Abstract Submitted
for the DFD19 Meeting of
The American Physical Society

Airfoil control with Proximal Policy Optimization DENIS DUMOULIN, PHILIPPE CHATELAIN, UCLouvain — Airfoil control generally relies on techniques based on a dynamical model of the actual system one wants to control. Such methods are in some cases limited by the expressiveness of the model and its linearization around equilibrium points. This is the case in aerodynamics where non-linearities (e.g., turbulent structures or dynamic stall) strongly affect the flight conditions and more specifically aerodynamic loads undergone by the flying body. To avoid such limitations, we propose the use of a model-free control method based on Reinforcement Learning (RL) algorithms. This study relies on a low-order inviscid planar solver which accounts for separation; this model controls its computational cost through lumping of the shed vortices in the wake. The complexity of the model is critical for RL methods which require large amounts of data. We then use Proximal Policy Optimization (PPO) to control the dynamics of a NACA0012 airfoil. This planar body goes through a random train of vortical dipoles while trying to keep its aerodynamic coefficients constant. Within a few thousands of episodes of training, the controller is able to stabilize the airfoil regardless of the incoming vortices using pressure measurements distributed on the body and its kinematics.

Denis Dumoulin
UCLouvain

Date submitted: 31 Jul 2019 Electronic form version 1.4