

## Full Length Article



## A reinforcement learning strategy for p-adaptation in high order solvers

David Huergo<sup>a,\*</sup>, Gonzalo Rubio<sup>a,b</sup>, Esteban Ferrer<sup>a,b</sup><sup>a</sup> ETSIAE-UPM-School of Aeronautics, Universidad Politécnica de Madrid, Plaza Cardenal Cisneros 3, E-28040 Madrid, Spain<sup>b</sup> Center for Computational Simulation, Universidad Politécnica de Madrid, Campus de Montegancedo, Boadilla del Monte, 28660 Madrid, Spain

## ARTICLE INFO

## Keywords:

Reinforcement learning  
 Proximal policy optimization  
 High-order discontinuous Galerkin  
 p-adaptation  
 Mesh adaptation

## ABSTRACT

Reinforcement learning (RL) has emerged as a promising approach to automating decision processes. This paper explores the application of RL techniques to optimise the polynomial order in the computational mesh when using high-order solvers. Mesh adaptation plays a crucial role in improving the efficiency of numerical simulations by increasing accuracy while reducing the cost. Here, actor-critic RL models based on Proximal Policy Optimization offer an approach for agents to learn optimal mesh modifications based on evolving conditions. The paper provides a strategy for p-adaptation in high-order solvers and includes insights into the main aspects of RL-based mesh adaptation, including the formulation of appropriate reward structures and the interaction between the RL agent and the simulation environment. The proposed strategy does not require a high-fidelity solution during the training process and the formulation is general for any computational mesh and partial differential equation (PDE), solved in a discontinuous Galerkin solver. We discuss the impact of RL-based mesh p-adaptation on computational efficiency and accuracy. We apply the RL p-adaptation strategy to a one-dimensional inviscid Burgers' equation, focusing our analysis on smooth solutions of the equation to showcase its effectiveness. The RL strategy reduces the computational cost and improves accuracy over uniform adaptation, while minimising human intervention.

## 1. Introduction

In recent years, the field of machine learning has witnessed remarkable advancements, enabling machines to learn and adapt to complex tasks with increasing autonomy. Of particular interest are advances related to fluid mechanics and the related solution of partial differential equations [3,12,31,2,5]. Reinforcement learning (RL), a prominent subfield of machine learning, has emerged as a powerful framework to enable intelligent decision making in dynamic and uncertain environments. RL encompasses algorithms and methodologies that enable agents to learn optimal strategies by interacting with an environment and receiving feedback in the form of rewards or penalties.

Reinforcement learning has garnered significant attention across various domains, ranging from robotics and game-playing to finance and healthcare. Its success lies in its ability to handle sequential decision-making problems, where an agent learns to make optimal decisions by taking into account the current state of the environment, selecting appropriate actions, and observing the resulting rewards or consequences. This iterative process of learning from interactions al-

lows RL agents to navigate through complex environments, optimise their behaviours, and achieve predefined objectives.

In the field of computational fluid dynamics (CFD), which deals with the numerical simulation of fluid flow, RL techniques have shown promise for flow control (e.g., drag reduction), shape optimisation, and turbulence modelling [12,32]. RL is still in its infancy when combined with numerical techniques. In particular, only few publications have studied RL to control mesh adaptation in CFD [34,33,35,22,11] and concluded that RL policies have the potential to outperform adaptation based on error estimators. Let us note that all the previous works have considered low-order numerical methods. Even though there are some new RL approaches for high-order h-adaptation [10], to the authors' knowledge there is no RL strategy for high-order p-adaptation, which is the topic of this work. Mesh adaptation plays a crucial role in numerical methods, as it allows for the refinement or coarsening of the computational mesh based on the solution smoothness and the computational cost. Traditionally, mesh adaptation has relied on heuristics or manual intervention, which can be time consuming and limited in their ability to capture complex flow phenomena. Automatic

\* Corresponding author.

E-mail address: [david.huergo.perea@upm.es](mailto:david.huergo.perea@upm.es) (D. Huergo).

mesh adaptation algorithms require specific criteria to determine the regions that require adjustment. These criteria can be classified into three distinct categories, as discussed in [24]: feature-based indicators [25,26,21,13], local error-based indicators [28,6,1], and goal-oriented indicators [20,7,4]. Feature-based indicators are derived from the physical properties of the flow (e.g., norm of the gradient), and although these methods are relatively inexpensive and straightforward to implement, they often lack robustness and their effectiveness can vary depending on the specific case. The second category, local error-based indicators, identify regions that require refinement based on high values of quantified error. Implementing this approach involves obtaining solutions at different levels of mesh refinement, which can be computationally expensive until an appropriate numerical setup is achieved. Finally goal-oriented indicators evaluate the error contribution associated with a specific target functional, often relying on adjoint-based techniques. However, this approach is costly as it necessitates solving a new system of equations.

RL offers a novel approach to automate and optimise the mesh adaptation process by enabling the solver to learn the most suitable mesh modifications based on the time-evolving conditions. By integrating reinforcement learning algorithms into simulations, it is possible to train agents to adapt the computational mesh (here increasing/decreasing the polynomial order in a discontinuous Galerkin solver) in response to changes in the solution. The RL agent interacts with the numerical solver, observing the state of the flow, making decisions on mesh modifications (e.g., increasing the polynomial order in certain mesh elements), and receiving rewards based on the resulting accuracy and efficiency of the simulation. In fact, it is possible to decouple the training so that a more efficient decision process is incorporated into the numerical solver. This training considers high-order polynomials and is an essential part of the proposed RL strategy studied in this work. Through iterative learning, the RL agent improves its mesh adaptation strategies, ultimately leading to more accurate and efficient simulations.

The application of reinforcement learning in mesh adaptation brings several advantages. First, it reduces the reliance on manual intervention, allowing for automated and adaptive mesh refinement. Secondly, RL-based mesh adaptation can optimise the computational resources by dynamically allocating degrees of freedom (here high order polynomials) in regions of interest and reducing unnecessary refinement in areas with low flow variations. This resource optimisation leads to computational cost savings without compromising the accuracy of the solution. Moreover, reinforcement learning enables the exploration of unconventional mesh adaptations that might not be intuitively considered by traditional methods. RL agents have the capability to discover non-obvious mesh modifications that enhance the accuracy and robustness of the CFD simulations.

In this work, we design a strategy for polynomial adaptation in high-order methods using RL. Without loss of generality we restrict the methodology to 1D cases and focus on the design of appropriate rewards and state definitions that lead to efficient RL-adaptation. This approach could also be applied to 3D cases within the nodal high-order discontinuous Galerkin spectral element method (DGSEM) framework, where the 3D solution is obtained by the tensor product of 1D solutions inside each element. First, we describe the methodology in section 2, including the Proximal Policy Optimisation algorithm, the numerical solution of the PDE and the RL strategy (detailing agent, states, rewards, training and testing). Second, we provide results in section 3 and finalise with conclusions in section 4.

## 2. Methodology

### 2.1. Mathematical model and numerical discretisation

We use a one-dimensional nonlinear partial differential equation to test our strategy for mesh adaptation. In particular, we use the 1D inviscid Burgers' equation:

$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial u^2}{\partial x} = 0. \quad (1)$$

We solve this equation in the domain  $x \in [0, 1]$ , with periodic boundary conditions and the initial condition:

$$u(x, 0) = 2 + \sin(2\pi x). \quad (2)$$

The analytical solution to the previous equation can be obtained through the method of characteristics [19], leading to:

$$u(x, t) = h(x - ut), \quad (3)$$

where  $h$  is a function that can be computed from the initial condition. In this case, we use the condition defined in Eq. (2), the final solution is:

$$u(x, t) = 2 + \sin(2\pi(x - ut)). \quad (4)$$

Equation (4) will be used to measure the performance of our methodology in Section 3. The solutions of Eq. (4) for every pair  $(x, t)$  are obtained using a Newton-Raphson method.

We discretise the equation using a nodal high-order discontinuous Galerkin spectral element method (DGSEM), see details in Appendix A. The method allows for the selection of an arbitrary polynomial order,  $p$ , in each mesh element, which can be different from the ones selected in the neighbouring elements. This is a key advantage of DG since it combines flexibility and accuracy. The discontinuous nature of the method allows for local refinement and adaptivity, enabling precise resolution of features with high gradients or regions of interest. This adaptivity is particularly beneficial when solving PDEs in complex geometries or when dealing with problems in which the solution undergoes abrupt changes.

### 2.2. Reinforcement learning strategy for $p$ -adaptation

Reinforcement learning is generally considered a semi-supervised approach, as the agent learns by itself the optimal policy through the interaction with an environment. However, the user must define an objective function to reward the agent when it shows a positive behaviour. Every RL method is based on the scheme shown in Fig. 1.

This framework highlights several points that must be defined for the correct understanding of the proposed methodology:

1. **Agent:** The model which decides the best action to take, in our case the best polynomial order,  $p$ , in each mesh element. It is made up of two neural networks: the actor and the critic. See more details in section 2.2.1.
2. **Environment:** Each element and polynomial of the computational mesh and the DG scheme. It allows to compute a solution in one element given the polynomial order,  $p$ .
3. **Action:** It is an output for the agent and an input for the environment. Three different actions will be considered: to increase the polynomial order one unit, to decrease the polynomial order one unit, or to keep the polynomial order constant.
4. **State:** It is an output for the environment and an input for the agent. It must contain enough information to define the current scenario. See more details in section 2.2.2.
5. **Reward:** The user-defined objective function. It should have a higher value when the agent is performing better. See more details in section 2.2.3.

While the definitions of the environment and the actions are straightforward, the agent, the state, and the reward have to be cautiously designed to obtain good performance.

#### 2.2.1. Agent definition

The reinforcement learning model has been trained using the Proximal Policy Optimisation (PPO) algorithm, which allows the agent to




Fig. 1. Schematic definition of the RL framework.

learn the optimum policy that maximises the expected cumulative reward for the current state. PPO is a reinforcement learning algorithm that has gained significant attention for its effectiveness and stability in training complex decision-making agents [29]. PPO belongs to the class of policy optimisation algorithms and is designed to strike a balance between exploration and exploitation, enabling agents to learn robust policies. At its core, PPO aims to maximise the expected cumulative rewards by iteratively updating the policy of an RL agent. The key idea behind PPO is to define a surrogate objective function that measures the policy improvement. This surrogate objective is optimised through multiple iterations, with each iteration consisting of two steps: sampling trajectories and updating the policy. In the trajectory sampling step, the agent interacts with the environment to collect a batch of experiences. These experiences are then used to estimate the advantages of different actions based on the rewards received and the estimated value function. The policy update step in PPO involves computing the surrogate objective, which quantifies the improvement in the policy. PPO introduces a clip function that limits the policy update to a specified range, preventing large deviations from the previous policy. By constraining the policy updates, PPO ensures stability during training and mitigates the risk of policy divergence.

The methodology is based on the actor-critic approach, in which two neural networks interact with each other. In PPO, the actor is responsible for generating actions based on the current policy, while the critic estimates the value function and provides feedback on the quality of the chosen actions. By combining both models, the actor-critic approach enables the model to learn in an autonomous way from the critic's feedback to update the current policy, leading to a better decision-making actor. Overall, the PPO algorithm provides a robust and stable framework for training reinforcement learning agents. Its conservative policy updates, combined with the use of surrogate objectives and entropy regularisation, enable efficient exploration and exploitation, leading to the discovery of optimal policies in complex environments, such as polynomial adaptation in high order methods.

The PPO agent has been trained through the Tensorforce Python library [30], which provides a wide set of reinforcement learning algorithms to train the agent. Furthermore, it offers a flexible framework and the user has a complete control over the definitions of the environment, the state and the reward, as well as the value of the different hyperparameters involved in the process. Some of the most important hyperparameters are the learning rate,  $lr = 10^{-3}$ , the exploration factor,  $\epsilon = 10^{-2}$ , the discount factor,  $\gamma = 0.99$ , and the likelihood ratio clipping,  $rc = 0.1$ .

The actor-critic approach used to train the agent requires two neural networks, both with the same architecture for this case. In addition to the input layer, with the shape of the state vector (see section 2.2.2), and the output layer, which provides a single output, each network has two hidden layers with 64 neurons each and an activation function of tanh. The output of the actor network discriminates among the three

possible actions through a linear function. Furthermore, the Adam optimiser [14] has been used to train the networks.

### 2.2.2. State definition

We call state the set of variables that provides enough information to characterise different situations. Based on this state, the agent can learn the best course of action for each scenario. In this case, we have defined a state based on two variables: the current polynomial degree,  $p$ , and an error,  $e$ . The polynomial degree will range as  $p \in [2, 10]$ , which is reasonable for solving PDEs with DGSEM. The error is a key point of the RL strategy and must provide an estimate of the accuracy that can be achieved when selecting a polynomial degree  $p$ .

When using a nodal high order approach (e.g., DGSEM in our case), we pose the equation at  $p + 1$  nodes (in our case Gauss nodes) in each element, given a polynomial of order  $p$ . Having computed the numerical value at the nodes, it is possible to compute/reconstruct the continuous solution inside each element using Lagrange polynomials. Then, the polynomial can be used to interpolate the solution for the nodes associated with a lower degree  $p - 1$ . The low-degree solution is then compared with the previous solution to compute  $N_s = 3$  errors,  $e_i$  for  $i = 1, 2$  and  $3$ , which are obtained by subtracting the low-order solution from the high-order solution in the middle and at the boundaries of the element, as shown in Fig. 2. The value of  $N_s = 3$  is selected, independently of the current polynomial order,  $p$ , because it is fast and considers the boundary values of the elements. This is interesting because the inter-element jumps are related to the error in the solution in DG methods [9,18]. However, another value could be used for  $N_s$ , specially when employing a solver other than the discontinuous Galerkin approach to address the problem. If the error between two consecutive polynomials is small, then the shape of the true solution is accurately captured, and hence that polynomial order is high enough to obtain a precise numerical approximation. Finally, these three errors are combined to obtain a single scalar value  $mse$ :

$$mse = \frac{\sum_{i=1}^{N_s} e_i^2}{N_s}, \quad (5)$$

$$e = -\log_{10}(\min(mse + 10^{-10}, 1)). \quad (6)$$

The error,  $e$ , has been mapped using the logarithm to separate values of different orders of magnitude. The image of the error is  $e \in [0, 10]$ . In this way, the values of  $e$  and  $p$ , have similar ranges, which is advantageous for the training process. We favour this estimation of the error because of its efficiency (only 3 points need to be evaluated) and its locality (all information is contained in the element), but of course other error estimations could be used (e.g., [27,28,20]). As the error,  $e$ , is a continuous variable, an infinite number of states ( $p, e$ ) are possible; and hence, the selected RL model must be capable of handling continuous states, as is the case for the PPO model which is used in this work.




Fig. 2. Error calculation between two consecutive polynomials of order  $p$  and  $p - 1$ .

### 2.2.3. Reward definition

We define the reward as the objective function that the agent should learn to decide the best course of action based on the current state. This function should have a higher value as the error decreases. In addition, it should have a higher value as the polynomial order decreases, because the objective is to use the minimum polynomial order possible to reduce the computational time while preserving the accuracy.

The error used to calculate the reward should not be the same as that used for the state, as the reward is only computed during the training phase and it must provide a measurement of the real accuracy of the model. In this case, we require the error between our polynomial approximation and the analytical solution, which is a reliable value for evaluating the performance of the model. Therefore, the analytical solution (or an estimation) must be known during the training, as will be explained later in section 2.2.4. We use the root mean squared error,  $rmse$ , as the difference between the analytical and approximate ( $p - 1$ ) solutions sampled at  $N_r$  Gauss points. In this work, we use a constant value of  $N_r = 2p_{max} + 1$  which is enough to accurately approximate Eq. (8). In our proposed strategy, we calculate an approximate solution using the polynomial order  $p - 1$  instead of  $p$ . By doing so, we reward the accuracy of the solution  $p - 1$ , which is the lowest order required to compute the state (see the previous section 2.2.2), leading to a more robust method. We finally define the reward function as a smooth Gaussian-like distribution, which provides a trade-off between cost (low polynomial order) and accuracy (low error):

$$reward = \underbrace{\frac{p_{max}^2 + 1}{p^2 + 1}}_{\text{COST}} \underbrace{\exp\left(-\frac{rmse^2}{2\sigma^2}\right)}_{\text{ERROR}}, \quad (7)$$

$$rmse = \sqrt{\frac{\sum_{i=1}^{N_r} (y_i - y_i^*)^2}{N_r}}, \quad (8)$$

with  $p_{max} = 10$  as the maximum polynomial order available,  $\sigma$  as the standard deviation,  $y$  as the analytical solution and  $y^*$  as the polynomial approximation. The main contribution of the reward is the exponential part (the error), which will drop the reward if the error is big compared to the standard deviation. The second contribution is the quadratic part  $\frac{p_{max}^2 + 1}{p^2 + 1}$ , it relates to the computational cost, and only becomes relevant if the error  $rmse$  is small enough. The cost part increments the reward if the polynomial order is small. Consequently, the reward enforces the error to be small, while keeping the polynomial degree as small as possible (but without increasing the error significantly). We illustrate the behaviour of the reward as a function of  $rmse$ ,  $p$  and  $\sigma$  in Figs. 3 and 4. Fig. 3a shows the exponential decay of the reward (the error part) with the  $rmse$  and Fig. 3b shows the quadratic part (the cost) of the reward with the polynomial order. The value of  $\sigma$  is important, as it provides an estimation of the value of  $rmse$  that must be achieved before the polynomial order can be optimised. If the value of  $\sigma$  is small, the agent will

try to significantly reduce the  $rmse$ , which will lead to an adaptation with high polynomials, and vice versa. The effect of  $\sigma$  on the reward is shown in Fig. 4. The final error  $rmse$  (after the adaptation process) will have an approximate value defined by the region where the reward is increasing with its maximum slope (in Fig. 4). This way, an agent can be trained to achieve a specific threshold error, while minimising the cost. Based on preliminary tests (not shown), we select  $\sigma = 0.01$ , which provides an appropriate trade-off between accuracy and computational cost of the solution. This is what defines the efficiency and accuracy of the  $p$ -adaptation strategy.

### 2.2.4. Training

During the training process, the agent interacts with the environment and learns which action provides the highest reward for each state. In our training, the agent learns to modify the polynomial order of one element of the mesh. Within the DGSEM framework, every element of the mesh is mapped into the computational space where  $x \rightarrow \xi \in [-1, 1]$ , see Appendix A for details. We can simplify and optimise the training process if the agent is trained to optimise the polynomial order in computational space. This approach allows to use a common agent for every element of the mesh.

One of the main problems of using RL for numerical applications is that the training is very time consuming. Furthermore, the reward, Eq. (7), is based on an error that is computed in relation to the analytical solution, which is not known for most partial differential equations (e.g., Navier-Stokes equations). To tackle both problems, in this work we propose to decouple the training process of the RL agent from the simulation. The training is performed using a random function generator, which provides random sinusoidal-type functions in one element. This random function generator can create functions,  $g : [-1, 1] \rightarrow \mathbb{R}$ , with the following structure:

$$g(\xi) = \frac{1 + a \cdot \sin(2\pi f (\xi - c))}{2} \quad (9)$$

where  $a = \{-1, 1\}$ ,  $c \in [-1, 1]$  are random parameters and  $f \in [0, 1]$  is a random frequency. More complex functions can be created to capture specific behaviours, but this simple option shows promising results for smooth solutions. Let us note that when correlating polynomials and sinusoidal functions, we are, in fact, performing a non-linear von Neumann (or eigensolution analysis) type analysis [23,17,16] and consequently we are finding the optimal polynomials that capture wavelike solutions.

The training process is divided in episodes, each one formed by the following steps:

1. A random function is generated using equation (9) and the initial polynomial order  $p \in [2, 10]$  is chosen randomly.
2. The initial state is computed.
3. The agent is called to select an action based on the current state.
4. The action is performed, and a new state and a reward are generated.
5. The agent learns whether the chosen action was good or not based on the reward.
6. Steps 3) to 5) are repeated  $N$  times.

The value of  $N$  has to be big enough to allow the agent to reach the optimum polynomial order from any state. In the worst case scenario, the initial polynomial order is the minimum,  $p = 2$ , and the optimum polynomial is the maximum,  $p_{opt} = 10$ . In this case, the agent needs at least  $N = 8$  steps to reach the optimum. However, at the beginning of the training process the performance of the agent is low and it is very unlikely to reach the optimum in 8 steps. Furthermore, if  $N$  is too big the computational cost of the training is significantly increased. Therefore, we have decided to use  $N = 20$  as a trade-off, allowing the agent to reach the optimum easily during the early phase of the training, while keeping a low computational cost. The whole process is repeated  $N_{ep} = 25000$  episodes until the model is trained.




Fig. 3. Reward distribution as a function of a) the error  $rmse$  and b) the polynomial order  $p$ , both with a value of  $\sigma = 0.01$ .

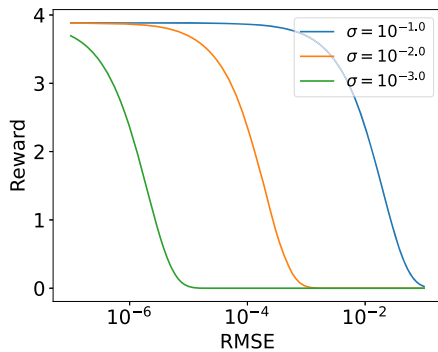


Fig. 4. Reward distribution as a function of the error  $rmse$  for three different values of  $\sigma$  and a constant polynomial order  $p = 5$ .

Let us clarify the methodology using an example. First, a initial polynomial order,  $p$ , is selected randomly ( $p = 6$  in this example) and Eq. (9) is used to generate a random analytical sinusoidal function. This analytical solution is sampled at the Gauss nodes corresponding to the selected polynomial, as represented in Fig. 5a. Then, the value at the nodes is used to construct a Lagrange interpolating polynomial, which approximates the solution. Knowing the approximate sampled solution (for  $p = 6$  in this example), the initial state can be calculated, following section 2.2.2.

The initial state is the input for our agent that provides the course of action. As explained above, the agent can choose among three possible actions: to increase the polynomial order by one ( $p = 7$ ), to decrease the polynomial order by one ( $p = 5$ ) or to keep the current polynomial order ( $p = 6$ ). For this example, let us consider that the action decided by the agent is to increase the polynomial order; that is, to change the current  $p = 6$  to  $p_{new} = 7$ .

Given this new polynomial, the value of the solution at the Gauss nodes is calculated from the analytical solution (Eq. (9)), which is the same function as before. The new sampled Gauss nodes are represented in Fig. 5b. These new values are then used to generate a new Lagrange polynomial to approximate the solution. For this condition, a new state (see section 2.2.2) and a reward (see section 2.2.3) are computed.

Finally, the agent stores the initial state, the action chosen and the reward to update its policy at the end of the episode. This update is the core of the training process, and it allows the agent to learn the best course of action for a given state.

The previous steps are repeated several times (we have selected  $N = 20$  times per episode). After  $N$  iterations, the episode finishes, the agent updates the internal weights of the networks, and a new episode begins with a new random function as the analytical solution.

Table 1

Accuracy of the rewards in relation to the maximum possible value and mean error of the polynomial order for the RL agent compared with the optimal polynomial based on the reward function. Each test set is formed by 100 simulations of random functions.

| Test set | Reward accuracy | Mean error $ p - p_{opt} $ |
|----------|-----------------|----------------------------|
| 1        | 0.875           | 0.48                       |
| 2        | 0.864           | 0.59                       |
| 3        | 0.875           | 0.56                       |
| 4        | 0.862           | 0.53                       |
| 5        | 0.889           | 0.49                       |
| 6        | 0.874           | 0.54                       |
| 7        | 0.880           | 0.51                       |
| 8        | 0.869           | 0.55                       |
| 9        | 0.848           | 0.59                       |
| 10       | 0.864           | 0.49                       |

### 2.2.5. Testing

Once the model has been trained, the agent can be tested by generating a new set of random functions and comparing the maximum reward obtained for each case using RL with the expected optimal reward. Each test set is formed by 100 simulations of random functions, and the accuracy of the rewards and the mean error of the polynomial order for the RL agent are computed. The results are summarised in Table 1. In each test set, the accuracy obtained by the model is close to 90%; that is, the mean reward obtained in the simulations is 90% of the maximum possible value, which represents very good performance. The mean error between the optimal polynomial order and the one selected by the agent is close to 0.5, which means that most of the simulations (50% at least) have been perfectly performed by the agent. This result is represented in Fig. 6, which shows the difference between the optimum polynomial order (based on the reward function) and the one selected by the RL agent in 100 simulations. It can be seen that the agent can select the best polynomial order the 60% of the time. Of course, the agent sometimes selects one degree less than or more from the optimum. In these cases, when the optimum is not achieved, the agent tends to overestimate the polynomial order; that is, it chooses a higher polynomial degree, which leads to a higher accuracy, at the cost of a higher computational cost. In general, the trained model is accurate and capable of performing a competitive  $p$ -adaptation over a wide set of functions, and it can be used to improve the performance of numerical simulations.

## 3. Results

To demonstrate the effectiveness of the proposed methodology, the trained RL agent is used to improve the accuracy of a Burgers' equation, Eq. (1). In a separate phase, the PPO agent has learnt how to increase or decrease the polynomial order of each element of a given mesh to

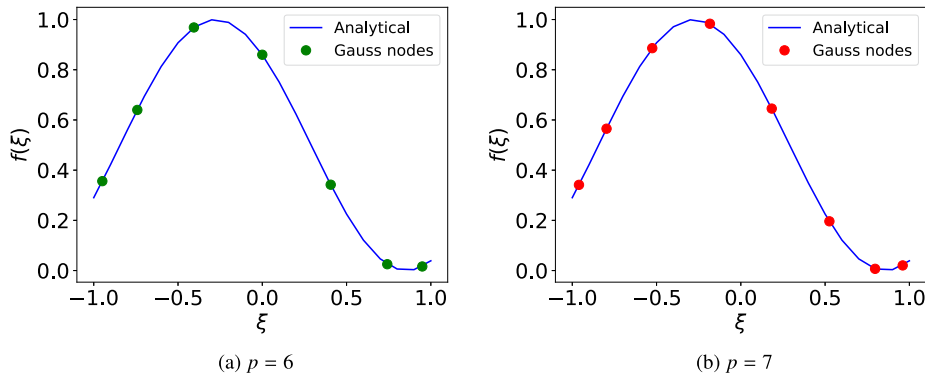


Fig. 5. Random function generated during the training and the Gauss nodes a) for the initial value of  $p = 6$  and b) for its value after the action is executed  $p_{new} = 7$ .




Fig. 6. Difference between the optimum polynomial order (based on the reward function) and the one selected by the RL agent in 100 simulations.

obtain an optimum performance; that is, a very low error at the lowest computational cost.

The computational mesh is made up of 8 elements and the initial polynomial degree is  $p = 4$  for every element, as shown in Fig. 7, where the initial solution is shown at Gauss points. A time step of  $\Delta t = 10^{-5}$  s is used to perform the temporal integration.

The solution of the problem without p-adaptation (uniform  $p = 4$ ) for  $t \in [0, 0.12]$  s is represented in Fig. 8, where we have used colours to represent the different elements. Note that we stop the simulation before the shock forms, since we are interested in the smooth phase of this problem. Two main regions can be highlighted in the solution at the final time. On the one hand, we observe a left region, where the solution is almost linear, and the right region, where a strong gradient is present. The left region could be captured with a polynomial order  $p = 2$ , because the velocity changes progressively and its gradient is almost constant. On the other hand, the right region requires a higher polynomial degree to capture the abrupt change of the solution. Therefore, the use of a constant polynomial order  $p = 4$  is far from ideal, as some elements use a polynomial order higher than required, which increases the computational cost, while other elements use a polynomial order lower than required, which deteriorates the precision of the solution. Both problems can be solved using the proposed RL strategy.

We now simulate the same problem using the RL p-adaptation strategy. The solution for  $t \in [0, 0.12]$  s is represented in Fig. 9. The adaptation process is performed each 0.01 s of simulation time, that is each 1000 iterations, for every element in the mesh. To improve the adaptation process, the solution within each element has been normalised between 0 and 1 before the computation of the state. In this way, the input will be more similar to the states used during the training.

In this case, the RL model is able to dynamically modify the polynomial order, reducing the order when possible and increasing the order near regions with strong gradients. When the agent decides to modify the polynomial order inside one element, the interpolated solution is




Fig. 7. Initial condition to solve the 1D inviscid Burgers' equation with periodic BC. Different colours are used to represent different elements of the computational mesh. Dots are used to represent the Gauss nodes within each element.

used to continue the simulation. The final solution is more accurate and less time consuming than the previous solution without p-adaptation (computational costs are compared in section 3.1).


We perform two additional simulations where we change the number of elements in the mesh (to 4 and 16 elements), and perform RL adaptation. Results are shown in Fig. 10. We observe that since the elements are transformed into the computational domain, where  $\xi \in [-1, 1]$ , the model can perform the adaptation for an arbitrary size of the elements. If the number of elements is increased, the local gradients within each element have smaller fluctuations, and the RL agent selects a lower polynomial order. As expected, if a very coarse mesh is used, the agent selects a very high polynomial order to capture the solution, and if the mesh is fine, the agent decreases the polynomial order to reduce the cost of the simulation, while providing an accurate result.

Finally, for a medium mesh of 8 elements, we use the analytical solution (Eq. (4)) to compute the point-wise errors across the domain, for the RL p-adaptation and the uniform  $p = 4$  simulations. These errors are calculated by subtracting the analytical solution from the numerical approximation at the Gauss nodes, where the solution is known. The results are shown in Fig. 11. We see that the error is significantly smaller when using RL adaptation than when using a constant polynomial order  $p = 4$ . The solution not only shows improved accuracy, but also presents a more cost-effective approach, as evidenced by the data provided in Table 3, which can be found in the next section.


In summary, we have shown that the RL adaptation is able to select the correct polynomial order in each element in the mesh. In the next section, we analyse the computational cost associated to the RL strategy.

### 3.1. Trade-off between accuracy and computational time


Previous results have shown that the RL p-adaptation algorithm is effective and capable of performing dynamic mesh adaptation, increas-




**Fig. 8.** 1D inviscid Burgers' equation solved with a DGSEM spatial scheme and an explicit Euler temporal scheme, using a constant polynomial order,  $p = 4$ , in each element of the computational mesh. The solution is shown in four different timestamps. Different colours are used to represent different elements of the computational mesh. Dots are used to represent the Gauss nodes within each element.



**Fig. 9.** 1D inviscid Burgers' equation solved with a DGSEM spatial scheme and an explicit Euler temporal scheme, using a RL p-adaptation algorithm. The polynomial of each element of the computational mesh has been adapted once every 0.01 s of simulation time, starting with a polynomial order  $p = 4$  in every element. The solution is shown in four different timestamps.



**Fig. 10.** 1D inviscid Burgers' equation solved with a DGSEM spatial scheme and an explicit Euler temporal scheme, using a RL p-adaptation algorithm. The polynomial of each element of the computational mesh has been adapted once every 0.01 s of simulation time, starting with a polynomial order  $p = 4$  in every element. The solution is shown at the timestamp  $t = 0.12$  s for three different meshes.



**Fig. 11.** Point-wise error between the analytical solution and the numerical simulation with and without p-adaptation at the simulation time  $t = 0.12$  s. The computational mesh has 8 elements. These errors are calculated by subtracting the analytical solution,  $y$ , from the numerical approximation,  $y^*$ , at the Gauss nodes.

ing the accuracy in problematic regions while reducing the polynomials (and the cost) in very smooth regions.

However, there is no universal rule that establishes the best polynomial order for each element. This decision lies on the users, who should be able to adapt the algorithm to their own needs. In some cases, a better accuracy could be desired to capture very small fluctuations. In other cases, if the resources are more limited, it could be desirable to obtain a solution accurate enough, but prioritising the reduction of the computational time.

To have a trade-off between these two key points, a parameter,  $\delta$ , is introduced. This parameter is independent of the training of the RL agent and is included afterward to offer the user additional control.

The control parameter,  $\delta$ , is used to modify the state and make the agent believe that the error is higher or lower than the real value. Given a possible scenario, it is possible to compute the state  $[p, e]$  (see Section 2.2.2), which is the input for the trained agent and, hence, it completely defines the best action to take. However, this control parameter modifies the state by multiplying the error,  $e$ , resulting in a


modified state  $[p, \delta e]$ . This way, the agent believes that the accuracy is higher (if  $0 < \delta < 1$ ) or smaller (if  $\delta > 1$ ) than the real value. Therefore, the agent adapts the output to this new state to obtain a different set of solutions depending on the value of  $\delta$ .

We can discriminate among three different cases:

1. If  $\delta = 1$ , the model shows the standard behaviour.
2. If  $0 < \delta < 1$ , then the model tries to reduce the computational time by selecting a lower polynomial order during the adaptation process.
3. If  $\delta > 1$ , then the model tries to improve the accuracy by selecting a higher polynomial order during the adaptation process.

We simulate again using the RL adaptation and show the solution at the time  $t = 0.12$  s in Fig. 12, where three different values of  $\delta$  are reported. In this case, it is clear that the general polynomial order has its value increased when  $\delta$  has a higher value.

To check the effect of the p-adaptation algorithm and the value of  $\delta$  on the computational time and the accuracy of the solution, several simulations have been carried out and the results are summarised in Tables 3, 5 and 6. Table 3 shows that, when using 8 and 16 elements, the RL adaptation provides clear accelerations when compared to a uniform polynomial with  $p = 4$ , for all values of  $\delta$ . Even for high values of  $\delta$  (e.g.  $\delta = 3.0$ ), the computational time using p-adaptation is similar or smaller than the time required to compute the simulation with a constant polynomial order. For completeness, Table 5 shows a comparison of the  $rmse$  of the approximate solution (in relation to the analytical solution provided in Eq. (4)) with and without p-adaptation, where we see that improved accuracy is obtained in most cases when RL adaptation is used. We perform additional simulations without p-adaptation for different polynomial orders, and include the errors and computational times in Tables 2 and 4. These results have been compared with those obtained using p-adaptation, as shown in Fig. 13, where it is clear that the RL p-adaptation approach improves the accuracy and reduces the computational cost in all cases. The  $rmse$  shown in this figure is a size-normalized version of the  $L_2$  norm:



**Fig. 12.** 1D inviscid Burgers' equation solved with a DGSEM spatial scheme and an explicit Euler temporal scheme, using a RL p-adaptation algorithm. The polynomial of each element of the computational mesh has been adapted once every 0.01 s of simulation time, starting with a polynomial order  $p = 4$  in every element. The solution is shown at the timestamp  $t = 0.12$  s for three different values of the control parameter  $\delta$ .

**Table 2**

Comparison between the computational time without p-adaptation for three meshes and four values of the polynomial order  $p$ .

| Elements | No adaptation $p = 3$ | No adaptation $p = 4$ | No adaptation $p = 5$ | No adaptation $p = 6$ |
|----------|-----------------------|-----------------------|-----------------------|-----------------------|
| 4        | 1.89 s                | 2.36 s                | 3.11 s                | 3.60 s                |
| 8        | 3.70 s                | 4.76 s                | 5.81 s                | 6.93 s                |
| 16       | 7.17 s                | 9.24 s                | 11.33 s               | 14.27 s               |

**Table 3**

Comparison between the computational time with and without p-adaptation for three meshes and three values of the control parameter  $\delta$ . The time required to adapt the mesh for each 0.01 s is included in these values. The average polynomial order,  $p_{av}$ , represents the mean of the polynomials among the elements of the computational mesh.

| Elements | No adaptation $p = 4$ | Adaptation $\delta = 0.3$  | Adaptation $\delta = 1.0$  | Adaptation $\delta = 3.0$  |
|----------|-----------------------|----------------------------|----------------------------|----------------------------|
| 4        | 2.36 s                | 2.15 s ( $p_{av} = 4.75$ ) | 2.76 s ( $p_{av} = 5.75$ ) | 2.91 s ( $p_{av} = 6.0$ )  |
| 8        | 4.76 s                | 3.56 s ( $p_{av} = 3.38$ ) | 3.79 s ( $p_{av} = 3.63$ ) | 4.46 s ( $p_{av} = 4.0$ )  |
| 16       | 9.24 s                | 6.25 s ( $p_{av} = 2.63$ ) | 6.58 s ( $p_{av} = 2.88$ ) | 7.18 s ( $p_{av} = 3.13$ ) |

$$rmse = \sqrt{\frac{\sum_{i=1}^N (y_i - y_i^*)^2}{N}} = \frac{L_2}{\sqrt{N}}, \tag{10}$$

where  $N$  is the number of DOFs (Gauss nodes) of the mesh,  $y$  is the analytical solution and  $y^*$  is the polynomial approximation. The use of the  $rmse$  instead of the  $L_2$  norm allows to perform a fair comparison among different meshes with variable polynomial order.

Finally and for completeness, Table 6 shows a measurement of the trade-off between accuracy and efficiency. This quantity, which has a smaller value if the performance is better, has been obtained by multiplying the computational time (see Table 3) by the error (see Table 5). In the table, we have highlighted in bold the best results for each mesh. Based on these results, the value of the control parameter,  $\delta$ , should be increased if the mesh is refined to obtain accurate and fast solutions. In general, the proposed method provides a good trade-off between efficiency and accuracy, and both are improved in most cases.

The p-adaptation strategy shows good performance and appears competitive for speeding up numerical simulations. It allows to obtain a significant reduction of the computational cost while maintaining a high level of accuracy.

#### 4. Conclusions

In conclusion, reinforcement learning (RL) is a powerful approach to automate and optimise mesh adaptation in the field of high-order numerical simulations. By integrating RL algorithms into high-order simulations, we have demonstrated the ability of RL agents to learn and adapt the computational mesh based on evolving solutions.

The application of RL in mesh p-adaptation brings numerous benefits. First, it reduces the reliance on manual intervention, enabling automated and adaptive mesh refinement. RL agents can discover non-obvious mesh modifications that enhance the accuracy while reducing cost. The agent's ability to learn optimal mesh modifications based on

**Table 4**

Comparison between the *rmse* of the approximated solution in relation to the analytical solution without p-adaptation for three meshes and four values of the polynomial order *p*.

| Elements | No adaptation $p = 3$ | No adaptation $p = 4$ | No adaptation $p = 5$ | No adaptation $p = 6$ |
|----------|-----------------------|-----------------------|-----------------------|-----------------------|
| 4        | $1.74 \cdot 10^{-2}$  | $1.47 \cdot 10^{-2}$  | $1.12 \cdot 10^{-2}$  | $7.22 \cdot 10^{-3}$  |
| 8        | $9.33 \cdot 10^{-3}$  | $5.15 \cdot 10^{-3}$  | $1.99 \cdot 10^{-3}$  | $6.59 \cdot 10^{-4}$  |
| 16       | $1.98 \cdot 10^{-3}$  | $3.80 \cdot 10^{-4}$  | $2.98 \cdot 10^{-4}$  | $2.25 \cdot 10^{-4}$  |

**Table 5**


Comparison between the *rmse* of the approximated solution in relation to the analytical solution with and without p-adaptation for three meshes and three values of the control parameter  $\delta$ . The average polynomial order,  $p_{av}$ , represents the mean of the polynomials among the elements of the computational mesh.

| Elements | No adaptation $p = 4$ | Adaptation $\delta = 0.3$                | Adaptation $\delta = 1.0$                | Adaptation $\delta = 3.0$                |
|----------|-----------------------|--|--|--|
| 4        | $1.47 \cdot 10^{-2}$  | $9.73 \cdot 10^{-3}$ ( $p_{av} = 4.75$ ) | $2.11 \cdot 10^{-3}$ ( $p_{av} = 5.75$ ) | $2.02 \cdot 10^{-3}$ ( $p_{av} = 6.0$ )  |
| 8        | $5.15 \cdot 10^{-3}$  | $4.13 \cdot 10^{-3}$ ( $p_{av} = 3.38$ ) | $1.39 \cdot 10^{-3}$ ( $p_{av} = 3.63$ ) | $5.72 \cdot 10^{-4}$ ( $p_{av} = 4.0$ )  |
| 16       | $3.80 \cdot 10^{-4}$  | $9.79 \cdot 10^{-4}$ ( $p_{av} = 2.63$ ) | $4.86 \cdot 10^{-4}$ ( $p_{av} = 2.88$ ) | $4.03 \cdot 10^{-4}$ ( $p_{av} = 3.13$ ) |

**Table 6**

Measurement of the trade-off between efficiency (the computational time) and accuracy (the error *rmse*), by multiplying both quantities, with and without p-adaptation for three meshes and three values of the control parameter  $\delta$ . The best result (the smallest) for each mesh is highlighted in bold. The average polynomial order,  $p_{av}$ , represents the mean of the polynomials among the elements of the computational mesh.

| Elements | No adaptation $p = 4$ | Adaptation $\delta = 0.3$                | Adaptation $\delta = 1.0$                                  | Adaptation $\delta = 3.0$                                  |
|----------|-----------------------|--|--|--|
| 4        | $3.47 \cdot 10^{-2}$  | $2.09 \cdot 10^{-2}$ ( $p_{av} = 4.75$ ) | <b><math>5.82 \cdot 10^{-3}</math></b> ( $p_{av} = 5.75$ ) | $5.88 \cdot 10^{-3}$ ( $p_{av} = 6.0$ )                    |
| 8        | $2.45 \cdot 10^{-2}$  | $1.47 \cdot 10^{-2}$ ( $p_{av} = 3.38$ ) | $5.27 \cdot 10^{-3}$ ( $p_{av} = 3.63$ )                   | <b><math>2.55 \cdot 10^{-3}</math></b> ( $p_{av} = 4.0$ )  |
| 16       | $3.51 \cdot 10^{-3}$  | $6.12 \cdot 10^{-3}$ ( $p_{av} = 2.63$ ) | $3.20 \cdot 10^{-3}$ ( $p_{av} = 2.88$ )                   | <b><math>2.89 \cdot 10^{-3}</math></b> ( $p_{av} = 3.13$ ) |



**Fig. 13.** Error *rmse* and CPU time convergence with and without p-adaptation for three different meshes, as a function of the polynomial order *p*. When p-adaptation is used, the average polynomial order,  $p_{av}$ , is used, which represents the mean of the polynomials among the elements of the computational mesh.

the solution characteristics allows for the discovery of innovative strategies that enhance simulation accuracy and efficiency. Despite these advances, challenges remain in the application of RL for mesh adaptation. Sample inefficiency, exploration in high-dimensional spaces, handling discontinuous solutions and generalisation to unseen scenarios are among the key hurdles that must be addressed. By leveraging RL’s ability to learn optimal mesh modifications, we can enhance the efficiency of numerical simulations while reducing human intervention.

**CRedit authorship contribution statement**

**David Huergo:** Conceptualization, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Gonzalo Rubio:** Supervision, Writing – original draft, Writing – review & editing. **Esteban Ferrer:** Funding acquisition, Resources, Supervision, Writing – original draft.

**Declaration of competing interest**

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Esteban Ferrer reports financial support was provided by the Comunidad de Madrid. Esteban Ferrer reports financial support was provided by Agencia Estatal de Investigación.

**Data availability**

Data will be made available on request.

**Acknowledgements**

EF would like to thank the support of Agencia Estatal de Investigación (for the grant “Europa Excelencia 2022” Proyecto EUR2022-134041/AEI/10.13039/501100011033) y del Mecanismo de Recu-

peración y Resiliencia de la Unión Europea, and the Comunidad de Madrid and Universidad Politécnica de Madrid for the Young Investigators award: APOYO-JOVENES-21-53NYUB-19-RRX1A0. GR and EF acknowledge the funding received by the Grant DeepCFD (Project No. PID2022-137899OB-I00) funded by ERDF A way of making Europe and MCIN/AEI/10.13039/501100011033/. Finally, all authors gratefully acknowledge the Universidad Politécnica de Madrid ([www.upm.es](http://www.upm.es)) for providing computing resources on Magerit Supercomputer.

## Appendix A. DGSEM

To derive the DGSEM, the physical domain is tessellated with non-overlapping elements,  $e$ , which are geometrically transformed into a reference element  $el = [-1, 1]$ . This transformation is carried out using an affine mapping that relates the physical coordinates  $x$  and the local reference coordinates  $\xi$ . The transformation is applied to the Burgers' equation (1), resulting in the following:

$$J u_t + \frac{\partial f}{\partial \xi} = 0, \quad (\text{A.1})$$

where  $J$  is the Jacobian of the mapping, and  $f = \frac{u^2}{2}$  is the flux. Next, we multiply Eq. (A.1) by a element-wise smooth test function  $\phi_j$  (i.e., in our case a polynomial), for  $0 \leq j \leq p$ , where  $p$  is the polynomial degree, and integrate over an element  $el$  to obtain the weak form

$$\int_{-1}^1 J u_t \phi_j + \int_{-1}^1 \frac{\partial f}{\partial \xi} \phi_j = 0. \quad (\text{A.2})$$

We can now integrate by parts the second integral, to obtain a local weak form of the equations (one per element) with the boundary fluxes separated from the interior

$$\int_{-1}^1 J u_t \phi_j + f \phi_j \Big|_{-1}^1 - \int_{-1}^1 f \frac{\partial \phi_j}{\partial \xi} = 0. \quad (\text{A.3})$$

We replace discontinuous fluxes at inter-element faces by a numerical flux,  $f^*$ , to couple neighbouring elements,

$$\int_{-1}^1 J u_t \phi_j + f^* \phi_j \Big|_{\partial el} - \int_{-1}^1 f \frac{\partial \phi_j}{\partial \xi} = 0. \quad (\text{A.4})$$

In this work, we have chosen the Roe numerical flux for  $f^*$ . The final step to obtain a usable numerical scheme is to approximate the numerical solution, fluxes, and test functions by polynomials of order  $p$ , and to use Gaussian quadrature rules to numerically approximate the integrals. Here, we use for Gauss-Legendre quadrature points. The resulting system of ordinary differential equations can be solved by any explicit or implicit time-stepping method; in this work we use an explicit Euler scheme. For more details, see [15,8].

## References

- [1] Francesca Basile, Jean-Baptiste Chapelier, Marta de la Llave Plata, Romain Laraufie, Pascal Frey, Unstructured h- and hp-adaptive strategies for discontinuous Galerkin methods based on a posteriori error estimation for compressible flows, *Comput. Fluids* 233 (2022) 105245.
- [2] M.P. Brenner, J.D. Eldredge, J.B. Freund, Perspective on machine learning for advancing fluid mechanics, *Phys. Rev. Fluids* 4 (10) (2019) 100501.
- [3] Steven L. Brunton, J. Nathan Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Cambridge University Press, 2019.
- [4] Nicholas Burgess, Dimitri Mavriplis, An hp-adaptive discontinuous Galerkin solver for aerodynamic flows on mixed-element meshes, in: 49th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2011, p. 490.
- [5] S. Le Clainche, E. Ferrer, S. Gibson, E. Cross, A. Parente, R. Vinuesa, Improving aircraft performance using machine learning: a review, [arXiv:2210.11481](https://arxiv.org/abs/2210.11481), 2022.
- [6] Marta de la Llave Plata, Fabio Naddei, Vincent Couaillier, Les of the flow past a circular cylinder using a multiscale discontinuous Galerkin method, in: 5th International Conference on Turbulence and Interactions, TI 2018, 2018.
- [7] Richard P. Dwight, Goal-oriented mesh adaptation for finite volume methods using a dissipation-based error indicator, *Int. J. Numer. Methods Fluids* 56 (8) (2008) 1193–1200.
- [8] E. Ferrer, G. Rubio, G. Ntoukas, W. Laskowski, O.A. Mariño, S. Colombo, A. Mateo-Gabín, H. Marbona, F. Manrique de Lara, D. Huergo, J. Manzanero, A.M. Rueda-Ramírez, D.A. Kopriva, E. Valero, Horses3d: a high-order discontinuous Galerkin solver for flow simulations and multi-physics applications, *Comput. Phys. Commun.* 287 (2023) 108700.
- [9] Esteban Ferrer, An interior penalty stabilised incompressible discontinuous Galerkin–Fourier solver for implicit large eddy simulations, *J. Comput. Phys.* 348 (2017) 754–775.
- [10] Corbin Foucart, Aaron Charous, Pierre F.J. Lermusiaux, Deep reinforcement learning for adaptive mesh refinement, *J. Comput. Phys.* 491 (2023) 112381.
- [11] Niklas Freymuth, Philipp Dahlinger, Tobias Würth, Luise Kärger, Gerhard Neumann, Swarm reinforcement learning for adaptive mesh refinement, 2023.
- [12] Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, Elie Hachem, A review on deep reinforcement learning for fluid mechanics, *Comput. Fluids* 225 (2021) 104973.
- [13] S.J. Kamkar, Andrew M. Wissink, Venkateswaran Sankaran, A. Jameson, Feature-driven Cartesian adaptive mesh refinement for vortex-dominated flows, *J. Comput. Phys.* 230 (16) (2011) 6271–6298.
- [14] Diederik P. Kingma, Jimmy Ba Adam, A method for stochastic optimization, [arXiv preprint, arXiv:1412.6980](https://arxiv.org/abs/1412.6980), 2014.
- [15] David A. Kopriva, *Implementing Spectral Methods for Partial Differential Equations: Algorithms for Scientists and Engineers*, Springer Science & Business Media, May 2009.
- [16] Jiaqing Kou, Soledad Le Clainche, Esteban Ferrer, Data-driven eigensolution analysis based on a spatio-temporal Koopman decomposition, with applications to high-order methods, *J. Comput. Phys.* 449 (2022) 110798.
- [17] Jiaqing Kou, Aurelio Hurtado-de Mendoza, Saumitra Joshi, Soledad Le Clainche, Esteban Ferrer, Eigensolution analysis of immersed boundary method based on volume penalization: applications to high-order schemes, *J. Comput. Phys.* 449 (2022) 110817.
- [18] Jiaqing Kou, Oscar A. Marino, Esteban Ferrer, Jump penalty stabilisation techniques for under-resolved turbulence in discontinuous Galerkin schemes, [arXiv preprint, arXiv:2208.11426](https://arxiv.org/abs/2208.11426), 2022.
- [19] Hung-Chi Kuo, R.T. Williams, Semi-Lagrangian solutions to the inviscid Burgers equation, *Mon. Weather Rev.* 118 (6) (1990) 1278–1288.
- [20] Wojciech Laskowski, Gonzalo Rubio, Eusebio Valero, Esteban Ferrer, A functional oriented truncation error adaptation method, *J. Comput. Phys.* 451 (2022) 110883.
- [21] Weizhao Li, Hong Luo, Aditya Pandare, Jozsef Bakosi, A p-adaptive discontinuous Galerkin method for compressible flows using charm++, in: AIAA Scitech 2020 Forum, 2020, p. 1565.
- [22] Cooper Lorsche, Amir Barati Farimani, Mesh deep q network: a deep reinforcement learning framework for improving meshes in computational fluid dynamics, *AIP Adv.* 13 (1) (2023) 015026.
- [23] Juan Manzanero, Gonzalo Rubio, Esteban Ferrer, Eusebio Valero, Dispersion-dissipation analysis for advection problems with nonconstant coefficients: applications to discontinuous Galerkin formulations, *SIAM J. Sci. Comput.* 40 (2) (2018) A747–A768.
- [24] Fabio Naddei, Marta de la Llave Plata, Vincent Couaillier, Frédéric Coquel, A comparison of refinement indicators for p-adaptive simulations of steady and unsteady flows using discontinuous Galerkin methods, *J. Comput. Phys.* 376 (January 2019) 508–533.
- [25] Gerasimos Ntoukas, Juan Manzanero, Gonzalo Rubio, Eusebio Valero, Esteban Ferrer, A free-energy stable p-adaptive nodal discontinuous Galerkin for the Cahn–Hilliard equation, *J. Comput. Phys.* 442 (2021) 110409.
- [26] Gerasimos Ntoukas, Juan Manzanero, Gonzalo Rubio, Eusebio Valero, Esteban Ferrer, An entropy-stable p-adaptive nodal discontinuous Galerkin for the coupled Navier–Stokes/Cahn–Hilliard system, *J. Comput. Phys.* 458 (2022) 111093.
- [27] Per-Olof Persson, Jaime Peraire, *Sub-Cell Shock Capturing for Discontinuous Galerkin Methods*, 2012.
- [28] Andrés M. Rueda-Ramírez, Juan Manzanero, Esteban Ferrer, Gonzalo Rubio, Eusebio Valero, A p-multigrid strategy with anisotropic p-adaptation based on truncation errors for high-order discontinuous Galerkin methods, *J. Comput. Phys.* 378 (2019) 209–233.
- [29] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, Oleg Klimov, Proximal policy optimization algorithms, [arXiv preprint, arXiv:1707.06347](https://arxiv.org/abs/1707.06347), 2017.
- [30] Tensorforce Team, *Tensorforce*, <https://tensorforce.readthedocs.io/en/latest/>.
- [31] Ricardo Vinuesa, Steven L. Brunton, Enhancing computational fluid dynamics with machine learning, *Nat. Comput. Sci.* 2 (6) (2022) 358–366.
- [32] J. Viquerat, P. Meliga, A. Larcher, E. Hachem, A review on deep reinforcement learning for fluid mechanics: an update, *Phys. Fluids* 34 (11) (2022) 111301.
- [33] Tailin Wu, Takashi Maruyama, Qingqing Zhao, Gordon Wetzstein, Jure Leskovec, Learning controllable adaptive simulation for multi-resolution physics, in: The Eleventh International Conference on Learning Representations, 2023.
- [34] Jiachen Yang, Tarik Dzanic, Brenden Petersen, Jun Kudo, Ketan Mittal, Vladimir Tomov, Jean-Sylvain Camier, Tuo Zhao, Hongyuan Zha, Tzanio Kolev, Robert Anderson, Daniel Faissol, Reinforcement Learning for Adaptive Mesh Refinement, 2023.

- [35] Jiachen Yang, Ketan Mittal, Tarik Dzanic, Socratis Petrides, Brendan Keith, Brenden Petersen, Daniel Faissol, Robert Anderson, Multi-agent reinforcement learning for adaptive mesh refinement, in: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '23, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems, 2023, pp. 14–22.