HydroGym: A Platform for Reinforcement Learning in Fluid Dynamics

Samuel Ahnert[*,†], Ludger Paehler[+], Jared Callaham[*], Sajeda Mokbel[*,†],
Nikolaus Adams[+], Christian Lagemann[†], Esther Lagemann[†], and Steven L. Brunton[*,†]

The modelling and control of fluid flows is a canonically challenging problem due to the strong non-linearity, high-dimensionality, and multi-scale physics present. Interest in this problem is further driven by the opportunity to lower the consumption of resources in industry. For example, improved flow control may enable drag reduction, mixing enhancement, and noise reduction for industrial-scale flows, offering the potential to save immensely on costs across domains as diverse as transportation, medicine, and energy. Reinforcement learning has rapidly emerged as a highly effective control paradigm in environments as complex as robotics, games, and protein folding. A key accelerant of these advances is the foundation of scalable reinforcement learning frameworks and standardized environments that allow for direct comparison between learned policies. In contrast, progress in reinforcement learning for flow control has been limited by the scarcity of such platforms, and it is further hindered by the burden of running and validating computationally intensive CFD. To this end, we present HydroGym: a scalable and extensible platform that closes the loop between efficient flow solvers, sophisticated flow control benchmark problems, and state-of-the-art reinforcement learning.

The four initial flow environments presented are all evaluated on a broad set of reinforcement learning algorithms with run-times which scale seamlessly from laptops to high performance computing resources. The included environments represent a diverse set of challenging dynamics commonly found in complex fluid flows, and the framework is intentionally designed to be solver-flexible. This allows future benchmarks to encompass even more complicated dynamics than what is currently implemented, enabling researchers to push the applications and limits of modern reinforcement learning for flow control. Additionally, we emphasize our choice to ensure HydroGym is compatible with a standardized reinforcement learning interface, thereby promoting participation in flow control research by domain experts outside of fluid mechanics and facilitating the interdisciplinary advancement of scientific machine learning.

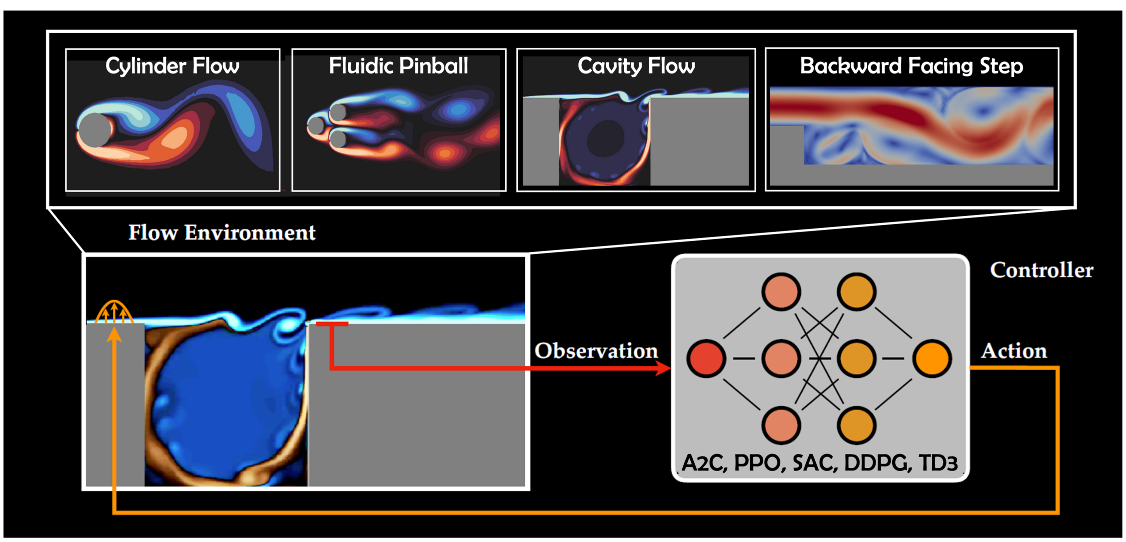

Figure 1: Flow configurations and benchmark algorithms included in the initial release of HydroGym.

[*]Department of Mechanical Engineering, University of Washington, Seattle, USA
[†]AI Institute in Dynamic Systems, University of Washington, Stevens Way NE, Seattle, United States
[+]School of Engineering and Design, Technical University of Munich, Munich, GER