



Model-based deep reinforcement learning for accelerated learning from flow simulations

Andre Weiner · Janis Geise

Received: 26 February 2024 / Accepted: 17 April 2024
© The Author(s) 2024

Abstract In recent years, deep reinforcement learning has emerged as a technique to solve closed-loop flow control problems. Employing simulation-based environments in reinforcement learning enables a priori end-to-end optimization of the control system, provides a virtual testbed for safety-critical control applications, and allows to gain a deep understanding of the control mechanisms. While reinforcement learning has been applied successfully in a number of rather simple flow control benchmarks, a major bottleneck toward real-world applications is the high computational cost and turnaround time of flow simulations. In this contribution, we demonstrate the benefits of model-based reinforcement learning for flow control applications. Specifically, we optimize the policy by alternating between trajectories sampled from flow simulations and trajectories sampled from an ensemble of environment models. The model-based learning reduces the overall training time by up to 85% for the fluidic pinball test case. Even larger savings are expected for more demanding flow simulations.

Andre Weiner and Janis Geise have contributed equally to this work.

A. Weiner (✉) · J. Geise
Institute of Fluid Mechanics, TU Dresden,
George-Bähr-Str. 3c, Dresden 01069, Saxony, Germany
e-mail: andre.weiner@tu-dresden.de

J. Geise
e-mail: janis.geise@tu-dresden.de

Keywords Closed-loop flow control · Model ensemble proximal policy optimization · Fluidic pinball

Abbreviations

AFC	Active flow control
CFD	Computational fluid dynamics
DRL	Deep reinforcement learning
IDDES	Improved delayed detached eddy simulation
MB	Model-based
MEPPO	Model ensemble proximal policy optimization
METRPO	Model ensemble trust region policy optimization
MF	Model-free
RL	Reinforcement learning

1 Introduction

According to the 2022 IPCC report, smart control technologies are a key enabler in mitigating the impact of global warming and climate change [1]. Control tasks in which the manipulation of fluid flows can significantly reduce the carbon dioxide footprint are ubiquitous in society and industries, e.g., reducing aerodynamic forces acting on vehicles [2], maximizing the coefficient of performance of heat pumps [3], or adjusting the operating conditions of process engineering systems in response to the availability

of renewable energy resources [4], to name a few. To achieve optimal control in the aforementioned example applications, it is necessary that the controller can respond to changes in the ambient conditions. However, the design and implementation of closed-loop active flow control (AFC) systems is highly complex. As an example, consider the flow past a truck. At highway speeds, such a flow is fully 3D and comprises turbulent boundary layers, flow separation, reattachment, and large turbulent wakes [5, 6]. To implement closed-loop AFC, this complexity must be captured with a limited amount of sensors placed on the vehicle's surface. Similarly, the actuators must be designed and placed sensibly to control the dominant flow structures. Moreover, a suitable control law must be derived, i.e., a mapping from the sensor reading to the optimal actuation. Finally, the nonlinear flow dynamics create a tight coupling between sensors, actuators, and control law. Hence, the AFC system should be designed and optimized as a whole.

Recently, reinforcement learning (RL), and in particular deep RL (DRL), started to emerge as a technology in fluid mechanics with the potential to tackle the complexity of closed-loop AFC. Viquerat et al. [7] provide an extensive account of DRL-based applications in fluid mechanics. DRL solves control problems through trial-and-error interactions with an environment, where *environment* is the DRL term for the system to be controlled. The unique advantage of learning from simulation-based environments is the possibility of end-to-end optimization of the full control system, i.e., sensor locations, control law, and actuator positions [8–10]. However, a critical challenge of simulation-based RL is the high computational cost and turnaround time of flow simulations. Even for relatively simple flow control benchmark problems, state-of-the-art RL algorithms require $O(100)$ episodes to converge [7], which is far from practical for real-world AFC applications.

To emphasize the need for sample-efficient learning, consider again the application of DRL-based closed-loop AFC to vehicles. The current gold standard for accurate and fast simulations of vehicle aerodynamics are hybrid approaches like improved delayed detached eddy simulations (IDDES) [11]. Computing converged mean force coefficients for the *DrivAer* test case at $Re = 768000$, required approximately two days on 700 CPU cores in 2016 [11]. It is conceivable that roughly five sequential IDDES can

be performed per day because it is not necessary to obtain converged statistics in every run. Moreover, hardware and scalability have improved significantly over the past years. Assuming that 10 simulations are performed in parallel per episode to generate more data, each simulation being run on 1000 CPU cores, the computational cost for 100 episodes is approximately 5 M CPU hours, and the training lasts 20 days. Cloud providers charge 0.01–0.05 euros per CPU hour, which puts a price tag of 0.5–2.5M euros on a single training at a single flow condition. Additional hyperparameter optimization, sensor/actuator variants, and an extension to multiple flow conditions can easily increase the cost by a factor of 10–100. This cost, combined with a turnaround time of 20 days, is far from practical for the majority of potential users.

Several remedies exist to reduce the computational cost and turnaround time, e.g., by exploiting invariances [12, 13] or by pre-training on simulations with very coarse meshes [14]. While effective, the applicability of the aforementioned techniques strongly depends on the control problem, i.e., the problem must exhibit invariances, and the mesh coarsening must not change the flow characteristics. A more general approach to improve data efficiency is model-based DRL (MBDRL) [15]. The main idea is to substitute the expensive simulation-based environment with one or more surrogate models. The surrogate models are optimized based on data coming from the high-fidelity environment. Additional data to optimize the control law can be created with little effort by querying the optimized surrogate model. A plethora of MBDRL algorithms exist that differ in the type of surrogate model and how the control law is derived from the model. Moerland et al. [15] provide an extensive overview.

Key challenges in MBDRL are (i) the efficient creation of accurate surrogate models and (ii) dealing with the presence of model error. Due to their flexibility, neural networks are a common choice for building data-driven surrogate models. However, the models must be created on the fly based on data that is unavailable for prior tests, so the optimization pipeline must be robust and efficient. Moreover, the models are auto-regressive. Hence, long-term stability is a persistent challenge. If the prediction accuracy is not sufficiently high, querying the models repeatedly leads to vanishing or exploding solutions. Finally, even if the model creation workflow is robust

and efficient, there is always epistemic uncertainty due to the alternating iterative updates of environment models and control law. With each update of the control law, previously unknown actuations and states will emerge, and consequently, the environment models will eventually become too inaccurate. Estimating whether the surrogate models are still reliable or should be updated with high-fidelity data from the true environment is essential to making MBDRL robust.

In this contribution, we present a modified version of the model ensemble trust region policy optimization (METRPO) [16] algorithm and demonstrate the benefits of MBDRL for fluid-mechanical applications. Specifically, we compare model-free (MF) and model-based (MB) learning on the flow past a circular cylinder [17, 18] and the fluidic pinball configuration [9, 19]. The appendix describes various implementation details and serves as a reference for the abbreviations used throughout the text. Moreover, we provide a complimentary code repository with instructions to reproduce the numerical experiments and data analyses presented hereafter [20]. A persistent record of all scientific results, including a snapshot of the code repository, is publicly accessible [21].

2 Theory

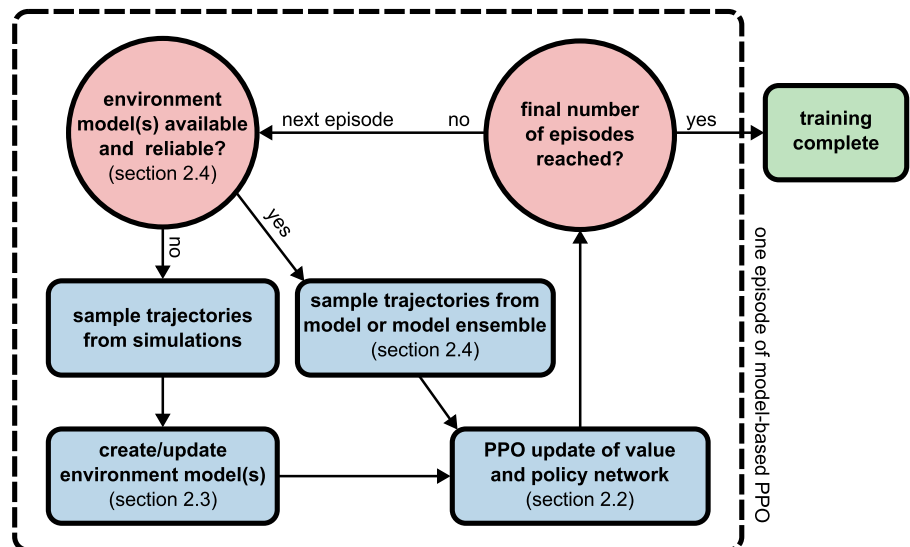
This section starts with a very short introduction to essential RL concepts and then describes the MF

algorithm, proximal policy optimization (PPO), employed to solve the flow control problems. Thereafter, we explain the creation of model ensembles, which emulate the simulation outputs, and how these ensembles are employed in an MBDRL variant of PPO. Figure 1 shows how the different algorithmic components are connected within one episode of model-based PPO.

2.1 Reinforcement learning

There are two main entities in RL: the agent and the environment. The agent contains the control logic, while the environment represents the simulation or experiment that shall be controlled. The environment at a discrete time step $n \in \mathbb{N}$ is characterized by a state $S_n \in \mathcal{S}$, where \mathcal{S} is the space of all possible states. Formally, there is a difference between the system's state and a partial observation of the state, since the full state is often inaccessible. Nonetheless, we follow standard practice and refer to partial observations and full states alike simply as states. In fluid mechanics, S_n could represent instantaneous pressure or velocity values at one or more sensor locations in the region of interest. The agent can manipulate the state via an action $A_n \in \mathcal{A}$, where \mathcal{A} is the space of all possible actions, upon which the environment transitions to a new state S_{n+1} . Common actions (means of actuation) for flows are suction, blowing, rotation, or heating. The control objective in RL is expressed by the reward $R_{n+1} \in \mathcal{R}$,

Fig. 1 High-level overview of one model-based PPO episode



where \mathcal{R} is the space of all possible rewards, and is assessed based on the transition from S_n to S_{n+1} . For example, if the control objective is drag reduction, the reward would be defined such that large drag forces yield low rewards while small drag forces yield high rewards. The combination of S_n , A_n , R_{n+1} , and S_{n+1} is called an experience tuple [22]. The aggregation of experience tuples over N interactions between agent and environment forms a so-called trajectory [22], $\tau = [(S_0, A_0, R_1, S_1), \dots, (S_{N-1}, A_{N-1}, R_N, S_N)]$. Trajectories form the basis for optimizing the policy $\pi(A_n = a | S_n = s)$ (the control law), which predicts the probability of taking some action a given that the current state is s . Optimality in RL is defined in terms of cumulative rewards, the so-called return $G_n = \sum_{i=n+1}^N \gamma^i R_i$, because the system's response to an action may be delayed. The discount factor $\gamma \in (0, 1]$ introduces a notion of urgency, i.e., the same reward counts more if achieved early in a trajectory. Given the same policy π and the same state S_n , the return G_n might vary between trajectories. Incomplete knowledge of the state is one common source of uncertainty. To consider uncertainty in the optimization, RL ultimately aims to find the policy that maximizes the expected return $v_\pi(s)$ under policy π over all possible states $s \in \mathcal{S}$ [22]:

$$v_\pi^*(s) = \max_{\pi} \underbrace{\mathbb{E}_\pi[G_n | S_n = s]}_{v_\pi(s)} \tag{1}$$

where $\mathbb{E}_\pi[\cdot]$ denotes the expected value under policy π , e.g., the mean value for a normally distributed random variable, and $v_\pi^*(s)$ is the optimal state-value function. Equation (1) forms the basis for all RL algorithms.

2.2 Policy learning

Expression (1) implicitly defines the optimal policy $\pi^*(a|s)$, which is the policy associated with the optimal value function $v_\pi^*(s)$. However, additional algorithmic building blocks are required to derive a workable policy. PPO is the most common algorithm to solve problem (1) for fluid-mechanical systems [7]. PPO employs deep neural networks as ansatz for the value function $v_\pi(s)$ and the policy $\pi(a|s)$. Hence, PPO belongs to the group of actor-critic DRL algorithms. PPO is relatively straightforward

to implement and enables accelerated learning from multiple trajectories, i.e., multiple trajectories may be sampled and processed in parallel rather than sequentially to produce new training data $T = [\tau_1, \tau_2, \dots, \tau_K]$. It is important to note that there are many PPO details that vary across different implementations. Many of these details are not thoroughly described in the standard PPO reference article [23], but they are nonetheless of utmost importance. Andrychowicz et al. [24] provide guidelines based on comprehensive test results for a multitude of algorithmic variants. Our implementation is based on the textbook by Morales [25]. We outline only those elements needed for the MBDRL version.

Based on the trajectories, the free parameters θ_v of the parametrized value network $v_{\theta_v}(s)$ are optimized to predict the expected return over all trajectories. The value function's subscript π has been dropped to simplify the notation. The loss function includes some additional clipping to limit the change between two updates [25]:

$$L_{\text{val}} = \frac{1}{KN} \sum_{k=1}^K \sum_{n=0}^{N-1} \max \left\{ \left(G_n^k - v_{\theta_v}(S_n^k) \right)^2, \left(G_n^k - V_{\text{clip},n}^k \right)^2 \right\},$$

$$V_{\text{clip},n}^k = v_{\theta_v}^{\text{old}}(S_n^k) + \text{clip} \left\{ v_{\theta_v}(S_n^k) - v_{\theta_v}^{\text{old}}(S_n^k), -\delta, \delta \right\}, \tag{2}$$

where $v_{\theta_v}^{\text{old}}$ is the value network's state before optimization, and the clip function is defined as:

$$\text{clip}(x, -\delta, \delta) = \begin{cases} x, & \text{for } -\delta \leq x \leq \delta \\ -\delta, & \text{for } x < -\delta \\ \delta, & \text{for } x > \delta \end{cases}. \tag{3}$$

Similar to the value function network, the policy network $\pi_{\theta_\pi}(a|s)$ parametrizes a multivariate Beta distribution over possible actions. During exploration, actions are sampled at random from this distribution, i.e., $A_n \sim \pi_{\theta_\pi}(S_n)$. We do not consider a potential cross-correlation between actions for exploration. To evaluate if a sampled action performs better than expected, the reward is compared to the expected reward, i.e. [23, 25]:

$$\delta_n = R_n - (v_{\theta_v}(S_n) - \gamma v_{\theta_v}(S_{n+1})), \tag{4}$$

where δ_n is called the 1-step advantage estimate at time step n . Actions with a positive advantage estimate should be favored. Rather than looking one step ahead, one could also look l steps ahead to consider delayed effects. The generalized advantage estimate \hat{A} balances between short-term and long-term effects using the hyperparameter $\lambda \in (0, 1]$ [23, 25]:

$$\hat{A}_n = \sum_{l=0}^{N-n} (\gamma \lambda)^l \delta_{n+l}. \tag{5}$$

Finally, the free parameters of the policy network, θ_π , are optimized to make favorable actions ($\hat{A}_n > 0$) more likely and poor actions ($\hat{A}_n \leq 0$) less likely. This concept is combined with a clipping mechanism that restricts policy changes [25]:

$$L_{\text{pol}} = -\frac{1}{KN} \sum_{k=1}^K \sum_{n=0}^{N-1} \min \left\{ \frac{\pi_{\theta_\pi}(A_n^k | S_n^k)}{\pi_{\theta_\pi^{\text{old}}}(A_n^k | S_n^k)} \hat{A}_n^k, \text{clip} \left\{ \frac{\pi_{\theta_\pi}(A_n^k | S_n^k)}{\pi_{\theta_\pi^{\text{old}}}(A_n^k | S_n^k)}, -\epsilon, \epsilon \right\} \hat{A}_n^k \right\} - \frac{\beta}{KN} \sum_{k=1}^K \sum_{n=0}^{N-1} E[\pi_{\theta_\pi}(S_n^k)] \tag{6}$$

where $\pi_{\theta_\pi} / \pi_{\theta_\pi^{\text{old}}}$ is the probability ratio between the current and old (before optimization) policy, ϵ is the clipping parameter, $E(\pi)$ is the policy’s entropy [25], and β is the entropy’s weighting factor. The additional entropy term avoids a premature stop of the exploration [15].

One sequence of generating trajectories, updating the value network according to Eq. (2), and updating the policy according to Eq. (6) is called an episode. Since the optimization starts with randomly initialized value and policy networks, multiple episodes are required to find the optimal policy.

2.3 Model learning

The type of environment model employed here is a simple feed-forward neural network with weights θ_m that maps from the last $d + 1$ states and a given action to the next state and the received reward:

$$m_{\theta_m} : (S_{n-d}, \dots, S_{n-1}, S_n, A_n) \rightarrow (S_{n+1}, R_{n+1}). \tag{7}$$

To simplify the notation, we introduce the current extended state \hat{S}_n , which comprises the last $d + 1$ visited states:

$$\hat{S}_n = [S_{n-d}, \dots, S_{n-1}, S_n]. \tag{8}$$

Arranging the current extended state and the current action into a feature vector, $\mathbf{x}_n = [\hat{S}_n, A_n]$, and the next state and the received reward into a label vector, $\mathbf{y}_n = [S_{n+1}, R_{n+1}]$, the model weights θ_m can be optimized employing a mean squared error (MSE) loss of the form:

$$L_m = \frac{1}{|D|} \sum_i^{|D|} (\mathbf{y}_i - m_{\theta_m}(\mathbf{x}_i))^2, \tag{9}$$

where D is a set of feature-label-pairs constructed from high-fidelity trajectories, and $|D|$ is the overall number of feature-label-pairs. A single trajectory of length N yields $|D| = N - (d + 1)$ such pairs. To obtain the model weights based on equation (9), we employ state-of-the-art deep learning techniques, e.g., normalization of features and labels, layer normalization, batch training, learning rate schedule, and early stopping. More details about the model training are provided in appendix A.

2.4 Model-based proximal policy optimization

Given a trained model of the form (7), it becomes straightforward to sample new (*fictitious*) trajectories from the model by employing the model recursively. This process is described in detail in algorithm 1. The initial extended state \hat{S}_0 is selected from the existing high-fidelity trajectories. Note that the action sampling makes the trajectory sampling stochastic, so even when starting from the same \hat{S}_0 , executing algorithm 1 K times generates K different trajectories. For efficiency, the sampling of multiple trajectories should be performed in parallel.

In principle, MBDRl can work by sampling trajectories from a single model. However, the policy changes with each episode, which leads to newly explored states and actions. Naturally, the environment model’s prediction error increases with each episode, and it becomes increasingly challenging to sample meaningful trajectories from the model. Therefore, it is vital to monitor the model’s confidence in a prediction to decide at which point new

Algorithm 1 Sampling of a trajectory from an environment model.

Input: $\hat{S}_0, \pi_{\theta_\pi}, m_{\theta_m}$	▷ initial history, policy, env. model
Result: τ	▷ sampled model trajectory
$\tau \leftarrow []$	▷ initialize empty trajectory
$\hat{S}_n \leftarrow \hat{S}_0$	▷ initialize current extended state
while $i < N$ do	
$A_n \sim \pi_{\theta_\pi}(S_n)$	▷ sample an action from the current policy
$[S_{n+1}, R_{n+1}] \leftarrow m_{\theta_m}(\hat{S}_n, A_n)$	▷ predict the next state and the reward
$\tau \leftarrow \tau \cup [[S_n, A_n, R_{n+1}, S_{n+1}]]$	▷ append experience tuple to trajectory
$\hat{S}_n \leftarrow [S_{n-d+1}, \dots, S_n, S_{n+1}]$	▷ overwrite current extended state
end while	

high-fidelity episodes should be generated. There are two pathways to include model uncertainty: one could train a fully probabilistic environment model, i.e., a model that predicts a distribution over the next possible states, or one could train an ensemble of regular models [15]. Here, we follow the METRPO algorithm [16] and train an ensemble of N_m simple environment models. Each member in the ensemble $M = [m_1, m_2, \dots, m_{N_m}]$ has a different set of weights θ_i , which we drop from the index to simplify the notation. The ensemble approach introduces a new hyperparameter, namely the number of models, but training the ensemble is typically easier than training a fully probabilistic network. Moreover, it is straightforward to optimize the models in parallel. Each model is trained as described before, i.e., employing loss function (9). Even though the training procedure stays the same, the optimized models differ from one another for several reasons:

1. The dataset D is split randomly into N_m subsets, and each model is trained on a different subset.
2. Each model has a different set of randomly initialized weights. Since the optimization is non-linear, the optimized weights likely differ, even if the same dataset is used for training.

3. The gradient descent algorithm updates the model weights multiple times per epoch (iteration) based on batch gradients. The mini-batches are chosen at random from the training data (batch training).

There are several options to generate new trajectories from the model ensemble. The most obvious one is to repeat algorithm 1 once per model. In contrast, the original METRPO algorithm mixes the models within the trajectory sampling to improve robustness [16]. The authors state that different models are likely to have different model biases due to the training on a different subset of the data. When mixing the models, the different biases can partially cancel out. The alternation between models is achieved by sampling a different model for each step from a categorical distribution P_M over N_m classes, where each class i has an equal selection probability of $p_i = 1/N_m$. The modified sampling strategy, which we also employ here, is depicted in algorithm 2.

Besides the robust generation of trajectories, the ensemble allows to quantify prediction uncertainty. More precisely, we would like to determine, at which point the ensemble becomes unsuitable to generate new trajectories. Therefore, in each

Algorithm 2 Sampling of a trajectory from an ensemble of environment models. Differences compared to algorithm 1 are marked in bold.

Input: $\hat{S}_0, \pi_{\theta_\pi}, M$	▷ initial history, policy, model ensemble
Result: τ	▷ sampled model trajectory
$\tau \leftarrow []$	▷ initialize empty trajectory
$\hat{S}_n \leftarrow \hat{S}_0$	▷ initialize current extended state
while $i < N$ do	
$A_n \sim \pi_{\theta_\pi}(S_n)$	▷ sample an action from the current policy
$m_n \sim P_M$	▷ randomly select a model from the ensemble
$[S_{n+1}, R_{n+1}] \leftarrow m_n(\hat{S}_n, A_n)$	▷ predict the next state and the reward
$\tau \leftarrow \tau \cup [[S_n, A_n, R_{n+1}, S_{n+1}]]$	▷ append experience tuple to trajectory
$\hat{S}_n \leftarrow [S_{n-d+1}, \dots, S_n, S_{n+1}]$	▷ overwrite current extended state
end while	

Algorithm 3 Model-based proximal policy optimization (MEPPO) algorithm.

```

Input:  $v_{\theta_v}, \pi_{\theta_\pi}, M$  ▷ initial value and policy networks, model ensemble
Result:  $\theta_\pi^*$  ▷ optimal policy weights
 $e \leftarrow 0$  ▷ initialize episode counter
while  $e < e_{\max}$  do
  if  $N_{\text{pos}} < N_{\text{thr}}$  or  $e < 2$  then ▷ sample high-fidelity data
     $T \leftarrow \text{GenerateCFDTrajectories}(\pi_{\theta_\pi})$  ▷ run  $K$  simulations
     $M \leftarrow \text{UpdateModelEnsemble}(M, T)$  ▷ loss function (9)
     $N_{\text{pos}} \leftarrow N_m$  ▷ all models trustworthy
  else ▷ sample from model ensemble
     $T \leftarrow \text{GenerateEnsembleTrajectories}(\pi_{\theta_\pi}, M)$  ▷ use algorithm 2  $K$  times
     $N_{\text{pos}} \leftarrow \text{EvaluateModels}(\pi_{\theta_\pi}, M)$  ▷ algorithm 1, equation (10)
  end if
   $\pi_{\theta_\pi} \leftarrow \text{UpdatePolicyNet}(\pi_{\theta_\pi}, v_{\theta_v}, T)$  ▷ loss function (6)
   $v_{\theta_v} \leftarrow \text{UpdateValueNet}(v_{\theta_v}, T)$  ▷ loss function (2)
end while

```

model-based episode, K trajectories are generated from the ensemble employing algorithm 2, and N_m additional trajectories, one for each model, are generated employing algorithm 1. The ensemble trajectories are used to update policy and value networks. The model-specific trajectories are used to evaluate the policy loss (6) individually for each model to obtain a scalar value expressing the model’s quality (not for gradient descent). Comparing the i th model’s loss values of the current episode, i.g., $L_{\text{pol},i}^{\text{new}}$, with the loss of the same model obtained in the previous episode, i.e., $L_{\text{pol},i}^{\text{old}}$, allows assessing the quality of the policy update in the current episode. The loss trends for all models in the ensemble are combined by evaluating:

$$N_{\text{pos}} = \sum_{i=1}^{N_m} H\left(L_{\text{pol},i}^{\text{old}} - L_{\text{pol},i}^{\text{new}}\right), \tag{10}$$

where H is the Heaviside step function. In simple terms, N_{pos} is the number of models in the ensemble for which the policy loss improves after an update of the policy. Values of N_{pos}/N_m close to zero indicate that the models’ generalization capabilities are exhausted. Consequently, new high-fidelity data should be generated to update the ensemble. We define the condition $N_{\text{pos}} \geq N_{\text{thr}}$ to switch between model-based and simulation-based trajectory sampling. Values of N_{thr}/N_m close to zero encourage learning from the models at the risk of decreased robustness, whereas values close to one ensure high model quality at the risk of decreased computational efficiency. It is conceivable that $N_{\text{thr}}/N_m \approx 0.5$ might be a suitable compromise between efficiency and

control performance. The final model ensemble PPO (MEPPO) workflow is depicted in algorithm 3.

3 Results

We demonstrate the MEPPO algorithm on two flow configurations, namely a rotating cylinder in channel flow and the fluidic pinball. The numerical setups are described in Sect. 3.1. In Sect. 3.2, we compare MF and MB trainings and investigate the influence of ensemble size N_m and switching criterion N_{thr} . Section 3.3 presents a short discussion of the optimal policies found for each test case. The numerical simulations are performed with the *OpenFOAM-v2206* toolbox [26]. The orchestration of simulations and the DRL logic are implemented in the *drlFoam* package [27]. Further implementation details are available in the complementary code repository [20].

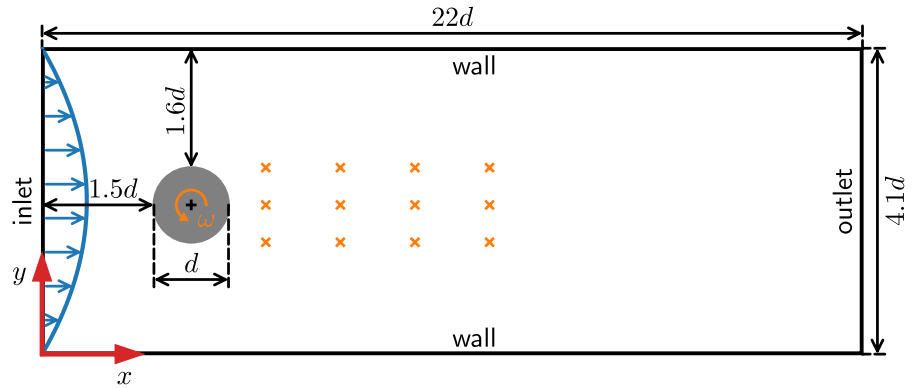
3.1 Flow control problems

3.1.1 Cylinder flow

The flow past a circular cylinder has become an established AFC benchmark. Originally, the setup was introduced by Schäfer et al. [28] as a benchmark for numerical methods. Rabault et al. [17] were the first to adopt the setup for DRL-based AFC. Since then, many variants with different means of actuation, sensor positions, and Reynolds numbers have been investigated [7].

The setup used in the present study is depicted in Fig. 2. The inlet velocity profile is parabolic [28],

Fig. 2 AFC setup of the flow past a cylinder; based on [17, 18, 28]



while the velocity vector at the upper and lower domain boundary is zero. At the outlet boundary, the velocity gradient is set to zero. For the pressure, a fixed reference value is applied at the outlet, while the gradient is set to zero on all other boundaries. The Reynolds number based on the mean inlet velocity U_{in} , the cylinder diameter d and the kinematic viscosity ν is $Re = U_{in}d/\nu = 100$. To solve the incompressible Navier Stokes equations, we employ the *pimpleFoam* solver with residual control. The pressure–velocity coupling is stopped once the initial residuals for pressure and momentum equations drop below 10^{-4} . Since the flow is laminar, no additional turbulence modeling is necessary. The mesh consists of approximately 5.3×10^3 hexahedral cells and is created using the *blockMesh* utility. For completeness, we note that the mesh has one cell layer in the third spatial direction of depth $\Delta z = 0.1d$, even though the simulation is 2D. The extension in the third direction is a technical requirement of 2D simulations in *OpenFOAM*. To discretize convective and diffusive fluxes, we employ pure linear interpolation. An implicit first-order Euler method is used for the temporal discretization.

The state (observation) used as input for control is formed by 12 pressure probes placed in the cylinder’s wake. The flow is actuated by rotating the cylinder. The control aims to minimize the forces acting on the cylinder. Given the density-normalized integral force vector $\mathbf{F} = [F_x, F_y]^T$, the corresponding force coefficients along each coordinate are:

$$c_x = \frac{2F_x}{U_{in}^2 A_{ref}}, \quad c_y = \frac{2F_y}{U_{in}^2 A_{ref}}, \quad (11)$$

where the reference area is $A_{ref} = d\Delta z$. The instantaneous reward at time step n is computed as [18]:

$$R_n = 3 - (c_{x,n} + 0.1|c_{y,n}|). \quad (12)$$

Note that we do not use time or window-averaged coefficients but the instantaneous values $c_{i,n}$. The *magical constants* in Eq. (12) simply scale the reward to a value range near zero, which avoids additional tuning of PPO hyperparameters.


Given the convective time scale $t_{conv} = d/U_{in}$, the control starts once the quasi-steady state is reached, which is at $40t_{conv}$. The policy is queried once every 20 numerical time steps, which corresponds to a control time step of $\Delta t_c = 0.1t_{conv}$. Within the control interval, the angular velocity linearly transitions from ω_n to ω_{n+1} [18]. The value range of the normalized angular velocity, $\omega^* = \omega d/U_{in}$, is limited to $\omega^* \in [-0.5, 0.5]$. Overall, $N = 400$ experience tuples are generated within one trajectory, which corresponds to a duration of $40t_{conv}$. The generation of a single trajectory takes approximately 4min on two MPI ranks.

3.1.2 Fluidic pinball

Noack et al. [29] introduced the fluidic pinball, which is a triangular configuration of three rotating cylinders. The unforced flow undergoes several different vortex shedding regimes, namely symmetric, asymmetric, and chaotic vortex shedding. The cylinder-based Reynolds number of the present setup is $Re = 100$, which places the flow in the asymmetric vortex shedding regime.

The setup is depicted in Fig. 3. To reduce the initial transient simulation phase, we apply a step function

Fig. 3 AFC setup of the fluidic pinball; based on [29]



velocity profile at the inlet with $\epsilon = 0.01U_{in}$. A constant velocity vector is set on the lower and upper domain boundaries, i.e., $1.01U_{in}$ and $0.99U_{in}$, respectively. The velocity gradient at the outlet boundary is set to zero. Similar to the cylinder setup, the pressure is fixed at the outlet, and the pressure gradient on all other boundaries is set to zero. The remaining details of the setup are largely the same as those presented in Sect. 3.1.1. Due to the more complex geometry, the mesh consists of approximately 5.3×10^4 control volumes. For the convective term of the momentum equation, we employ a *limitedLinear* scheme with a coefficient of 1.0.

A total of 14 pressure sensors form the state. The positions depicted in Fig. 3 are loosely inspired by a preliminary optimal sensor placement study [9]. The precise numerical coordinates of each sensor can be inferred from the setup files in the complementary code repository. The policy parametrizes a trivariate Beta distribution over the angular velocities ω_i , $i \in \{1, 2, 3\}$. As the control objective, we aim to minimize the cumulative forces acting on all three cylinders. Given the density-normalized integral force $\mathbf{F}_i = [F_{x,i}, F_{y,i}]^T$ of the i th cylinder, the corresponding force coefficients are defined as:

$$c_{x,i} = \frac{2F_{x,i}}{U_{in}^2 A_{ref}}, \quad c_{y,i} = \frac{2F_{y,i}}{U_{in}^2 A_{ref}}, \quad (13)$$

where the reference area is the projected area of all three cylinders in x -direction, namely $A_{ref} = 2.5d\Delta z$. Definition (13) ensures that the sum over all cylinders recovers the correct total force coefficients:

$$c_x = \sum_{i=1}^3 c_{x,i}, \quad c_y = \sum_{i=1}^3 c_{y,i}. \quad (14)$$

Based on the cumulative force coefficients, the instantaneous reward is computed as:

$$R_n = 1.5 - (c_{x,n} + 0.5|c_{y,n}|). \quad (15)$$

The reward definition is very similar to the one used for the single cylinder. Note that there is no particular reason for this definition other than that it was employed in a previous study [19]. Alternatively, we could have also summed up the magnitudes of the individual force coefficients. However, the focus of this work is on the demonstration of MBDRL. The *magical constants* in Eq. 15 have the same purpose as before, namely normalizing the reward values.

Keeping the same definition of the convective time scale as before, the control starts at the end of the initial transient regime, which is at $200t_{conv}$. The policy is queried every 50 numerical time steps, corresponding to a control interval of $\Delta t_c = 0.5t_{conv}$. The trajectory extends over $100t_{conv}$ such that $N = 200$ experience tuples are generated. The dimensionless angular velocities are limited to

Table 1 Characteristic parameters of the two AFC setups; *ranks* refers to the number of MPI ranks; *CV* is the number of control volumes, and T_{tr} is the time required to execute one simulation (to sample one trajectory)

case	Re	t_{conv}	$\Delta t_c/t_{conv}$	ω^* range	N	ranks	CV	T_{tr}
cylinder	100	0.1s	0.1	$[-0.5, 0.5]$	400	2	5.3×10^3	$\approx 4min$
pinball	100	1s	0.5	$[-5, 5]$	200	8	5.3×10^4	$\approx 40min$

the range $\omega_i^* \in [-5, 5]$. Generating a single trajectory with eight MPI ranks requires approximately 40min. The main parameters relevant to AFC are summarized in Table 1 for both test cases.

3.2 Training performance

3.2.1 General remarks about the training


We compare the control performance and the computational cost between MF and MB trainings. Each training is executed in isolation on a dedicated compute node with 96 CPU cores. To discuss the schematics presented later on, it is necessary to outline several implementation details. In each episode, exactly ten trajectories are generated in parallel before updating the policy and value networks. The *buffer size* and the parallel trajectory generation are the same for CFD and MB trajectories. We perform a very rough check of the force coefficients in the MB trajectories and discard trajectories that are obviously extremely inaccurate. Specifically, we only keep an MB trajectory for the cylinder flow if all coefficients fulfill the criteria $2.85 \leq c_x \leq 3.5$ and $|c_y| \leq 1.3$. The analogous bounds for pinball trajectories are $-3.5 \leq c_{x,i} \leq 3.5$ and $|c_{y,i}| \leq 3.5$. Note that these bounds are really generous and filter out only extreme prediction errors. If one or more trajectories are discarded, the sampling is repeated until ten valid trajectories are available. The training is stopped after 200 and 150 episodes for the cylinder and pinball cases, respectively. Finally, we note that all parameter configurations for the cylinder flow are repeated with five different seed values. All results presented in Sect. 3.2.2 are averaged over these five runs. Due to the computational cost of the fluidic pinball, we execute only a single training run per parameter configuration.

3.2.2 Cylinder flow

Figure 4 shows the episode-wise mean rewards received with different training configurations. Within 200 episodes almost all configurations achieve a similar maximum reward. The maximum reward in the MF training is slightly lower, but the reward still keeps increasing moderately toward the end of the training. Surprisingly, all MB trainings reach the maximum reward earlier than the MF training. We noticed that the MB trajectories display less small-scale variance than the CFD trajectories. Presumably, this indirect *filtering* performed by the environment models affects the policy and value network updates positively. In general, the MB trainings display a higher variance in the mean reward, especially within the first 50 episodes, where the changes in the reward are the strongest. This behavior is related to individual MB trajectories with poor prediction quality. Luckily, these trajectories seem to have a limited impact on the learning.

The markers in the MB training in Fig. 4 indicate CFD trajectories. For the MB training with one model, we switch back to simulation-based sampling every fourth episode, so 75% of all episodes use the environment model. The same ratio results in the MEPPPO training with $N_m = 5$ and $N_{thr} = 3$. Interestingly, the automated switching between CFD and model sampling leads to a relatively regular alternation. The same number of models with a lowered threshold value of $N_{thr} = 2$ reduces the amount of CFD episodes to 15%, notably without performance degradation. The training with $N_m = 10$ and $N_{thr} = 5$ also employs the model ensemble in 75% of the episodes. Reducing the threshold value in the latter case to $N_{thr} = 3$ reduces the number of CFD episodes to 6%. However, the optimization becomes significantly less stable and does not reach optimal control performance. Note that the amount of training data per model decreases as the number of models increases. Therefore, increasing the number of

Fig. 4 Cylinder flow: episode-wise mean reward R for different training configurations; the shaded area encloses one standard deviation below and above the mean; mean and standard deviation are computed over all trajectories of all seeds; for the MB training, the markers indicate CFD-based trajectory sampling



sampled trajectories per episode or allowing overlapping training data could potentially stabilize trainings with a lowered threshold value. The increased amount of training data when employing only a single model could explain the quick convergence of the corresponding training after approximately 100 episodes.

Assuming that the creation of environment models is reliable and much less costly than a CFD simulation, the more CFD episodes can be replaced with MB episodes, the lower the overall training time. This trend is reflected by the normalized MB training time depicted in Fig. 5. Remarkably, all configurations reduce the training time by at least 65%, even though the computational cost to simulate the cylinder flow is fairly small. To better understand differences in the overall training time, Fig. 5 also shows the time spent on the most relevant parts of the training. For all configurations, the simulation runs remain the most

time-consuming element. The time spent updating the networks for environment, policy, and value function remains fairly constant.

A noticeable increase can be observed in the time spent to sample trajectories from a single environment model. This increase, as well as other variations in the sampling time, are related to the removal of obviously corrupted trajectories, as explained in Sect. 3.2.1. Figure 6 shows the number of discarded trajectories per training. Clearly, the ensemble models lead to a significantly more robust trajectory sampling. Moreover, ensembles with more models produce fewer invalid samples. It is noteworthy to mention that the majority of invalid trajectories result in the initial training phase, $e < 50$, presumably because of the strong exploration and policy changes. In our implementation, trajectories are sampled in parallel. However, the constraint of executing the

Fig. 5 Cylinder flow: composition of the total MB training time T_{MB} normalized with the MF training time $T_{MF} \approx 14h$

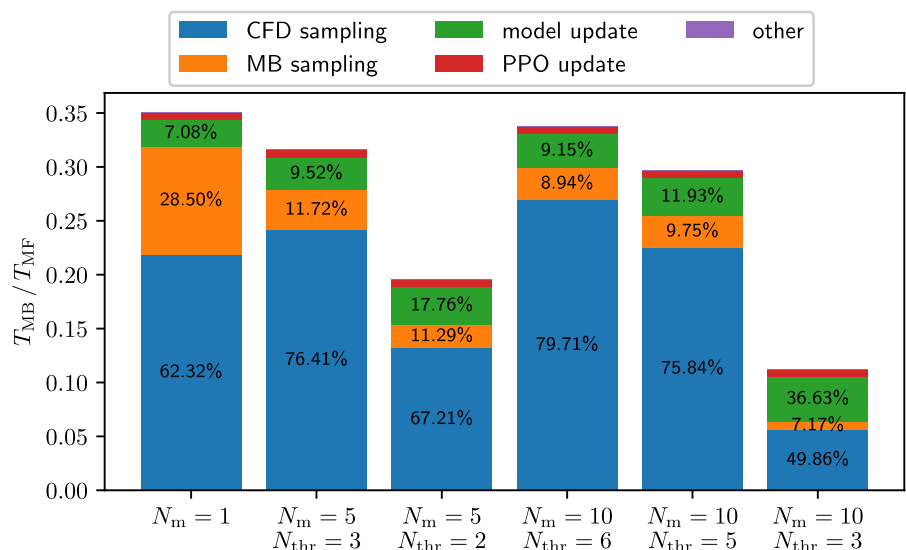



Fig. 6 Cylinder flow: number of discarded trajectories N_d per training; the box plot shows the outcome of five independent runs (seed values); the orange line indicates the median value, and the markers indicate extreme outcomes



PPO update with exactly ten valid trajectories leads to repeated sampling in the case of discarded trajectories. Of course, this extra time could be avoided by sampling more trajectories than required in parallel or by loosening the constraint.

We note that there are also small variations in the measured execution times that are not straightforward to explain. One means to reduce these variations would be executing additional training runs with varying seeds. However, the main trends in terms of control performance and training time are clearly visible and explainable. Moreover, there are nontrivial coupling effects between simulations, environment models, and PPO updates. As discussed before, policy updates differ between MF and MB trainings and also depend to some degree on ensemble size and switching criterion. The policy employed in a simulation impacts the execution time, e.g., by subsequent changes in the pressure–velocity-coupling iterations, or the linear solver iterations.

3.2.3 Fluidic pinball

Due to the computational cost of the pinball setup, we do not perform multiple training runs and investigate

fewer hyperparameter configurations. Figure 7 shows the episode-wise mean rewards of all trainings. The final reward is similar for all tested configurations. The MF training reaches the maximum reward after approximately 60 episodes. Thereafter, the control performance drops slightly. The highest overall reward is achieved by the MB training with $N_m = 10$. For the MB training with a single model we switch every fourth episode to CFD trajectories. Also the automated switching criterion in the training with $N_m = 10$ and $N_{thr} = 5$ leads to 25% CFD episodes. The smaller ensemble with five models switches approximately every fifth episode back to the high-fidelity samples.

Figure 8 shows the main contributions to the normalized training time. Even though the relative amount of model trajectories is similar to the cylinder flow, the runtime reduction is even more significant due to the higher computational cost of performing the pinball simulation. Remarkably, all configurations reduce the training time by at least 80%. The remaining CFD episodes dominate the overall time consumption. As for the cylinder test case, the time spent on sampling model trajectories increases when using only a single model. Ten trajectories were discarded

Fig. 7 Fluidic pinball: episode-wise mean reward R for different training configurations; the shaded area encloses one standard deviation below and above the mean; mean and standard deviation are computed over all trajectories; for the MB training, the markers indicate CFD-based trajectory sampling


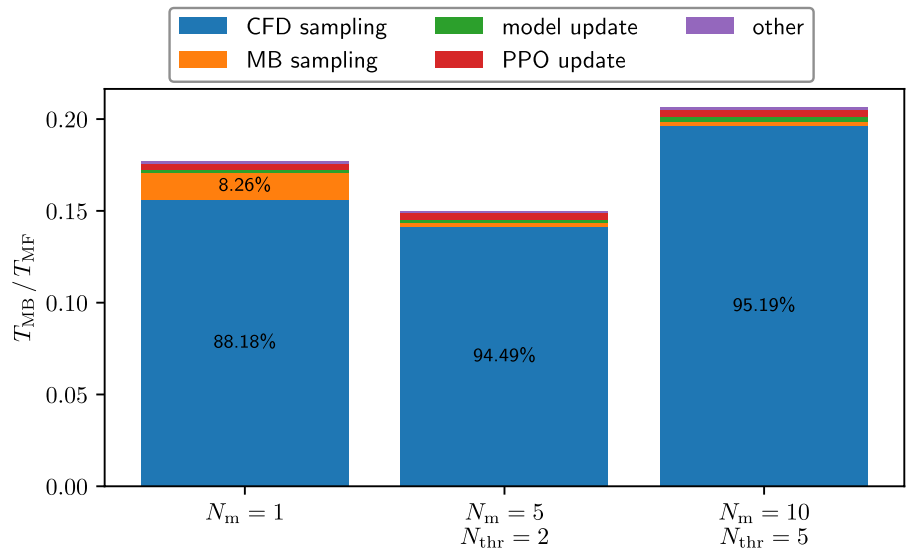


Fig. 8 Fluidic pinball: composition of the total MB training time T_{MB} normalized with the MF training time $T_{MF} \approx 130h$



during the training and had to be re-sampled. In the MEPPPO runs, no invalid trajectories occurred. The increased time spent on simulations in the training with ten models is a result of the non-trivial interaction between policy learning and the numerical properties of the simulation, as discussed before.

3.3 Analysis of the best policies

3.3.1 Cylinder flow

In this section, we compare the final policies of the MF and MB trainings that achieved the highest mean reward. Note that the achieved force reduction and the corresponding control are by no means new [7]. However, we would like to point out a few subtle differences in the final MF and MB policies. It is noteworthy that all MB policies achieved a similar force reduction, even though we present only results for the MEPPPO training with ten models.

Figure 9 shows the optimal control and the corresponding force coefficients. Both MF and MB policies reduce the drag force by approximately 7%. The second force coefficient remains near zero after 20 convective time units. Both the MF and MB policies perform a very similar actuation within the first five convective time units. Thereafter, the MB policy rotates the cylinder slightly less and archives an almost perfectly steady force balance. Small


fluctuations remain in the angular velocity and the side force when employing the MF policy.

The cylinder rotation steadies and extends the wake. The pressure drop over the cylinder decreases, as can be inferred visually in Fig. 10. The figure also shows the significant suppression of vortex shedding. While small pressure fluctuations remain when employing the MF policy, the MB policies leads to an almost perfectly steady flow field. It is conceivable that the MF training could achieve a similar control performance with more training episodes and additional PPO hyperparameter tuning. However, the MB runs achieve the presented control performance with even less than 200 episodes.

3.3.2 Fluidic pinball

Several established control mechanisms exist for the fluidic pinball, e.g., boat tailing and base bleeding [30]. The strategy learned by the policies in the present work is referred to as boat tailing. Figure 11 compares the control and the resulting force coefficients of the best-performing MF and MB policies. Both policies follow a similar strategy, i.e., rotating cylinders 2 and 3 with maximum speed in opposite directions while keeping the cylinder in the front nearly still. This actuation leads to boat tailing and reduces the drag coefficient by approximately 87%. Of course, the reduction strongly depends on the allowed maximum rate of rotation and would be smaller for smaller rates of rotation.

Fig. 9 Cylinder flow: angular velocity and force coefficients resulting from the best MF and MB policies; the dimensionless time is $t^* = t/t_{conv}$



While the drag reduction is fairly similar, the MF policy introduces a nonzero net force in y -direction. This asymmetry is caused by a subtle difference in the absolute values of ω_2^* and ω_3^* , i.e., the absolute value of ω_2^* is marginally larger. The same small difference between ω_2^* and ω_3^* is present in the MB policy. However, the MB policy applies a small positive spin to the cylinder in the front, which balances the side forces to net zero. The reader might be wondering if the same drag reduction and net zero force in y -direction could be achieved by setting $\omega_1^* = 0$ and $\omega_2^* = -\omega_3^* = 5$. The short answer is yes. However, the latter control leads to relatively strong c_y fluctuations whose amplitude decays only slowly. Approximately 150 convective time units

are necessary to reach the steady state with open-loop control, which is longer than the prescribed trajectory length employed in the training. The asymmetric control performed by the MB policy achieves the same force reduction in about 20 convective time units.

Finally, we want to shed some light on the local flow dynamics created by the spinning cylinders. Figure 12 shows the velocity field in the vicinity of the cylinders. As mentioned before, the uncontrolled state is asymmetric, as can be observed by looking at the wakes of the two cylinders in the back. In the controlled flow, the rotation of cylinders two and three pushes fluid between them in upstream direction. This suction significantly reduces the extent of

Fig. 10 Temporal mean and standard deviation of the pressure fields with and without control; both mean and standard deviation are normalized with the minimum (blue) and maximum (yellow) values of the uncontrolled case; the coordinates are normalized with the diameter

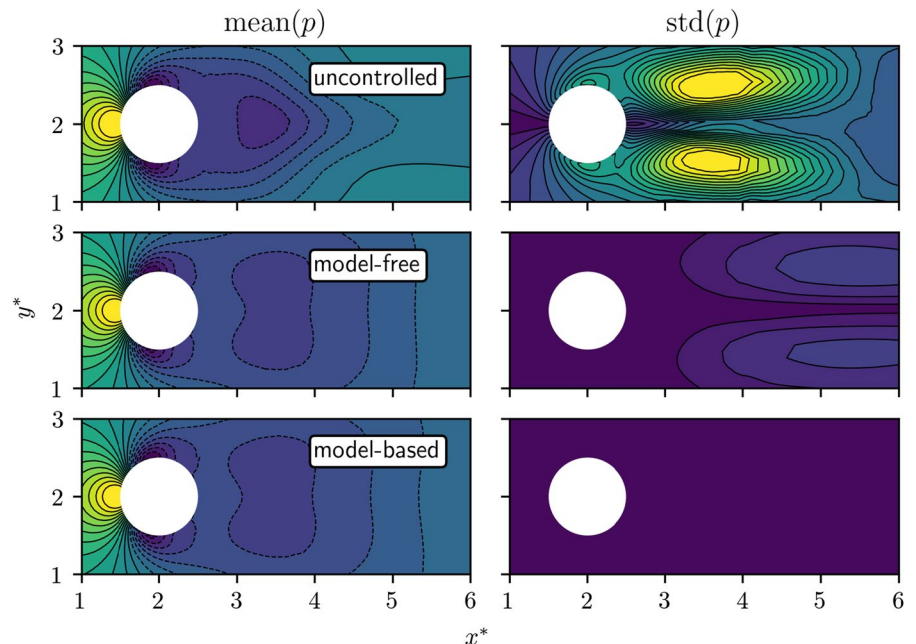
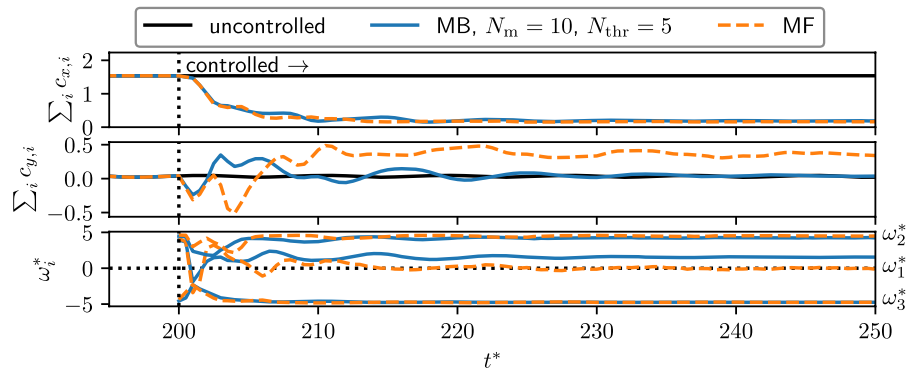


Fig. 11 Fluidic pinball: angular velocities and summed force coefficients resulting from the best MF and MB policies; the dimensionless time is $t^* = t/t_{conv}$



the wake and the cumulative pressure drop over the cylinders. The fluid then passes through the gaps between the rear cylinders and the front cylinder. For the MF policy, the fluxes in positive and negative y -direction are fairly balanced. Instead, the MB policy hinders the flux between cylinders one and two and diverts more fluid to the gap between cylinders one and three. This small imbalance is enough to compensate for the different rotation speeds of the rear cylinders.

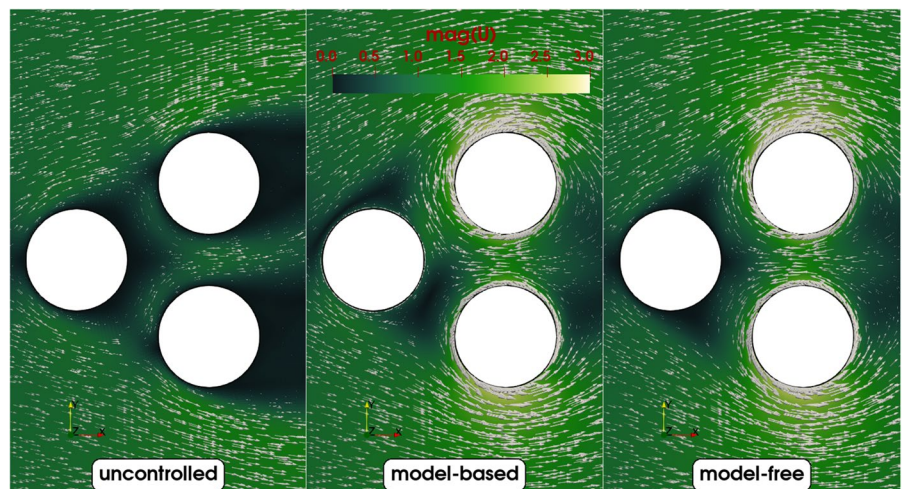
4 Conclusion

DRL-based AFC is a promising approach to enable smart control technology. A current limitation of simulation-based DRL is the high computational cost and turnaround time associated with the optimization. The idea of MBDR is the creation of low-cost

environment models that partially replace the costly sampling of high-fidelity trajectories from simulations. On two common benchmark AFC problems, we demonstrate that the MEPPPO algorithm adopted here can tremendously reduce the training cost and time while achieving optimal control performance. The relative reduction of training time increases with the cost of the CFD simulation. Therefore, we expect the MEPPPO algorithm, or variants thereof, to be a key enabler in performing DRL-based AFC on realistic, 3D simulations at an industrial scale. Of course, further tests on more complex flow configurations are required to consolidate this hypothesis.

There are a number of promising options to reduce the training cost even further. Creating accurate autoregressive environment models on the fly posed a significant challenge. Automating the model creation, e.g., by means of Bayesian hyperparameter optimization, could be beneficial for the models' accuracy

Fig. 12 Comparison of instantaneous velocity fields with and without control; the velocity field is normalized with the inlet velocity



and adaptability to new control problems. The more accurate the models can extrapolate, the more savings are possible. More advanced recurrent network architectures or transformers [31] might be capable of achieving higher accuracy than the simple, fully connected networks employed here. However, such advanced networks are typically also harder to train. Alternatively, more straightforward approaches to creating reduced-order models might be worthwhile to explore, too. For example, dynamic mode decomposition with control [32] could be a drop-in replacement for the models employed here.

We also noticed a strong dependency of the learning progress on the PPO hyperparameters. Due to the high training cost, automated tuning of these parameters is infeasible for complex simulations. However, recent improvements in the PPO algorithm, e.g., the PPO-CMA variant [33], reduce the number of hyperparameters, improve robustness, and accelerate learning. Such improvements would be extremely beneficial both for MF and MB learning.

Acknowledgements The authors gratefully acknowledge the Deutsche Forschungsgemeinschaft DFG (German Research Foundation) for funding this work in the framework of the research unit FOR 2895 under the grant WE 6948/1-1. The computational resources to conduct the numerical experiments were kindly provided by Amazon Web Services. Finally, the authors gratefully acknowledge the organizing committee of

the 18th OpenFOAM Workshop for the support provided during the preparation of this manuscript.

Author's contribution AW acquired the funding, planned the research project, and wrote the manuscript. JG conducted the numerical experiments and wrote the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. This work has been funded by the German Research Foundation (DFG) under the grant WE 6948/1-1.

Data availability The full research data and instructions to reproduce the numerical experiments and analyses are available at: <https://doi.org/10.23728/b2share.85ab8f3f68724372b83babbdaca85910>

Declarations

Conflict of interest The authors have no Conflict of interest to declare.

Ethical approval No experiments on humans or animals have been conducted.

Appendix A: Environment models

We use a standard deep learning workflow to train the environment models. The different optimization techniques are described in detail by Raff [31]. The activation function, the number of hidden layers, and the number of neurons per layer have been determined by employing a simple grid search. Table 2 provides

Table 2 Hyperparameters for the training of the environment models

Parameter	Cylinder	Pinball
Hidden layers	3	3
Layer normalization	LayerNorm	LayerNorm
Neurons per hidden layer	100	100
Time delays d	30	30
Input neurons	450	690
Output neurons	14	20
Activation function	leaky ReLU	leaky ReLU
Max. number of epochs	2500	2500
Batch size	25	25
Train/validation split	75/25%	75/25%
Absolute stopping loss	10^{-6}	10^{-6}
Relative stopping loss	10^{-7}	10^{-7}
Loss function	MSELoss	MSELoss
Optimizer	AdamW	AdamW
Learning rate schedule	ReduceLROnPlateau	ReduceLROnPlateau
Init./min. learning rate	$10^{-2}/10^{-4}$	$10^{-2}/10^{-4}$

an overview of the most important hyperparameters. Since the amount of training data is fairly small, we employ 8 CPU cores to train a single model. Note that the models predict the individual force coefficients rather than the full reward. Consequently, the output layer has 12+2 and 14+6 neurons for the cylinder and pinball test cases, respectively. The training is stopped once the maximum number of epochs (gradient descent steps) is reached, the absolute loss value falls below a threshold value, or the relative change of the loss averaged over 40 epochs becomes too small.

Appendix B: PPO hyperparameters

Most PPO hyperparameters employed here are standard values suggested in the literature [23, 25]. In contrast to the original implementation, we use two separate networks for policy and value function. The two networks are optimized sequentially by two different optimizers. Since the optimization is inexpensive, we allocate only 5 CPU cores. As suggested in [25], the policy optimization is stopped if the difference between old and new policy, measured in terms of KL divergence, exceeds a prescribed threshold. However, we noticed that this criterion is rarely triggered. All hyperparameters are summarized in Table 3.

Table 3 PPO hyperparameters

Parameter	Cylinder	Pinball
Hidden layers	2	2
Neurons per hidden layer	64	512
Input neurons	12	14
Output neurons	2	6
Activation function	ReLU	ReLU
Discount factor γ	0.99	0.99
Advantage smoothing λ	0.97	0.97
Learning rate policy net	10^{-3}	10^{-5}
Learning rate value net	5×10^{-4}	10^{-5}
Max. number of epochs	100	100
Batch size	full	full
Value and policy clipping	0.1	0.1
Optimizer	AdamW	AdamW
KL divergence	0.2	0.2
Entropy weight β	0.01	0.01

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. IPCC: Climate Change 2022 (2022) Impacts, Adaptation and Vulnerability. Summary for Policymakers, pp. 3–33. Cambridge University Press, Cambridge, UK and New York, USA
2. Seifert A, Shtendel T, Dolgopyat D (2015) From lab to full scale active flow control drag reduction: How to bridge the gap? *J Wind Eng Ind Aerodyn* 147:262–272. <https://doi.org/10.1016/j.jweia.2015.09.012>
3. Rohrer T, Frison L, Kaupenjohann L, Scharf K, Hergenröther E (2023) Deep reinforcement learning for heat pump control. In: Arai K (ed) *Intelligent Computing*, pp 459–471. Springer, Cham. https://doi.org/10.1007/978-3-031-37717-4_29
4. Esche E, Repke J-U (2020) Dynamic process operation under demand response - a review of methods and tools. *Chem Ing Tech* 92(12):1898–1909. <https://doi.org/10.1002/cite.202000091>
5. Hucho W, Sovran G (1993) Aerodynamics of road vehicles. *Annu Rev Fluid Mech* 25(1):485–537. <https://doi.org/10.1146/annurev.fl.25.010193.002413>
6. Choi H, Lee J, Park H (2014) Aerodynamics of heavy vehicles. *Annu Rev Fluid Mech* 46(1):441–468. <https://doi.org/10.1146/annurev-fluid-011212-140616>
7. Viquerat J, Meliga P, Larcher A, Hachem E (2022) A review on deep reinforcement learning for fluid mechanics: an update. *Phys Fluids* 34(11):111301. <https://doi.org/10.1063/5.0128446>
8. Paris R, Beneddine S, Dandois J (2021) Robust flow control and optimal sensor placement using deep reinforcement learning. *J Fluid Mech* 913:25. <https://doi.org/10.1017/jfm.2020.1170>
9. Krogmann T (2023) Optimal sensor placement for active flow control with deep reinforcement learning. <https://doi.org/10.5281/zenodo.7636959>
10. Paris R, Beneddine S, Dandois J (2023) Reinforcement-learning-based actuator selection method for active flow control. *J Fluid Mech* 955:8. <https://doi.org/10.1017/jfm.2022.1043>

11. Ashton N, West A, Lardeau S, Revell A (2016) Assessment of rans and des methods for realistic automotive models. *Comput Fluids* 128:1–15. <https://doi.org/10.1016/j.compfluid.2016.01.008>
12. Belus V, Rabault J, Viquerat J, Che Z, Hachem E, Reglade U (2019) Exploiting locality and translational invariance to design effective deep reinforcement learning control of the 1-dimensional unstable falling liquid film. *AIP Adv* 9(12):125014. <https://doi.org/10.1063/1.5132378>
13. Vignon C, Rabault J, Vasanth J, Alcántara-Ávila F, Mortensen M, Vinuesa R (2023) Effective control of two-dimensional Rayleigh-Bénard convection: invariant multi-agent reinforcement learning is all you need. *Phys Fluids* 35(6):065146. <https://doi.org/10.1063/5.0153181>
14. Dixit A, Elsheikh AH (2023) Robust optimal well control using an adaptive multigrid reinforcement learning framework. *Math Geosci* 55(3):345–375. <https://doi.org/10.1007/s11004-022-10033-x>
15. Moerland TM, Broekens J, Jonker CM (2020) Model-based reinforcement learning: a survey. *CoRR* [arXiv:2006.16712](https://arxiv.org/abs/2006.16712)
16. Kurutach T, Clavera I, Duan Y, Tamar A, Abbeel P (2018) Model-ensemble trust-region policy optimization. *CoRR* [arXiv:1802.10592](https://arxiv.org/abs/1802.10592)
17. Rabault J, Kuchta M, Jensen A, Réglade U, Cerardi N (2019) Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *J Fluid Mech* 865:281–302. <https://doi.org/10.1017/jfm.2019.62>
18. Tokarev M, Palkin E, Mullyadzhyanov R (2020) Deep reinforcement learning control of cylinder flow using rotary oscillations at low reynolds number. *Energies* 13(22):5920. <https://doi.org/10.3390/en13225920>
19. Holm M (2020) Using reinforcement learning for active flow control. *DUO Research Archive*. <http://hdl.handle.net/10852/79212>
20. Geise J, Weiner A (2024) Git repository accompanying the article. https://github.com/JanisGeise/MB_DRL_for_accelerated_learning_from_CFD
21. Weiner A, Geise J (2024) Model-based deep reinforcement learning for accelerated learning from flow simulations. <https://b2share.eudat.eu>. <https://doi.org/10.23728/B2SHARE.85AB8F3F68724372B83BABBDACA85910>
22. Sutton RS, Barto AG (2018) Reinforcement Learning, Second Edition: An Introduction. Adaptive Computation and Machine Learning series. MIT Press, ???
23. Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O (2017) Proximal policy optimization algorithms. *CoRR* [arXiv:1707.06347](https://arxiv.org/abs/1707.06347)
24. Andrychowicz M, Raichuk A, Stanczyk PM, Orsini M, Girgin S, Marinier R, Hussenot L, Geist M, Pietquin O, Michalski M, Gelly S, Bachem OF (2021) What matters for on-policy deep actor-critic methods? A large-scale study. In: *ICLR*. <https://openreview.net/pdf?id=nIAxjsniDzg>
25. Morales M (2020) Grokking Deep Reinforcement Learning. Manning Publications
26. OpenFOAM version 2206 (2022). <https://www.openfoam.com/news/main-news/openfoam-v2206>
27. drlFoam - deep reinforcement learning with OpenFOAM. <https://github.com/OFDataCommittee/drlfoam> (2022)
28. Schäfer M, Turek S, Durst F, Krause E, Rannacher R (2022) In: Hirschel EH (ed) Benchmark Computations of Laminar Flow Around a Cylinder, pp. 547–566. Vieweg+Teubner Verlag, Wiesbaden x
29. Noack BR, Stankiewicz W, Morzyński M, Schmid PJ (2016) Recursive dynamic mode decomposition of transient and post-transient wake flows. *J Fluid Mech* 809:843–872. <https://doi.org/10.1017/jfm.2016.678>
30. Raibaudo C, Zhong P, Noack BR, Martinuzzi RJ (2020) Machine learning strategies applied to the control of a fluidic pinball. *Phys Fluids* 32(1):015108. <https://doi.org/10.1063/1.5127202>
31. Raff E (2022) Inside deep learning: math, algorithms, models. Manning
32. Proctor JL, Brunton SL, Kutz JN (2016) Dynamic mode decomposition with control. *SIAM J Appl Dyn Syst* 15(1):142–161. <https://doi.org/10.1137/15M1013857>
33. Hämmäläinen P, Babadi A, Ma X, Lehtinen J (2020) Ppocma: proximal policy optimization with covariance matrix adaptation. In: 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP), pp 1–6. <https://doi.org/10.1109/MLSP49062.2020.9231618>