2.922	-0.0441	0.0330
	Present study	3D	0.113	2.949	-0.0467	0.0297
	Wang et al.[154]	2D	0.118	2.960	-0.043	0.0319

Table 3: Validation of the fluidic pinball flow results at different Reynolds numbers. Dim: Dimensionality (2D/3D); f_0 : Natural frequency; $\overline{C_D}$: Time-averaged drag coefficient; $\overline{C_{L,i}}$: Time-averaged lift coefficient for cylinder i

Environment setup. The fluidic pinball environment features a multi-cylinder flow control setup comprising three circular cylinders positioned in an equilateral triangular arrangement. The primary reward function r for the fluidic pinball targets total drag reduction across all three cylinders and is defined as $r = -|\sum_{i=1}^3 C_{D,i}| - \omega|\sum_{i=1}^3 C_{L,i}|$, where $C_{D,i}$ represents the normalized drag coefficient of cylinder i and $\omega = 1.0$ serves as the scaling parameter consistently applied across all Reynolds numbers to avoid control policies that introduce




Figure 2: Computational domain setup, sensor probe distribution and actuation strategy for the fluidic pinball environment: Three-dimensional view showing the probe distribution around the fluidic pinball with vortical structures visualized through Q-criterion isosurfaces. The yellow dots in the zoom-in highlight sensor probe locations in a single spanwise plane. Active flow control is performed by individual surface rotation of each cylinder.

a significant lift generation. Users can readily adapt this reward structure to explore different objectives including thrust control or independent cylinder optimization.

The control mechanism employs surface rotation actuation, where the RL agent manages the angular velocity of each cylinder individually (see Fig. 2). This creates a three-dimensional continuous action space with components bounded within $[-1, 1]$, facilitating exploration of cooperative multi-cylinder control approaches. Although jet-based actuation is not part of the standard setup, the framework design accommodates user integration of blowing/suction mechanisms for extended research applications or comparative analysis.

The framework encompasses both two-dimensional and three-dimensional fluidic pinball flow scenarios across various Reynolds numbers ($Re = 30, 75, 100, 150$), covering flow regimes from steady symmetric configurations through periodic vortex shedding to bistable asymmetric behavior and quasi-periodic dynamics. This Reynolds number spectrum encompasses the distinctive bifurcation characteristics of the fluidic pinball, including the initial Hopf bifurcation at $Re \approx 18$, the symmetry-breaking pitchfork bifurcation at $Re \approx 68$, and progressive transitions toward chaotic flow patterns. The setup utilizes a single-agent configuration, where one agent manages all three cylinder rotations concurrently, promoting research into unified control methodologies for multi-body flow applications.

The observation space utilizes the same adaptable probe-based methodology as established in other environments, featuring user-definable point measurements of velocity components, pressure, density, and body forces positioned throughout the computational domain. Training episodes encompass 200 control actions, equivalent to roughly 10 vortex shedding cycles, offering an adequate temporal scope for agents

to perceive and react to the intricate multi-frequency behavior inherent in the fluidic pinball. The CFD-RL coupling utilizes Reynolds number-dependent action intervals tailored to match evolving flow characteristics: 230, 190, 170, and 100 CFD timesteps per control action for 2D scenarios at $Re = 30, 75, 100, 150$ respectively, with corresponding 3D cases using 250, 235, 225, and 200 timesteps. These intervals adapt to the changing temporal scales and flow complexity encountered as the system progresses through distinct dynamical states, maintaining suitable control responsiveness while preserving computational tractability. All outlined specifications serve as baseline configurations that users can readily adjust to meet particular research needs, encompassing modified reward structures, alternative control mechanisms, customized sensor layouts, and tailored temporal coupling settings.

5C Open cavity flow

Characteristic physics. The flow around open cavities in the Reynolds number range 4,000-7,500 based on the cavity height and the freestream velocity exhibits complex transitional behavior characterized by distinct regime evolutions and sophisticated three-dimensional instabilities that have profound implications for both fundamental aeroacoustics research and practical flow control applications. This transitional range encompasses the evolution from predominantly steady two-dimensional recirculation patterns to complex three-dimensional unsteady dynamics with strong acoustic coupling, presenting unique challenges for modeling, prediction, and control that distinguish it from both lower and higher Reynolds number regimes.

Within this Reynolds number range, the flow physics is governed by several critical transitions that fundamentally alter the cavity dynamics and acoustic characteristics. At Reynolds numbers around 4,000-5,000, the initially stable two-dimensional recirculation undergoes its first significant three-dimensional instability through centrifugal mechanisms arising from curved streamlines in the recirculation region [155, 156]. This transition manifests as weakly unsteady flow behavior with emerging spanwise variations that mark the onset of complex temporal dynamics. As the Reynolds number continues to increase, a second critical transition occurs around $Re = 7,000-8,000$, where strong temporal instabilities emerge through Hopf bifurcation characteristics, introducing multiple competing temporal modes and mode-switching phenomena where different frequencies compete for energy [157]. These successive transitions create a regime where simplified two-dimensional approximations become progressively inadequate for an accurate prediction of both aerodynamic and acoustic behavior.

The emergence of these transitional instabilities creates fundamental differences between 2D and 3D cavity flow configurations that extend far beyond geometric considerations. While two-dimensional simulations remain computationally efficient and provide valuable insights into fundamental feedback mechanisms, they systematically overpredict cavity tone amplitudes by up to 15 dB compared to three-dimensional cases and miss critical mixing processes that govern realistic flow physics [158]. The three-dimensional effects introduce end-wall influences that cause flow spillage over cavity sides, development of Taylor-Görtler-like vortices, and breakdown of coherent shear layer structures due to spanwise instabilities. In compressible flows within this Reynolds number range, the acoustic feedback mechanisms become particularly complex, with the classical Rossiter formula requiring depth-sensitized modifications to account for three-dimensional cavity effects and modified convection velocities [159, 160].

This rich transitional behavior makes cavity flows particularly valuable as canonical test cases for developing and validating advanced flow control strategies, especially those employing machine learning approaches. The geometric simplicity combined with complex physics provides an ideal benchmark that bridges fundamental research and practical aerodynamic applications. Various frameworks have provided theoretical foundations for understanding the underlying bifurcation mechanisms and sensitivity characteristics that govern cavity flow transitions [155]. However, the practical implementation of flow control in this Reynolds number range faces several specific challenges that stem from the complex transitional flow physics and multi-frequency acoustic coupling. The strong nonlinear coupling between Rossiter modes creates frequency

cross-talk and triadic interactions that complicate traditional linear control approaches, while mode-switching regimes introduce time-varying dynamics that challenge adaptive control systems [161]. Multiple time scales ranging from fast Kelvin-Helmholtz instabilities to slow mode-switching phenomena create complex control requirements, with real-time implementation demanding computational capabilities supporting control frequencies up to 10 kHz. Energy efficiency considerations become particularly critical in practical applications, as the spatially disjoint measurement and control locations inherent in cavity flows introduce time-delay effects that can destabilize feedback control systems if not properly addressed [162]. As a result, the rich transitional physics and computational accessibility make cavity flows in this Reynolds number range an enduring benchmark for flow control innovation, particularly as machine learning approaches mature and computational resources continue to expand.

Numerical setup and validation. The open cavity benchmark comprises 2D and 3D low Mach number flows ($Ma < 0.2$) over square cavities at Reynolds numbers $Re_D \in [4, 200, 7, 500]$, based on the cavity depth D and free-stream velocity U_∞ , following the configuration of Barbagallo et al. [157].

For the 2D case, the computational domain is defined by $(x/D, y/D) \in [0.0, 2.5] \times [0, 1.5]$ with the square cavity located at $(x/D, y/D) \in [0.27, 1.27] \times [0, 1]$. The grid spacing is determined by the grid refinement level n via $\Delta_n = \Delta_0/2^n$ with $\Delta_0^{(2D)} = 2.5D$. Multiple refinement patches are employed with maximum refinement levels of $n = 12$ near the cavity edges and shear layer regions, corresponding to approximately 227 cells per cavity depth (see Fig. 4). Special attention is given to resolving the boundary layer upstream of the cavity, targeting a displacement thickness of $\delta_1 \approx 0.012D$. This numerical setup demonstrates good agreement with established literature [157] as shown in Fig. 3. In 3D, the domain extends to $(x/D, y/D, z/D) \in [0.0, 2.5] \times [0, 1.5] \times [0, 2.25]$ with periodic boundary conditions applied in the spanwise z -direction. A similar grid compared to the 2D test case is used to maintain adequate resolution of three-dimensional instabilities while keeping computational costs manageable.

At the inlet, a Dirichlet boundary condition implements a uniform velocity profile $(u, v, w) = (U_\infty, 0, 0)$. The top boundary implements symmetry conditions ($\partial_y u = 0, v = 0, \partial_y w = 0$). No-slip conditions ($u = v = w = 0$) are applied on the cavity walls and the bottom surface from $x/D = 0.1777$ onwards, while symmetry conditions are used for $x/D < 0.1777$ to generate appropriate boundary layer development.

Environment setup. The open cavity environment features a flow control setup designed to stabilize the inherent shear layer instability of cavity flows. This configuration provides a consistent framework across both 2D and 3D flow scenarios, enabling comprehensive investigation of the cavity flow dynamics and active control strategies.

The primary objective centers on minimizing flow fluctuations through targeted stabilization of the shear




Figure 3: Time-averaged leading edge velocity profiles for 2D cavity flow at $Re = 7,500$. Present numerical results (solid line) compared with Barbagallo et al. (2009) [157] reference data (crosses). Velocities normalized by free-stream velocity; vertical coordinate normalized by cavity length.




Figure 4: Computational domain setup, sensor probe distribution and actuation strategy for the open cavity flow: Three-dimensional view showing the probe distribution in the open cavity environment with vortical structures visualized through Q-criterion isosurfaces. The yellow dots highlight sensor probe locations in a single spanwise plane. Active flow control is performed using up to three independent jet actuators positioned at the leading edge and inside the cavity.

layer that develops across the cavity opening. The reward function r quantifies this stabilization goal as:

$$r = - \sum_i \left(\frac{\text{obs}_i - \text{target}_{\text{mean},i}}{\text{target}_{\text{std},i}} \right)^2 \quad (36)$$

where obs_i represents the observed quantity at measurement location i , $\text{target}_{\text{mean},i}$ denotes the target reference value, and $\text{target}_{\text{std},i}$ provides normalization scaling. This formulation encourages the control system to maintain observed flow quantities close to their target states, thereby reducing the amplitude of shear layer oscillations and associated cavity resonance phenomena. For simplicity, the target system state is approximated by a temporal mean over 1,000 instability cycles.

The control mechanism employs jet-based actuation strategically positioned to influence the shear layer development and cavity flow dynamics. The standard configuration utilizes a single jet oriented in the y -direction and located at the cavity leading edge, providing direct interaction with the separating boundary layer. For enhanced control authority, an extended configuration incorporates additional jets positioned within the cavity at coordinates $(1.277, 0.75)$ and $(0.277, 0.5)$ that act in streamwise direction, enabling multi-point flow manipulation and improved stabilization capabilities.

The jet actuation creates a continuous action space where the RL agent controls the jet momentum coefficients or velocity ratios. This setup facilitates exploration of both passive flow stabilization through steady jet injection and active control strategies involving time-varying actuation patterns. The framework

design readily accommodates user modifications to jet positioning, orientation, and actuation parameters to investigate alternative control approaches or cavity geometries.

The CFD-RL coupling employs 400 CFD timesteps per control action for 2D test cases and 375 timesteps for 3D counterparts, providing sufficient temporal resolution to capture the essential flow physics while maintaining computational efficiency. This interval balances the need to resolve the characteristic timescales of cavity flow instabilities with practical training considerations, allowing the control agent adequate time to observe the effects of its actions on the flow field evolution. Training episodes are designed to encompass multiple characteristic flow timescales, enabling agents to learn both short-term stabilization responses and longer-term flow management strategies. The environment supports the investigation of various Reynolds numbers and cavity aspect ratios, accommodating research into the Reynolds number dependence of control effectiveness and the geometric sensitivity of stabilization approaches.

5D Square cylinder flow

Characteristic physics. The flow around a square cylinder represents a canonical configuration for studying bluff body aerodynamics, providing fundamental insights relevant to both scholarly inquiry and a wide range of engineering applications, including the design of buildings, bridges, and offshore structures. The interaction of the fluid with the sharp corners of the cylinder forces flow separation, leading to the formation of an unsteady wake dominated by the periodic shedding of large-scale vortices [119]. This alternating vortex shedding generates fluctuating aerodynamic forces on the body, which can induce structural vibrations—a phenomenon termed vortex-induced vibration—posing a significant risk of material fatigue and creating acoustic noise.

The dynamics of the flow are highly dependent on the Reynolds number, and within the range of 200 to 3,900, the wake undergoes a series of complex transitions from a two-dimensional laminar state to a fully three-dimensional turbulent one. The initial transition from a purely two-dimensional to a three-dimensional wake occurs at a critical Reynolds number between approximately 150 and 200 [163, 164]. This transition is marked by the emergence of spanwise instabilities, akin to those observed in circular cylinder wakes. The first of these is the Mode A instability, which appears around $Re \approx 165$ and is characterized by the formation of large-scale streamwise vortex loops with a spanwise wavelength of about four to five times the cylinder diameter [164, 165]. As the Reynolds number increases further, a transition to the Mode B instability occurs around $Re \approx 210$, which involves smaller-scale, rib-like streamwise vortices with a much shorter spanwise wavelength of approximately one cylinder diameter [165]. Beyond these initial instabilities, for Reynolds numbers approaching and exceeding 1,000, the separated shear layers themselves become unstable, leading to the roll-up of Kelvin-Helmholtz vortices which subsequently break down into finer turbulent structures [166]. At higher Reynolds numbers the wake is fully turbulent and characterized by a broad spectrum of interacting eddies [166, 167].

Significant distinctions exist between the flow around square and circular cylinders, primarily stemming from the nature of the flow separation. The sharp corners of the square cylinder fix the separation points, whereas for a circular cylinder, these points are mobile and their location depends on the Reynolds number [168]. This fixed separation generally leads to a wider, more pronounced wake and a higher overall drag coefficient for the square cylinder under identical inflow conditions [169]. While the wake instability modes are qualitatively similar, their onset and development occur at different critical Reynolds numbers. The differences between 2D and 3D configurations are also profound; two-dimensional simulations are incapable of capturing the physics of spanwise instabilities and the resulting energy cascade to smaller turbulent scales. Consequently, 2D models often yield inaccurate predictions of aerodynamic forces and wake dynamics for Reynolds numbers beyond the onset of three-dimensionality [163].

The complex, multi-scale nature of this flow presents considerable challenges for numerical modeling. Reynolds-Averaged Navier-Stokes (RANS) models often fail to accurately predict the highly unsteady and

(a) 2D, $Re = 200$				(b) 2D, $Re = 1,000$				(c) 3D, $Re = 300$			
Reference	St	$\overline{C_D}$	$C_{L'}$	Reference	St	$\overline{C_D}$	$std(C_L)$	Reference	St	$\overline{C_D}$	$C_{L'}$
Present study	0.151	1.469	0.391	Present study	0.121	2.302	1.405	Present study	0.146	1.441	0.222
Sohankar et al.[163]	0.170	1.46	0.32	Yan et al.[23]	0.120	2.34	1.49	Sohankar et al.[163]	0.153	1.47	0.2
Sohankar et al.[171]	0.15	1.462	0.377	Okajima[173]	0.124	–	–	Yoon et al.[177]	0.145	1.431	0.205
Cheng et al.[172]	0.15	1.45	0.372	Norberg[175]	0.124	2.02–2.33	–	Jiang & Cheng[174]	0.146	1.434	0.183
Okajima[173]	0.148	1.45	–	Cao & Tamura[176]	0.123	2.07	–				
Jiang & Cheng[174]	0.152	1.443	0.412	Bai & Alam[164]	0.120	2.10–2.33	–				

Table 4: Comparison of 2D and 3D square cylinder flow results at different Reynolds numbers. St: Strouhal number; $\overline{C_D}$: Time-averaged drag coefficient; $C_{L'}$: RMS lift coefficient; $std(C_L)$: Standard deviation of lift coefficient; –: Data not reported or not applicable

separated flow, particularly the subtle dynamics of vortex shedding and turbulent transition [170]. Scale-resolving approaches, such as Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS), are better suited but are computationally intensive. Achieving accurate results with these methods requires not only extremely fine computational meshes but also large domain sizes and long integration times to ensure that the statistics are fully converged and free from boundary condition influences [166, 171].

These well-defined yet complex flow characteristics make the square cylinder an important test case for the development and validation of flow control strategies, including those leveraging advanced RL techniques. Key challenges for applying RL in this context include the high computational cost of simulations required for training, which can limit the exploration of the solution space, and the persistent difficulty of transferring control policies learned in simulation to real-world experimental setups.

Numerical setup and validation. The square cylinder benchmark follows the same computational setup as the circular cylinder case described above, with the key difference being the bluff body geometry. The square cylinder has a side length D and the Reynolds number is based on this characteristic length, i.e. $Re \in [200, 1,000, 3,900]$.

The computational domains, grid spacing definitions, and refinement strategies remain identical to the circular cylinder configuration. For 2D cases, the domain spans $(x/D, y/D) \in [0, 40] \times [-10, 10]$ with the square cylinder centered at $(x/D, y/D) = (10, 0)$. The 3D domain is $(x/D, y/D, z/D) \in [0, 32] \times [0, 16] \times [-2, 2]$ with the square cylinder centered at $(x/D, y/D) = (8, 8)$ and periodic boundary conditions in the z -direction as shown in Fig. 5.

All boundary condition implementations—inlet Dirichlet conditions with prescribed velocity profiles, non-reflecting outlet conditions, and periodic lateral boundaries—are applied identically to the circular cylinder setup. The grid refinement levels up to $n = 14$ and the resolution criteria also follow the same specifications as detailed in the previous subsection resulting in 512 cells per side length D .

Our validation results for the square cylinder demonstrate comparable accuracy to the circular cylinder cases, with computed Strouhal numbers and drag coefficients showing good agreement with established literature values across available Reynolds numbers in literature (see Tab. 4).

Environment setup. The square cylinder environment follows the same fundamental framework as the circular cylinder configuration described above, with identical reward formulation ($r = -|C_D| - \omega|C_L|$), state space representation through configurable point probes, and episode structure spanning 200 control actions over approximately 10 vortex shedding periods.

The key distinction lies in the actuation strategy, where jet-based control is exclusively employed. Unlike the circular cylinder’s zero-net mass-flux configuration, the square cylinder utilizes individual jets that can




Figure 5: Computational domain setup, sensor probe distribution and actuation strategy: Three-dimensional view showing the probe distribution around the square cylinder with vortical structures visualized through Q-criterion isosurfaces. The yellow dots in zoom-in highlight sensor probe locations in a single spanwise plane. Active flow control is performed through jet actuators positioned at the leading and/or trailing edge.

be positioned at four discrete locations: leading edge upper, leading edge lower, trailing edge upper, and trailing edge lower surfaces. The RL agent controls the mass flux of each active jet independently, with actions normalized to the range $[-1, 1]$ for consistent algorithm performance.

The temporal coupling between the CFD solver and RL agent employs Reynolds number-specific action frequencies: 225, 750, and 375 CFD timesteps per control action for 2D cases at $Re = 200, 1,000, 3,900$ respectively, and 700, 900, and 875 timesteps for corresponding 3D cases. This environment configuration enables the investigation of asymmetric control strategies and the comparative effectiveness of leading versus trailing edge actuation for square cylinder wake management, while maintaining the same comprehensive sensing and reward framework established for the circular cylinder testbed.

5E Flow around a NACA 0012 airfoil

Characteristic physics. The flow around a NACA 0012 airfoil in the Reynolds number range of 100 to 200,000 presents a rich variety of fluid dynamics phenomena, transitioning from steady, laminar flow to complex unsteady, transitional, and separated regimes. At the lower end of this spectrum, for Reynolds numbers on the order of 10^2 , the flow is characterized by thick laminar boundary layers that are susceptible to separation, even at high angles of attack [178]. As the angle of attack is increased, the wake is often dominated by unsteady vortex shedding.

With an increase in the Reynolds number into the range of 10^3 , the flow dynamics are frequently governed by the formation of a laminar separation bubble (LSB) on the suction side of the airfoil [179, 180]. This phenomenon occurs when the laminar boundary layer separates due to an adverse pressure gradient, transitions to a turbulent state within the separated shear layer, and then reattaches to the airfoil surface. The

formation, size, and behavior of the LSB are highly sensitive to both the Reynolds number and the angle of attack, which can lead to significant nonlinearities in the aerodynamic characteristics of the airfoil [179]. Near-stall conditions can induce self-sustained, low-frequency oscillations in the flow field, which are linked to the quasi-periodic bursting and reforming of the laminar separation bubble [180]. The various flow regimes can be categorized based on these dynamics, from steady attached flow at very low angles of attack, to the formation of a stable LSB, and ultimately to large-scale separation and stall, which is characterized by the shedding of large vortices [181].

A primary challenge in the computational modeling of these flow regimes is the accurate prediction of the laminar-to-turbulent transition. This transition is a critical element as it determines whether the separated flow will reattach. Standard RANS models often struggle to accurately capture the complex physics within the separation bubble. Consequently, higher-fidelity methods such as LES or DNS are necessary to resolve the fine-scale turbulent structures and transition mechanisms, albeit at a much greater computational expense. The intricate interplay between the Kelvin-Helmholtz instability of the separated shear layer, the amplification of disturbances, and the eventual turbulent breakdown and reattachment complicates the fundamental understanding of the physics.

Significant differences are observed between 2D and 3D configurations. While 2D simulations can capture many of the fundamental aspects of LSB formation and vortex shedding, they do not account for spanwise instabilities. In a 3D flow, the separated shear layer is subject to three-dimensional instabilities that can lead to a more rapid and complex transition to turbulence [182]. These three-dimensional effects can alter the size and shape of the separation bubble and, as a result, the overall aerodynamic forces [183]. This generally leads to a decrease in aerodynamic performance when compared to idealized two-dimensional predictions.

The NACA 0012 airfoil within the considered Reynolds number range is a crucial benchmark case for the development and testing of flow control applications, including those that utilize RL. The simple, symmetric geometry of the airfoil produces fundamentally important and challenging flow phenomena, which are directly relevant to the performance of micro air vehicles (MAVs), drones, and small wind turbine blades. The high sensitivity of the flow to small perturbations makes it an excellent platform for evaluating active flow control strategies aimed at suppressing separation, reducing drag, and enhancing lift. To this end, RL is a particularly promising approach as it can discover complex, non-intuitive, time-dependent control strategies, such as using blowing and suction jets, without requiring a precise analytical model of the flow physics. Nonetheless, the inherent instabilities of the shear layer and the broadband nature of turbulence necessitate control systems with high bandwidth and precise actuation. The effectiveness of a control strategy can be highly dependent on the specific flow state, which poses a challenge for the development of robust controllers that can operate effectively across a range of angles of attack and Reynolds numbers.

Characteristic physics in gust-airfoil interactions. Understanding the aerodynamic consequences of an airfoil encountering an atmospheric gust is a pivotal challenge in aeronautics, with direct implications for aircraft structural integrity, safety, and flight stability. The issue is especially pronounced for smaller aerial platforms, such as drones and MAVs, whose lower operating speeds make them highly susceptible to disturbances. When the flight speed of such a vehicle is on the same order as the gust velocity, a condition of near-resonance can occur, which dramatically magnifies the gust’s aerodynamic impact [184]. The increasing deployment of autonomous aircraft in unpredictable environments, such as urban settings or mountainous terrain, has intensified the need for robust technologies capable of mitigating these adverse gust-induced effects.

The nature of this aerodynamic interaction can be broadly classified based on the gust ratio, G , which compares the gust’s velocity to that of the aircraft’s freestream. For large commercial aircraft, their high cruise speeds ensure the gust ratio is typically low ($G < 1$). In these scenarios, provided the gust amplitude is not excessive, the aerodynamic response can often be approximated using linear theoretical models [185].

However, the validity of these linear models degrades as nonlinearities become more prominent. In sharp contrast, smaller vehicles frequently operate in conditions where the gust velocity matches or exceeds their flight speed ($G \geq 1$). This high- G regime, often termed 'extreme aerodynamics', introduces a set of physical phenomena far more complex than those seen in milder encounters, rendering linear theories inadequate [186].

Physically, a high- G interaction is dominated by highly unsteady and often violent aerodynamic events. As the gust alters the oncoming flow, the airfoil experiences a rapid and substantial change in its effective angle of attack. This can trigger the formation of a large, coherent leading-edge vortex. Concurrently, the rapid adjustment in the airfoil's overall circulation leads to the shedding of vorticity from the trailing edge, forming a trailing-edge vortex. If the effective angle of attack surpasses a critical threshold, the flow can undergo dynamic stall, a process distinct from its steady-state counterpart. During dynamic stall, large-scale vortical structures detach from the airfoil's surface and convect downstream, causing dramatic and often unpredictable shifts in lift and moment characteristics [187]. The resulting aerodynamic loads can momentarily spike to levels 200 – 300% greater than the maximum steady-state loads. This behavior is also accompanied by significant hysteresis in the force and moment curves, meaning the aerodynamic response depends on the history of the motion, which presents a major complication for control system design.

To summarize, the development of effective flow control strategies for gust load alleviation is the ultimate engineering goal. Conventional control systems, which rely on the deflection of large control surfaces, often have actuation delays and bandwidth limitations that make them too slow to counteract the rapid onset of high-frequency gusts. This has spurred innovation in active flow control technologies. The inherent complexity and unpredictability of gust encounters make them an ideal application area for reinforcement learning as the agent learns an optimal control strategy through direct trial-and-error interaction with its environment, making it well-suited to mastering the nonlinear dynamics of gust-airfoil interactions. However, the transient nature of the gust event demands a control policy that can respond almost instantaneously. To systematically develop and validate these advanced data-driven models and control algorithms, the aerospace community urgently needs well-documented benchmark test cases, such as flow over the NACA 0012 airfoil.




Figure 6: Comparison of (a) drag coefficient and (b) lift coefficient as a function of the angle of attack for a NACA 0012 airfoil at $Re = 1,000$. Results from the present study (2D and 3D simulations) are compared with experimental and numerical data from the literature.

Numerical setup and validation - m-AIA. The NACA 0012 airfoil simulations are conducted in both 2D and 3D configurations to capture the full spectrum of flow physics relevant to gust-airfoil interactions. For the 2D cases, the computational domain is defined by $(x/c, y/c) \in [0, 40] \times [-10, 10]$ with the NACA 0012 airfoil centered at $(x/c, y/c) = (10, 0)$. This 2D configuration allows for efficient exploration of parameter spaces and validation against classical airfoil theory, while serving as a baseline for understanding the fundamental flow mechanisms before extending to 3D effects. The reduced computational cost of 2D simulations enables systematic studies across wider ranges of angles of attack and Reynolds numbers.

For the 3D cases, the computational domain spans $[32 \times 16 \times 4]$ chord lengths $c = 1.0$ for $Re \leq 1,000$ and $[32 \times 16 \times 0.5]$ chord lengths for $Re \in [10,000, 50,000]$, with the NACA 0012 airfoil positioned $8c$ from the inflow boundary. The domain extent is designed to minimize boundary reflections while maintaining computational efficiency. The airfoil surface is treated as a no-slip boundary condition implemented through an interpolated bounce-back scheme. Periodic boundary conditions are imposed in the spanwise direction to simulate a nominally two-dimensional configuration while capturing three-dimensional flow instabilities.

The grid resolution is chosen to ensure that the smallest turbulent scales are adequately captured, particularly in the separated flow regions. The computational mesh employs an adaptive refinement approach with 14 refinement levels, with the grid spacing determined by $\Delta_n = \Delta_0/2^n$ where $\Delta_0^{(3D)} = 32c$. For 2D cases, the base grid spacing is $\Delta_0^{(2D)} = 40c$ with refinement levels up to $n = 12$, providing approximately 200 grid points per chord length in the highest resolution regions around the airfoil. Similar to [188], multiple refinement patches are strategically positioned to capture the boundary layer, wake region, and areas of expected flow separation with sufficient resolution. The numerical setup is validated against benchmark data from the literature [186, 189–193] showing good agreement in time-averaged lift and drag coefficients across various angles of attack $0^\circ < \alpha < 30^\circ$ for both 2D and 3D configurations (see Fig. 6).

Numerical setup and validation - Nek5000. High-resolution large-eddy simulation (LES) is employed to resolve the turbulent boundary layers (TBLs) developing over the NACA0012 wing sections at a chord-based Reynolds number of $Re_c = 200,000$ and an angle of attack of 0° . The present computations adopt a high-resolution LES strategy, with a grid density approaching that typically associated with direct numerical simulation (DNS) [194].

The computational domain is constructed using a C-type mesh, with streamwise and vertical extents of $L_x = 6c$ and $L_y = 4c$, respectively. The spanwise width of the NACA0012 configuration is set to $L_z = 0.1c$, resulting in a discretization comprising 220,000 spectral elements [195]. The leading and trailing edges of the airfoil are positioned at distances of $2c$ and $3c$ from the inflow boundary, respectively. This configuration has been validated in previous investigations [194], demonstrating that the selected spanwise extent is sufficient to capture both the energetically relevant turbulent scales and the largest coherent structures characterizing the wing boundary layer.

The near-wall spatial resolution, expressed in viscous units, satisfies $\Delta x_t^+ < 18.0$, $\Delta y_n^+ < (0.64, 11.0)$, and $\Delta z^+ < 9.0$ in the wall-tangential, wall-normal, and spanwise directions, respectively. The viscous length scale is defined as $l^* = \nu/u_\tau$, where the friction velocity is given by $u_\tau = \sqrt{\tau_w/\rho}$. The mean wall-shear stress is expressed as $\tau_w = \rho\nu (dU_t/dy_n)|_{y_n=0}$. To achieve the prescribed resolution, a polynomial order of $N = 7$ is adopted, resulting in a total number of grid points of approximately $N_{\text{grid}} \approx 1.1 \times 10^8$. In the wake region, the mesh resolution satisfies $\Delta x/\eta < 9$, where $\eta = (\nu^3/\epsilon)^{1/4}$ denotes the Kolmogorov length scale and ϵ represents the local isotropic dissipation rate.

The smallest unresolved turbulent scales are represented through a subgrid-scale (SGS) model based on a time-independent relaxation-term filter [196]. This filtering operation is implemented implicitly via a volume-force formulation that dissipates unresolved fluctuations while preserving mass conservation. The LES filter is a high-pass filter that restricts the LES dynamics to a subset of spectral modes within each element. The implementation of this approach in Nek5000 has been extensively validated in previous studies [194].




Figure 7: Inner-scaled mean component of wall-tangential velocity and the Reynolds-stress tensors from the 3D NACA4412 simulation at $Re_c = 400,000$ with DNS data. Figure reproduced from Ref [194] with permission of the publisher (Elsevier)

Dirichlet boundary conditions are imposed at the inflow, upper, and lower boundaries using a far-field velocity distribution obtained from an auxiliary Reynolds-averaged Navier–Stokes (RANS) computation. The RANS solution is performed with the $k-\omega$ shear-stress transport (SST) model [197] in a circular domain of radius $200c$. At the outflow boundary, the formulation proposed by [198] is employed to prevent the unphysical influx of kinetic energy. Boundary-layer transition is enforced at $x/c = 0.1$ on both the suction and pressure sides through a localized wall-normal body force [199], which generates intense, time-dependent streaks that subsequently break down and promote transition to turbulence [200].

One flow-over time (T_{FO}) is defined as the duration required for a fluid particle convecting at the free-stream velocity U_∞ to travel one chord length c . The procedure adopted to generate the uncontrolled reference simulations is described in Ref. [201]. Controlled simulations are initialized from fully developed turbulent fields obtained from the corresponding uncontrolled cases. Statistical convergence of the flow quantities is achieved after approximately 20 T_{FO} for the NACA0012 configuration.

The adequacy of the present mesh for the selected SGS modelling approach has been demonstrated through excellent agreement between DNS and LES data for a NACA4412 airfoil at $Re_c = 400,000$ and $AoA = 5^\circ$ (see figure 7). Specifically, the inner-scaled mean velocity profile and selected components of the Reynolds-stress tensor at $x_{ss}/c = 0.7$ (where ss denotes the suction side), expressed in the local tangential and wall-normal reference frame, are compared between both simulations. The comparison reveals excellent agreement in the mean velocity distribution, as well as in the Reynolds shear stress and in the outer-layer behaviour of the tangential velocity fluctuations. These results provide a validation of the numerical fidelity of the present configuration.

Environment implementation - Drag reduction setup. The control objective is to reduce skin-friction drag using the zero-shot deployment strategy. Although the implementation closely follows that adopted for turbulent channel flows described in Sec. 5K, several critical modifications are required. First, owing to the spatial development of TBLs, the friction velocity becomes a function of the chordwise coordinate, i.e. $u_\tau = u_\tau(x_{ss}/c)$. This spatial variation is represented using a ninth-order polynomial fit of the u_τ distribution




Figure 8: Exemplary visualization of the transversal gust scenario for 2D (a) and 3D (b) environments

available in the database reported in [195], yielding an exact reconstruction of the reference profile.

As a consequence, (1) the bounds of the actuation amplitude, (2) the wall-normal position of the sensing plane at $y^+ = 15$, and (3) the control update interval ($\Delta t^+ = 0.6$) all become functions of x_{ss}/c . The formulation of (1) and (2) follows directly once $u_\tau(x_{ss}/c)$ is known. Ensuring an exact update interval at each streamwise location, however, is considerably more challenging. To address this issue, the update interval is enforced to be uniform within each control block Ω_i , based on a viscous time scale computed using the block-averaged friction velocity, i.e. $t_{\Omega_i}^* = \nu / \langle u\tau \rangle_{\Omega_i}^2$.

Furthermore, a coordinate transformation between the wall-normal and Cartesian reference frames is required. In the wing configuration, actuation is defined in terms of the wall-normal velocity, whereas the numerical solver operates in Cartesian coordinates. Consequently, the actuation is imposed through modifications of the horizontal and vertical velocity components such that their projection along the wall-normal direction yields the desired control intensity. Conversely, the observable states are defined in the wall-normal frame. Both horizontal and vertical velocity components therefore contribute to the wall-tangential and wall-normal velocities, with their relative contribution depending on the local curvature of the surface. Owing to the streamwise inhomogeneity of TBLs, the spatial means of both actuation and state variables are obtained by averaging instantaneous quantities exclusively in the spanwise direction.

Environment implementation - Transverse gust setup. The transverse gust is implemented through a prescribed velocity perturbation imposed at the inflow boundary. The gust profile follows a 1-cosine approach with a gust factor $G = 2.0$, representing the ratio of peak gust velocity to freestream velocity (see Fig. 8). The gust temporal development is controlled to achieve the desired interaction dynamics with the airfoil. The state space representation consists of user-configurable point probes distributed throughout the computational domain, providing measurements of velocity components, pressure, density, and body forces at specified locations. To facilitate interactions between the RL agent and the flow environment, three jet actuators are distributed across the leading edge of the airfoil (see Fig. 9). Each actuator can be controlled independently, covers 3% of the chord length, and extends over the entire spanwise location. The action spaces are continuous and normalized in the range $[-1.0, 1.0]$.

Depending on the investigated Reynolds number, the number of CFD timesteps varies in the range of 225 – 1,000, delivering the necessary temporal granularity to represent critical flow physics while preserving computational tractability. The reward function is formulated to minimize gust-induced force fluctuations while maintaining aerodynamic efficiency:

$$R(t) = -|C_L(t) - \overline{C_{L,ref}}| - \omega |C_D(t) - \overline{C_{D,ref}}| \quad (37)$$

where $C_L(t)$ and $C_D(t)$ are the instantaneous drag and lift coefficients while $\overline{C_{L,ref}}$ and $\overline{C_{D,ref}}$ represent




Figure 9: Computational domain setup, sensor probe distribution and actuation strategy for the gust-airfoil interaction environments: Three-dimensional view showing the probe distribution around the airfoil with vortical structures visualized through Q-criterion isosurfaces. The yellow dots in lower right zoom-in highlight sensor probe locations in a single spanwise plane. Active flow control is performed using three independent jet actuators positioned along the leading edge of the NACA 0012 airfoil, each covering 3% of the chord length and extending across the full span (see lower left zoom-in).

time-averaged coefficients of the unperturbed reference case. The weighting coefficient $\omega = 0.25$ is tuned to balance the competing objectives of gust mitigation and drag reduction.

5F Flow around a cube

Characteristic physics. The flow around a suspended cube represents a canonical problem in fluid dynamics, offering a fundamental case study for three-dimensional bluff body wakes with fixed separation points. Unlike a sphere, where the separation location is Reynolds number dependent, the sharp leading edges of a cube dictate the points of flow separation, leading to a unique and complex evolution of the wake structure as the Reynolds number increases. For Reynolds numbers in the range of 300 to 3,700, the flow transitions through a sequence of distinct regimes, each characterized by specific vortical dynamics and stability properties.

Within the lower end of this Reynolds number range, several key bifurcations govern the wake's topology. The flow is initially steady and symmetric up to a critical Reynolds number of approximately 215 [202]. As the Reynolds number increases beyond this point, the flow undergoes a first bifurcation, losing its orthogonal symmetry to become steady but asymmetric, a state that persists up to a Reynolds number of about 265. This is followed by a second Hopf bifurcation around a Reynolds number of 270, where the flow becomes unsteady and periodic, characterized by the shedding of hairpin-like vortices [202, 203]. These hairpin vortices are initially shed in a regular, periodic manner. Experimental investigations using techniques such as Particle Image Velocimetry (PIV) and laser-induced fluorescence have confirmed this sequence of bifurcations,

aligning well with numerical simulations [203]. As the Reynolds number further increases towards 400, the hairpin vortex shedding becomes more complex, eventually transitioning to a chaotic vortex shedding regime around a Reynolds number of 310, where the coherence of the shed structures begins to break down [204].

Modeling the flow around a cube presents significant challenges, particularly as the Reynolds number increases into the turbulent regime. Accurately capturing the precise Reynolds numbers for the initial bifurcations requires high-fidelity numerical methods with sufficient spatial resolution to resolve the thin shear layers separating from the leading edges [202]. The interaction between these shear layers from the top, bottom, and side faces of the cube creates a highly three-dimensional and inherently unstable wake. As the Reynolds number climbs towards 3,700, the wake becomes fully turbulent, characterized by a broad range of spatial and temporal scales. The breakdown of larger coherent structures into smaller-scale turbulence and the complex interactions between different vortical structures are computationally expensive to resolve. Furthermore, ensuring the numerical schemes accurately capture the sharp-corner-induced separation without introducing excessive numerical diffusion is a persistent challenge. Understanding the physics is complicated by the fully three-dimensional nature of the vortex dislocations and the transition to turbulence in the separated shear layers.

Therefore, the flow around a cube serves as an important test case for the development and benchmarking of flow control strategies, including those based on reinforcement learning. The well-defined sequence of flow regimes offers a varied landscape to test the robustness and adaptability of control algorithms. The objective of such control could be drag reduction, lift suppression, or wake stabilization. The sharp edges provide fixed locations for the placement of actuators, such as synthetic jets or plasma actuators, simplifying one aspect of the control problem. However, the inherent three-dimensionality of the wake presents a significant challenge for flow control. Actuation strategies must be able to influence a complex, three-dimensional flow field, which necessitates a high-dimensional action space for an RL agent. For instance, controlling the interaction and instability of the multiple shear layers simultaneously is a non-trivial task. A key challenge for RL is the development of a low-dimensional state representation from achievable sensor measurements (e.g., surface pressure sensors) that can adequately describe the high-dimensional state of the flow. The agent must learn a control policy that can effectively manipulate the different instabilities present in the various flow regimes, from the steady asymmetry to the chaotic vortex shedding. The time delay between an actuation and its effect on the global flow properties, such as drag, also poses a significant hurdle for the learning process, requiring the RL algorithm to handle delayed rewards and long-term consequences of its actions. The complexity of the cube wake with its multiple interacting shear layers and instabilities therefore provides a challenging and realistic benchmark for advancing the capabilities of reinforcement learning in active flow control.

Numerical setup and validation. The cube benchmark follows the same computational setup as the square cylinder case, with the key difference being the three-dimensional bluff body geometry. The cube has a side length D and the Reynolds numbers are $Re \in [300, 1,000, 3,700]$. The 3D domain is $(x/D, y/D, z/D) \in [0, 32] \times [0, 16] \times [0, 16]$ with the cube centered at $(x/D, y/D, z/D) = (8, 8, 8)$ and non-periodic boundary conditions in all directions. All other boundary conditions, grid refinement levels, and resolution criteria follow the square cylinder specifications. Validation results shown in Tab. 5 demonstrate comparable accuracy to the square cylinder cases, with good agreement to established literature values where available.

Environment implementation. The cube environment follows the same fundamental framework as the square cylinder configuration, with identical reward formulation, state space representation, and episode structure. The key distinction is the three-dimensional jet actuation strategy, where jets can be positioned across multiple cube faces: leading/trailing edges with upper, lower, and side face locations (see Fig. 10). The RL agent controls each jet’s mass flux independently, with actions normalized to $[-1, 1]$. Temporal coupling follows the square cylinder approach with Reynolds number-specific action frequencies in the range of 1,000

(a) $Re = 300$			(b) $Re = 1,000$	
Reference	St	$\overline{C_D}$	Reference	$\overline{C_D}$
Present study	0.102	0.855	Present study	0.984
Saha et al.[202]	0.095	0.804	Holzer & Sommerfeld[207]	0.815
Saha[205]	0.097	0.833	Khan et al.[209]	0.883
Haider & Levenspiel [206]	–	0.985	Khan et al. [208]	0.968
Holzer & Sommerfeld[207]	–	0.923	Haider & Levenspiel [206]	1.121
Khan et al. [208]	–	0.843		

Table 5: Comparison of 3D cube flow results at different Reynolds numbers. St: Strouhal number; $\overline{C_D}$: Time-averaged drag coefficient; –: Data not reported or not applicable




Figure 10: Computational domain setup, sensor probe distribution and actuation strategy for the flow around a cube: Three-dimensional view showing the probe distribution around the cube with vortical structures visualized through Q-criterion isosurfaces. The yellow dots highlight sensor probe locations in a single spanwise plane. Active flow control is performed using independent jet actuators positioned along the leading and trailing edge of the cube.

to 1,250 CFD timesteps per control action. This configuration enables the investigation of multi-face actuation strategies for three-dimensional wake control.

5G Flow around a sphere

Characteristic physics. Flow past spheres within the Reynolds number regime 300-3,700 demonstrates intricate transitional phenomena marked by sequential symmetry-breaking bifurcations and progressive

wake complexity that establishes this configuration as fundamental for understanding three-dimensional bluff body aerodynamics and developing advanced control methodologies. This intermediate Reynolds number range captures the essential transition from organized periodic vortex structures to chaotic subcritical turbulent behavior, presenting distinctive opportunities and obstacles for both theoretical analysis and practical intervention strategies.

The underlying flow mechanics throughout this range are dictated by a systematic bifurcation sequence that progressively destroys flow symmetries and introduces increasingly complex temporal dynamics. The initial critical transition manifests around $Re = 212$ through a supercritical pitchfork bifurcation that eliminates axisymmetry while preserving planar symmetry, establishing steady non-axisymmetric flow patterns [210]. Subsequently, at approximately $Re = 280$, a supercritical Hopf bifurcation introduces periodic temporal behavior characterized by regular hairpin vortex shedding with Strouhal numbers near 0.18-0.20 [211]. The final primary transition occurs through a Neimark-Sacker bifurcation around $Re = 330$, generating quasiperiodic dynamics with two incommensurate frequencies that create complex modulation patterns while maintaining overall wake organization [212].

These successive instability mechanisms establish fundamental distinctions between lower Reynolds number steady flows and the complex transitional behavior observed in the $Re = 300 - 3,700$ range. While computational approaches utilizing steady-state assumptions remain viable below $Re = 212$, the emergence of temporal instabilities necessitates fully unsteady three-dimensional simulation methodologies that capture both large-scale wake evolution and fine-scale vortical structures [211, 213]. The hairpin vortex morphology characteristic of this regime creates a distinctive flow topology involving interconnected vortex loops and induced secondary structures that significantly influence drag, heat transfer, and mixing properties. Beyond $Re = 800$, the wake transitions toward subcritical turbulent characteristics while retaining coherent large-scale organization, establishing a complex multiscale environment where deterministic and stochastic elements coexist [214].

Consequently, this rich transitional behavior positions spherical wake flows as exceptionally valuable canonical benchmark for advancing flow control science and validating emerging computational strategies. The geometric simplicity eliminates confounding effects associated with complex shapes while preserving all essential three-dimensional bluff body physics, creating optimal conditions for isolating and understanding control mechanisms. The Reynolds number window of $Re \in [300 - 3,700]$ provides particularly favorable conditions for these developments because the flow complexity challenges algorithm sophistication while maintaining computational feasibility for extensive training iterations.

In summary, the well-established theoretical foundation surrounding spherical wake physics renders this configuration invaluable as advancing computational methodologies and experimental capabilities continue expanding flow control horizons for purely three-dimensional bluff body aerodynamics. The thorough documentation of fundamental mechanisms enables confident evaluation of whether innovative approaches genuinely capture essential physics or merely exploit numerical artifacts. Consequently, the extensively characterized sphere wake provides an indispensable validation platform for novel developments prior to extension toward more complex geometries and flow physics.

Numerical setup and validation. The sphere setup follows the cylinder benchmark scenario with 3D low Mach number flows ($Ma < 0.2$) at Reynolds numbers $Re \in [300, 3,700]$ based on the sphere diameter D . The computational domain differs from the cylinder configuration with $(x/D, y/D, z/D) \in [0, 32] \times [0, 16] \times [0, 16]$, featuring an extended z-direction domain. The sphere is centered at $(x/D, y/D, z/D) = (8, 8, 8)$. Grid spacing and refinement follow the cylinder setup with $\Delta_0^{(3Dsphere)} = 32D$ and refinement levels up to $n = 14$, leading up to 512 cells per sphere diameter D . Inlet and outlet boundary conditions are identical to the cylinder configuration. However, outlet boundary conditions replace the periodic conditions on all lateral boundaries due to the extended z-domain. Validation results demonstrate similarly good accuracy compared to the

(a) $Re = 300$			(b) $Re = 3,700$		
Reference	St	$\overline{C_D}$	Reference	St	$\overline{C_D}$
Present study	0.142	0.636	Present study	0.212	0.391
Johnson & Patel[210]	0.137	0.656	Almedeij[216]	–	0.409
Roos & Willmarth[215]	–	0.629	Clift et al.[217]	–	0.378
Almedeij[216]	–	0.602	Haider & Levenspiel[206]	–	0.392
Clift et al.[217]	–	0.674	Turton & Levenspiel[218]	–	0.386
			Rodriguez et al.[213]	0.215	0.394
			Yun et al.[219]	0.21	0.355

Table 6: Comparison of 3D sphere flow results at different Reynolds numbers. St: Strouhal number; $\overline{C_D}$: Time-averaged drag coefficient; –: Data not reported or not applicable

cylinder setup, with Strouhal numbers and drag coefficients matching within established literature ranges for sphere flows (see Tab. 6), where benchmark results are available.

Environment implementation. The sphere flow environment uses the same default reward objective as the cylinder setup, targeting drag minimization and lift force oscillation reduction. The actuation strategy differs from the cylinder configuration through eight cylindrical jets positioned on a circle at the sphere’s backside, located at 60° around the streamwise x-axis. Unlike the cylinder’s slot-like jets, these jets have a circular shape in the z-direction (see Fig. 11) and allow independent mass flux control by the RL agent.

The environment supports the Reynolds numbers $Re = 300, 3,700$, with an identical state space representation through user-configurable point probes as the cylinder counterpart. Similarly, training episodes and temporal coupling parameters follow the cylinder configuration, featuring 300 and 750 CFD timesteps per control action, respectively.

5H Zero-pressure gradient turbulent boundary layer

Characteristic physics. Turbulent boundary layer flows (TBLs) over flat plates represent a canonical flow problem for fundamental turbulence studies and the investigation of viscous drag reduction mechanisms. The flow is dominated by chaotic near-wall coherent structures, such as quasi-streamwise vortices and velocity streaks, which are primarily responsible for high skin-friction drag. Specifically, the buffer layer ($y^+ \approx 10\text{--}50$) hosts quasi-streamwise vortices with a characteristic diameter of approximately 30 wall units and alternating low- and high-speed velocity streaks with a robust mean spanwise spacing of $\lambda_z^+ \approx 100$ [220, 221]. These structures participate in an autonomous self-sustaining cycle comprising streak formation via the lift-up effect, inflectional instability and breakdown of the streaks, and nonlinear regeneration of streamwise vortices [222, 223]. Disrupting any phase of this cycle, for instance through external forcing, can suppress turbulence production and reduce skin friction. While this near-wall cycle is the primary driver of skin-friction drag at moderate Reynolds numbers, its dynamics do not evolve in isolation at higher Re_τ . The inner-scaled peak of the streamwise velocity variance at $y^+ \approx 15$ grows logarithmically with Re_τ , driven by the superposition and amplitude modulation of energy from outer-layer large-scale motions onto the near-wall region [224, 225]. At $Re_\tau \gtrsim 1,000$, very-large-scale motions (superstructures) with streamwise coherence exceeding 6δ emerge in the logarithmic layer and exert a measurable top-down modulation on the near-wall small-scale activity [226]. Critically, the wall-normal footprint of these outer-layer structures onto the near-wall region intensifies significantly with increasing Re_τ : the superimposed low-frequency signature of the superstructures progressively dominates the near-wall velocity and shear-stress fields, meaning that the inner-layer cycle alone becomes an incomplete picture of the drag-generating dynamics. The Fukagata–




Figure 11: Computational domain setup, sensor probe distribution and actuation strategy for the flow around a sphere: Three-dimensional view showing the probe distribution around the sphere with vortical structures visualized through Q-criterion isosurfaces. The zoom-in highlights sensor probe locations in a single spanwise plane. Independent jet actuators positioned in the sphere’s rear half enable active flow control.

Iwamoto–Kasagi identity [227] directly links the skin-friction coefficient to a weighted wall-normal integral of the Reynolds shear stress, providing a quantitative framework for assessing how modifications to the turbulent stress field translate into drag changes. Because the scale separation between inner and outer motions grows linearly with Re_τ , and because the outer-layer footprint on the near-wall region intensifies accordingly, control strategies that are effective at low Reynolds numbers may lose efficacy at higher Re_τ if they target only the near-wall cycle—motivating actuation methods capable of simultaneously addressing both the inner-layer regeneration mechanism and the outer-layer footprint. Spanning this Reynolds number range with a single benchmark platform therefore enables the systematic investigation of how learned control policies adapt—or fail to adapt—to the evolving multi-scale structure of the turbulent boundary layer.

Numerical setup and validation. The computational domain is a rectangular box spanning $L_x \times L_y \times L_z$ in the streamwise, wall-normal, and spanwise directions, respectively. To validate the baseline numerical setup, an unactuated flat plate boundary layer is simulated at a momentum-thickness Reynolds number of $Re_\theta = 1000$. The grid utilizes a DNS-like near-wall resolution ($\Delta x^+ \approx 12, \Delta y_{wall}^+ \approx 1.0, \Delta z^+ \approx 4.0$) to accurately resolve the small-scale turbulent motions. A reformulated synthetic turbulence generation (RSTG) method prescribes the inflow conditions, and a sponge layer is applied at the outflow to prevent the spurious reflection of acoustic waves. Validation of the unactuated baseline demonstrates excellent agreement with reference direct numerical simulation data for both the mean streamwise velocity profile and the Reynolds stress tensor components (see Fig 12). To facilitate the study of Reynolds number scaling on control strategies, the HydroGym environment offers pre-configured setups at friction Reynolds numbers $Re_\tau \in [200, 1550, 2200]$.




Figure 12: ZPGTBL Validation for $Re_\theta = 1000$: Wall-normal distributions of (a) the streamwise velocity and (b) the symmetric components of the Reynolds stress tensor and the Reynolds shear stress of the non-actuated reference case at $x = 165$. Results are compared to an incompressible TBL from [228].

Environment implementation. We investigate two distinct actuation setups: a streamwise traveling surface wave and a jet actuation on the surface. For both setups, the primary reward function for the RL agent is the net power saving (ΔP_{net}), which evaluates the energetic benefit of the drag reduction minus the power required to drive the respective actuation (i.e., the wall deformation or the jets). The continuous action spaces are defined according to the setup. For the surface wave, the agent controls the wave geometry and speed, parameterized in outer flow units by the wave amplitude $A \in [0.001, 0.40]$, the wavelength $\lambda \in [1, 8]$, and the phase speed $c \in [0.1, 0.40]$. For the jet actuation, the action space controls the jet velocity, which is bounded within $[-0.1U_\infty, 0.1U_\infty]$. Because macroscopic modifications to the turbulent boundary layer require time to develop, the duration of a control action depends on the setup: one control action comprises 5000 CFD time steps for the surface wave actuation and 1000 CFD time steps for the jet actuation. Across both configurations, a single training trajectory consists of 200 actuation steps. The state observation space is highly customizable; it relies on user-




Figure 13: Schematic of the computational setup for zero-pressure gradient turbulent boundary layer environment leveraging a streamwise traveling surface wave.

defined probe locations distributed throughout the flow field that can sample instantaneous velocity components (u, v, w), static pressure (p), density (ρ), and local surface forces. Episodes run for the predefined control interactions or terminate early if the solver diverges due to the exploration of non-physical actuation parameters.

5I Supercritical DRA2303 airfoil

Characteristic physics. Moving beyond canonical flat plates, the flow around a realistic lifting surface like the DRA2303 airfoil introduces the complexities of strong streamwise pressure gradients and surface curvature. Operating at a chord-based Reynolds number of $Re_c = 400,000$ and Mach numbers spanning $Ma = [0.2, 0.7]$, the aerodynamic efficiency is dictated by a delicate balance between friction drag and pressure (form) drag. At higher subsonic Mach numbers, compressibility effects become highly pronounced, and the interaction of active flow control mechanisms—such as a dynamically deforming wall or surface jets—with the boundary layer can induce localized supersonic regions or weak shock waves. Applying flow control in this regime poses a distinct control challenge: minimizing skin-friction drag might inadvertently increase pressure drag or penalize the lift coefficient. RL agents are uniquely positioned to discover spatially and temporally varying actuation parameters that navigate these competing aerodynamic forces to optimize global efficiency across different flight regimes.

Numerical setup and validation. The flow domain is discretized using a boundary-conforming C-type curvilinear mesh that extends approximately 24 chord lengths upstream and 25 chord lengths in both the downstream and vertical directions. Spanwise periodicity is enforced with a domain width of $L_z = 0.1c$, which is sufficient to ensure the decorrelation of turbulent flow structures along the wingspan. The boundary layer is tripped near the leading edge ($x/c = 0.1$) to ensure stable, fully developed turbulence over the actuated regions of the airfoil. The baseline numerical framework has been rigorously validated against high-resolution large-eddy simulations for similar airfoil geometries under identical flow conditions, showing excellent agreement for chordwise pressure (C_p) and skin-friction (C_f) distributions (see Fig. 14). The active flow control (traveling wave or jet actuation) is applied over the majority of the suction and pressure sides (typically $0.2 \leq x/c \leq 0.8$). For the surface wave setup, a smooth spatial blending function is utilized to transition the surface from static to dynamically wavy.




Figure 14: Validation results of numerical framework for high Re-number airfoil flows: Comparison of the present method applied to the flow around a NACA4412 airfoil with DNS results by [229] for the (a) pressure coefficient; (b) skin-friction coefficient; and (c) momentum thickness based Reynolds number.

Environment implementation. Similar to the flat plate case, the reward function is formulated to maximize the global net power saving, inherently balancing the reduction in total airfoil drag against the energetic cost of the respective actuation (i.e., surface deformation or jet blowing). Both actuation modes expose continuous action spaces tailored to the aerodynamic sensitivities of the curved geometry. For the traveling surface wave, the agent manipulates three parameters in outer units (Fig. 15): the amplitude is restricted to $A \in [0.00001, 0.0005]$ to account for the heightened sensitivity of the curved boundary layer, while the wavelength and phase speed are bounded by $\lambda \in [0.005, 0.02]$ and $c \in [0.001, 0.25]$, respectively. The jet actuator exposes a single degree of freedom—the jet exit velocity—constrained to $[-0.1 U_\infty, 0.1 U_\infty]$. To allow aerodynamic forces and wake dynamics to physically respond to updated wall boundary conditions, the action interval is set to 5000 CFD iterations for the surface wave and 1000 iterations for the jet. Each training trajectory comprises 200 actuation steps across both configurations. Flow observations are recorded through user-defined probe arrays positioned within the boundary layer or downstream wake, capturing instantaneous velocities, pressure, density, and local aerodynamic forces. The environment accommodates a range of freestream Mach numbers ($Ma = 0.2\text{--}0.7$), supporting the development of control policies that generalise across compressible flow regimes. Episodes conclude either upon completion of the prescribed actuation sequence or upon early termination triggered by solver divergence under non-physical actuation parameters.




Figure 15: Schematic of the computational setup for zero-pressure gradient turbulent boundary layer environments leveraging a streamwise traveling surface wave.

5J Flow through a stenotic pipe

Characteristic physics. Flow through a stenotic pipe has been extensively studied as a canonical configuration to investigate complex hemodynamic and biofluidic phenomena. The presence of a localized narrowing introduces strong acceleration of the fluid jet, followed by an adverse pressure gradient downstream that promotes flow separation, shear layer instabilities, and, at sufficiently high Reynolds numbers, transition to turbulence [230]. These dynamics create characteristic features such as recirculation zones, vortex shedding, and highly fluctuating wall shear stresses, which are of particular relevance in cardiovascular applications, for instance, in the progression of atherosclerosis and thrombosis. The flow field is highly sensitive to the severity and geometry of the stenosis as well as the Reynolds number. As such, stenotic pipe flow serves as a benchmark problem for studying transitional shear flows as well as fluid–structure interaction in arterial geometries.

Beyond its relevance to vascular hemodynamics, the stenotic pipe is also a valuable generic case for studying respiratory flows under obstructed conditions. The localized constriction mimics the airflow limitations encountered in the human nasal cavity during pathological states such as septal deviation or turbinate hypertrophy. In this context, a level-set approach has been developed for monitoring the change in pressure and temperature distributions when iteratively modifying the surface from a stenosed to an unstricted pipe [231]. After successfully testing this on the stenotic pipe, the authors have extended this framework to capture surface modifications from pre- to post-surgical nasal geometries. Building on this methodological foundation, subsequent studies integrated RL with CFD by letting an agent control interpolation factors that modify the geometry via the same level-set representation [31]. The agent’s goal was to find a compromise between reducing pressure loss and improving the heating capability. In the context of nasal obstruction surgery, this approach has been extended to nasal cavity geometries to design optimal




Figure 16: Computational domain for validating the stenotic pipe flow.

surgical intervention strategies [53]. In this sense, the stenotic pipe serves as a bridge between canonical fluid mechanics problems and clinically relevant respiratory applications, making it a suitable benchmark for developing and validating RL-CFD platforms.

Beyond steady inflow conditions, stenotic pipe flow has also been investigated under pulsatile forcing, which more closely resembles the physiological environment in cardiovascular and respiratory systems. The periodic acceleration and deceleration of the fluid introduce additional instabilities that interact with the stenosis-induced shear layers, often amplifying vortex shedding and enhancing the complexity of the downstream flow structures [232]. These effects can result in intermittent turbulence, increased wall shear stress fluctuations, and significant cycle-to-cycle variability in recirculation zones.

In addition to pulsatile inflow, oscillatory flow through stenotic geometries has been examined as a canonical setup for studying shear reversal and resonance phenomena under zero-mean periodic forcing. Unlike pulsatile flow, which superimposes forward and reverse phases on a mean flow, purely oscillatory forcing accentuates the interaction between unsteady shear layers and flow separation in both directions, leading to enhanced mixing and strongly asymmetric vortex dynamics. Under such conditions it has been demonstrated how the frequency of oscillation relative to the stenosis geometry governs vortex strength and flow reversal patterns [233].

Building on this canonical framework, the present study explores a novel RL-CFD coupling in which an RL agent interacts with oscillatory stenotic pipe flow by controlling the inlet temperature. The agent’s objective is to maintain the temperature field downstream of the stenosis as close as possible to a prescribed target value. This setup not only leverages the complex unsteady dynamics of oscillatory stenotic flows as a challenging benchmark for flow control, but also provides a direct link to biomedical applications. Specifically, the stenotic pipe serves as a generic surrogate for obstructed respiratory passages, where maintaining controlled temperature and flow conditions is critical. In this context, the methodology holds promise for extension to clinically relevant scenarios such as the pre-calibration of mechanical ventilation machines, where RL-driven control strategies could help adapt inlet conditions to patient-specific airway geometries and pathologies.

Implementation validation. The computational domain used for validating the stenotic pipe flow and its coordinate system are illustrated in Fig. 16. Details about the curvature of the stenosis can be found in [230]. The Reynolds number based on the pipe diameter at the inlet D is set to $Re = (U_{ref} \cdot D)/\nu = 500$ to match the conditions in [230], where $D = 1m$, $\nu = 1.63 \cdot 10^{-5} m^2/s$ is the kinematic viscosity of air, and U_{ref} the spatially averaged velocity at the inlet. The domain length is $L = 38D$.

At the inlet, the velocity profile of a fully developed laminar pipe flow is prescribed and the density of the fluid is linearly extrapolated from the neighboring inner cells. At the outlet, a constant pressure p_{out} is set, and the velocity is extrapolated from the inner cells. To satisfy the no-slip condition at the pipe walls, an interpolated bounce-back scheme is used [80]. These boundary conditions match the boundary conditions of the reference simulation in [230].

Figure 17a shows results of a mesh refinement study. Profiles of the normalized streamwise velocity component u/U_{ref} at $x = 5D$ with various mesh resolutions are compared to the reference solution from [230] (black). The coarse mesh (blue) has a resolution of $\Delta x = D/50$ and $2.5 \cdot 10^6$ cells. Results with the coarse mesh indicate a smaller recirculation zone between $r/D = 0.4$ and $r/D = 0.5$ and an underprediction of




Figure 17: Profiles of u/U_{ref} for different mesh resolutions (a) or at different locations with the highest resolution (b, c) compared to the reference case from the DNS in [230].

the maximum velocity at $r/D = 0$ by 5.2%. The solution of the medium mesh (red) with a resolution of $\Delta x = D/100$ and $20 \cdot 10^6$ cells agrees better with the reference solution near the recirculation zone, but still shows a slight underprediction of the maximum velocity by 2.5%. Finally, the fine grid (green) has a resolution of $\Delta x = D/200$ and features $160 \cdot 10^6$ cells. The corresponding velocity profile matches well with the reference solution and has only a deviation in the maximum velocity of 1.4%.

Figure 17b shows that the inlet profile of the simulation with the LB method matches the inlet profile in [230]. The velocity profiles at $x = 2.5D$ in Fig. 17c reveal that the LB simulation agrees well with the reference solution in the near-wall region and recirculation zone. Deviations below 5% are observed near $r/D = 0$.

The comparisons of velocity profiles and recirculation dynamics across different grid resolutions demonstrate that the fine mesh solution agrees closely with the established DNS benchmark [230]. Together with the previously verified thermal model [77], these results confirm that the present LB solver reliably reproduces the key flow features of the stenotic pipe.

Environment implementation. The stenotic pipe control environment is likewise formulated as a discrete-time Markov Decision Process. The underlying flow is driven by an oscillatory inflow with a Womersley number of $\alpha = \frac{1}{2}D\sqrt{\frac{\omega}{\nu}} = 4$, where ω denotes the angular frequency. To introduce variability and mimic physiological fluctuations, a Gaussian noise term with mean 0, standard deviation 1, and amplitude 0.05 is superimposed on the sinusoidal signal. The agent interacts with the system by modulating the inlet temperature. Actions are continuous and normalized, with the initial condition corresponding to a normalized inlet-to-wall temperature ratio of $T_{inlet}/T_{wall} = 0.9$. At each time step, the agent can adjust the inlet temperature within the bound $|\Delta T_{inlet}/T_{wall}| \leq 0.05$.

The environment advances in intervals of 4,000 simulation time steps between two consecutive control decisions. This duration ensures that thermal perturbations introduced at the inlet have sufficient time to advect downstream and alter the velocity and temperature fields in the vicinity of the stenotic region. The state space is defined by 10 velocity, pressure, and temperature probes placed along the pipe centerline, distributed between $x = -3.5D$ upstream of the stenosis and $x = D$ downstream (see Fig. 18). These probes capture the temporal development of the flow response to the applied control, providing the agent with the necessary feedback for learning effective strategies. The reward function is formulated as $R(t) =$




Figure 18: Normalized velocity and temperature fields of states I-IV of the uncontrolled 3D stenotic pipe flow.

$-100 \cdot |(T_{target}(t)/T_{wall}) - 0.95|$, where T_{target} is the temperature on the centerline at $x = D$.

Figure 18 presents the normalized velocity and temperature fields of the uncontrolled 3D stenotic pipe flow for four characteristic states. In state I, the inflow velocity reaches its peak and cold fluid enters the pipe. In state II, the inflow velocity transitions from positive to negative, resulting in near-zero velocities throughout most of the domain. Under these conditions, wall-to-fluid heat transfer is enhanced, leading to a rise in fluid temperature. In state III, the reverse flow attains its maximum, driving warm fluid into the pipe from the outlet. Finally, state IV marks the transition from negative back to positive inflow velocity, again associated with low velocities and elevated heat transfer.

5K Turbulent channel flow

Characteristic physics. Turbulent channel flow (TCF) represents a canonical configuration for investigating wall-bounded turbulence. Owing to its periodicity in two spatial directions, TCF provides an ideal training environment for multi-agent reinforcement learning (MARL), as this feature enables the generation of numerous pseudo-environments that can be exploited for parallel training. Consequently, global invariance is

preserved while learning efficiency is maximized. This distinctive advantage is absent in turbulent boundary layers (TBLs), where inherent spatial development substantially increases computational cost. Nevertheless, at comparable Reynolds numbers, TCF exhibits near-wall dynamics and coherent structures that closely resemble those observed in TBLs.

These considerations motivate the zero-shot MARL methodology, whereby policies trained in a channel configuration are subsequently deployed on TBLs without additional training. The TCF setup is therefore designed to reproduce the mean characteristics of the TBL within the targeted control region, thereby serving as an analogue of the flow environment encountered in the boundary layer. In the present study, the reference TBL corresponds to the suction side of a NACA0012 wing at $Re_c = 200,000$ and $AoA = 0^\circ$, as described in Sec. 5E, representing a boundary layer subjected to a mild adverse pressure gradient (APG). This strategy facilitates high-Reynolds-number deep reinforcement learning (DRL) control in complex flow configurations while substantially reducing computational expense.

Following the zero-shot framework, the principal objective is to match the momentum-thickness-based Reynolds number (Re_θ) as the target flow characteristic. This quantity is defined as $Re_\theta = U_0\theta/\nu$, where momentum thickness (θ) is defined as

$$\theta = \int_0^\infty (U/U_0) (1 - (U/U_0)) dy.$$

Here, U_0 corresponds to the boundary-layer edge velocity (U_e) in TBLs and to the centerline velocity (U_{cl}) in TCFs. The resulting configuration yields $Re_\theta = 365$, which aligns closely with the chordwise-averaged value $\langle Re_\theta \rangle_x = 382$ on the wing, corresponding to a deviation of only 4.5% from the target. This setup consequently produces a friction Reynolds number of $Re_\tau = 206$.

Numerical setup and validation. The three-dimensional open channel flow is simulated using direct numerical simulation (DNS) based on the spectral-element method (SEM). The computational domain has dimensions $(L_x, L_y, L_z) = (2\pi, 1, \pi)h$ in the streamwise, wall-normal, and spanwise directions, respectively, and is discretized into $(N_x, N_y, N_z) = (10, 8, 9)$ spectral elements. A polynomial order of $N = 7$ is adopted, yielding $(80, 64, 72)$ grid points in the streamwise, wall-normal, and spanwise directions, respectively. The resulting inner-scaled spatial resolutions are $\Delta x^+ < 18.8$, $\Delta y^+|_{\min, \max} < (0.5, 11.3)$, and $\Delta z^+ < 10.1$ in the streamwise, wall-normal, and spanwise directions, respectively. This resolution is selected to ensure that the effective size of the control actuators closely matches that employed in the wing simulations expressed in viscous units, thereby providing the grid-level consistency required to justify zero-shot transfer. A no-slip boundary condition is imposed at the lower wall, whereas a symmetry condition is applied at the upper boundary to avoid additional complications associated with wall actuation. Periodic boundary conditions are enforced in both the streamwise and spanwise directions, and the flow is driven by a prescribed constant bulk velocity to maintain statistically stationary turbulence.

To validate the numerical framework, a DNS of a full channel flow (i.e. with no-slip conditions imposed at both walls) is performed in Nek5000 at $Re_\tau = 180$, following the configuration detailed in [201]. The inner-scaled mean streamwise velocity profile and Reynolds-stress components are compared against the reference DNS database of [234] at $Re_\tau = 180$ as functions of the inner-scaled wall-normal distance y^+ , presented in logarithmic scale in figure 19. The excellent agreement observed across all quantities confirms the accuracy of the present numerical setup.

Environment implementation. The control objective is to reduce the wall-shear stress $\tau_w = \rho\nu (dU/dy)_{y=0}$ in the channel flow. Accordingly, the reward (R) is defined as the relative reduction in τ_w :

$$R = 1 - \tau_w^{\text{ctrl}} / \tau_w^{\text{ref}} \quad (38)$$




Figure 19: Inner-scaled (left) mean component of streamwise velocity and (right) Reynolds stresses as a function of inner-scaled wall-normal distance y^+ compared to Hoyas et al. DNS data of a full channel at $Re_\tau = 180$.

where ctrl and ref denote the controlled and uncontrolled flow states, respectively. For a fully developed TCF, this quantity is spatially uniform at the wall, thereby constituting a global invariant well suited for MARL. The actuation consists of zero-mean wall-normal blowing and suction at the wall, as detailed in Sec. 3C, with amplitudes constrained within the range $[-u_\tau, u_\tau]$, where $u_\tau = \sqrt{\tau_w/\rho}$. Control actions are updated at intervals of $\Delta t^+ = 0.6$, scaled using the viscous time unit $t^* = \nu/u_\tau^2$. Each training episode spans a duration of $t^+ = 1,500$. The observable states available to each agent consist of the streamwise and wall-normal velocity components sampled at $y^+ = 15$ (with $y^+ = yu_\tau/\nu$). For each component, the spatial mean computed as the ensemble average over all homogeneous directions is removed. In fully developed TCF, this mean value is identically zero for all velocity components. Further implementation details can be found in [235], as the present setup closely follows the methodology established in that work.

5L Differentiable turbulent channel flow

Characteristic physics. The turbulent channel flow at $Re_\tau = 180$ is an ideal example of a fully developed turbulent flow. It exhibits the phenomena of wall-bounded flows in general, serving as an ideal case for understanding flow over surfaces. In HydroGym, the channel flow is solved on a 3D domain of size $(2\pi \times 2 \times \pi)$ discretized into $(72 \times 72 \times 72)$ cells. The streamwise and spanwise directions contain uniform spacing, while a tunable hyperbolic tangent function is applied to the spacing in the wall-normal direction resulting in a finer grid close to the wall. A no-slip boundary condition is implemented at the walls, and periodic boundary conditions are applied in the streamwise and spanwise directions. A Neumann boundary condition is applied on the walls for pressure, $\frac{\partial p}{\partial z_{wall}} = 0$. The channel flow is driven by a constant, external body force in the x-direction, $f = (f_x, 0, 0)$, to ensure a statistically stationary turbulent flow [236]. The flow is initialized with a laminar flow plus a random perturbation to generate turbulence. The forcing term sets the friction Reynolds number, Re_τ , through the following relation:

$$f_x = \frac{8\nu^2}{H^3} Re_\tau^2, \quad (39)$$




Figure 20: Dimensionless velocity fluctuations and mean velocity from the 3D channel flow compared to Moser et al. DNS data.

where H is the height of the channel. The velocity and pressure are both solved at the cell-centroids of the grid. The timestep, Δt , of the simulation is dynamically updated to ensure a maximum Courant number of 0.9 [237].

Numerical setup and validation. The 3D channel flow problem is solved with a finite-difference direct numerical solver (DNS) inspired by work of Vuorinen et al. [237]. A compact way to numerically solve the derivatives in equation 24 is with differentiation matrices. Consider a derivative operator $\frac{\partial}{\partial x}$ and the velocity u and pressure p at the cell-centroid points. For a grid of size (N, N, N) , the velocity and pressure can be expressed as column vectors of size $(N^3, 1)$, and the differentiation operators, say $\frac{\partial}{\partial x}$, as a sparse differentiation matrix $\mathbf{D}_x$ of size (N^3, N^3) . All together, there are six differentiation matrices for velocity and six for pressure, $(\mathbf{D}_x, \mathbf{D}_{xx}, \mathbf{D}_y, \mathbf{D}_{yy}, \mathbf{D}_z, \mathbf{D}_{zz})$. The derivatives in the Navier-Stokes equations can then be computed with basic linear algebra $\frac{\partial u}{\partial x} \approx \mathbf{D}_x u$. For instance, the diffusive term in x would be computed as: $\nu(\mathbf{D}_{xx}u + \mathbf{D}_{yy}u + \mathbf{D}_{zz}u)$. The Poisson equation is used to solve for the pressure, using the formula: $\mathbf{M}p = \mathbf{D}_xu + \mathbf{D}_yu + \mathbf{D}_zu$ where $\mathbf{M}$ is the Poisson operator. This is a linear system, which is solved using the Biconjugate gradient solver from `lineax BiCGStab` [238]. Once the pressure is solved for, the velocity is corrected using the pressure gradient: $u^* = u - \mathbf{D}_{x,p}p$. To validate the solver, the dimensionless velocity fluctuations and mean velocity profile were plotted against the well-known Moser et al. 1999 DNS study [239], as shown in Figure 20.

Environment implementation. The control task for the 3D channel flow is to locally minimize the wall-shear stress with synthetic jets. Specifically, zero-mass-flux jets with periodic blowing and suction are implemented at the wall of the channel. This is essentially a wall-normal velocity distribution with a parabolic velocity profile, ensuring it is maximal at the centerline and smoothly decays to zero at the jet's edges. The actuation time is determined by the travel time of the velocity streaks, and ranges from $50\Delta t$ to $150\Delta t$ depending on the amount of jets implemented and the friction Reynolds number. The environment observation is the point value of the velocity in the x -direction, u , sampled from an off-wall location parallel to the jet locations. The sampling height was chosen as $y_s^+ \approx 15$, where there exists a strong correlation between wall-normal velocity and near-wall vortical structures [240]. Similar works have also used this sampling plane to ensure the best results [241]. The objective was to reduce the skin friction, or wall-shear stress, of the channel. The wall shear stress was computed as $\tau_w = \nu u(y)/y$, where y was chosen within the viscous sub-layer. Mathematically,

this is equivalent to computing the velocity gradient with respect to the wall, since $u(y=0) = 0$ [242, 243]. To decrease τ_w , reward signal was calculated as the ratio between the magnitude of the uncontrolled wall shear stress and the controlled wall shear stress. Both GPPO and PPO used the same observations and rewards, with the main distinction being that the GPPO training uses the reward signal in the loss function to differentiate through the environment dynamics when updating the policy.

5M Kolmogorov flow

Characteristic physics. Chaotic dynamical systems often exhibit extreme events in which the system significantly deviates from its expected behavior. These extreme events exist in both nature and engineered systems, such as oceanic rogue waves, earthquakes, and shocks in power grids, often leading to adverse financial and societal impacts [244]. As such, it is important to be able to predict and control these events. In turbulence, extreme events manifest as energy bursts, which are random and intermittent in both space and time [245]. To demonstrate this phenomena, a canonical two-dimensional turbulent flow driven by a sinusoidal forcing term is developed. This flow is known as the Kolmogorov flow, and in a certain parameter regime, it exhibits sporadic energy dissipation bursts due to a non-linear energy transfer between scales [246]. This often happens between Reynolds numbers 40 to 80 [247], as depicted in Figure 21 (b). For this implementation, it was found that the energy transfer that occurs is between a larger fluid structure with wavenumber $(1,0)$, which transfers energy to smaller structures with wavenumber $(0,k_f)$ [246]. This rapidly increases the energy dissipation, as turbulent energy dissipates at smaller scales.

Numerical implementation. Earlier in this section, the vector form of the incompressible Navier-Stokes equation 24 was introduced. This can be reformulated in vorticity space by taking the curl of the momentum equation, which can then be solved with Fourier methods. To demonstrate this, recall that $\nabla \times \mathbf{u} = \omega$, $\nabla \cdot \mathbf{u} = 0$ for an incompressible fluid, and $\nabla \times \nabla \cdot \mathbf{q} = 0$ where q is a scalar quantity. Equation 24 can then be simplified to:

$$\frac{\partial \omega}{\partial t} + \mathbf{u} \cdot \nabla \omega = \nu \nabla^2 \omega + \mathbf{f}, \quad \omega = -\nabla^2 \Psi; \quad (40)$$

where the stream function Ψ , defined as $\mathbf{u} = \nabla \times \Psi$, ensures continuity. Equation 40 is also known as the vorticity equation [248], and is approximated numerically on a doubly-periodic domain of $[(0, 2\pi) \times (0, 2\pi)]$ with a sinusoidal forcing term of $\mathbf{f} = (\sin k_f y, 0)$, where k_f is the forcing wavenumber. A pseudo-spectral method with 2/3 aliasing was developed to solve the equation [249]. For the time integration, a fourth order Runge Kutta and Crank Nicolson method from [250] was implemented.

Environment implementation. The Kolmogorov flow contained two control objectives: (1) mitigating energy dissipation, ϵ , and (2) increasing turbulent kinetic energy (TKE). While the obvious control objective is to predict and mitigate the energy burst, it can instead be leveraged to drive the system into a more turbulent regime, thereby increasing TKE and mixing. For the first objective, the energy dissipation rate is computed as $\epsilon = 2\nu \langle S_{ij} S_{ij} \rangle$, where S_{ij} is the strain rate tensor [251]. The reward function penalizes both energy dissipation magnitude and events characterized as extreme. In this case, any energy dissipation value that is greater than three standard deviations is penalized. For the second objective, the average turbulent kinetic energy (TKE) is computed as: $\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n [\frac{1}{2}(u_i - U)^2 + (v_i - V)^2]$, where (u, v) are the x and y velocity components with mean fields of U and V , respectively. In this case, the reward is proportional to the TKE. The observations are the same for both objectives - equally spaced velocity point values across the domain. The actuation is formulated as amplitudes of a forcing wavenumber, where the amplitudes are the actions of the agent. Four wavenumbers are chosen to manipulate the system, and the control function took the form $c = a_1 f(k_{f_1}) + a_2 f(k_{f_2}) + a_3 f(k_{f_3}) + a_4 f(k_{f_4})$, where f is the sinusoidal forcing function with wavenumber k_f . The wavenumbers selected are greater than the forcing wavenumber, as any additional




Figure 21: (a) A snapshot of the vorticity field for a $Re \approx 100$ case of the Kolmogorov flow, and (b) the energy dissipation rate over time for $Re \approx 40$. Highlighted with the red circle is the stable state, and the red star indicates one of the few extreme energy dissipation bursts.

energy in the large-scale structures dominated the dynamics. In the case presented in this work, the specific forcing wavenumbers of $k_1, k_2, k_3, k_4 = 4, 5, 6, 7$ are chosen. For both objectives, the actions are added as a term to the reward function to penalize large actions and promote more efficient controllers.

5N Shock vector control in a single divergent nozzle

Characteristic physics. Shock vector control (SVC) is a fluidic thrust vectoring technique in which a secondary jet is injected into the divergent section of a nozzle, deflecting the primary supersonic flow and generating a controllable lateral force. This approach offers a lightweight and mechanically simple alternative to conventional systems based on movable flaps or gimbaled nozzles, thereby reducing structural complexity, weight, and maintenance requirements [252]. The underlying physics are governed by compressible flow interactions: the injected secondary jet generates a bow shock that induces an adverse pressure gradient, forcing the primary flow to turn. This results in an asymmetric wall pressure distribution, which is the primary source of the lateral force and the resulting thrust vector deflection. The effectiveness of SVC is highly sensitive to operating conditions such as the pressure ratios of the primary and the secondary flow. Developing control strategies is therefore essential to ensure stable and predictable performance across a wide operating envelope.

The nozzle geometry and flow conditions considered in this study are based on the experimental work of Waithe et al. [253], who examined the flow in a single divergent planar nozzle. A schematic of the geometry is shown in Fig. 22. In addition to the planar (2D) configuration, we also investigate the corresponding three-dimensional bell nozzle. The 3D geometry is constructed by revolving the 2D nozzle contour about the x -axis. The nozzle contour consists of linear and circular segments, with a linear convergent section. The throat diameter is $D_t = 27.48\text{mm}$ and the injector width is 2.032mm . The injection direction is aligned with the planar injector normal. The flow conditions are parameterised by the nozzle pressure ratio $NPR = P_0/P_\infty$ and the secondary pressure ratio $SPR = P_s/P_0$. Here, P_0 and P_s denote the total pressure of the primary and secondary flow, respectively. The ambient pressure is denoted by P_∞ . Key performance metrics are the thrust

vector angles δ_0 and δ_1 . Given the thrust vector $\mathbf{F} = [F_x, F_y, F_z]$, they are defined as

$$\delta_0 = \tan^{-1}(F_y/F_x), \quad \delta_1 = \tan^{-1}(F_z/F_x). \quad (41)$$

Numerical setup and validation. The simulations are performed using the JAX-Fluids CFD solver [102, 103]. JAX-Fluids solves the compressible Navier-Stokes equations on a Cartesian mesh using a high-order Godunov-type method. In particular, a fifth-order WENO-Z scheme [254] is employed for cell-face reconstruction. The numerical cell-face flux is then computed with the HLLC approximate Riemann solver [255]. Temporal integration is performed using the third-order total-variation-diminishing Runge-Kutta scheme [256]. Solids are modeled with a conservative cut-cell immersed boundary method [257], where the interface is represented as the zero level-set of a signed-distance function.

Figure 22 shows a schematic of the computational domain. The domain extends over $(x/D_t, y/D_t, z/D_t) \in [0, 250] \times [-75, 75] \times [-75, 75]$, where the nozzle centerline coincides with the x -axis. The mesh consists of a highly resolved uniform region around the nozzle and a stretched region, increasing cell size towards the domain boundaries to reduce boundary effects. Two grid resolutions are considered: a coarse grid with $D_t/\Delta x = 50$, resulting in $(N_x, N_y, N_z) = (500, 300, 300)$ cells, and a fine grid with $D_t/\Delta x = 100$, resulting in $(N_x, N_y, N_z) = (900, 500, 500)$ cells.

At the primary and secondary inlets, stagnation pressure and a total temperature of $T_0 = 300$ K are prescribed. At the secondary inlet, no assumption is made a priori regarding choking. Instead, the local static pressure at the boundary is compared to the injector total pressure P_s . If the corresponding pressure ratio indicates choked conditions, the boundary state is imposed assuming isentropic expansion from P_s to sonic conditions ($M = 1$). Otherwise, the local static pressure is imposed, and the remaining flow variables are obtained from isentropic relations between P_s and the local pressure. The ambient pressure is set to $P_\infty = 1$ bar. The fluid is an ideal gas with specific heat ratio $\gamma = 1.4$ and specific gas constant $R = 287.14$ Jkg⁻¹K⁻¹.

In the present work, viscous effects are neglected, eliminating the need to resolve the boundary layer and thereby significantly reducing computational cost. As a consequence, boundary layer separation is not captured. Figure 23 shows the resulting Mach number field for the planar nozzle at $NPR = 4.6$ and $SPR = 0.7$. The injected secondary jet generates a bow shock that penetrates into the supersonic core flow, inducing a strong flow deflection and forming an asymmetric shock structure responsible for thrust vectoring. In the absence of viscous effects, the adverse pressure gradient does not lead to boundary layer separation. In viscous flows, the resulting separation bubble modifies the shock system such that, instead of a strong bow shock located close to the injector, a comparatively weaker oblique (separation) shock forms further upstream [252, 253]. This mechanism is absent in the inviscid case, which is also reflected in the wall pressure distributions in Fig. 23. Here, the overall trend is well captured, but the inviscid solution exhibits a more pronounced pressure peak closer to the injector. Nevertheless, the predicted thrust vector angles show good agreement with RANS and experimental data, indicating that the essential flow dynamics governing the shock-vectoring mechanism are well captured despite the absence of boundary layer effects. Future work will extend the present approach to wall-modeled large-eddy simulation to account for these near-wall effects.

Environment implementation. Two environments are considered: a planar (2D) and an axisymmetric (3D) nozzle configuration, each simulated at two grid resolutions as described above. The flow is characterized by a fixed nozzle pressure ratio $NPR = 4.6$, while the secondary pressure ratio is user-defined within $SPR \in [0.7, 0.9]$. In the 2D configuration, two injectors are present, one located on the upper and one on the lower wall, resulting in two control inputs. In the 3D configuration, the number of injectors N_{inj} is user-defined, and the injectors are uniformly distributed along the nozzle circumference, yielding N_{inj} independent control inputs, as illustrated in Fig. 24. The control input is defined through normalized actuator signals $a_i \in [0, 1]$, $i = 1, \dots, N_{inj}$, for each injector. Each actuator controls the total pressure of the corresponding secondary jet




Figure 22: Schematic of the nozzle geometry and the computational domain.

according to

$$P_{s,i} = p_{\text{local},i} + a_i (P_{s,\text{max}} - p_{\text{local},i}), \quad (42)$$

where $p_{\text{local},i}$ denotes the instantaneous local static pressure at injector i , and $P_{s,\text{max}} = P_{\infty} \cdot NPR \cdot SPR$.

The observations are the instantaneous thrust vector angles (δ_0 and δ_1), the instantaneous target thrust vector angles (δ_0^{ref} and δ_1^{ref}), and wall pressure measurements downstream of the injectors. In particular, two pressure probes are placed along the wall between each injection location and the nozzle exit, resulting in a total of $2N_{\text{inj}}$ pressure probes. The reward r is defined based on tracking a reference trajectory of the thrust vector angles. Using the instantaneous thrust vector angles and their corresponding reference values, the reward is formulated as

$$r = -\sqrt{(\delta_0 - \delta_0^{\text{ref}})^2 + (\delta_1 - \delta_1^{\text{ref}})^2}, \quad (43)$$

which penalizes deviations from the desired thrust vector direction. Note that the environment admits time-dependent reference values for the thrust vector direction.




Figure 23: Results for the planar nozzle with $SPR = 0.7$. Visualized is an instantaneous snapshot of the Mach number field for $NPR = 4.6$ and the corresponding mean wall pressure p_w/P_0 over the centerline x/x_t , where x_t denotes the throat location. Additionally, the mean thrust vector angle δ_0 over a variation of the nozzle pressure ratio NPR is shown.




Figure 24: Illustration of the 3D bell nozzle environment with $N_{\text{inj}} = 8$. The injector locations are indicated by white patches on the nozzle surface, which is colored by the static pressure. A 2D xy -slice at $z = 0$ shows the corresponding Mach number field. The computational grid is also depicted, highlighting the highly resolved uniform region and the stretched region towards the domain boundaries.

References

- [1] S. L. Brunton and B. R. Noack. Closed-loop turbulence control: Progress and challenges. *Appl. Mech. Rev.*, 67:050801–1–050801–48, 2015.
- [2] S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *Ann. Rev. Fluid Mech.*, 52:477–508, 2020.
- [3] Russ Tedrake and the Drake Development Team. Drake: Model-based design and verification for robotics, 2019.
- [4] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023.
- [5] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.
- [6] J. Degraeve, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- [7] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [8] R. Vinuesa and S. L. Brunton. Enhancing computational fluid dynamics with machine learning. *Nature Computational Science*, 2(6):358–366, 2022.
- [9] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer. Machine learning–accelerated computational fluid dynamics. *Proc. Nat. Acad. Sci.*, 118(21):e2101784118, 2021.
- [10] Eliot Xing, Vernon Luk, and Jean Oh. Stabilizing reinforcement learning in differentiable multiphysics simulation. *arXiv preprint arXiv:2412.12089*, 2024.
- [11] Rishabh Puri, Junya Onishi, Mario Rüttgers, Rakesh Sarma, Makoto Tsubokura, and Andreas Lintermann. On the choice of physical constraints in artificial neural networks for predicting flow fields. *Future Generation Computer Systems*, 161:361–375, 2024.
- [12] Mario Rüttgers, Junshin Park, and Donghyun You. Large-eddy simulation of turbulent flow over the driver fastback vehicle model. *Journal of Wind Engineering and Industrial Aerodynamics*, 186:123–138, 2019.
- [13] E. Lagemann, S.L. Brunton, W. Schröder, and C. Lagemann. Towards extending the aircraft flight envelope by mitigating transonic airfoil buffet. *Nat. Comm.*, 15(1):5020, 2024.
- [14] Esther Lagemann, Marian Albers, Christian Lagemann, and Wolfgang Schröder. Impact of reynolds number on the drag reduction mechanism of spanwise travelling surface waves. *Flow, Turbulence and Combustion*, 113(1):27–40, 2024.
- [15] 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. *J. Fluid Mech.*, 865:281–302, 2019.
- [16] M. Chatzimanolakis, P. Weber, and P. Koumoutsakos. Learning in two dimensions and controlling in three: Generalizable drag reduction strategies for flows past circular cylinders through deep reinforcement learning. *Phys. Rev. Fluids*, 9(4):043902, 2024.
- [17] Mikhail Tokarev, Egor Palkin, and Rustam Mullyadzhanov. Deep reinforcement learning control of cylinder flow using rotary oscillations at low reynolds number. *Energies*, 13(22):5920, 2020.
- [18] Hui Xu, Wei Zhang, Jian Deng, and J. Rabault. Active flow control with rotating cylinders by an artificial neural network trained by deep reinforcement learning. *Journal of Hydrodynamics*, 32(2):254–258, 2020.
- [19] J. Rabault and Alexander Kuhnle. Accelerating deep reinforcement learning strategies of flow control through a multi-environment approach. *Physics of Fluids*, 31(9), 2019.
- [20] Qiulei Wang, Lei Yan, Gang Hu, Wenli Chen, J. Rabault, and B. R. Noack. Dynamic feature-based deep reinforcement learning for flow control of circular cylinder with sparse surface pressure sensing. *J. Fluid*

Mech., 988:A4, 2024.

- [21] D. Fan, L. Yang, Z. Wang, M. S. Triantafyllou, and G. E. Karniadakis. Reinforcement learning for bluff body active flow control in experiments and simulations. *Proc. Nat. Acad. Sci.*, 117(42), 2020.
- [22] Wenjie Chen, Qiulei Wang, Lei Yan, Gang Hu, and B. R Noack. Deep reinforcement learning-based active flow control of vortex-induced vibration of a square cylinder. *Physics of Fluids*, 35(5), 2023.
- [23] Lei Yan, Yuerong Li, Gang Hu, Wen-li Chen, Wei Zhong, and B. R Noack. Stabilizing the square cylinder wake using deep reinforcement learning for different jet locations. *Physics of Fluids*, 35(11), 2023.
- [24] Mustafa Z Yousif, Paraskovia Kolesova, Yifan Yang, Meng Zhang, Linqi Yu, J. Rabault, R. Vinuesa, and Hee-Chang Lim. Optimizing flow control with deep reinforcement learning: Plasma actuator placement around a square cylinder. *Physics of Fluids*, 35(12), 2023.
- [25] Lei Yan, Gang Hu, Wenli Chen, and B. R Noack. Deep reinforcement transfer learning for active flow control of a 3D square cylinder under state dimension mismatch. *arXiv preprint arXiv:2401.12543*, 2024.
- [26] Chengwei Xia, Junjie Zhang, Eric C Kerrigan, and Georgios Rigas. Active flow control for bluff body drag reduction using reinforcement learning with partial measurements. *Journal of Fluid Mechanics*, 981:A17, 2024.
- [27] Wang Jia and Hang Xu. Robust and adaptive deep reinforcement learning for enhancing flow control around a square cylinder with varying reynolds numbers. *Physics of Fluids*, 36(5), 2024.
- [28] 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 $Re_D = 3900$. *Flow, Turb. & Combust.*, 2025.
- [29] Pol Suárez, Francisco Alcántara-Ávila, J. Rabault, Arnau Miró, Bernat Font, O. Lehmkuhl, and R. Vinuesa. Flow control of three-dimensional cylinders transitioning to turbulence via multi-agent reinforcement learning. *Communications Engineering*, 4(1):113, 2025.
- [30] Jonathan Viquerat, J. Rabault, Alexander Kuhnle, Hassan Ghraieb, Aurélien Larcher, and Elie Hachem. Direct shape optimization through deep reinforcement learning. *J. Comp. Phys.*, 428:110080, 2021.
- [31] M. Rüttgers, M. Waldmann, W. Schröder, and A. Lintermann. Machine learning-based control of perturbed and heated channel flows. In *High Performance Computing*, pages 7–22, 2021.
- [32] Jinhua Lou, Rongqian Chen, Jiaqi Liu, Yue Bao, Yancheng You, and Zhengwu Chen. Aerodynamic optimization of airfoil based on deep reinforcement learning. *Physics of Fluids*, 35(3), 2023.
- [33] S. Bhola, S. Pawar, P. Balaprakash, and R. Maulik. Multi-fidelity reinforcement learning framework for shape optimization. *J. Comp. Phys.*, 482:112018, 2023.
- [34] Thomas P Dussauge, Woong Je Sung, Olivia J Pinon Fischer, and Dimitri N Mavris. A reinforcement learning approach to airfoil shape optimization. *Scientific Reports*, 13(1):9753, 2023.
- [35] Jignesh Patel, Yannis Spyridis, and Vasileios Argyriou. Enhancing vehicle aerodynamics with deep reinforcement learning in voxelised models. In *2024 14th International Conference on Pattern Recognition Systems (ICPRS)*, pages 1–7. IEEE, 2024.
- [36] Jonathan Tran, Kai Fukami, Kenta Inada, Daisuke Umehara, Yoshimichi Ono, Kenta Ogawa, and Kunihiro Taira. Aerodynamics-guided machine learning for design optimization of electric vehicles. *Communications Engineering*, 3(1):174, 2024.
- [37] Sang-Jae Yeo, Suk-Yoon Hong, and Jee-Hun Song. Deep-reinforcement-learning-based hull form optimization method for stealth submarine design. *International Journal of Naval Architecture and Ocean Engineering*, 16:100595, 2024.
- [38] S.-J. Oh, M.-J. Oh, and E.-Y. Son. Reinforcement learning-based optimal hull form design with variations in fore and aft parts. *J. Comp. Design Eng.*, 11(6):1–19, 2024.
- [39] B. Font, F. Alcántara-Ávila, J. Rabault, R. Vinuesa, and O. Lehmkuhl. Deep reinforcement learning for active flow control in a turbulent separation bubble. *Nat. Comm.*, 16(1):1422, 2025.
- [40] J. Xiang, H. Zong, Y. Wu, J. Li, and H. Liang. Experimental control of the flow separation behind a backward facing step using deep reinforcement learning. *Physics of Fluids*, 36(10), 2024.
- [41] Satoshi Shimomura, Satoshi Sekimoto, Akira Oyama, Kozo Fujii, and Hiroyuki Nishida. Closed-loop

- flow separation control using the deep Q network over airfoil. *AIAA Journal*, 58(10):4260–4270, 2020.
- [42] X. Garcia, A. Miró, P. Suárez, F. Alcántara-Ávila, J. Rabault, B. Font, O. Lehmkuhl, and R. Vinuesa. Deep-reinforcement-learning-based separation control in a two-dimensional airfoil. *Int. J. Heat & Fluid Flow*, 116:109913, 2025.
 - [43] L. Guastoni, J. Rabault, Philipp Schlatter, Hossein Azizpour, and R. Vinuesa. Deep reinforcement learning for turbulent drag reduction in channel flows. *The European Physical Journal E*, 46(4):27, 2023.
 - [44] Taehyuk Lee, Junhyuk Kim, and Changhoon Lee. Turbulence control for drag reduction through deep reinforcement learning. *Phys. Rev. Fluids*, 8(2):024604, 2023.
 - [45] Zisong Zhou, Mengqi Zhang, and Xiaojue Zhu. Reinforcement-learning-based control of turbulent channel flows at high Reynolds numbers. *J. Fluid Mech.*, 1006:A12, 2025.
 - [46] Takahiro Sonoda, Zhuchen Liu, Toshitaka Itoh, and Yosuke Hasegawa. Reinforcement learning of control strategies for reducing skin friction drag in a fully developed turbulent channel flow. *J. Fluid Mech.*, 960:A30, 2023.
 - [47] G. Novati, S. Verma, D. Alexeev, D. Rossinelli, W. M. Van Rees, and P. Koumoutsakos. Synchronisation through learning for two self-propelled swimmers. *Bioinspiration & biomimetics*, 12(3):036001, 2017.
 - [48] S. Verma, G. Novati, and P. Koumoutsakos. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Nat. Acad. Sci.*, 115(23):5849–5854, 2018.
 - [49] Ke Qin, Zonghao Zou, Lailai Zhu, and On Shun Pak. Reinforcement learning of a multi-link swimmer at low Reynolds numbers. *Physics of Fluids*, 35(3), 2023.
 - [50] Tongzhao Xiong, Zhaorong Liu, Yufei Wang, Chong Jin Ong, and Lailai Zhu. Chemotactic navigation in robotic swimmers via reset-free hierarchical reinforcement learning. *Nat. Comm.*, 16(1):5441, 2025.
 - [51] Yi Zhu, Jian-Hua Pang, Tong Gao, and Fang-Bao Tian. Learning to school in dense configurations with multi-agent deep reinforcement learning. *Bioinspiration & Biomimetics*, 18(1):015003, 2022.
 - [52] Daniel Wälchli, Pascal Weber, Michail Chatzimanolakis, Robert Katzschmann, and P. Koumoutsakos. Inverse reinforcement learning for objective discovery in collective behavior of artificial swimmers. *Phys. Rev. Fluids*, 10(6):064901, 2025.
 - [53] M. Rüttgers, M. Waldmann, K. Vogt, J. Ilgner, W. Schröder, and A. Lintermann. Automated surgery planning for an obstructed nose by combining computational fluid dynamics with reinforcement learning. *Computers in Biology and Medicine*, 173:108383, 2024.
 - [54] M. Rüttgers, F. Hübenthal, M. Tsubokura, and A. Lintermann. Parallel Reinforcement Learning and Gaussian Process Regression for Improved Physics-Based Nasal Surgery Planning. In *Parallel Processing and Applied Mathematics*, pages 79–96. Springer Nature Switzerland, 2025.
 - [55] Parisa Sarikhani, Hao-Lun Hsu, Mahmoud Zeydabadinezhad, Yuyu Yao, Mayuresh Kothare, and Babak Mahmoudi. Reinforcement learning for closed-loop regulation of cardiovascular system with vagus nerve stimulation: a computational study. *Journal of Neural Engineering*, 21, 06 2024.
 - [56] Yuyang Shi, Zhike Xu, Chenghan Chen, Feng He, Pengfei Hao, and Xiwen Zhang. Enhancing cardiac hemodynamic and pulsatility in heart failure via deep reinforcement learning: An in-silico and in-vitro validation study of percutaneous ventricular assist devices. *Computer Methods and Programs in Biomedicine*, 270:108975, 2025.
 - [57] E. Hachem, P. Meliga, A. Goetz, P. Rico, J. Viquerat, A. Larcher, R. Valette, A. Sanches, V. Lannelongue, H. Ghraieb, R. Nemer, Y. Ozpeynirci, and T. Liebig. Reinforcement learning for patient-specific optimal stenting of intracranial aneurysms. *Scientific Reports*, 13:7147, 05 2023.
 - [58] J. Yang, T. Dzanic, B. Petersen, J. Kudo, K. Mittal, V. Tomov, J.-S. Camier, T. Zhao, H. Zha, T. Kolev, et al. Reinforcement learning for adaptive mesh refinement. In *International conference on artificial intelligence and statistics*, pages 5997–6014. PMLR, 2023.
 - [59] Corbin Foucart, Aaron Charous, and Pierre FJ Lermusiaux. Deep reinforcement learning for adaptive mesh refinement. *J. Comp. Phys.*, 491:112381, 2023.
 - [60] Y. Zhu, S. Zhao, Y. Zhou, H. Liang, and X. Bian. An unstructured adaptive mesh refinement for steady flows based on physics-informed neural networks. *J. Comp. Phys.*, page 114283, 2025.
 - [61] D. Huergo, Gonzalo Rubio, and Esteban Ferrer. A reinforcement learning strategy for p-adaptation in

- high order solvers. *Results in Engineering*, 21:101693, 2024.
- [62] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Ann. Rev. Fluid Mech.*, 51(1):357–377, 2019.
 - [63] Marius Kurz, Philipp Offenhäuser, and A. Beck. Deep reinforcement learning for turbulence modeling in large eddy simulations. *Int. J. Heat & Fluid Flow*, 99:109094, 2023.
 - [64] G. Novati, H. L. de Laroussilhe, and P. Koumoutsakos. Automating turbulence modelling by multi-agent reinforcement learning. *Nat. Mach. Intell.*, 3(1):87–96, 2021.
 - [65] Zhan Zhang. Application of deep reinforcement learning in parameter optimization and refinement of turbulence models. *Scientific Reports*, 15(1):25236, 2025.
 - [66] H. J. Bae and P. Koumoutsakos. Scientific multi-agent reinforcement learning for wall-models of turbulent flows. *Nat. Comm.*, 13(1):1443, 2022.
 - [67] D. Zhou and H. J. Bae. Wall modeling of turbulent flows with varying pressure gradients using multi-agent reinforcement learning. *AIAA Journal*, 62(10):3713–3727, 2024.
 - [68] Wulong Hu, Zhangze Jiang, Mingyang Xu, and Hanyu Hu. Efficient deep reinforcement learning strategies for active flow control based on physics-informed neural networks. *Physics of Fluids*, 36(7), 2024.
 - [69] Colin Rodwell and Phanindra Tallapragada. Physics-informed reinforcement learning for motion control of a fish-like swimming robot. *Scientific Reports*, 13(1):10754, 2023.
 - [70] Minghui Ye and Ahmed H Elsheikh. Model-based reinforcement learning for active flow control. *Physics of Fluids*, 37(9), 2025.
 - [71] Andre Weiner and Janis Geise. Model-based deep reinforcement learning for accelerated learning from flow simulations. *Meccanica*, 60(6):1771–1788, 2025.
 - [72] Nicholas Zolman, Christian Lagemann, Urban Fasel, J Nathan Kutz, and Steven L Brunton. Sindy-rl for interpretable and efficient model-based reinforcement learning. *Nature Communications*, 16(1):10714, 2025.
 - [73] Institute of Aerodynamics. m-AIA, 2024. <https://doi.org/10.5281/zenodo.13350586>.
 - [74] P. L. Bhatnagar, E. P. Gross, and M. Krook. A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems. *Physical Review*, 94(3):511, 1954.
 - [75] Martin Geier, Martin Schönherr, A. Pasquali, and Manfred Krafczyk. The cumulant lattice boltzmann equation in three dimensions: Theory and validation. *Computers & Mathematics with Applications*, 70(4):507–547, 2015.
 - [76] Andreas Lintermann, Matthias Meinke, and Wolfgang Schröder. Fluid mechanics based classification of the respiratory efficiency of several nasal cavities. *Computers in Biology and Medicine*, 43(11):1833–1852, 2013.
 - [77] M. Rüttgers, M. Waldmann, W. Schröder, and A. Lintermann. A machine-learning-based method for automatizing lattice-boltzmann simulations of respiratory flows. *Applied Intelligence*, 52:9080–9100, 2022.
 - [78] G. Eitel-Amor, M. Meinke, and W. Schröder. A lattice-boltzmann method with hierarchically refined meshes. *Computers & Fluids*, 75:127–139, April 2013.
 - [79] Alexandre Dupuis and Bastien Chopard. Theory and applications of an alternative lattice boltzmann grid refinement algorithm. *Physical Review E*, 67(6):066707, 2003.
 - [80] Mohamed Bouzidi, Mouaouia Firdaouss, and Pierre Lallemand. Momentum transfer of a boltzmann-lattice fluid with boundaries. *Physics of Fluids*, 13(11):3452–3459, 2001.
 - [81] C. Lagemann, M. Rüttgers, M. Gondrum, M. Meinke, W. Schröder, A. Lintermann, and S.L. Brunton. HydroGym-GPU: From 2D to 3D Benchmark Environments for Reinforcement Learning in Fluid Flows. In *Proc. 35th Int. Conf. on Parallel Comp. Fluid Dyn.*, 2024.
 - [82] Peter D. Thomas and Charles K. Lombard. Geometric conservation law and its application to flow computations on moving grids. *AIAA Journal*, 17(10):1030–1037, 1979.
 - [83] Meng-Sing Liou and Christopher J. Steffen, Jr. A new flux splitting scheme. *Journal of Computational Physics*, 107(1):23–39, 1993.

- [84] Bram van Leer. Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method. *Journal of Computational Physics*, 32(1):101–136, 1979.
- [85] S Stolz and Nikolaus A Adams. An approximate deconvolution procedure for large-eddy simulation. *Physics of Fluids*, 11(7):1699–1701, 1999.
- [86] Philipp Schlatter, Steffen Stolz, and Leonhard Kleiser. Les of transitional flows using the approximate deconvolution model. *International journal of heat and fluid flow*, 25(3):549–558, 2004.
- [87] Matthias Meinke, Wolfgang Schröder, Egon Krause, and Thomas Rister. A comparison of second- and sixth-order methods for large-eddy simulations. *Computers & Fluids*, 31(4–7):695–718, 2002.
- [88] David L. Whitfield. Three-dimensional unsteady Euler equations solution using flux vector splitting. In *AIAA Paper No. 84-0122*, 1984.
- [89] Benedikt Roidl, Matthias Meinke, and Wolfgang Schröder. A reformulated synthetic turbulence generation method for a zonal RANS–LES method and its application to zero-pressure gradient boundary layers. *International Journal of Heat and Fluid Flow*, 44:28–40, 2013.
- [90] Marian Albers, Xiao Shao, and Wolfgang Schröder. Energy efficient actuated drag reduced compressible turbulent flat plate flow. *International Journal of Heat and Fluid Flow*, 106:109314, 2024.
- [91] Esther Mäteling and Wolfgang Schröder. Analysis of spatiotemporal inner-outer large-scale interactions in turbulent channel flow by multivariate empirical mode decomposition. *Physical Review Fluids*, 7(3):034603, 2022.
- [92] E. Mäteling, M. Albers, and W. Schröder. How spanwise travelling transversal surface waves change the near-wall flow. *J. Fluid Mech.*, 957:A30, 2023.
- [93] Esther Mäteling, Michael Klaas, and Wolfgang Schröder. Detection of small-scale/large-scale interactions in turbulent wall-bounded flows. *Physical Review Fluids*, 5(11):114610, 2020.
- [94] Paul F. Fischer, James W. Lottes, and Stefan G. Kerkemeier. Nek5000 Web page, 2008. <http://nek5000.mcs.anl.gov>.
- [95] BB De Moura, MR Machado, S Dey, and T Mukhopadhyay. Manipulating flexural waves to enhance the broadband vibration mitigation through inducing programmed disorder on smart rainbow metamaterials. *Appl. Math. Modelling*, 125:650–671, 2024.
- [96] Yvon Maday and Anthony T Patera. Spectral element methods for the incompressible navier-stokes equations. IN: *State-of-the-art surveys on computational mechanics (A90-47176 21-64)*. New York, pages 71–143, 1989.
- [97] S. Son, L. Zheng, R. Sullivan, Y.-L. Qiao, and M. Lin. Gradient informed proximal policy optimization. *Advances in Neural Information Processing Systems*, 36:8788–8814, 2023.
- [98] Sajeda Mokbel, C. Lagemann, E. Lagemann, and S. L. Brunton. Controlling chaotic energy events in fluids with reinforcement learning. In *Proceedings of the 16th ACM International Conference on Future and Sustainable Energy Systems*, pages 954–958, 2025.
- [99] Robert Tjarko Lange. gymjax: A JAX-based Reinforcement Learning Environment Library, 2022.
- [100] Chris Lu. PureJaxRL: End-to-End RL Training in Pure Jax, 2023.
- [101] S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. G. M. Araújo. CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
- [102] Deniz A Bezgin, Aaron B Buhendwa, and Nikolaus A Adams. Jax-fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Computer Physics Communications*, 282:108527, 2023.
- [103] Deniz A Bezgin, Aaron B Buhendwa, and Nikolaus A Adams. Jax-fluids 2.0: Towards hpc for differentiable cfd of compressible two-phase flows. *Computer Physics Communications*, 308:109433, 2025.
- [104] David A. Ham, Paul H. J. Kelly, Lawrence Mitchell, Colin J. Cotter, Robert C. Kirby, Koki Sagiya, Nacime Bouziani, Sophia Vorderwuelbecke, Thomas J. Gregory, Jack Betteridge, Daniel R. Shapero, Reuben W. Nixon-Hill, Connor J. Ward, Patrick E. Farrell, Pablo D. Brubeck, India Marsden, Thomas H. Gibson, Miklós Homolya, Tianjiao Sun, Andrew T. T. McRae, Fabio Luporini, Alastair Gregory, Michael Lange, Simon W. Funke, Florian Rathgeber, Gheorghe-Teodor Bercea, and Graham R. Markall. *Firedrake*

User Manual. Imperial College London and University of Oxford and Baylor University and University of Washington, first edition edition, 5 2023.

- [105] Albert Bou, Matteo Bettini, S. Dittert, Vikash Kumar, Shagun Sodhani, Xiaomeng Yang, Gianni De Fabritiis, and Vincent Moens. TorchRL: A data-driven decision-making library for PyTorch, 2023.
- [106] James Bradbury, Roy Frostig, M. James Hawkins, P. and Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.
- [107] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [108] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [109] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [110] Volodymyr Mnih, Koray Kavukcuoglu, D. Silver, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [111] D. Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning*, pages 387–395, 2014.
- [112] George E Uhlenbeck and L. S Ornstein. On the theory of the Brownian motion. *Physical Review*, 36(5):823–841, 1930.
- [113] Sergey Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 448–456, 2015.
- [114] Scott Fujimoto, Herke Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. *Proceedings of the 35th International Conference on Machine Learning*, pages 1587–1596, 2018.
- [115] H. van Hasselt. Double Q-learning. *Advances in Neural Information Processing Systems*, 23:2613–2621, 2010.
- [116] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [117] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*, volume 105. SIAM, 2008.
- [118] Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, A. Steiner, and Marc van Zee. Flax: A neural network library and ecosystem for JAX, 2020.
- [119] C.H.K. Williamson. Vortex dynamics in the cylinder wake. *Ann. Rev. Fluid Mech.*, 28:477–539, 1996.
- [120] CHK Williamson. Three-dimensional wake transition. *J. Fluid Mech.*, 328:345–407, 1996.
- [121] A.I. Aleksyuk and M. Heil. On the short-wavelength three-dimensional instability in the cylinder wake. *J. Fluid Mech.*, 999, 2024.
- [122] S. Bhattacharya and J.W. Gregory. Investigation of the cylinder wake under spanwise periodic forcing with a segmented plasma actuator. *Physics of Fluids*, 27(1):014102, 2015.
- [123] F. Ren, L. Cheng, Q. Xiong, F. Tong, and C. Chen. Three-dimensional wake transitions of steady flow past two side-by-side cylinders. *J. Fluid Mech.*, 969, 2023.
- [124] P. Parnaudeau, J. Carlier, D. Heitz, and E. Lamballais. Experimental and numerical studies of the flow over a circular cylinder at Reynolds number 3900. *Physics of Fluids*, 20:085101, 2008.
- [125] Arthur G Kravchenko and P. Moin. Numerical studies of flow over a circular cylinder at $Re_D = 3900$. *Physics of Fluids*, 12(2):403–417, 2000.
- [126] Bing-Zheng Han, Wei-Xi Huang, and Chun-Xiao Xu. Deep reinforcement learning for active control of flow over a circular cylinder with rotational oscillations. *Int. J. Heat & Fluid Flow*, 96:109008, 2022.
- [127] S.L Brunton and B. R Noack. Closed-loop turbulence control: progress and challenges. *Appl. Mech. Rev.*, 67(5):050801, 2015.

- [128] S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2nd edition, 2022.
- [129] Kai Lagemann, Christian Lagemann, and Sach Mukherjee. Invariance-based learning of latent dynamics. In *International Conference on Learning Representations*, pages 38843–38887, 2024.
- [130] Julio R Meneghini, Fábio Saltara, Cesareo de La Rosa Siqueira, and JA Ferrari Jr. Numerical simulation of flow interference between two circular cylinders in tandem and side-by-side arrangements. *J. Fluid Struct.*, 15(2):327–350, 2001.
- [131] H Ding, C Shu, KS Yeo, and D Xu. Numerical simulation of flows around two circular cylinders by mesh-free least square-based finite difference methods. *International journal for numerical methods in fluids*, 53(2):305–332, 2007.
- [132] M Braza, PHHM Chassaing, and H Ha Minh. Numerical study and physical analysis of the pressure and velocity fields in the near wake of a circular cylinder. *J. Fluid Mech.*, 165:79–130, 1986.
- [133] O Posdziech and R Grundmann. A systematic approach to the numerical calculation of fundamental quantities of the two-dimensional flow over a circular cylinder. *J. Fluid Struct.*, 23(3):479–499, 2007.
- [134] A. B. Harichandan and A. Roy. Numerical investigation of low reynolds number flow past two and three circular cylinders using unstructured grid cfr scheme. *Int. J. Heat & Fluid Flow*, 31(2):154–171, 2010.
- [135] A. Colagrossi, G. Graziani, and M. Pulvirenti. Particles for fluids: Sph versus vortex methods. *Mathematics and Mechanics of complex systems*, 2(1):45–70, 2013.
- [136] Danilo Durante, Emanuele Rossi, A. Colagrossi, and Giorgio Graziani. Numerical simulations of the transition from laminar to chaotic behaviour of the planar vortex flow past a circular cylinder. *Communications in Nonlinear Science and Numerical Simulation*, 48:18–38, 2017.
- [137] L Ridgway Scott and Rebecca Durst. Chaotic dynamics of two-dimensional flows around a cylinder. *Physics of Fluids*, 36(2), 2024.
- [138] Henry von Wahl, Leo G Rebholz, and L Ridgway Scott. Benchmark stress tests for flow past a cylinder at higher reynolds numbers using emac. *arXiv preprint arXiv:2507.09274*, 2025.
- [139] Lixia Qu, Christoffer Norberg, Lars Davidson, Shia-Hui Peng, and Fujun Wang. Quantitative numerical analysis of flow past a circular cylinder at Reynolds number between 50 and 200. *J. Fluid Struct.*, 39:347–370, 2013.
- [140] O. Posdziech and R. Grundmann. Numerical simulation of the flow around an infinitely long circular cylinder in the transition regime. *Theor. Comp. Fluid Dyn.*, 15(2):121–141, 2001.
- [141] Bruno Souza Carmo and Julio Romano Meneghini. Numerical investigation of the flow around two circular cylinders in tandem. *J. Fluid Struct.*, 22(6-7):979–988, 2006.
- [142] DFL Labbé and PA Wilson. A numerical investigation of the effects of the spanwise length on the 3-D wake of a circular cylinder. *J. Fluid Struct.*, 23(8):1168–1188, 2007.
- [143] Suresh Behara and Sanjay Mittal. Parallel finite element computation of incompressible flows. *Parallel Computing*, 35(4):195–212, 2009.
- [144] BN Rajani, A Kandasamy, and Sekhar Majumdar. Numerical simulation of laminar flow past a circular cylinder. *Applied Mathematical Modelling*, 33(3):1228–1247, 2009.
- [145] Feifei Tong, Liang Cheng, and Ming Zhao. Numerical simulations of steady flow past two cylinders in staggered arrangements. *J. Fluid Mech.*, 765:114–149, 2015.
- [146] Chengwang Lei, Liang Cheng, and Ken Kavanagh. Spanwise length effects on three-dimensional modelling of flow over a circular cylinder. *Computer methods in applied mechanics and engineering*, 190(22-23):2909–2923, 2001.
- [147] Ming Zhao, Liang Cheng, and Tongming Zhou. Direct numerical simulation of three-dimensional flow past a yawed circular cylinder of infinite length. *J. Fluid Struct.*, 25(5):831–847, 2009.
- [148] O Lehmkuhl, I Rodríguez, R Borrell, and A Oliva. Low-frequency unsteadiness in the vortex formation region of a circular cylinder. *Physics of Fluids*, 25(8), 2013.
- [149] F Tremblay, M Manhart, and R Friedrich. LES of flow around a circular cylinder at a subcritical Reynolds number with cartesian grids. In *Advances in LES of Complex Flows: Proceedings of the Euromech Colloquium 412*, pages 133–150. Springer, 2002.

- [150] B. R. Noack and Marek Morzyński. The Fluidic PinBall - a Toolkit for Multiple-Input Multiple-Output Flow Control. *Technical Report 02/2016*, 2017.
- [151] Nan Deng, B. R. Noack, Marek Morzyński, and Luc R Pastur. Low-order model for successive bifurcations of the fluidic pinball. *J. Fluid Mech.*, 884:A37, 2020.
- [152] Luc R. Pastur, Nan Deng, Marek Morzyński, and B. R. Noack. Reduced-order modeling of the fluidic pinball. In *11th Chaotic Modeling and Simulation International Conference*, pages 205–213. Springer International Publishing, 2019.
- [153] Nan Deng, Luc R Pastur, B. R. Noack, and Marek Morzyński. Galerkin force model for transient and post-transient dynamics of the fluidic pinball. *J. Fluid Mech.*, 918:A4, 2021.
- [154] Xin Wang, Nan Deng, Guy Y Cornejo Maceda, and B. R Noack. Cluster-based control for net drag reduction of the fluidic pinball. *Physics of Fluids*, 35(2), 2023.
- [155] D. Sipp and A. Lebedev. Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows. *J. Fluid Mech.*, 593:333–358, 2007.
- [156] Thierry M Faure, Panayotis Adrianos, Flavio Lusseyran, and Luc Pastur. Visualization of the flow inside an open cavity at medium range reynolds numbers. *Experiments in Fluids*, 42(2):169–184, 2007.
- [157] A. Barbagallo, D. Sipp, and P. J. Schmid. Closed-loop control of an open cavity flow using reduced-order models. *J. Fluid Mech.*, 641:1–50, 2009.
- [158] Jesús de Vicente, Eusebio Valero, L. M González, and Vassilis Theofilis. Linear three-dimensional global and asymptotic stability analysis of incompressible open cavity flow. *J. Fluid Mech.*, 768:113–140, 2015.
- [159] M. F Barone and Srinivasan Arunajatesan. A cavity depth sensitized Rossiter mode formula. *Applied Mathematics and Computation*, 340:221–226, 2018.
- [160] JE Rossiter. Wind tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. *Aeronautical Research Council Reports and Memoranda*, 1964.
- [161] C. W Rowley and D. R Williams. Dynamics and Control of High-Reynolds-Number Flow over Open Cavities. *Ann. Rev. Fluid Mech.*, 38:251–276, 2006.
- [162] L. N Cattafesta III, Q. Song, D. R Williams, C. W Rowley, and F. S. Alvi. Active control of flow-induced cavity oscillations. *Prog. Aerosp. Sci.*, 44(7-8):479–502, 2008.
- [163] Ahmad Sohankar, Christoffer Norberg, and Lars Davidson. Simulation of three-dimensional flow around a square cylinder at moderate reynolds numbers. *Physics of Fluids*, 11(2):288–306, 1999.
- [164] Honglei Bai and Md Mahbub Alam. Dependence of square cylinder wake on Reynolds number. *Physics of Fluids*, 30(1), 2018.
- [165] Hongyi Jiang, Liang Cheng, and Hongwei An. Three-dimensional wake transition of a square cylinder. *J. Fluid Mech.*, 842:102–127, 2018.
- [166] F. X. Trias, A. Gorobets, and A. Oliva. Turbulent flow around a square cylinder at Reynolds number 22000: A DNS study. *Computers & Fluids*, 123:87–98, 2015.
- [167] Y. Srinivas, G. Biswas, A. S. Parihar, and R. Ranjan. Large-eddy simulation of high Reynolds number turbulent flow past a square cylinder. *Journal of Engineering Mechanics*, 132(3):327–336, 2006.
- [168] Zihao Zhu, Pau Fradera-Soler, William Jo, and Feng Liu. Numerical simulation of the flow around a square cylinder under plasma actuator control. *Physics of Fluids*, 33(12), 2021.
- [169] M. I. Yuce and D. A. Kareem. A numerical analysis of fluid flow around circular and square cylinders. *Journal - American Water Works Association*, 108(10):E546–E554, 2016.
- [170] W Rodi. Comparison of LES and RANS calculations of the flow around bluff bodies. *Journal of wind engineering and industrial aerodynamics*, 69:55–75, 1997.
- [171] Ahmad Sohankar, Christoffer Norberg, and Lars Davidson. Low-Reynolds-number flow around a square cylinder at incidence: study of blockage, onset of vortex shedding and outlet boundary condition. *International journal for numerical methods in fluids*, 26(1):39–56, 1998.
- [172] M Cheng, DS Whyte, and J Lou. Numerical simulation of flow around a square cylinder in uniform-shear flow. *J. Fluid Struct.*, 23(2):207–226, 2007.
- [173] Atsushi Okajima. Strouhal numbers of rectangular cylinders. *J. Fluid Mech.*, 123:379–398, 1982.

- [174] Hongyi Jiang and Liang Cheng. Hydrodynamic characteristics of flow past a square cylinder at moderate Reynolds numbers. *Physics of Fluids*, 30(10), 2018.
- [175] Christoffer Norberg. Flow around rectangular cylinders: pressure forces and wake frequencies. *Journal of wind engineering and industrial aerodynamics*, 49(1-3):187–196, 1993.
- [176] Shuyang Cao, Yaojun Ge, and Yukio Tamura. Shear effects on flow past a square cylinder at moderate Reynolds numbers. *Journal of engineering mechanics*, 138(1):116–123, 2012.
- [177] Dong-Hyeog Yoon, Kyung-Soo Yang, and Choon-Bum Choi. Three-dimensional wake structures and aerodynamic coefficients for flow past an inclined square cylinder. *Journal of Wind Engineering and Industrial Aerodynamics*, 101:34–42, 2012.
- [178] D Mateescu and M Abdo. Analysis of flows past airfoils at very low Reynolds numbers. *Proc. Inst. Mech. Eng. G*, 224(7):757–775, 2010.
- [179] Eltayeb Eljack, Julio Soria, Yasir Elawad, and Tomohisa Ohtake. Simulation and characterization of the laminar separation bubble over a NACA-0012 airfoil as a function of angle of attack. *Phys. Rev. Fluids*, 6(3):034701, 2021.
- [180] Jaber AlMutairi, Eltayeb ElJack, and Ibraheem AlQadi. Dynamics of laminar separation bubble over naca-0012 airfoil near stall conditions. *Aerospace Science and Technology*, 68:193–203, 2017.
- [181] Ivette Rodríguez, O. Lehmkuhl, R Borrell, and A Oliva. Direct numerical simulation of a NACA0012 in full stall. *Int. J. Heat & Fluid Flow*, 43:194–203, 2013.
- [182] Hua Shan, Li Jiang, and Chaoqun Liu. Direct numerical simulation of flow separation around a NACA 0012 airfoil. *Computers & fluids*, 34(9):1096–1114, 2005.
- [183] Sergio Martínez-Aranda, AL García-González, Luis Parras, JF Velázquez-Navarro, and C Del Pino. Comparison of the aerodynamic characteristics of the NACA0012 airfoil at low-to-moderate Reynolds numbers for any aspect ratio. *International Journal of Aerospace Sciences*, 4(1):1–8, 2016.
- [184] Sam Zarovy, Mark Costello, Ankur Mehta, Greg Gremillion, Derek Miller, Badri Ranganathan, J Sean Humbert, and Paul Samuel. Experimental study of gust effects on micro air vehicles. In *AIAA atmospheric flight mechanics conference*, page 7818, 2010.
- [185] Herbert G Küssner. Zusammenfassender Bericht über den instationären Auftrieb von Flügeln. *Luftfahrtforschung*, 13:410, 1936.
- [186] Kai Fukami, Hiroya Nakao, and Kunihiro Taira. Data-driven transient lift attenuation for extreme vortex gust–airfoil interactions. *J. Fluid Mech.*, 992:A17, 2024.
- [187] Kenneth Granlund, Bruno Monnier, M Ol, and D. Williams. Airfoil longitudinal gust response in separated vs. attached flows. *Physics of Fluids*, 26(2), 2014.
- [188] M. Gondrum, M. Meinke, W. Schroeder, F. Avallone, and D. Ragni. Porous fairings for landing gear noise mitigation. In *30th AIAA/CEAS Aeroacoustics Conference (2024)*, page 3174, 2024.
- [189] Giovanni Di Ilio, Daniele Chiappini, Stefano Ubertini, Gino Bella, and Sauro Succi. Fluid flow around NACA 0012 airfoil at low-Reynolds numbers with hybrid lattice Boltzmann method. *Computers & Fluids*, 166:200–208, 2018.
- [190] Yan Liu, Kailun Li, Jiazhong Zhang, Hang Wang, and Liguang Liu. Numerical bifurcation analysis of static stall of airfoil and dynamic stall under unsteady perturbation. *Communications in Nonlinear Science and Numerical Simulation*, 17(8):3427–3434, 2012.
- [191] Dilek Funda Kurtulus. On the unsteady behavior of the flow around NACA 0012 airfoil with steady external conditions at $Re = 1000$. *International journal of micro air vehicles*, 7(3):301–326, 2015.
- [192] Taiba Kouser, Yongliang Xiong, Dan Yang, and Sai Peng. Direct Numerical Simulations on the three-dimensional wake transition of flows over NACA0012 airfoil at $Re = 1000$. *International Journal of Micro Air Vehicles*, 13:17568293211055656, 2021.
- [193] Y. Hoarau, M. Braza, Y. Ventikos, and D. Faghani. First stages of the transition to turbulence and control in the incompressible detached flow around a NACA0012 wing. *Int. J. Heat & Fluid Flow*, 27(5):878–886, 2006.
- [194] Ricardo Vinuesa, Prabal Singh Negi, Marco Atzori, Ardeshtir Hanifi, Dan S Henningson, and Philipp Schlatter. Turbulent boundary layers around wing sections up to $Re_c = 1,000,000$. *Int. J. Heat and*

Fluid Flow, 72:86–99, August 2018.

- [195] Álvaro Tanarro, Ricardo Vinuesa, and Philipp Schlatter. Effect of adverse pressure gradients on turbulent wing boundary layers. *J. Fluid Mech.*, 883:A8, January 2020.
- [196] Prabal Singh Negi, Ricardo Vinuesa, Ardeshir Hanifi, Philipp Schlatter, and Dan S Henningson. Unsteady aerodynamic effects in small-amplitude pitch oscillations of an airfoil. *Int. J. Heat and Fluid Flow*, 71:378–391, 2018.
- [197] Florian R Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA J.*, 32(8):1598–1605, 1994.
- [198] Suchuan Dong, George E Karniadakis, and Chrysostomos Chrysostomidis. A robust and accurate outflow boundary condition for incompressible flow simulations on severely-truncated unbounded domains. *J. Comput. Phys.*, 261:83–105, 2014.
- [199] Philipp Schlatter and Ramis Örlü. Turbulent boundary layers at moderate Reynolds numbers: inflow length and tripping effects. *J. Fluid Mech.*, 710:5–34, 2012.
- [200] Ricardo Vinuesa, Seyed M Hosseini, Ardeshir Hanifi, Dan S Henningson, and Philipp Schlatter. Pressure-gradient turbulent boundary layers developing around a wing section. *Flow, turbulence and combustion*, 99:613–641, 2017.
- [201] Yuning Wang, Marco Atzori, and Ricardo Vinuesa. Opposition control applied to turbulent wing sections. *J. Fluid Mech.*, 1010:A29, May 2025.
- [202] Arun K Saha. Three-dimensional numerical simulations of the transition of flow past a cube. *Physics of Fluids*, 16(5):1630–1646, 2004.
- [203] Lukasz Klotz, S Goujon-Durand, Jacek Rokicki, and JE Wesfreid. Experimental investigation of flow behind a cube for moderate Reynolds numbers. *J. Fluid Mech.*, 750:73–98, 2014.
- [204] Qingrui Meng, Hongwei An, Liang Cheng, and Mehrdad Kimiaei. Wake transitions behind a cube at low and moderate Reynolds numbers. *J. Fluid Mech.*, 919:A44, 2021.
- [205] AK Saha. Three-dimensional numerical study of flow and heat transfer from a cube placed in a uniform flow. *Int. J. Heat & Fluid Flow*, 27(1):80–94, 2006.
- [206] A Haider and O Levenspiel. Drag coefficient and terminal velocity of spherical and nonspherical particles. *Powder technology*, 58(1):63–70, 1989.
- [207] A. Hölzer and Martin Sommerfeld. New simple correlation formula for the drag coefficient of non-spherical particles. *Powder Technology*, 184(3):361–365, 2008.
- [208] Majid Hassan Khan, Atul Sharma, and Amit Agrawal. Simulation of flow around a cube at moderate reynolds numbers using the lattice boltzmann method. *Journal of Fluids Engineering*, 142(1):011301, 2020.
- [209] Majid Hassan Khan, P Sooraj, Atul Sharma, and Amit Agrawal. Flow around a cube for Reynolds numbers between 500 and 55,000. *Experimental thermal and fluid science*, 93:257–271, 2018.
- [210] TA Johnson and VC Patel. Flow past a sphere up to a Reynolds number of 300. *J. Fluid Mech.*, 378:19–70, 1999.
- [211] Ananias G Tomboulides and S.A Orszag. Numerical investigation of transitional and weak turbulent flow past a sphere. *J. Fluid Mech.*, 416:45–73, 2000.
- [212] Vincenzo Citro, Lorenzo Siconolfi, D. Fabre, Flavio Giannetti, and Paolo Luchini. Stability and sensitivity analysis of the secondary instability in the sphere wake. *AIAA Journal*, 55(11):3661–3668, 2017.
- [213] Ivette Rodriguez, Ricard Borell, O. Lehmkuhl, Carlos D Perez Segarra, and Assensi Oliva. Direct numerical simulation of the flow over a sphere at $Re = 3700$. *J. Fluid Mech.*, 679:263–287, 2011.
- [214] E. K. W. Poon, A. S. H. Ooi, M. Giacobello, G. Iaccarino, and D. Chung. Flow past a transversely rotating sphere at Reynolds numbers above the laminar regime. *J. Fluid Mech.*, 759:751–781, 2014.
- [215] Frederick W Roos and William W Willmarth. Some experimental results on sphere and disk drag. *AIAA Journal*, 9(2):285–291, 1971.
- [216] Jaber Almedeij. Drag coefficient of flow around a sphere: Matching asymptotically the wide trend. *Powder Technology*, 186(3):218–223, 2008.
- [217] R. Clift, J. R. Grace, and M. E. Weber. *Bubbles, drops, and particles*. Courier Corporation, 2005.

- [218] R Turton and O Levenspiel. A short note on the drag correlation for spheres. *Powder technology*, 47(1):83–86, 1986.
- [219] Giwoong Yun, Dongjoo Kim, and Haeccheon Choi. Vortical structures behind a sphere at subcritical Reynolds numbers. *Physics of Fluids*, 18(1), 2006.
- [220] Stephen J. Kline, William C. Reynolds, Fred A. Schraub, and Peter W. Runstadler. The structure of turbulent boundary layers. *Journal of Fluid Mechanics*, 30(4):741–773, 1967.
- [221] Jinhee Jeong, Fazle Hussain, Wade Schoppa, and John Kim. Coherent structures near the wall in a turbulent channel flow. *Journal of Fluid Mechanics*, 332:185–214, 1997.
- [222] James M. Hamilton, John Kim, and Fabian Waleffe. Regeneration mechanisms of near-wall turbulence structures. *Journal of Fluid Mechanics*, 287:317–348, 1995.
- [223] Javier Jiménez and Alfredo Pinelli. The autonomous cycle of near-wall turbulence. *Journal of Fluid Mechanics*, 389:335–359, 1999.
- [224] I Marusic, R Mathis, and N Hutchins. Predictive model for wall-bounded turbulent flow. *Science*, 329(5988):193–196, 2010.
- [225] Alexander J. Smits, Beverley J. McKeon, and Ivan Marusic. High-Reynolds number wall turbulence. *Annual Review of Fluid Mechanics*, 43:353–375, 2011.
- [226] Nicholas Hutchins and Ivan Marusic. Evidence of very long meandering features in the logarithmic region of turbulent boundary layers. *Journal of Fluid Mechanics*, 579:1–28, 2007.
- [227] Koji Fukagata, Kaoru Iwamoto, and Nobuhide Kasagi. Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Physics of Fluids*, 14(11):L73–L76, 2002.
- [228] Philipp Schlatter and Ramis Örlü. Assessment of direct numerical simulation data of turbulent boundary layers. *Journal of fluid mechanics*, 659:116–126, 2010.
- [229] Seyed Mohammad Hosseini, Ricardo Vinuesa, Philipp Schlatter, Ardeshir Hanifi, and Dan S Henningson. Direct numerical simulation of the flow around a wing section at moderate reynolds number. *International Journal of Heat and Fluid Flow*, 61:117–128, 2016.
- [230] Sonu S. Varghese, S.H. Frankel, and Paul F. Fischer. Direct numerical simulation of stenotic flows. part 1. steady flow. *J. Fluid Mech.*, 582:253–280, 2007.
- [231] M. Waldmann, M. Rüttgers, A. Lintermann, and W. Schröder. Virtual Surgeries of Nasal Cavities Using a Coupled Lattice-Boltzmann–Level-Set Approach. *Journal of Engineering and Science in Medical Diagnostics and Therapy*, 5(3), 2022.
- [232] Sonu S. Varghese, S.H. Frankel, and Paul F. Fischer. Direct numerical simulation of stenotic flows. part 2. pulsatile flow. *J. Fluid Mech.*, 582:281–318, 2007.
- [233] Kartik Jain. Transition to turbulence in an oscillatory flow through stenosis. *Biomechanics and Modeling in Mechanobiology*, pages 1–19, 02 2020.
- [234] Sergio Hoyas and Javier Jiménez. Reynolds number effects on the reynolds-stress budgets in turbulent channels. *Physics of Fluids*, 20(10):101511, 10 2008.
- [235] L. Guastoni et al. Multi-agent reinforcement learning for wall-bounded turbulence control. *Nat. Comm.*, 14, 2023.
- [236] K. S. Nelson and O. B. Fringer. Reducing spin-up time for simulations of turbulent channel flow. *Physics of Fluids*, 29(10):105101, 10 2017.
- [237] Ville Vuorinen and K Keskinen. Dnslab: A gateway to turbulent flow simulation in matlab. *Computer Physics Communications*, 203:278–289, 2016.
- [238] Jason Rader, Terry Lyons, and Patrick Kidger. Lineax: unified linear solves and linear least-squares in jax and equinox, 2023.
- [239] Robert D Moser, John Kim, Nagi N Mansour, et al. Direct numerical simulation of turbulent channel flow up to $Re = 590$. *Phys. fluids*, 11(4):943–945, 1999.
- [240] E. P. Hammond, T. R. Bewley, and P. Moin. Observed mechanisms for turbulence attenuation and enhancement in opposition-controlled wall-bounded flows. *Physics of Fluids*, 10(9):2421–2423, 09 1998.
- [241] A. Stroh, B. Frohnäpfel, P. Schlatter, and Y. Hasegawa. A comparison of opposition control in turbulent

- boundary layer and turbulent channel flow. *Physics of Fluids*, 27(7):075101, 07 2015.
- [242] E. Lagemann, S. L. Brunton, and C. Lagemann. Uncovering wall-shear stress dynamics from neural-network enhanced fluid flow measurements. *Proc. Roy. Soc. A*, 480(2292):20230798, 2024.
 - [243] Esther Lagemann, Julia Roeb, Steven L Brunton, and Christian Lagemann. A deep learning approach to wall-shear stress quantification: from numerical training to zero-shot experimental application. *Journal of Fluid Mechanics*, 1014:A1, 2025.
 - [244] Mohammad Farazmand and Themistoklis P. Sapsis. Closed-loop adaptive control of extreme events in a turbulent flow. *Phys. Rev. E*, 100:033110, Sep 2019.
 - [245] H.K. Moffatt. Extreme events in turbulent flow. *J. Fluid Mech.*, 914:F1, 2021.
 - [246] A. Blanchard and T. P. Sapsis. Stabilization of unsteady flows by reduced-order control with optimally time-dependent modes. *Phys. Rev. Fluids*, 4:053902, 2019.
 - [247] Gaurav Kumar and Aditya Nair. Stabilizing two-dimensional turbulent kolmogorov flow via selective modification of inviscid invariants. *Physics of Fluids*, 35(12), 2023.
 - [248] Guido Boffetta and Robert E. Ecke. Two-dimensional turbulence. *Ann. Rev. Fluid Mech.*, 44, 2012.
 - [249] Z. Yin, H.J.H. Clercx, and D.C. Montgomery. An easily implemented task-based parallel scheme for the Fourier pseudospectral solver applied to 2D Navier–Stokes turbulence. *Computers & Fluids*, 33(4):509–520, 2004.
 - [250] Gideon Dresdner, Dmitrii Kochkov, P.C. Norgaard, L. Zepeda-Nunez, Jamie Smith, Michael Brenner, and Stephan Hoyer. Learning to correct spectral methods for simulating turbulent flows. *Transactions on Machine Learning Research*, 2023.
 - [251] R. Byron Bird, Warren E. Stewart, and Edwin N. Lightfoot. *Transport Phenomena*. John Wiley & Sons, New York, 2nd edition, 2002.
 - [252] Arnab K. Das, Tapan K. Mankodi, and Ujjwal K. Saha. Thrust vectoring of double divergent nozzle using secondary fluid injection. *Journal of Propulsion and Power*, 41:486–502, 2025.
 - [253] Kenrick A. Waithe and Karen A. Deere. Experimental and computational investigation of multiple injection ports in a convergent-divergent nozzle for fluidic thrust vectoring. In *21st AIAA Applied Aerodynamics Conference*. American Institute of Aeronautics and Astronautics Inc., 2003.
 - [254] Rafael Borges, Monique Carmona, Bruno Costa, and Wai Sun Don. An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws. *Journal of Computational Physics*, 227:3191–3211, 2008.
 - [255] Eleuterio F Toro, M Spruce, and W Speares. Restoration of the contact surface in the hll-riemann solver. *Shock waves*, 4:25–34, 1994.
 - [256] Sigal Gottlieb and Chi-Wang Shu. Total variation diminishing runge-kutta schemes. *Mathematics of computation of the American Mathematical Society*, 67:73–85, 1998.
 - [257] X Y Hu, B C Khoo, N A Adams, and F L Huang. A conservative interface method for compressible flows. *J. Comput. Phys.*, 219:553–578, 2006.