



HAL
open science

A review on deep reinforcement learning for fluid mechanics

Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, Elie Hachem

► **To cite this version:**

Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, et al.. A review on deep reinforcement learning for fluid mechanics. *Computers and Fluids*, 2021, 225, pp.104973. 10.1016/j.compfluid.2021.104973 . hal-02401491v2

HAL Id: hal-02401491

<https://hal.science/hal-02401491v2>

Submitted on 17 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A REVIEW ON DEEP REINFORCEMENT LEARNING FOR FLUID MECHANICS

A PREPRINT

Paul Garnier

MINES Paristech , PSL - Research University
paul.garnier@mines-paristech.fr

Jonathan Viquerat

MINES Paristech , PSL - Research University
CEMEF
jonathan.viquerat@mines-paristech.fr

Jean Rabault

Department of Mathematics, University of Oslo
and
MINES Paristech , PSL - Research University
CEMEF
jean.rblt@gmail.com

Aurélien Larcher

MINES Paristech , PSL - Research University
CEMEF
aurelien.larcherm@mines-paristech.fr

Alexander Kuhnle

University of Cambridge
alexkuhnle@t-online.de

Elie Hachem*

MINES Paristech , PSL - Research University
CEMEF
elie.hachem@mines-paristech.fr

February 25, 2021

ABSTRACT

Deep reinforcement learning (DRL) has recently been adopted in a wide range of physics and engineering domains for its ability to solve decision-making problems that were previously out of reach due to a combination of non-linearity and high dimensionality. In the last few years, it has spread in the field of computational mechanics, and particularly in fluid dynamics, with recent applications in flow control and shape optimization. In this work, we conduct a detailed review of existing DRL applications to fluid mechanics problems. In addition, we present recent results that further illustrate the potential of DRL in Fluid Mechanics. The coupling methods used in each case are covered, detailing their advantages and limitations. Our review also focuses on the comparison with classical methods for optimal control and optimization. Finally, several test cases

*Corresponding author

are described that illustrate recent progress made in this field. The goal of this publication is to provide an understanding of DRL capabilities along with state-of-the-art applications in fluid dynamics to researchers wishing to address new problems with these methods.

Keywords Deep Reinforcement Learning · Fluid Mechanics

1 Introduction

The process of learning and how the brain understands and classifies information has always been a source of fascination for humankind. Since the beginning of research in the field of artificial intelligence (AI), an algorithm able to learn to make decisions on its own has been one of the ultimate goals. Going back to the end of the 1980s (Sutton, 1988), reinforcement learning (RL) proposed a formal framework in which an agent learns by interacting with an environment through the gathering of experience (François-lavet *et al.*, 2018). Significant results were obtained during the following decade (Tesauro, 1995), although they were limited to low-dimensional problems.

In recent years, the rise of deep neural networks (DNN) has provided reinforcement learning with new powerful tools Goodfellow *et al.* (2017) Strang (2019) Bottou *et al.* (2018). The combination of deep learning with RL, called deep reinforcement learning (DRL), lifted several major obstacles that hindered classical RL by allowing the use of high-dimensional state spaces and exploiting the feature extraction capabilities of DNNs. Recently, DRL managed to perform tasks with unprecedented efficiency in many domains such as robotics Pinto *et al.* (2017) or language processing Bahdanau *et al.* (2016), and even achieved superhuman levels in multiple games: Atari games Mnih *et al.* (2013), Go Silver *et al.* (2017), Dota II OpenAI (2018), Starcraft II Vinyals *et al.* (2019) and even Poker Brown & Sandholm (2019). Also in industrial application, DRL has proven itself useful. For example, Wayve trains autonomous cars both onboard (Kendall *et al.*, 2018) and through simulation (Bewley *et al.*, 2018), and Google uses DRL in order to control the cooling of its data centers (Knight, 2018).

In the field of fluid mechanics and mechanical engineering, one is also confronted with nonlinear problems of high dimensionality. For example, using computational simulations to test several different designs or configurations has proven a useful technique. However, the number of possibilities to explore can make such search difficult since it is often unfeasible to evaluate all configurations exhaustively. Therefore, the assistance of automatic optimization procedures is needed to help find optimal designs.

Possible applications range from multiphase flows in micro-fluidics to shape optimization in aerodynamics or conjugate heat transfer in heat exchangers, and many other fields can gain from such techniques such as biomechanics, energy, or marine technologies. This is the primary motivation for the combination of computational fluid dynamics (CFD) with numerical optimization methods. Most of the time, the main obstacle is the high computational cost of determining the sensitivity of the objective function to variations of the design parameters by repeated calculations of the flow. Besides, classical optimization techniques, like direct gradient-based methods, are known for their lack of robustness and for their tendency to fall into local optima. Generic and robust search methods, such as evolutionary or genetic algorithms, offer several attractive features and have been used widely for design shape optimization (Ali & Behdinan, 2003; Lee & Hajela, 1996; Mäkinen *et al.*, 1999). In particular, they can be used for multi-objective multi-parameter problems. They have been successfully tested for many practical cases, for example for design shape optimization in the aerospace (Matos *et al.*, 2004) and automotive industries (Muyl *et al.*, 2004). However, the use of a fully automatic evolutionary algorithms coupled with CFD for the optimization of multi-objective problems remains limited by the computing burden at stake and is still far from providing a practical tool for many engineering applications.

Despite great potential, the literature applying DRL to computational or experimental fluid dynamics remains limited. In recent years, only a handful of applications were proposed that take advantage of such learning methods. These applications include flow control problems (Rabault *et al.*, 2019; Guéniat *et al.*, 2016), process optimization (Novati *et al.*, 2017; Lee

Table 1: Notations regarding DRL

γ	discount factor
λ	learning rate
α, β	step-size
ϵ	probability of random action
s, s'	states
\mathcal{S}^+	set of all states
\mathcal{S}	set of non-termination states
a	action
\mathcal{A}	set of all actions
r	reward
\mathcal{R}	set of all rewards
t	time station
T	final time station
a_t	action at time t
s_t	state at time t
r_t	reward at time t
$R(\tau)$	discounted cumulative reward following trajectory τ
$R(s, a)$	reward received for taking action a in state s
π	policy
θ, θ'	parameterization vector of a policy
π_θ	policy parameterized by θ
$\pi(s)$	action probability distribution in state s following π
$\pi(a s)$	probability of taking action a in state s following π
$V^\pi(s)$	value of state s under policy π
$V^*(s)$	value of state s under the optimal policy
$Q^\pi(s, a)$	value of taking action a in state s under policy π
$Q^*(s, a)$	value of taking action a in state s under the optimal policy
$Q_\theta(s, a)$	estimated value of taking action a in state s with parameterization θ

et al., 2018), shape optimization (American Institute of Aeronautics and Astronautics, 2008), and our own results about laminar flows past a square cylinder presented later in this work, among others².

This paper presents an introduction to modern RL for the neophyte reader to get a grasp of the main concepts and methods existing in the domain, including the necessary basics of deep learning. Subsequently, the relevant fluid mechanics concepts are reiterated. We then go through a thorough review of the literature and describe the content of each contribution. In particular, we cover the choice of DRL algorithm, the problem complexity, and the concept of reward shaping. Finally, an outline is given on the future possibilities of DRL coupled with fluid dynamics.

2 Deep Reinforcement Learning

In this section, the basic concepts of DRL are introduced. For the sake of brevity, some topics are only minimally covered, and the reader is referred to more detailed introductions to (D)RL (Sutton & Barto, 2018; François-lavet *et al.*, 2018). First, RL is presented in its mathematical description as a Markov Decision Process. Then, its combination with deep learning is detailed, and an overview of several DRL algorithms is presented.

2.1 Reinforcement Learning

At its core, reinforcement learning models an agent interacting with an environment and receiving observations and rewards from it, as shown in figure 1. The goal of the agent is to determine the best action in any given state in order to maximize its cumulative reward over an episode, *i.e.* over one instance of the scenario in which the agent takes actions. Taking the illustrative example of a board game:

²These references and more are addressed in details in the remaining of this article.




Figure 1: **DRL agent and its environment**

- ◇ An episode is equivalent to one game;
- ◇ A state consists of the summary of all pieces and their positions;
- ◇ Possible actions correspond to moving a piece of the game;
- ◇ The reward may simply be +1 if the agent wins, -1 if it loses and 0 in case of a tie or a non-terminal game state.

Generally, the reward is a signal of how ‘good’ a board situation is, which helps the agent to learn distinguish more promising from less attractive decision trajectories. A trajectory is a sequence of states and actions experienced by the agent:

$$\tau = (s_0, a_0, s_1, a_1, \dots).$$

The cumulative reward (*i.e.* the quantity to maximize) is expressed along a trajectory, and includes a discount factor $\gamma \in [0, 1]$ that smoothes the impact of temporally distant rewards (it is then called *discounted* cumulative reward):

$$R(\tau) = \sum_{t=0}^T \gamma^t r_t.$$

RL methods are classically divided in two categories, namely *model-free* and *model-based* algorithms. Model-based method incorporates a model of the environment they interact with, and will not be considered in this paper (the reader is referred to Sutton & Barto (2018) and references therein for details about model-based methods). On the contrary, model-free algorithms directly interact with their environment, and are currently the most commonly used within the DRL community, mainly for their ease of application and implementation. Model-free methods are further distinguished between *value-based* methods and *policy-based* methods (Sutton & Barto, 1998). Although both approaches aim at maximizing their expected return, policy-based methods do so by directly optimizing the decision policy, while value-based methods learn to estimate the expected value of a state-action pair optimally, which in turn determines the best action to take in each state. Both families inherit from the formulation of Markov decision processes (MDP), which is detailed in the following section.

2.1.1 Markov decision processes

A MDP can be defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathbb{P}, R)$ (Bellman, 1957; Howard, 1960), where:

- ◇ \mathcal{S} , \mathcal{A} and R are defined in table 1,
- ◇ $\mathbb{P}: \mathcal{S} \times \mathcal{A} \times \mathcal{S}^+ \rightarrow [0, 1]$, where $\mathbb{P}(s'|s, a)$ is the probability of getting to state s' from state s following action a .

In the formalism of MDP, the goal is to find a policy $\pi(s)$ that maximizes the expected reward. However, in the context of RL, \mathbb{P} and R are unknown to the agent, which is expected to come up with an efficient decisional process by interacting with the environment. The way the agent induces this decisional process classifies it either in value-based or policy-based methods.

2.1.2 Value-based methods

In value-based methods, the agent learns to optimally estimate a *value function*, which in turn dictates the policy of the agent by selecting the action of the highest value. One usually defines the *state value function*:

$$V^\pi(s) = \mathbb{E}_{\tau \sim \pi} [R(\tau)|s],$$

and the *state-action value function*:

$$Q^\pi(s, a) = \mathbb{E}_{\tau \sim \pi} [R(\tau)|s, a],$$

which respectively denote the expected discounted cumulative reward starting in state s (resp. starting in state s and taking action a) and then follow trajectory τ according to policy π . It is fairly straightforward to see that these two concepts are linked as follows:

$$V^\pi(s) = \mathbb{E}_{a \sim \pi} [Q^\pi(s, a)],$$

meaning that in practice, $V^\pi(s)$ is the weighted average of $Q^\pi(s, a)$ over all possible actions by the probability of each action. Finally, the *state-action advantage function* can be defined as $A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s)$. One of the main value-based methods in use is called Q-learning, as it relies on the learning of the Q-function to find an optimal policy. In classical Q-learning, the Q-function is stored in a Q-table, which is a simple array representing the estimated value of the optimal Q-function $Q^*(s, a)$ for each pair $(s, a) \in \mathcal{S} \times \mathcal{A}$. The Q-table is initialized randomly, and its values are progressively updated as the agent explores the environment, until the Bellman optimality condition (Bellman & Dreyfus, 1962) is reached:

$$Q^*(s, a) = R(s, a) + \gamma \max_{a'} Q^*(s', a'). \quad (1)$$

The Bellman equation indicates that the Q-table estimate of the Q-value has converged and that systematically taking the action with the highest Q-value leads to the optimal policy. In practice, the expression of the Bellman equation (Bellman & Dreyfus, 1962) is used to update the Q-table estimates.

2.1.3 Policy-based methods

Another approach consists of directly optimizing a parameterized policy $\pi_\theta(a|s)$, instead of resorting to a value function estimate to determine the optimal policy. In contrast with value-based methods, policy-based methods offer three main advantages:

- ◇ They have better convergence properties, although they tend to be trapped in local minima;
- ◇ They naturally handle high dimensional action spaces;
- ◇ They can learn stochastic policies.

To determine how "good" a policy is, one defines an objective function based on the expected cumulative reward:

$$J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} [R(\tau)],$$

and seeks the optimal parameterization θ^* that maximizes $J(\theta)$:

$$\theta^* = \arg \max_{\theta} \mathbb{E}[R(\tau)].$$

To that end, the gradient of the cost function is needed. This is not an obvious task at first, as one is looking for the gradient with respect to the policy parameters θ , in a context where the effects of policy changes on the state distribution are unknown. Indeed, modifying the policy will most certainly modify the set of visited states, which could affect performance in an unknown manner. This derivation is a classical exercise that relies on the log-probability trick Williams (1992), which allows expressing $\nabla_{\theta} J(\theta)$ as an evaluable expected value:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \nabla_{\theta} \log(\pi_{\theta}(a_t | s_t)) R(\tau) \right]. \quad (2)$$

This gradient is then used to update the policy parameters:

$$\theta \leftarrow \theta + \lambda \nabla_{\theta} J(\theta). \quad (3)$$

As $\nabla_{\theta} J(\theta)$ takes the form of expected value, in practice, it is evaluated by averaging its argument over a set of trajectories. In such vanilla policy gradient methods, the actions are evaluated at the end of the episode (they belong to Monte-Carlo methods). If some low-quality actions are taken along the trajectory; their negative impact will be averaged by the high-quality actions and will remain undetected. This problem is overcome by actor-critic methods, in which a Q-function evaluation is used in conjunction with a policy optimization.

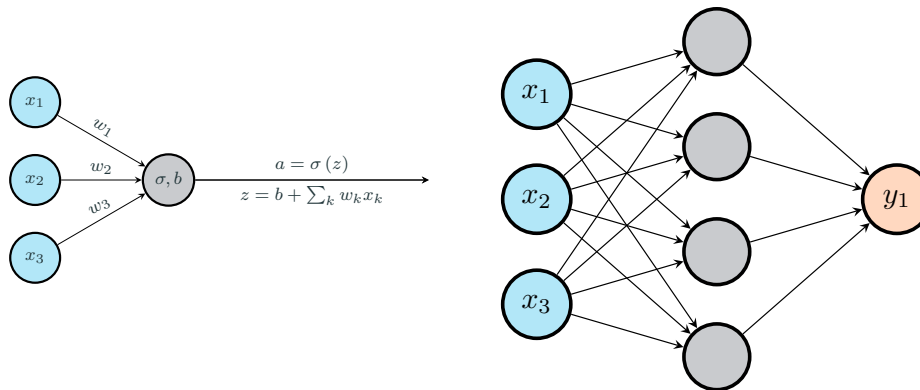
2.2 Deep Reinforcement Learning

In RL, the evaluation of value functions such as V^{π} or Q^{π} can be done in several ways. One possible way, as presented earlier, is to use tables that store the values for every state or action-state pair. However, this strategy does not scale with the sizes of state and action spaces. Another possibility is to use neural networks to estimate value functions, or, for policy-based methods, to output an action distribution given input states. Artificial neural networks can be seen as universal function approximators Siegelmann & Sontag (1995), which means that, if properly trained, they can represent arbitrarily complex mappings between spaces. This particularly suitable property of ANNs gave rise to deep reinforcement learning, *i.e.* reinforcement learning algorithms using deep neural networks as functions approximators. Regarding the notions introduced in previous sections, it is natural to think of θ as the parameterization of the neural network, *i.e.* as its set of weights and biases. In the following section, a basic introduction to neural networks is given. The reader is referred to Goodfellow *et al.* (2017) for additional details.

2.2.1 Artificial neurons and neural networks

The basic unit of a neural network (NN) is the neuron, whose representation is given in figure 2. An input vector x , associated with a set of weights w , is provided to the neuron. The neuron then computes the weighted sum $w \cdot x + b$, where b is called the bias and applies the activation function σ to this sum. This is the output of the neuron, hereafter noted z . In the neuron, the weights and the bias represent the degrees of freedom (*i.e.*, the parameters that can be adjusted to approximate the function f), while the activation function is a hyper-parameter, *i.e.*, it is part of the choices made during the network design.

In their purest form, neural networks consist of several layers of neurons connected to each other, as shown in the primary example of 2. Such a network is said to be fully connected (FC), in the sense that each neuron of a layer is connected to all the neurons of the following layer. As it was said in the last section, each connection is characterized by a weight, and in addition each neuron has a bias, except for the input layer. Indeed, as a convention, the input layer is usually drawn as a regular layer, but it does not hold biases, nor is an activation



(a) Representation of a single artificial neuron. The neuron receives an input vector x , assorted with a weight vector w . The output is computed as $w \cdot x$ corrected with the bias b , to which is applied the activation function σ .

(b) Simple example of neural network with an input vector $x \in \mathbb{R}^3$, a hidden layer composed of 4 neurons, and an output layer composed of a single neuron. As a convention, input variables are drawn using a neuron representation. However, it must be kept in mind that the input layer is not composed of neurons.

Figure 2: **Representation of a single artificial neuron and a basic 3-layer neural network.**

function applied at its output. The learning process in neural networks consists in adjusting all the biases and weights of the network in order to reduce the value of a well-chosen loss function that represents the quality of the network prediction. This update is usually performed by a stochastic gradient method, in which the gradients of the loss function with respect to the weights and biases are estimated using a back-propagation algorithm.

When working for instance with images as input, it is customary to exploit convolutional layers instead of FC ones. Rather than looking for patterns in their entire input space as FC layers, convolutional ones can extract local features. Additionally, they can build a hierarchy of increasingly sophisticated features. In such layers, a convolution kernel (*i.e.*, a tensor product with a weight matrix) is applied on a small patch of the input image and is used as input for a neuron of the next layer. The patch is then moved, and the operation repeated until the input image has been entirely covered. The size of the patch is usually known as the kernel size. The complete coverage of the image with this process generates a kernel, also called filter or feature map. In most cases, multiple kernels are generated at each layer, each encoding a specific feature of the input image. When used for regression applications, convolutional networks (or *convnets*) most often end with a fully connected layer, followed by the output layer, which size is determined by that of the sought quantity of interest. These considerations are discussed in details in a large variety of books and articles such as Goodfellow *et al.* (2017) Strang (2019) Bottou *et al.* (2018).

2.3 DRL algorithms

In this section, the main lines of the major DRL algorithms in use today are presented. For the sake of brevity, many details and discussions are absent from the following discussions. The reader is referred to the profuse literature on the topic.

2.3.1 Deep Q-networks

Instead of updating a table holding Q-values for each possible (s, a) pair, ANNs can be used to generate a map $\mathcal{S}^+ \times \mathcal{A} \rightarrow \mathbb{R}$ (called deep Q-networks (DQN)) that provides an estimate of the Q-value for each possible action given an input state. As it is usual for neural networks, the update of the Q-network requires a loss function for the gradient descent

algorithm, that will be used to optimize the network parameters θ . To do so, it is usual to exploit the Bellman optimality equation (1):

$$L(\theta) = \mathbb{E} \left[\frac{1}{2} \left(\underbrace{[R(s, a) + \gamma \max_{a'} Q_\theta(s', a')] - Q_\theta(s, a)}_{\text{Target}} \right)^2 \right]. \quad (4)$$

In the latter expression, $Q_\theta(s, a)$ represents the Q-value *estimate* provided by the DQN for action s and state a under network parameterization θ . It must be noted that the quantity denoted *target* is the same that appears in the Bellman optimality equation (1): when the optimal set of parameters θ^* is reached, $Q_\theta(s, a)$ is equal to the target, and $L(\theta)$ is equal to zero.

In vanilla deep Q-learning (and more generally in Q-learning), an exploration/exploitation trade-off must be implemented, to ensure a sufficient exploration of the environment. To do so, a parameter $\epsilon \in [0, 1]$ is defined, and a random value is drawn in the same range before each action. If this value is below ϵ , a random action is taken. Otherwise, the algorithm follows the action prescribed by $\max_a Q_\theta(s, a)$. Most often, ϵ follows a schedule, starting with high values at the beginning of learning, and progressively decreasing. The vanilla algorithm using DQN is shown below :

Algorithm 1 Vanilla deep Q-learning

```

1: Initialize action-value function parameters  $\theta_0$ 
2: for  $k = 0, M$  do                                     ▷ Loop over episodes
3:   for  $t = 0, T$  do                                       ▷ Loop over time stations
4:     Draw random value  $\omega \in [0, 1]$                        ▷  $\epsilon$ -greedy strategy
5:     if ( $\omega < \epsilon$ ) then
6:       Choose action  $a_t$  randomly
7:     else
8:       Choose action  $a_t = \max_a Q_\theta(s_t, a)$ 
9:     end if
10:    Execute action  $a_t$  and observe reward  $r_t$  and state  $s_{t+1}$ 
11:    Form target  $y_t = r_t + \max_a Q_\theta(s_{t+1}, a)$ 
12:    Perform gradient descent using  $(y_t - Q_\theta(s_t, a_t))^2$ 
13:    Update  $\theta$  and  $s_t$ 
14:   end for
15: end for

```

Still, the basic deep Q-learning algorithm presents several flaws, and many improvements can be added to it:

- ◇ **Experience replay:** To avoid forgetting about its past learning, a replay buffer can be created that contains random previous experiences. This buffer is regularly fed to the network as learning material, so previously acquired behaviors are not erased by new ones. This improvement also reduces the correlation between experiences: as the replay buffer is randomly shuffled, the network experiences past (s_t, a_t, r_t, s_{t+1}) tuples in a different order, and is therefore less prone to learn correlation between these (Tsitsiklis & Van Roy, 1997; Lin, 1993);
- ◇ **Fixed Q-targets:** When computing the loss (4), the same network is used to evaluate the target and $Q_\theta(s_t, a_t)$. Hence, at every learning step, both the Q-value and the target move, *i.e.* the gradient descent algorithm is chasing a moving target, implying big oscillations in the training. To overcome this, a separate network is used to evaluate the target (target network), and its weight parameters θ^- are updated less frequently than that of the Q-network. This way, the target remains fixed for several Q-network updates, making the learning easier;
- ◇ **Double DQN:** The risk of using a single Q-network is that it may over-estimate the Q-values, leading to sub-optimal policies, as π_θ is derived by systematically

choosing the actions of highest Q-values. In (van Hasselt *et al.*, 2015), two networks are used jointly: a DQN is used to select the best action, while the other one is used to estimate the target value;

- ◇ **Prioritized replay buffer:** In the exploration of the environment by the agent, some experiences might occur rarely while being of high importance. As the memory replay batch is sampled uniformly, these experiences may never be selected. In Schaul *et al.* (2015), the authors propose to sample the replay buffer according to the training loss value, in order to fill the buffer in priority with experiences where the Q-value evaluation was poor, meaning that the network will have better training on "unexpected events". To avoid overfitting on these events, the filling of the prioritized buffer follows a stochastic rule.

2.3.2 Deep policy gradient

When the policy π_θ is represented by a deep neural network with parameters θ , the evaluation of $\nabla_\theta J(\theta)$ can be delegated to the back-propagation algorithm (see section 2.2.1), as long as the gradient of the loss function equals the policy gradient (2). The loss can be expressed as:

$$L(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \log(\pi_\theta(a_t | s_t)) R(\tau) \right].$$

However, it can be shown that using the full history of discounted reward along trajectory τ is not necessary, as a given action a_t has no influence on the rewards obtained at previous time-stations. Several expressions are possible to use **instead** of $R(\tau)$, that do not modify the value of the computed policy gradient, while reducing its variance (and therefore the required amount of trajectories to correctly approximate the expected value) (Schulman *et al.*, 2015). Among those, the advantage function $A(s, a)$ is commonly chosen for its low variance:

$$A(s, a) = Q(s, a) - V(s).$$

The advantage function $A(s, a)$, represents the improvement obtained in the expected cumulative reward when taking action a in state s , compared to the average of all possible actions taken in state s . In practice, it is not readily available and must be estimated using the value function. In practice, this estimate is provided by a separately learned value function. As a result, the loss function writes as:

$$L(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \log(\pi_\theta(a_t | s_t)) A^{\pi_\theta}(s_t, a_t) \right].$$

The vanilla deep policy gradient algorithm is presented below:

Algorithm 2 Vanilla deep policy gradient

- 1: Initialize policy parameters θ_0
 - 2: **for** $k = 0, M$ **do** ▷ Loop over batches of trajectories
 - 3: **for** $i = 0, m$ **do** ▷ Loop over trajectories in batch k
 - 4: Execute policy π_{θ_k}
 - 5: Collect advantage estimates $A^{\pi_{\theta_k}}(s_t, a_t)$
 - 6: **end for**
 - 7: Estimate $L(\theta_k)$ over current batch of trajectories
 - 8: Update policy parameters θ_k using stochastic gradient ascent
 - 9: **end for**
-

2.3.3 Advantage actor-critic

As mentioned in section 2.1.3, actor-critic methods propose to simultaneously exploit two networks to improve the learning performance. One of them is policy-based, while the other one is value-based:

- ◇ The actor $\pi_\theta(s, a)$, controls the actions taken by the agent;
- ◇ The critic $Q_w(s, a)$, evaluates the quality of the action taken by the actor.

Both these networks are updated in parallel, in a time-difference (TD) fashion (one update at each time station) instead of the usual Monte-Carlo configuration (one update at the end of the episode). In practice, the advantage function $A(s, a)$ is preferred to the Q-value for its lower variability. In order to avoid having to evaluate two value functions, the advantage function is usually approximated, using the reward obtained by the actor after action a_t and the value evaluated by the critic in state s_{t+1} :

$$A_w(s_t, a_t) \sim R(s_t, a_t) + \gamma V_w(s_{t+1}) - V_w(s_t).$$

The actor-critic process then goes as follows:

Algorithm 3 Advantage actor-critic

```

1: Initialize policy parameters  $\theta_0$                                 ▷ Actor
2: Initialize action-value function parameters  $w_0$                 ▷ Critic
3: for  $k = 0, M$  do                                           ▷ Loop over episodes
4:   for  $t = 0, T$  do                                           ▷ Loop over time stations
5:     Execute action  $a_t$  following  $\pi_\theta$ , get reward  $r_t$  and state  $s_{t+1}$     ▷ Actor
6:     Evaluate target  $R(s_t, a_t) + \gamma V_w(s_{t+1})$             ▷ Critic
7:     Update  $w$  to  $w'$  using target                                ▷ Critic
8:     Evaluate advantage  $A_{w'}(s_t, a_t) = R(s_t, a_t) + \gamma V_{w'}(s_{t+1}) - V_{w'}(s_t)$  ▷ Critic
9:     Update  $\theta$  to  $\theta'$  using  $A_{w'}(s_t, a_t)$                 ▷ Actor
10:   end for
11: end for

```

In the latter algorithm, the actor and critic parameter updates follow the methods described in sections 2.3.1 and 2.3.2. The advantage actor critic algorithm has been developed in asynchronous (A3C Mnih *et al.* (2016)) and synchronous (A2C) versions. The version of the algorithm presented here corresponds to a time-difference A2C, although many possible variants exist, such as the n-step A2C.

2.3.4 Trust-region and proximal policy optimization

Trust region policy optimization (TRPO) was introduced in 2015 as a major improvement of vanilla policy gradient Schulman *et al.* (2015). In this method, the network update exploits a *surrogate advantage* functional:

$$\theta_{k+1} = \arg \max_{\theta} L(\theta_k, \theta),$$

with

$$L(\theta_k, \theta) = \mathbb{E}_{(s,a) \sim \pi_{\theta_k}} [\Pi(s, a, \theta, \theta_k) A^{\pi_{\theta_k}}(s, a)],$$

and

$$\Pi(s, a, \theta, \theta_k) = \frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}.$$

In latter expression, $L(\theta_k, \theta)$ measures how much better (or worse) the policy π_θ performs compared to the previous policy π_{θ_k} . In order to avoid too large policy updates that could collapse the policy performance, TRPO leverages second-order natural gradient optimization to update parameters within a trust-region of a fixed maximum Kullback-Leibler divergence between old and updated policy distribution. This relatively complex approach was replaced in the PPO method by simply clipping the maximized expression:

$$L(\theta_k, \theta) = \mathbb{E}_{(s,a) \sim \pi_{\theta_k}} \left[\min (\Pi(s, a, \theta, \theta_k) A^{\pi_{\theta_k}}(s, a), g(\varepsilon, A^{\pi_{\theta_k}}(s, a))) \right],$$

where

$$g(\varepsilon, A) = \begin{cases} (1 + \varepsilon)A & A \geq 0, \\ (1 - \varepsilon)A & A < 0, \end{cases}$$

where ε is a small, user-defined parameter. When $A^{\pi_{\theta_k}}(s, a)$ is positive, taking action a in state s is preferable to the average of all actions that could be taken in that state, and it is natural to update the policy to favor this action. Still, if the ratio $\Pi(s, a, \theta, \theta_k)$ is very large, stepping too far from the previous policy π_{θ_k} could damage performance. For that reason, $\Pi(s, a, \theta, \theta_k)$ is clipped to $1 + \varepsilon$ to avoid too large updates of the policy. If $A^{\pi_{\theta_k}}(s, a)$ is negative, taking action a in state s represents a poorer choice than the average of all actions that could be taken in that state, and it is natural to update the policy to decrease the probability of taking this action. In the same fashion, $\Pi(s, a, \theta, \theta_k)$ is clipped to $1 - \varepsilon$ if it happens to be lower than that value.

Many refinements can be added to these methods that are out of reach of this introduction, such as the use of generalized advantage estimator (Schulman *et al.*, 2016), the introduction of evolution strategies in the learning process (Houthoofd *et al.*, 2018), or the combination of DQN with PG methods for continuous action spaces (Lillicrap *et al.*, 2015).

2.4 A conclusion on DRL

Since the first edition of the book of Sutton Sutton & Barto (1998), RL has become a particularly active field of research. However, the domain really started striving after the paper from Mnih *et al.* (2013) and its groundbreaking results using DQN to play Atari games. As of today, the domain holds countless major achievements, and the ability of DRL to learn to achieve complex tasks is well established. It is now exploited in multiple fields of research, and is featured in an ever-increasing amount of publications, as shown in figure 3.

A major advantage of DRL is that it can be used as an agnostic tool for control and optimization tasks, both in continuous and discrete contexts. Its only requirement is a well-defined (s_t, a_t, r_t) interface from the environment, which can consist in a numerical simulation or a real-life experiment. As seen earlier, both DQN and DPG ensure a good exploration of the states space, making them able to unveil new optimal behaviors. Moreover, DRL is also robust to transfer learning, meaning that an agent trained on one problem will train much faster on a similar environment. This represents a great strength, especially for domains like fluid mechanics, where computational time represents a limiting factor (section 3.7 illustrates this point). Transfer learning also introduces a large contrast between DRL and regular optimization techniques, such as adjoint methods, where the product of optimization in a given case cannot be re-used to speed up optimization in a similar configuration. Finally, as illustrated in section 3.6, DRL libraries can exploit parallel learning with a close-to-perfect scaling on the available resources. Hence, for its ease of use and its numerous capabilities, DRL represents a promising tool for optimization and design processes involving computational or experimental environments.




Figure 3: **Number of publications mentioning "Deep reinforcement learning" per year** since 2010 (data from www.scholar.google.com). In blue, we indicate some of the most important algorithms or achievements of the field.

3 Applications

In this section, several applications combining DRL and fluid mechanics found in the literature are presented in details. For each case, the numerical experiments and the obtained results are detailed. The choice of DRL algorithm and the problem complexity are also considered. Table 2 presents an overview of the reviewed articles and their corresponding references.

Table 2: **Classification of the papers scanned in this review.** ADQN and DDPG respectively stand for asynchronous DQN and deep deterministic PG.

Main topic	Algorithm	References
Flow Control	QL	Guéniat <i>et al.</i> (2016)
	DQN	Gazzola <i>et al.</i> (2016), Novati <i>et al.</i> (2017)
	ADQN	Verma <i>et al.</i> (2018)
	TRPO	Ma <i>et al.</i> (2018)
	A3C	Garnier & Viquerat (2019)
Homogeneity optimization	PPO	Rabault <i>et al.</i> (2019), Rabault & Kuhnle (2019)
	QL	Colabrese <i>et al.</i> (2017), Gustavsson <i>et al.</i> (2017)
Shape Optimization	QL	Qiu <i>et al.</i> (2018), Hou Tsang <i>et al.</i> (2018)
	QL	American Institute of Aeronautics and Astronautics (2008)
Chaotic Systems	DDPG	Bucci <i>et al.</i> (2019)

3.1 Synchronised swimming of two fish - Novati *et al.* (2017)

In Novati *et al.* (2017), the authors study the swimming kinematics of two fish in a viscous incompressible flow. The first fish is the leader, and is affected a prescribed gait, while the second fish is a follower, which dynamics are unknown. An objective of the paper is to exploit DRL to derive a swimming strategy that reduces the energy expenditure of the follower. A two-dimensional Navier-Stokes equation is solved to simulate the viscous incompressible flow using a remeshed vortex method. To represent the fish and its undulation, a simplified physical model is used that describes the body curvature $k(s, t)$ of the fish:

$$k(s, t) = A(s) \sin \left[2\pi \left(\frac{t}{T_p} - \frac{s}{L} \right) + \phi \right], \quad (5)$$




Figure 4: **Leader and follower swimmer, reproduced from Novati *et al.* (2017)**, with the displacements Δx and Δy as well as the orientation θ between the leader and the follower.

where L is the length of the fish, T_p is the tail-beat frequency, ϕ is a phase-difference, t is the time variable, and s is the abscissa. The observation of the environment provided to the DRL agent is composed of:

- ◇ The displacements Δx and Δy (cf figure 4),
- ◇ The orientation θ between the follower and the leader (cf figure 4),
- ◇ The moment where the action is going to take place (during one period of a tail-beat undulation),
- ◇ The last two actions taken by the agent.

The agent provides actions that correspond to an additive curvature term for expression (5):

$$k'(s, t) = A(s)M\left(\frac{t}{T_p} - \frac{s}{L}\right).$$

The reward to maximize is defined as $r_t = 1 - 2\frac{|\Delta y|}{L}$, which penalizes the follower when it laterally strays away from the path of the leader. Finally, the state space is artificially restricted by terminating the current episode with an arbitrary reward $r_t = -1$ when the two fish get too far away from each other. In this application, a DQN algorithm is used to control the follower. For the best strategy found by the agent, swimming efficiency was increased by 20% compared to baseline of multiple solitary swimmers.

3.2 Efficient collective swimming by harnessing vortices through deep reinforcement learning Verma *et al.* (2018)

The authors in Verma *et al.* (2018) further investigated the swimming strategies of multiple fishes. They work with two and three-dimensional flows (see figure 5), and analyze with many details the movement of trained fishes. The DRL model is identical to what is used in Novati *et al.* (2017), but the neural network used was improved by using recurrent neural networks with Long-Short Term Memory (LSTM) layers (Hochreiter & Schmidhuber, 1997). The authors underline the importance of this feature, arguing that here, past observations hold informations that are relevant for future transitions (*i.e.* the process is not Markovian anymore). Two different rewards are used: the first one is identical to that of last section, while the second one is simply equal to the swimming efficiency of the follower (see Verma *et al.* (2018)). The result obtained show that collective energy-savings could be achieved with an appropriate use of the wake generated by other swimmers, when compared to the case of multiple solitary swimmers.




Figure 5: **Leader and follower swimmer, reproduced from Verma *et al.* (2018)**



Figure 6: **Fluid jets to control rigid body, reproduced from Ma *et al.* (2018).**

3.3 Fluid directed rigid body control using deep reinforcement learning Ma *et al.* (2018)

The control of complex environments that include two fluid flows interacting with multiple rigid bodies (cf figure 6) was first proposed by Ma *et al.* (2018). In this contribution, a multiphase Navier-Stokes solver is used to simulate the full characteristics of the interactions between fluids and rigid bodies at different scales. To provide a rich set of observations to the agent while limiting its input dimension, a convolutional autoencoder is used to extract low-dimensional features from the high-dimensional velocity field of the fluid (Liou *et al.*, 2014). This technique allows the use of a smaller neural network with faster convergence. The autoencoder is trained using a TRPO agent, which received two different observations: (i) extracted features from the fluid velocity field, and (ii) features from the rigid body.

In the environment, the agent is given control over several fluid jets (see figure 6), each of them described by 3 parameters: (i) the jet lateral acceleration \ddot{x}_{jet} , (ii) the jet angular acceleration $\ddot{\beta}_{jet}$, and (iii) the use of the jet δ_{jet} . Using lateral and angular accelerations instead of lateral and angular positions allows the jet to move smoothly even with a noisy policy. Different rewards were designed depending on the environment goals. A possible example consists in having a rigid body balanced at a fixed position. In this case the designed reward is:

$$r_t = w_c \exp(-\|c - c^*\|^2) + w_v \exp(-\|v\|^2) + w_e(1 - \delta_{jet}), \quad (6)$$

where the first term penalizes a position mismatch according to the goal c^* , the second penalizes the presence of a residual velocity, and the last one encourages the agent to use the smallest amount of control forces. w_c , w_v and w_e are user-defined parameters used to weight the three contributions.

Overall, impressive results were obtained on complex experiments, including playing music with the system or keeping a rotating object at a given position. An important feature of this contribution is the use of the autoencoder to extract features from the high-dimensional velocity fluid fields. Of particular interest is the simultaneous training of the two networks (autoencoder and policy network), leading to a high convergence rate of the overall method.




Figure 7: **Targets flow and results obtained from DRL agent**, reproduced from Lee *et al.* (2018).

3.4 Flow shape design for microfluidic devices using deep reinforcement learning Lee *et al.* (2018)

In Lee *et al.* (2018), the authors explore the capabilities of DRL for microfluidic flow sculpting using pillar-shaped obstacles to deform an input flow. Given a target flow, a double DQN is used to generate the corresponding optimal sequence of pillars. The observations provided to the agent consist in the current flow shape, represented as a matrix of black and white pixels. In return, the actions correspond to pillar configurations: the agent can add up to 32 pillars in order to reshape the input flow. A pixel match rate (PMR) function is used to evaluate the similarity between the target flow and the current flow. The reward is then defined using that PMR as:

$$r_t = -\left(1 - \frac{\text{PMR} - b}{b}\right),$$

where b is the PMR of the worst-case scenario performed. PMR over 90% were obtained for different target flows, showing that DRL represents a viable alternative to other optimization processes such as genetic algorithms (see figure 7). Transfer learning is also exploited, showing that an agent trained on a first flow could be retrained much more efficiently on a different flow than an untrained one, thus saving a lot of computational time.

3.5 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control (Rabault *et al.*, 2019)

In Rabault *et al.* (2019), the authors present the first application of DRL for performing active flow control in a simple CFD simulation. For this, a simple benchmark is used. More specifically, the flow configuration is based on the one presented in Schäfer *et al.* (1996) and features a cylinder immersed in a 2D incompressible flow injected following a parabolic velocity profile. The Reynolds number based on the cylinder diameter is moderate to keep the CFD affordable, and following Schäfer *et al.* (1996), $\text{Re}=100$ is used. The incremental pressure correction scheme is used together with a finite element approach to solve the discretized problem. A periodic Karman **vortex street** is then obtained, as visible in figure 8. Besides, small jets are added on the sides of the cylinder in order to allow controlling the separation of the wake. Therefore, the action performed by the DRL agent is to decide the instantaneous mass flow rate of the small jets, while the reward is chosen to encourage drag reduction. Finally, probes reporting the characteristics of the flow at several fixed points in the computational domain (either pressure or velocity can be used, both providing equivalent results) are added in the vicinity of the cylinder and are used as the observation provided as input to the network.

Good results are obtained, with training performed in around 1300 vortex shedding periods. The control strategy is more complicated than traditional harmonic forcing that had been performed in previous works (Bergmann *et al.*, 2005), which illustrates the value of using an




Figure 8: **Comparison of the velocity magnitude without (top) and with (bottom) active flow control**, reproduced from Rabault *et al.* (2019). A clear modification of the cylinder wake, similar to what would be obtained with boat-tailing, is visible.

ANN as the controller. Typically, around 93% of the drag induced by the vortex shedding effect (Bergmann *et al.*, 2005) is suppressed by the control law found, and the strength of the jets needed to reduce the drag is minimal. In the established regime, the magnitude of the mass flow rate injected by the jets normalized by the mass flow rate of the base flow intersecting the cylinder is only of typically 0.6%. The resulting flow, visible in figure 8, is similar to what would be obtained with boat tailing; more specifically, an apparent increase in the recirculation bubble is observed.

The work of Rabault *et al.* (2019) shows that active flow control is, at least on a simple flow configuration, achievable with small actuations if an adapted control law is used. While there are still many unanswered questions at the moment, such as the robustness of the control law, the application of the methodology to higher, more realistic (from an industrial point of view) Reynolds numbers, or the possibility to control 3D flows, this work establishes DRL as a new methodology that should be further investigated for the task of performing active flow control.

3.6 Accelerating deep reinforcement learning of active flow control strategies through a multi environment approach (Rabault & Kuhnle, 2019)

One of the main practical limitations of Rabault *et al.* (2019) lies within the time needed to perform the learning - around 24 hours on a modern CPU. This comes from the inherent cost of the CFD, and the fact that it is challenging to obtain computational speedups on small tasks. In particular, increasing the number of cores used for solving the CFD problem of Rabault *et al.* (2019) provides very little speedup, as the communication between the cores negates other speed gains. While this is particularly true on such a small computational problem, any CFD simulation will only speed up to a maximum point independently of the number of cores used. Therefore, a solution is needed to provide the data necessary for the ANN to perform learning in a reasonable amount of time.

In Rabault & Kuhnle (2019), the collection of data used for filling the memory replay buffer is parallelized by using an ensemble of simulations all running in parallel, independently of each other. In practice, this means that the knowledge obtained from several CFD simulations running in an embarrassingly parallel fashion can be focused on the same ANN to speed up its learning.

This technique results in large speedups for the learning process, as visible in figure 9. In particular, for a number of simulations that are a divider of the learning frequency of the ANN, perfect scaling is obtained, and the training is formally equivalent to the one obtained in serial (*i.e.*, with one single simulation). However, even more simulations can be used, and it is found empirically that the resulting off-policy data sampling does not sensibly reduce




Figure 9: **Illustration of the acceleration of ANN training through DRL using a multi environment approach** (reproduced from Rabault & Kuhnle (2019)). Both plots illustrate the convergence of the ANN learning as a function of UTC time, depending on the number of environments used in parallel. In both cases, 3 repetitions are performed for each number of environments. Thin lines indicate individual learnings. The average of all three learnings in each case is indicated by a thick line. In the case when the number of environments is a divider of the learning frequency of the ANN, the learning is formally equivalent to the serial case and perfect scaling is observed (left). If a larger number of environments is used, some stepping is observed in the learning process, but this does not affect the ability to perform learning, nor the general speedup of the process (right).

Table 3: **Summary of the six trained DRL agents** from Garnier & Viquerat (2019).

Agent	Re	Objective	Details
1.1	10	Control over 17 time steps	Trained from scratch
1.2	10	Direct optimization	Trained from scratch
2.1	40	Control over 17 time steps	Transfer learning from 1.1
2.2	40	Direct optimization	Transfer learning from 1.2
3.1	100	Control over 17 time steps	Transfer learning from 1.1
3.2	100	Direct optimization	Transfer learning from 1.2

the learning ability, as visible in figure 9. This means in practice that the authors achieve a training speedup factor of up to around 60.

This natural parallelism of the learning process of ANNs trained through DRL represents a significant advantage for this methodology, and a promising technical result that opens the way to the control of much more sophisticated situations. [It is worth noting that even larger speedups can be obtained by making use of surrogate or reduced-order models, either pre-trained or trained on-the-fly, although specific care must be taken regarding the validity range of these models \(as the agent could end up providing inputs to the surrogate model that would not match its training range\).](#)

3.7 Sensitivity of aerodynamic forces in laminar flows past a square cylinder

In this section, we propose an original comparison between a classical optimization approach and DRL. The test case is inspired by Meliga *et al.* (2014), where two cylinders are immersed in a moderate Reynolds flow. One is a square main cylinder, which position is fixed, while the other is a cylindrical control cylinder, which position can vary, and which radius is considered much smaller than that of the main cylinder. The goal of this application is to optimize the position of the small cylinder, such that the total drag of the two cylinders is inferior to that of the main cylinder alone. An incompressible Navier-Stokes solver Alnæs *et al.* (2015) is used to simulate the flow at different Reynolds number past the square cylinder. Moreover, the flow conditions are identical to that described in Rabault *et al.* (2019). Six different agents are trained in different conditions, as detailed in table 3.

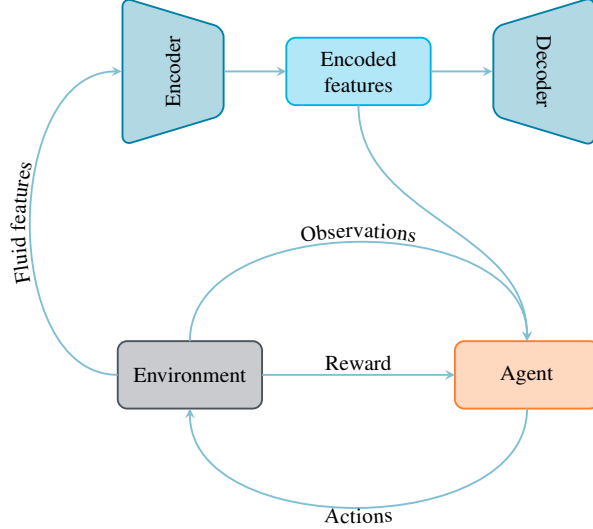


Figure 10: **Architecture for agents 1.1, 2.1 and 3.1.**

Two different architectures are proposed. The first one is a classical DRL architecture, similar to what is shown in figure 1, but with a difference that episodes were reduced to a single action. In that case, the agent directly attempts to propose an optimal (x, y) position from the same initial state, using a PPO algorithm. On the contrary, in the second architecture, successive displacements $(\Delta x, \Delta y)$ of the control cylinder are provided by an A3C agent, starting from a random initial position. The observations provided to the agent are made of the last two positions of the control cylinder, the Reynolds number and encoded features from the velocity fields. To that end, an autoencoder was trained simultaneously to the agent to reduce the state dimensionality (with a dimension reduction from around 8000 to 70), in a similar fashion to Ma *et al.* (2018) (see figure 10). The same reward function is used for both architectures:

$$r_t = C_D^0 - C_D,$$

where C_D^0 is the drag of the main cylinder alone, C_D is the total drag of the main and the control cylinders, and the drag coefficients are counted positively. With this simple expression, a positive reward is obtained when the combined drag of the cylinders is inferior to the reference drag value.

We compare the results obtained with those from Meliga *et al.* (2014), where a classical adjoint method is used. Overall, the same optimal positions were found for the control cylinder, at $Re = 40$ and $Re = 100$. Although the case $Re = 10$ was not present in the original paper, coherent results were obtained for the configuration too. For the first agents at $Re = 10$, training took approximately 5 hours using parallel learning with 17 CPUs. Regarding the last four agents, the transfer learning allowed them to learn at a much faster pace: only 10 hours were necessary on a single CPU. An overview of the results is shown in figure 11.

With this experiment, we showed that not only DRL could perform both control and direct optimization tasks, but also that transfer learning between different configurations represents a powerful feature to save computational time. We also confirmed that the use of an autoencoder is a safe and robust way to extract essential features from complex and high dimensional fluid fields. Details about this test case and the code used to reproduce these results can be found on the GitHub repository: <https://github.com/DonsetPG/fenics-DRL>. We recall that this code uses different library:

- ♦ FEniCS (Alnæs *et al.*, 2015) for the CFD solver,

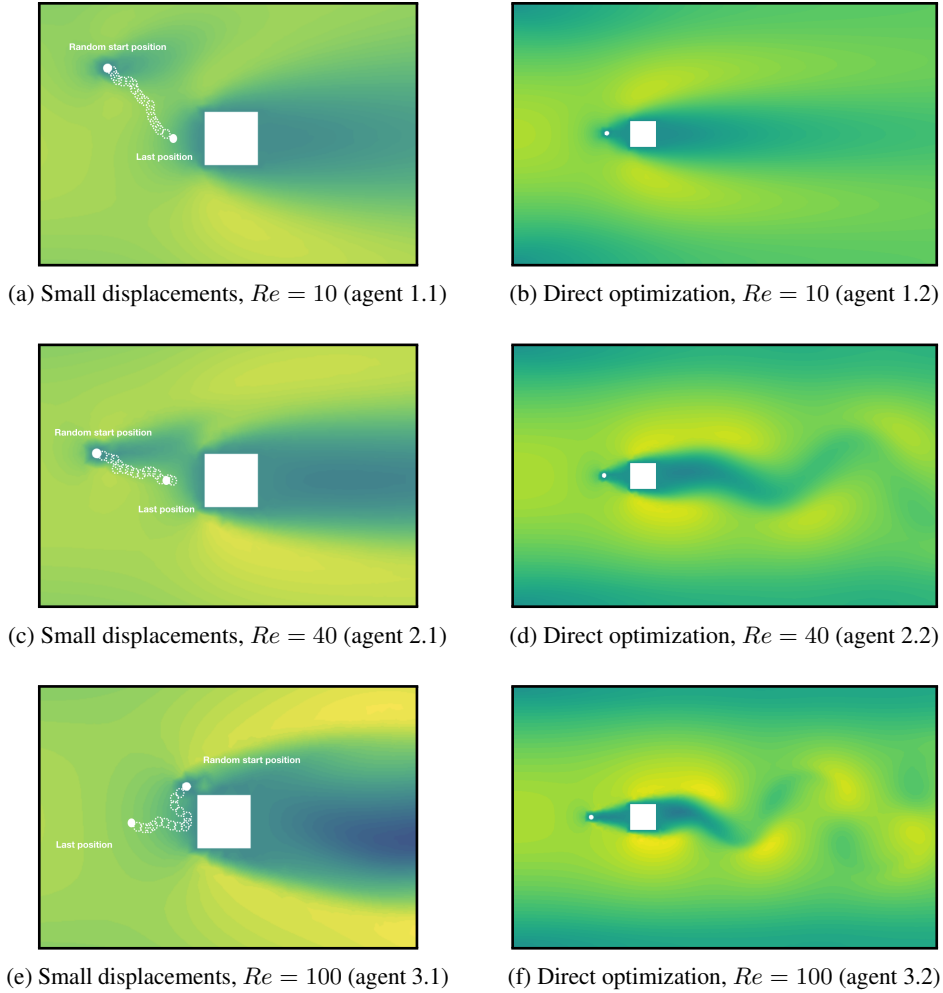


Figure 11: **Results obtained by the trained agents.** The environment consists of a square of unit lateral size centered in $(0, 0)$, immersed in a rectangular domain of dimension $[-5, 20] \times [-5, 5]$.

- ◇ Gym (Brockman *et al.*, 2016) for the DRL environment,
- ◇ Stable-baselines (Hill *et al.*, 2018) for the DRL algorithms.

Future work on this library will focus on generalizing the coupling of DRL with different open source CFD code.

4 Conclusion

Although DRL has already been applied to several cases of optimization and control in the context of fluid dynamics, the literature on the topic remains particularly shallow. In the present article, we reviewed the available contributions on the topic to provide the reader with a comprehensive state of play of the possibilities of DRL in fluid mechanics. In each of them, details were provided on the numerical context and the problem complexity. The choices of the DRL algorithm and the reward shaping were also described.

Given the high-level interfaces of existing libraries, the coupling of DRL algorithms with existing numerical CFD solvers can be achieved with a minimal investment, while opening a wide range of possibilities in terms of optimization and control tasks. The algorithms presented in this review proved to be robust when exposed to possible numerical noise,

although high Re applications remain to be achieved. Additionally, the parallel capabilities of the main DRL libraries represent a major asset in the context of time-expensive CFD computations. Transfer learning also proved to be a key feature in saving computational time, as re-training agents already exploited in similar situations led to a fast convergence of the agent policy. The use of autoencoders to feed the agent with both a compact and rich observation of the environment also showed to be beneficial, as it allows for a reduced size of the agent network, thus implying a faster convergence.

As of now, the capabilities and robustness of DRL algorithms in highly turbulent and non-linear flows remain to be explored. Also, their behavior and convergence speed in action spaces of high dimensionalities is unknown. It makes no doubt that the upcoming years will see the mastering of these obstacles, supported by the constant progress made in the DRL field and driven by the numerous industrial challenges that could benefit from it.

References

- ALI, N & BEHDINAN, K 2003 Optimal geometrical design of aircraft using genetic algorithms. *Transactions of the Canadian Society for Mechanical Engineering* **26**, 373–388.
- ALNÆS, M.S., BLETTCHA, J, HAKE, A, JOHANSSON, B, KEHLET, B, LOGG, A, RICHARDSON, C., RING, J., ROGNES, M.E. & WELLS, G.N. 2015 The FEniCS Project Version 1.5. *Archive of Numerical Software* **3**.
- AMERICAN INSTITUTE OF AERONAUTICS AND ASTRONAUTICS, ed. 2008 *Morphing Airfoils with Four Morphing Parameters*.
- BAHDANAU, DZMITRY, BRAKEL, PHILEMON, XU, KELVIN, GOYAL, ANIRUDH, LOWE, RYAN, PINEAU, JOELLE, COURVILLE, AARON & BENGIO, YOSHUA 2016 An Actor-Critic Algorithm for Sequence Prediction. *CoRR* **abs/1607.07086** (2015), 1–17.
- BELLMAN, R. 1957 A Markovian Decision Process. *Journal of Mathematics and Mechanics* **6** (5), 679–684.
- BELLMAN, R. & DREYFUS, S. E. 1962 *Applied dynamic programming*. Princeton University Press Princeton, N.J.
- BERGMANN, MICHEL, CORDIER, LAURENT & BRANCHER, JEAN-PIERRE 2005 Optimal rotary control of the cylinder wake using proper orthogonal decomposition reduced-order model. *Physics of fluids* **17** (9), 097101.
- BEWLEY, ALEX, RIGLEY, JESSICA, LIU, YUXUAN, HAWKE, JEFFREY, SHEN, RICHARD, LAM, VINH-DIEU & KENDALL, ALEX 2018 Learning to Drive from Simulation without Real World Labels. *arXiv e-prints* , arXiv: 1812.03823.
- BOTTOU, LÉON, CURTIS, FRANK E. & NOCEDAL, JORGE 2018 Optimization Methods for Large-Scale Machine Learning. *SIAM Rev.* **60**, 223–311.
- BROCKMAN, GREG, CHEUNG, VICKI, PETERSSON, LUDWIG, SCHNEIDER, JONAS, SCHULMAN, JOHN, TANG, JIE & ZAREMBA, WOJCIECH 2016 OpenAI Gym. *arXiv e-prints* .
- BROWN, NOAM & SANDHOLM, TUOMAS 2019 Superhuman AI for multiplayer poker. *Science* , arXiv: <https://science.sciencemag.org/content/early/2019/07/10/science.aay2400>.
- BUCCI, MICHELE ALESSANDRO, SEMERARO, ONOFRIO, ALLAUZEN, ALEXANDRE, WISNIEWSKI, GUILLAUME, CORDIER, LAURENT & MATHELIN, LIONEL 2019 Control of chaotic systems by Deep Reinforcement Learning. *arXiv p.* arXiv:1906.07672.
- COLABRESE, SIMONA, GUSTAVSSON, KRISTIAN, CELANI, ANTONIO & BIFERALE, LUCA 2017 Flow Navigation by Smart Microswimmers via Reinforcement Learning. *Physical Review Letters* **118** (15), 158004.
- FRANÇOIS-LAVET, VINCENT, HENDERSON, PETER, ISLAM, RIASHAT & BELLEMARE, MARC G 2018 An Introduction to Deep Reinforcement Learning. *Foundations and trends in machine learning* .
- GARNIER, PAUL & VIQUERAT, JONATHAN 2019 Position of a control cylinder to reduce drag. <https://github.com/DonsetPG/fenics-DRL>.

- GAZZOLA, M., TCHIEU, A. A., ALEXEEV, D., DE BRAUER, A. & KOUMOUTSAKOS, P. 2016 Learning to school in the presence of hydrodynamic interactions. *Journal of Fluid Mechanics* **789**, 726–749.
- GOODFELLOW, IAN, BENGIO, YOSHUA & COURVILLE, AARON 2017 *The Deep Learning Book*. MIT Press.
- GUÉNIAT, FLORIMOND, MATHELIN, LIONEL & HUSSAINI, M. YOUSUFF 2016 A statistical learning strategy for closed-loop control of fluid flows. *Theoretical and Computational Fluid Dynamics* **30** (6), 497–510.
- GUSTAVSSON, K., BIFERALE, L., CELANI, A. & COLABRESE, S. 2017 Finding Efficient Swimming Strategies in a Three Dimensional Chaotic Flow by Reinforcement Learning. *arXiv e-prints* , arXiv: 1711.05826.
- HILL, ASHLEY, RAFFIN, ANTONIN, ERNESTUS, MAXIMILIAN, GLEAVE, ADAM, TRAORE, RENE, DHARIWAL, PRAFULLA, HESSE, CHRISTOPHER, KLIMOV, OLEG, NICHOL, ALEX, PLAPPERT, MATTHIAS, RADFORD, ALEC, SCHULMAN, JOHN, SIDOR, SZYMON & WU, YUHUI 2018 Stable Baselines. <https://github.com/hill-a/stable-baselines>.
- HOCHREITER, SEPP & SCHMIDHUBER, JÜRGEN 1997 Long short-term memory. *Neural Comput* **9**, "1735–1780".
- HOU TSANG, ALAN CHENG, TONG, PUN WAI, NALLAN, SHREYES & PAK, ON SHUN 2018 Self-learning how to swim at low Reynolds number. *arXiv* p. arXiv:1808.07639.
- HOUTHOOFT, REIN, CHEN, RICHARD Y., ISOLA, PHILLIP, STADIE, BRADLY C., WOLSKI, FILIP, HO, JONATHAN & ABBEEL, PIETER 2018 Evolved Policy Gradients. *arXiv e-prints* .
- HOWARD, R. A. 1960 *Dynamic programming and Markov processes*. Technology Press and Wiley, New York.
- KENDALL, ALEX, HAWKE, JEFFREY, JANZ, DAVID, MAZUR, PRZEMYSŁAW, REDA, DANIELE, ALLEN, JOHN-MARK, LAM, VINH-DIEU, BEWLEY, ALEX & SHAH, AMAR 2018 Learning to Drive in a Day. *arXiv e-prints* , arXiv: 1807.00412.
- KNIGHT, WILL 2018 Google just gave control over data center cooling to an AI. <https://www.technologyreview.com/s/611902/google-just-gave-control-over-data-center-cooling-to-an-ai/>.
- LEE, J & HAJELA, P 1996 Parallel genetic algorithm implementation in multidisciplinary rotor blade design. *Journal of Aircraft* **33**, 962–969.
- LEE, XIAN YEOW, BALU, ADITYA, STOECKLEIN, DANIEL, GANAPATHYSUBRAMANIAN, BASKAR & SARKAR, SOUMIK 2018 Flow Shape Design for Microfluidic Devices Using Deep Reinforcement Learning. *CoRR abs/1811.12444*, 1–10.
- LILLICRAP, TIMOTHY P., HUNT, JONATHAN J., PRITZEL, ALEXANDER, HEESS, NICOLAS, EREZ, TOM, TASSA, YUVAL, SILVER, DAVID & WIERSTRA, DAAN 2015 Continuous control with deep reinforcement learning. *arXiv e-prints* .
- LIN, LONG-JI 1993 Reinforcement Learning for Robots Using Neural Networks. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, uMI Order No. GAX93-22750.
- LIU, CHENG-YUAN, CHENG, WEI-CHEN, LIU, JIUN-WEI & LIU, DAW-RAN 2014 Autoencoder for words. *Neurocomputing* **139**, 84–96.
- MA, PINGCHUAN, TIAN, YUNSHENG, PAN, ZHERONG, REN, BO & MANOCHA, DINESH 2018 Fluid directed rigid body control using deep reinforcement learning. *ACM Transactions on Graphics* **37**, 1–11.
- MÄKINEN, R.A.E, PERIAUX, J & TOIVANEN, J 1999 Multidisciplinary shape optimization in aerodynamics and electromagnetics using genetic algorithms. *International Journal for Numerical Methods in Fluids* **30**, 149–159.
- MATOS, R, LAURSEN, T, VARGAS, J & BEJAN, A 2004 Three-dimensional optimization of staggered finned circular and elliptic tubes in forced convection. *International Journal of Thermal Sciences* **43**, 477–487.

- MELIGA, PHILIPPE, BOUJO, EDOUARD, PUJALS, GREGORY & GALLAIRE, FRANÇOIS 2014 Sensitivity of aerodynamic forces in laminar and turbulent flow past a square cylinder. *Physics of Fluids* **26**, 26,104101.
- MNIH, VOLODYMYR, KAVUKCUOGLU, KORAY, SILVER, DAVID, GRAVES, ALEX, ANTONOGLU, IOANNIS, WIERSTRA, DAAN & RIEDMILLER, MARTIN 2013 Playing Atari with Deep Reinforcement Learning. *CoRR* **abs/1312.5602**.
- MNIH, VOLODYMYR, PUIGDOMÈNECH BADIA, ADRIÀ, MIRZA, MEHDI, GRAVES, ALEX, LILICRAP, TIMOTHY P., HARLEY, TIM, SILVER, DAVID & KAVUKCUOGLU, KORAY 2016 Asynchronous Methods for Deep Reinforcement Learning. *arXiv e-prints*, arXiv: 1602.01783.
- MUYL, F, DUMAS, L & HERBERT, V 2004 Hybrid method for aerodynamic shape optimization in automotive industry. *Computers & Fluids* **33**, 849–858.
- NOVATI, GUIDO, VERMA, SIDDHARTHA, ALEXEEV, DMITRY, ROSSINELLI, DIEGO, VAN REES, WIM M. & KOUMOUTSAKOS, PETROS 2017 Synchronised swimming of two fish. *Bioinspiration and Biomimetics* **12** (3), 036001.
- OPENAI 2018 OpenAI Five. <https://blog.openai.com/openai-five/>.
- PINTO, LERREL, ANDRYCHOWICZ, MARCIN, WELINDER, PETER, ZAREMBA, WOJCIECH & ABBEEL, PIETER 2017 Asymmetric Actor Critic for Image-Based Robot Learning. *CoRR* **abs/1710.06542**.
- QIU, JINGRAN, ZHAO, LIHAO, XU, CHUNXIAO & YAO, YICHEN 2018 Swimming strategy of settling elongated micro-swimmers by reinforcement learning. *arXiv* p. arXiv:1811.10880.
- RABAULT, JEAN, KUCHTA, MIROSLAV, JENSEN, ATLE, RÉGLADE, ULYSSE & CERARDI, NICOLAS 2019 Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of Fluid Mechanics* **865**, 281–302.
- RABAULT, JEAN & KUHNLE, ALEXANDER 2019 Accelerating Deep Reinforcement Learning of Active Flow Control strategies through a multi-environment approach. *arXiv e-prints*, arXiv: 1906.10382.
- SCHÄFER, MICHAEL, TUREK, STEFAN, DURST, FRANZ, KRAUSE, EGON & RANACHER, ROLF 1996 Benchmark computations of laminar flow around a cylinder. In *Flow simulation with High-Performance Computers II*, pp. 547–566. Springer.
- SCHAUL, TOM, QUAN, JOHN, ANTONOGLU, IOANNIS & SILVER, DAVID 2015 Prioritized Experience Replay. *arXiv e-prints*, arXiv: 1511.05952.
- SCHULMAN, JOHN, LEVINE, SERGEY, MORITZ, PHILIPP, JORDAN, MICHAEL I. & ABBEEL, PIETER 2015 Trust Region Policy Optimization. *arXiv e-prints*, arXiv: 1502.05477.
- SCHULMAN, JOHN, MORITZ, PHILIPP, LEVINE, SERGEY, JORDAN, MICHAEL & ABBEEL, PIETER 2016 High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv e-prints* pp. 1–14.
- SIEGELMANN, H.T. & SONTAG, E.D. 1995 On the Computational Power of Neural Nets. *Journal of Computer and System Sciences* **50** (1), 132–150.
- SILVER, DAVID, SCHRITTWIESER, JULIAN, SIMONYAN, KAREN, ANTONOGLU, IOANNIS, HUANG, AJA, GUEZ, ARTHUR, HUBERT, THOMAS, BAKER, LUCAS, LAI, MATTHEW, BOLTON, ADRIAN, CHEN, YUTIAN, LILICRAP, TIMOTHY, HUI, FAN, SIFRE, LAURENT, VAN DEN DRIESSCHE, GEORGE, GRAEPEL, THORE & HASSABIS, DEMIS 2017 Mastering the game of Go without human knowledge. *Nature* **550**.
- STRANG, GILBERT 2019 *Linear Algebra and Learning from Data*. Wellesley-Cambridge Press.
- SUTTON, R.S 1988 Learning to Predict by the Method of Temporal Differences. *Machine Learning* **3** (1), 9–44.
- SUTTON, R. S. & BARTO, A. G. 1998 *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.

- SUTTON, R. S. & BARTO, A. G. 2018 *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- TESAURO, GERALD 1995 Temporal Difference Learning and TD-Gammon. *Communications of the ACM* **38**.
- TSITSIKLIS, JOHN & VAN ROY, BENJAMIN 1997 An analysis of temporal-difference learning with function approximation. *IEEE Transactions* p. 674–690.
- VAN HASSELT, HADO, GUEZ, ARTHUR & SILVER, DAVID 2015 Deep Reinforcement Learning with Double Q-learning. *arXiv e-prints* , arXiv: 1509.06461.
- VERMA, SIDDHARTHA, NOVATI, GUIDO & KOUMOUTSAKOS, PETROS 2018 Efficient collective swimming by harnessing vortices through deep reinforcement learning. *arXiv e-prints* , arXiv: 1802.02674.
- VINYALS, ORIOL, BABUSCHKIN, IGOR, CHUNG, JUNYOUNG, MATHIEU, MICHAEL, JADERBERG, MAX, CZARNECKI, WOJCIECH M., DUDZIK, ANDREW, HUANG, AJA, GEORGIEV, PETKO, POWELL, RICHARD, EWALDS, TIMO, HORGAN, DAN, KROISS, MANUEL, DANIHELKA, IVO, AGAPIOU, JOHN, OH, JUNHYUK, DALIBARD, VALENTIN, CHOI, DAVID, SIFRE, LAURENT, SULSKY, YURY, VEZHNEVETS, SASHA, MOLLOY, JAMES, CAI, TREVOR, BUDDEN, DAVID, PAINE, TOM, GULCEHRE, CAGLAR, WANG, ZIYU, PFAFF, TOBIAS, POHLEN, TOBY, WU, YUHUI, YOGATAMA, DANI, COHEN, JULIA, MCKINNEY, KATRINA, SMITH, OLIVER, SCHAUL, TOM, LILICRAP, TIMOTHY, APPS, CHRIS, KAVUKCUOGLU, KORAY, HASSABIS, DEMIS & SILVER, DAVID 2019 AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>.
- WILLIAMS, RONALD J. 1992 Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* **8** (3), 229–256.