DFD19-2019-002016

Abstract for an Invited Paper
for the DFD19 Meeting of
the American Physical Society

Deep Reinforcement Learning for Flow Control.[1]

PETROS KOUMOUTSAKOS, ETH Zurich

Reinforcement Learning (RL) is a mathematical framework for problem solving that implies goal-directed interactions of an agent with its environment. The agent has a repertoire of actions, perceives states and learns a policy from its experiences in the form of rewards. Deep recurrent neural networks can be used to encode the state-action policy of the agent and effective training implies a selective choice of experiences. RL does not require a model of the dynamics, such as a Markov transition model, and relies instead on repeated interactions of the agent with the environment. I consider that this approximation makes it highly suitable for complex problems in fluid dynamics, albeit, presently, only when significant computing power or experimental automation is possible. Deep RL has been very successful in games (from backgammon to video games and the game of Go) and robotics, but remains rather unexplored in fluid mechanics. I hope to illustrate some of the strengths and limitations of Deep RL for flow control using simulations of perching and gliding bodies as well collective fish swimming.

[1]European Research Council Advanced Investigator Award 341117, Swiss National Science Foundation, Swiss National Supercomputing Centre (CSCS)