DOI: 10.61782/fa.2025.0737

FORUM ACUSTICUM
EURONOISE 2025

REINFORCEMENT LEARNING TO ENHANCE COMPUTATIONAL FLUID DYNAMIC SIMULATIONS WITH ACOUSTICS

David Huergo[1*] **Martín de Frutos**[1] **Óscar Mariño**[1]
Eduardo Jané[1] **Gonzalo Rubio**[1,2] **Esteban Ferrer**[1,2]
[1] ETSIAE-UPM-School of Aeronautics, Universidad Politécnica de Madrid,
Plaza Cardenal Cisneros 3, E-28040 Madrid, Spain.
[2] Center for Computational Simulation, Universidad Politécnica de Madrid,
Campus de Montegancedo, Boadilla del Monte, 28660 Madrid, Spain.

ABSTRACT

Reinforcement learning (RL) offers a powerful approach to automate decision-making in Computational Fluid Dynamic (CFD), improving accuracy and reducing human effort, leading to more efficient solvers.

In this work, we expand our RL-based polynomial adaptation technique for high-order solvers [1, 2] to capture acoustics. We improve aeroacoustic simulations employing RL to define a high-order acoustic path in the high-order solver HORSES3D [3], allowing precise wave capture from source to observer. Offline training eliminates the need for high-fidelity solutions, making the method versatile across various meshes and PDEs. This approach effectively addresses aeroacoustic challenges, focusing computational efforts where needed and ensuring accurate results at a reduced cost.

Keywords: *reinforcement learning, high-order discontinuous Galerkin, p-adaptation, adaptive mesh refinement, acoustics*

1. INTRODUCTION

The field of computational fluid dynamics (CFD) has undergone a paradigm change through the integration of machine learning techniques with classical numerical methods [4, 5]. In particular, reinforcement learning (RL) has been adopted as a valuable tool for a variety of fields, creating high-performing autonomous agents capable of making real-time decisions in complex environments [6]. In fluid dynamics, RL has been widely used to enable dynamic flow control [7–9], surpassing traditional approaches. Other applications in the field of numerical simulations for fluid dynamics include the selection of constants in turbulence models [10, 11], optimal parameters in high-order schemes [12, 13] or mesh adaptation.

Mesh adaptation or adaptive mesh refinement (AMR) allows to refine or coarsen the computational mesh based on sensors or metrics. Traditionally, this process has relied on manual intervention, heuristics or error estimators [14, 15], which is time-consuming and limits the ability to capture complex flow phenomena. The integration of RL for mesh adaptation offers a novel approach to automate the process, minimizing user intervention while providing a flexible and general framework applicable to a wide range of problems. RL has been successfully used for AMR (h-refinement) in recent works, where agents were trained to refine or coarsen the mesh, increasing the accuracy in critical regions and minimizing the computational cost [16]. Additionally, Yang et al. [17] consider new policy architectures which are trained from numerical simulations and are agnostic to the mesh size. These approaches provide accurate solutions and a reduced computational cost, surpassing traditional mesh refinement approaches.

This work extends our research presented in [1, 2], where

*__Corresponding author__: david.huergo.perea@upm.es.

we established the basis for applying reinforcement learning to p-adaptation in complex 3D turbulent problems. Furthermore, we provide an RL-based error estimator that can be used to estimate the order of magnitud of spatial errors inexpensively. Based on this original concept, we have developed a novel methodology to apply this approach to acoustic simulations. Here, we define an acoustic path, from a source to an observer, where the task of the p-adaptation algorithm is to accurately capture pressure waves; while outside of the acoustic path, the mesh is refined based on aerodynamic features. This hybrid adaptation strategy provides accurate solutions and results in a reduction of the required computational cost when running numerical simulations for acoustics.

2. METHODOLOGY

The methodology is implemented in the open source solver HORSES3D [3, 18, 19], developed at the ETSIAE-UPM School of Aeronautics in Madrid, and available on Github (`https://github.com/loganoz/horses3d`). This solver is a high-order discontinuous Galerkin (DG) tool capable of simulating compressible, incompressible, and multiphase flows with acoustics. HORSES3D allows for the use of high order polynomials in each element in the mesh, and shows exponential decay of the error when increasing the polynomial order but with only a linear increase in the cost, which results in very efficient computations.

The p-adaptation strategy [2] allows one to select the optimal polynomial order in each element of the mesh independently. As the proposed approach is anisotropic, three different polynomials are defined inside each element of the mesh. This methodology provides an automatic way of increasing the accuracy of the solver, by refining the mesh in regions with strong gradients, while keeping the computational cost down, as the polynomial order is reduced in the far field or other regions of less interest. The p-adaptation process is based on a set of *variables of interest* (e.g., momentum or pressure) that the user must provide. These variables will be used as a reference for the RL agent, resulting in an optimal adaptation for these specific variables.

When applied to acoustics, it is possible to define an *acoustic path* between a source and an observer. Within this acoustic path, the RL agent applies a pressure-based p-adaptation to accurately capture the acoustic waves; and outside the path, it applies a momentum/velocity-based adaptation to represent aerodynamic flow features. The acoustic path is defined along a straight line between the source and the observer. Then, the minimum distance from every high order Gauss node of the mesh to that line is calculated as depicted in Figure 1. All elements with at least one Gauss node whose distance d is below a threshold d_{th}, will be included inside the acoustic path. The definition of d_{th} is problem dependent and is an input provided by the user. Finally, in a real problem both the source and the observer could be defined as a surface (e.g., a boundary) instead of a single point. In this scenario, a different line is created from each node of the source to each node of the observer. Furthermore, since this step is performed only once, at the beginning of the simulation, the overhead in relation to the complete simulation time is negligible.

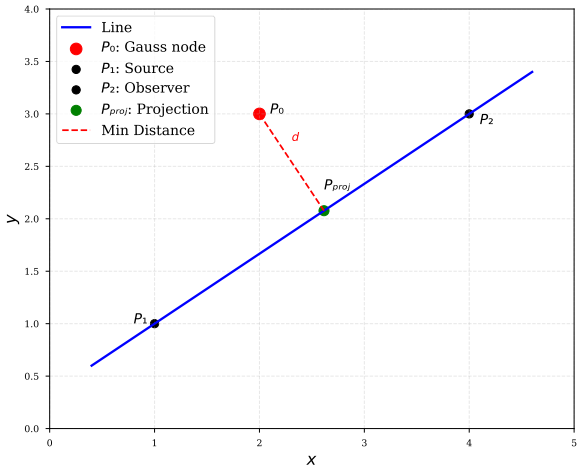

Figure 1. Example of the distance calculation from a high order Gauss node to the line that connects the acoustic source and the observer.

3. RESULTS

To test the proposed p-adaptation methodology for acoustics and the acoustic path, we simulate different problems, in increasing level of complexity. As the RL-based p-adaptation agent is agnostic to the computation mesh and the PDE being solved, only one agent is trained, which is used for every problem without modifications.

3.1 Monopole

First, to test the performance of the pressure-based p-adaptation against a momentum-based, a simple

monopole is simulated using the nonlinear Euler Equations (EE). The mesh is a square with 400×400 elements. In this test case, the acoustic path is not considered; instead, the polynomial order will be adapted by the RL agent in the whole mesh. A snapshot of the pressure field, once quasi-steady state conditions are achieved, is represented in Figure 2.

Additionally, to validate the proposed methodology, three simulations have been performed, with different p-adaptation strategies: an uniform $p = 6$ polynomial order (used as a reference), a pressure-based adaptation and a momentum-based adaptation. For the two p-adaptation strategies, the maximum polynomial order is limited to $p_{\max} = 6$; hence, the uniform $p = 6$ simulation will always be more accurate than the p-adaptation solutions, but also more costly. In Figure 3, we show the error between both p-adaptation strategies and the reference $p = 6$ solution for the pressure field. The solution computed with the pressure-based adaptation shows a significantly smaller error. Although momentum-based adaptation is appropriate to capture aerodynamic features [2], this problem highlights the importance of using pressure as the reference variable for acoustics.

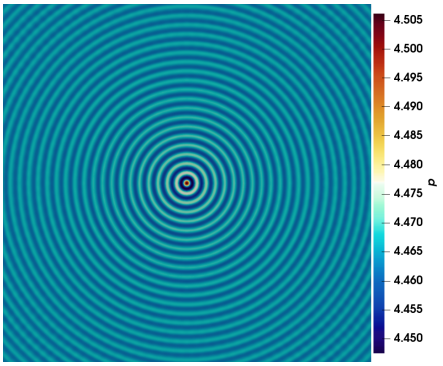

Figure 2. Snapshot of the pressure field.

Finally, in Figure 4 we show the rmse error of the pressure field, as a function of the degrees of freedom of the mesh, in a vertical line (from the source to the boundary at the top) and a diagonal line (from the source to the top-right corner) for both the momentum-based and the pressure-based simulations. As expected, the pressure-based adaptation shows a decreasing error with the number of degrees of freedom, while the momentum-based is not able to reduce the error below 10^{-5}, even if the mesh is refined. Therefore, the momentum field is not an appropriate reference variable to accurately capture acoustic waves.

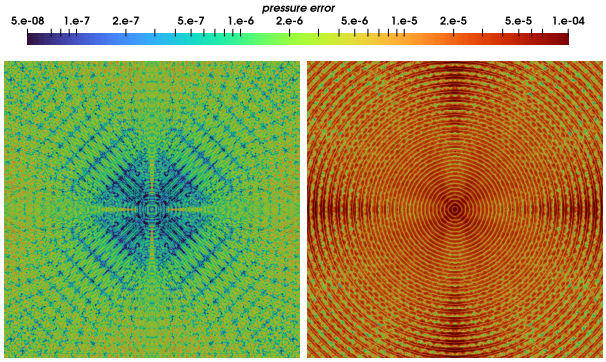

a) Pressure-based. b) Momentum-based.

Figure 3. Comparison of the pressure error, in relation to the reference $p = 6$ simulation, between the pressure-based and the momentum-based adaptations.

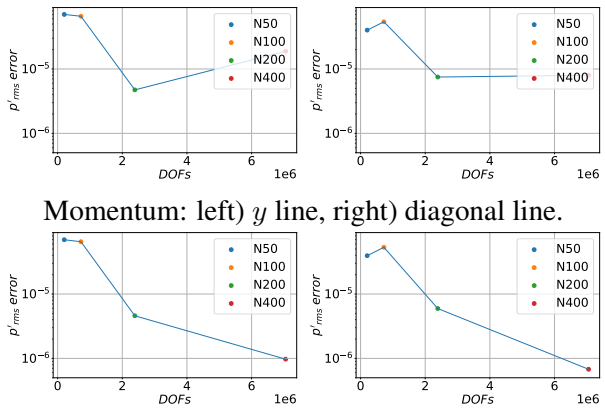

Momentum: left) y line, right) diagonal line.

Pressure: left) y line, right) diagonal line.

Figure 4. Comparison of the pressure rmse error for four different meshes, computed in relation to the reference $p = 6$ simulation, for the pressure-based adaptation and the momentum-based adaptation.

3.2 Multiphase monopole

Once it is demonstrated that the pressure is the appropriate variable to adapt the polynomial order in acoustic problems, we simulate a multiphase problem with two fluids: air and water. In this problem, an acoustic source generates pressure waves in the air region, which are propagated to the water region through an interface. The problem is solved inside a square with a 100×100 Cartesian mesh. To avoid reflections at the boundaries, a sponge is defined near the boundaries to dampen the acoustic waves. This setup is shown schematically in Figure 5. As the wavelength of the pressure waves are different in air and water, the p-adaptation agent should be able to select the optimal polynomial order in each region accordingly. Then, as we want to focus on the propagation of the acoustic waves from the air to the water, we define the acoustic path from the source to the boundary at the bottom, as shown in Figure 6, with a threshold distance of $d_{th} = 2.5$. This way, the RL agent will adapt the polynomial order inside the acoustic path only, reducing the computational cost of the overall simulation.

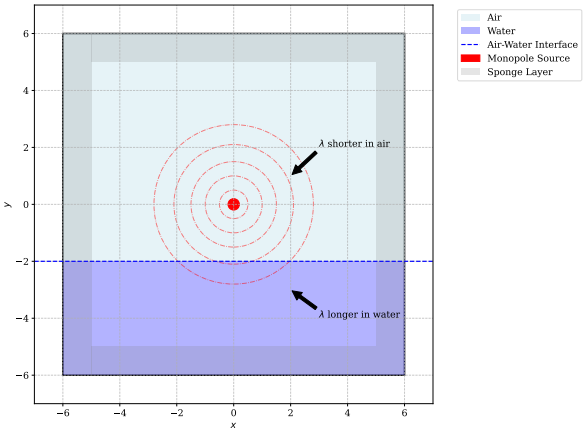

Figure 5. Setup of the multiphase problem.

The pressure field, once the simulation is stable, is represented in Figure 7 a). As expected, the wavelength of the pressure waves is significantly longer in the water and, therefore, the polynomial order should be smaller in that region. The average polynomial order (among the three anisotropic polynomials p_x, p_y and p_z) in each element of the mesh is shown in Figure 7 b). In this figure, there are four regions of interest that should be analyzed:

1. Near the boundaries, inside the sponge layer, the polynomial order has been manually fixed to a constant value to avoid undesired reflections.
2. Above the source, in the air region, but outside the acoustic path, the polynomial order is set to the minimum to reduce the computational cost, as we are not interested in the acoustics there.
3. Inside the acoustic path, in the air region, the wavelengths are shorter, and strong pressure gradients can be found inside some elements. Hence, the RL agent selects high-order polynomials to accurately capture the pressure waves.
4. Inside the acoustic path, in the water region, the wavelengths are longer and low order polynomials are enough to represent the pressure waves.

In conclusion, the RL agent is capable of selecting the optimum polynomial order inside the acoustic path, discriminating between different regions, and adapting the polynomial order accordingly.

Figure 6. Acoustic path of the multiphase problem.

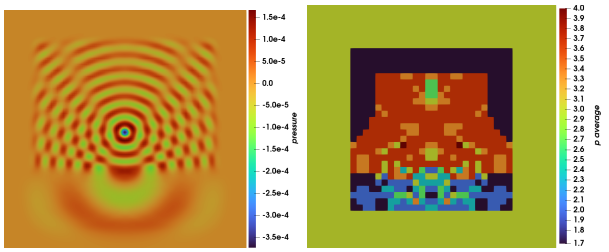

a) Pressure. b) Average polynomial.

Figure 7. Snapshot of the pressure field and the average polynomial order in each element of the mesh.

3.3 Cylinder at Reynolds 200

Finally, we propose a more challenging problem, the direct noise computation (DNC) produced by the flow around a cylinder at Reynolds 200, where the RL agent must consider two different adaptation strategies at the same time: inside the acoustic path, a pressure-based adaptation is performed to capture the acoustic waves, and outside the acoustic path, a momentum-based adaptation is used to represent aerodynamic flow features, such as vortices. For this problem we solve the full compressible Navier–Stokes (NS) equations, that accounts for both aerodynamics and acoustic physics. The acoustic path for this problem and the background mesh are depicted in Figure 8, where the acoustic source is the cylinder, the observer is defined 20 length units above the source and the threshold distance is $d_{th} = 1$. Also, it is necessary to define two tolerances to let the agent know when the fluctuations of the reference variable inside one element are negligible. In this case:

$$tol_m = 10^{-2} \qquad tol_p = 10^{-4}, \tag{1}$$

with tol_m as the momentum tolerance and tol_p as the pressure tolerance.

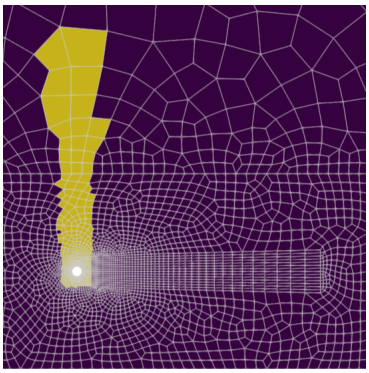

Figure 8. Acoustic path of the cylinder at $Re = 200$.

Once the flow field is stable, a snapshot of the average polynomial order in each element is represented in Figure 9, where there are three main regions:

1. The acoustic path: The polynomial order is increased to capture the pressure waves between the source and the observer.
2. The wake: The vortex shedding creates strong velocity gradients in this region, and high-order polynomials are required to accurately represent the flow field.
3. The far field: Although this region is also adapted based on momentums, the gradients are negligible, and the RL agent selects the minimum polynomial order available in each axis, as nothing of interest is happening.

Therefore, the RL agent, without additional information on this particular problem, is able to discriminate between different regions and to select the optimal polynomial order that provides an accurate solution while keeping the computational cost down.

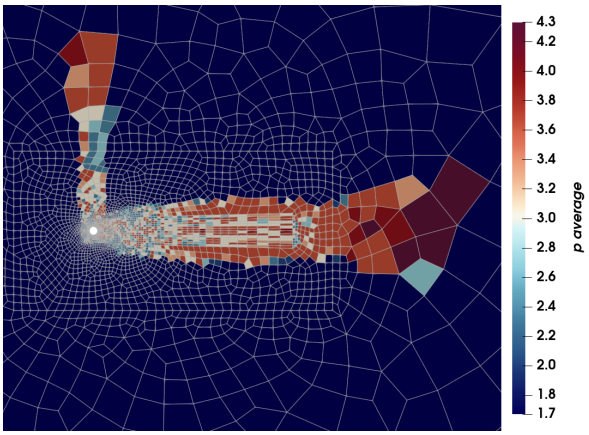

Figure 9. Snapshot of the average polynomial order of the cylinder at $Re = 200$.

Then, to check the accuracy of the proposed methodology, we compare the current results with two additional simulations with an uniform polynomial order of $p = 3$ and $p = 5$ in the whole mesh. First, in Figure 10 we show the root mean square pressure fluctuations along the acoustic path, from source to observer. Although the number of high-order elements in the p-adapted mesh is smaller, the results show a good agreement with the reference simulations. Furthermore, the reader may notice that the p-adapted solution is closer to the reference $p = 3$ simulation instead of $p = 5$, which is more accurate. The main reason behind this behavior is that, close to the cylinder, the aerodynamic pressure gradients are stronger than the acoustic fluctuations, but also very smooth. Therefore, near the cylinder a polynomial order $p = 2$ or $p = 3$ is enough to represent the pressure gradients, reducing the accuracy of the acoustic waves in that region, which are

propagated to the observer. To solve this issue, an uniform high-order polynomial order could be set around the body manually. Secondly, in Figure 11 we compare the SPL measured at the observer's location with different approaches, which highlights the accuracy of the current methodology.

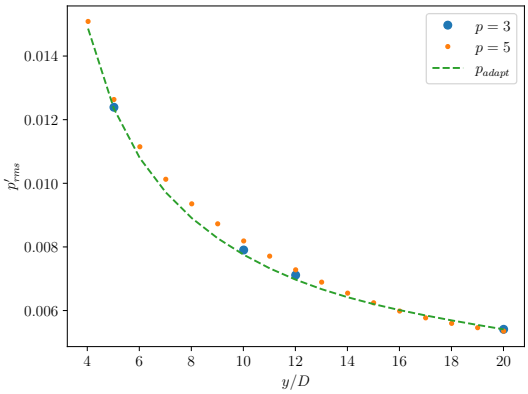

Figure 10. Comparison of the rms pressure fluctuations from the source to the observer.

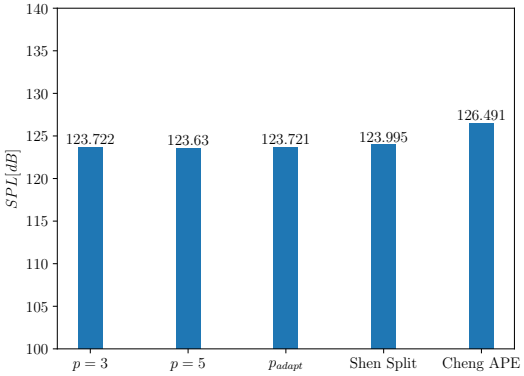

Figure 11. Comparison of the overall SPL at the observer's location.

Finally, for completeness, the u velocity field following the x axis is shown in Figure 12, where the vortex shedding is accurately represented, highlighting the effectiveness of the proposed p-adaptation strategy. In addition, this problem shows the flexibility of the p-adaptation RL agent, which can be applied in many different ways (to different problems, PDEs and meshes) and based on arbitrary reference variables, without retraining the agent. All these advantages make reinforcement learning for adaptive mesh refinement a good choice to solve aeroacoustic CFD problems, including DNC.

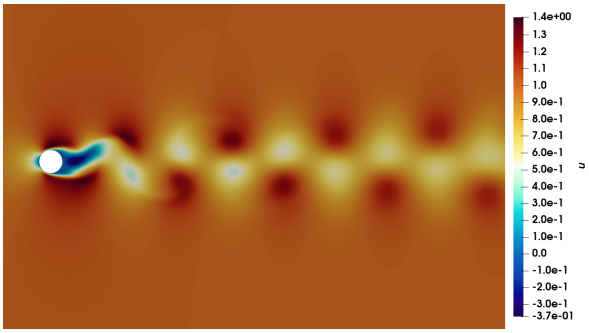

Figure 12. Snapshot of the velocity field of the cylinder at $Re = 200$.

4. CONCLUSIONS

In conclusion, this study presents a new approach to applying anisotropic p-adaptation for acoustics in high-order h/p solvers using reinforcement learning. The RL-based adaptation dynamically adjusts high-order polynomials based on the pressure field to capture acoustic fluctuations. Furthermore, it can handle different adaptation regions based on pressure, momentum, velocity, or other flow features. The proposed methodology is independent of the computational mesh and can be applied to any partial differential equation, illustrating its broad applicability and flexibility.

These findings open new avenues for further research in acoustics and other CFD applications using RL-based h/p-mesh adaptation.

5. ACKNOWLEDGMENTS

This research has received funding from the European Union (ERC, Off-coustics, project number 101086075). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

EF and GR acknowledge the funding received by the Grant DeepCFD (Project No. PID2022-137899OB-I00) funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU.
DH and EF acknowledge the funding received by the Comunidad de Madrid according to Orden 5067/2023, of December 27th, issued by the Consejero de Educación, Ciencia y Universidades, which announces grants for the hiring of predoctoral research personnel in training for the year 2023.
Finally, all authors gratefully acknowledge the Universidad Politécnica de Madrid (www.upm.es) for providing computing resources on Magerit Supercomputer and the computer resources at MareNostrum and the technical support provided by Barcelona Supercomputing Center (RES-IM-2024-1-0003).

6. REFERENCES

[1] D. Huergo, G. Rubio, and E. Ferrer, "A reinforcement learning strategy for p-adaptation in high order solvers," *Results in Engineering*, vol. 21, p. 101693, 2024.

[2] D. Huergo, M. de Frutos, E. Jané, O. A. Marino, G. Rubio, and E. Ferrer, "Reinforcement learning for anisotropic p-adaptation and error estimation in high-order solvers," *arXiv preprint arXiv:2407.19000*, 2024.

[3] E. Ferrer, G. Rubio, G. Ntoukas, W. Laskowski, O. Mariño, S. Colombo, A. Mateo-Gabín, H. Marbona, F. Manrique de Lara, D. Huergo, J. Manzanero, A. Rueda-Ramírez, D. Kopriva, and E. Valero, "Horses3d: A high-order discontinuous galerkin solver for flow simulations and multi-physics applications," *Computer Physics Communications*, vol. 287, p. 108700, 2023.

[4] S. Le Clainche, E. Ferrer, S. Gibson, E. Cross, A. Parente, and R. Vinuesa, "Improving aircraft performance using machine learning: A review," *Aerospace Science and Technology*, vol. 138, p. 108354, 2023.

[5] R. Vinuesa and S. L. Brunton, "Enhancing computational fluid dynamics with machine learning," *Nature Computational Science*, vol. 2, p. 358–366, June 2022.

[6] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. The MIT Press, 2^{nd} *ed.*, 2018.

[7] P. Garnier, J. Viquerat, J. Rabault, A. Larcher, A. Kuhnle, and E. Hachem, "A review on deep reinforcement learning for fluid mechanics," *Computers & Fluids*, vol. 225, p. 104973, 2021.

[8] C. Vignon, J. Rabault, and R. Vinuesa, "Recent advances in applying deep reinforcement learning for flow control: Perspectives and future directions," *Physics of Fluids*, vol. 35, no. 3, p. 031301, 2023.

[9] B. Font, F. Alcántara-Ávila, J. Rabault, R. Vinuesa, and O. Lehmkuhl, "Active flow control of a turbulent separation bubble through deep reinforcement learning," in *Journal of Physics: Conference Series*, vol. 2753, p. 012022, IOP Publishing, 2024.

[10] M. Kurz, P. Offenhäuser, and A. Beck, "Deep reinforcement learning for turbulence modeling in large eddy simulations," *International Journal of Heat and Fluid Flow*, vol. 99, p. 109094, 2023.

[11] A. Beck and M. Kurz, "Toward discretization-consistent closure schemes for large eddy simulation using reinforcement learning," *Physics of Fluids*, vol. 35, no. 12, 2023.

[12] Y. Feng, F. S. Schranner, J. Winter, and N. A. Adams, "A deep reinforcement learning framework for dynamic optimization of numerical schemes for compressible flow simulations," *Journal of Computational Physics*, vol. 493, p. 112436, 2023.

[13] D. Huergo, L. Alonso, S. Joshi, A. Juanicotena, G. Rubio, and E. Ferrer, "A reinforcement learning strategy to automate and accelerate h/p-multigrid solvers," *Results in Engineering*, p. 102949, 2024.

[14] M. Kompenhans, G. Rubio, E. Ferrer, and E. Valero, "Adaptation strategies for high order discontinuous Galerkin methods based on Tau-estimation," *Journal of Computational Physics*, vol. 306, pp. 216–236, 2016.

[15] M. Kompenhans, G. Rubio, E. Ferrer, and E. Valero, "Comparisons of p-adaptation strategies based on truncation- and discretisation-errors for high order discontinuous Galerkin methods," *Computers and Fluids*, vol. 139, pp. 36–46, 2016.

[16] C. Foucart, A. Charous, and P. F. Lermusiaux, "Deep reinforcement learning for adaptive mesh refinement," *Journal of Computational Physics*, vol. 491, p. 112381, 2023.

[17] J. Yang, T. Dzanic, B. Petersen, J. Kudo, K. Mittal, V. Tomov, J.-S. Camier, T. Zhao, H. Zha, T. Kolev, *et al.*, "Reinforcement learning for adaptive mesh refinement,"

in *International Conference on Artificial Intelligence and Statistics*, pp. 5997–6014, PMLR, 2023.

[18] L. Botero-Bolívar, O. A. Marino, C. H. Venner, L. D. de Santana, and E. Ferrer, "Low-cost wind turbine aeroacoustic predictions using actuator lines," *Renewable Energy*, vol. 227, p. 120476, 2024.

[19] O. A. Marino, R. Sanz, S. Colombo, A. Sivaramakrishnan, and E. Ferrer, "Modelling wind turbines via actuator lines in high-order h/p solvers," 2024.