Int. J. Adv. Eng. Comput. Sci. 2026; 2(3); 61-67

International Journal of Advances in Engineering and Computer Science

E-ISSN: 3121-6382

Volume 2, Issue 3, September, 2026
(An Academians Publishers of Research & Academic Resources)



Research Article

Physics-Informed Machine Learning for Enhanced Thermal-Fluid Predictions in Complex Engineering Systems: A Hybrid CFD-Data Driven Framework

Jonsh Uod*
[1]*Department of Enginnering, Adama Science and Technology University, Ethiopia*

Abstract: Computational fluid dynamics has become an indispensable tool for analysing complex thermal-fluid systems in engineering applications ranging from particle detector cooling to electronics thermal management. However, the computational cost of high-fidelity CFD simulations limits their use in real-time monitoring, design optimization, and uncertainty quantification. This investigation presents a physics-informed machine learning framework that combines the predictive accuracy of CFD with the computational efficiency of data-driven models. The proposed approach uses a K-nearest neighbours algorithm constrained by the governing conservation equations to reconstruct temperature and velocity fields from sparse sensor measurements. The framework is validated against three-dimensional CFD simulations of a simplified enclosure representative of particle detector thermal management systems. Results demonstrate that the physics-informed KNN model achieves a mean absolute error of 2.34 K for temperature predictions and 0.087 m/s for velocity magnitude, representing a 67% reduction in error compared to pure data-driven interpolation. The model successfully captures thermal stratification patterns and buoyancy-driven flow structures with a computational speedup of approximately 850 times compared to the full CFD solution. Sensitivity analysis reveals that the optimal number of neighbours varies between 3 and 7 depending on the local flow regime, with lower values performing better in regions of high gradient. The framework provides a practical methodology for developing fast surrogate models that preserve physical consistency while enabling real-time thermal monitoring and control in complex engineering systems.

Keywords: *Physics-informed machine learning, computational fluid dynamics, thermal management, surrogate modelling, K-nearest neighbours, particle detector cooling*

1. Introduction

The thermal management of complex engineering systems has become increasingly critical as power densities continue to rise while available installation space shrinks. In particle physics detectors, for example, silicon tracking sensors must operate at temperatures low enough to prevent thermal runaway while maintaining structural integrity under irradiation. Similarly, power electronics in electric vehicles and data centre servers require sophisticated cooling strategies to maintain junction temperatures within safe limits. The design and optimization of these thermal systems relies heavily on computational fluid dynamics simulations that resolve the detailed flow and temperature fields within the equipment enclosure.

The challenge with CFD-based design is the substantial computational cost of high-fidelity simulations. A single steady-state simulation of a moderately complex enclosure may require several hours of wall-clock time on a multi-core workstation, while transient analyses or parametric studies can consume days or weeks of computational time. This cost structure precludes the use of CFD for real-time monitoring, model predictive control, or the thousands of evaluations required for robust optimization under uncertainty. The need for fast, accurate surrogate models that can approximate the CFD solution with minimal computational overhead has therefore become a priority in thermal engineering research.

Physics-informed machine learning has emerged as a promising approach for developing such surrogates. The central idea is to incorporate physical constraints—governing equations, boundary conditions, or conservation laws—directly into the machine learning framework. This ensures that the resulting model not only fits the training data but also respects the fundamental physics of the system, improving generalization to conditions not represented in the

Received: 02 July 2026; Accepted: 27 August 2026; Published: 30 September 2026

© 2026 by the authors. This work is distributed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

training set and preventing non-physical predictions that violate conservation principles.

The present study develops and validates a physics-informed K-nearest neighbours framework for thermal-fluid field reconstruction in complex enclosures. The approach is motivated by the need for a computationally efficient yet physically consistent method for estimating temperature and velocity fields from limited sensor data. Such a capability would enable real-time thermal monitoring in particle detector systems, where fibre optic sensors provide sparse measurements but full-field reconstruction is required for control decisions.

The specific objectives of this investigation are threefold. First, to develop a physics-informed KNN algorithm that incorporates conservation constraints into the distance metric and prediction step. Second, to validate the approach against high-fidelity CFD simulations of a representative enclosure geometry with buoyancy-driven flow. Third, to characterize the performance of the method in terms of accuracy, computational cost, and robustness to sensor placement.

2. Literature Review

2.1 CFD for Thermal Management of Complex Systems

Computational fluid dynamics has been applied extensively to thermal management problems across diverse engineering domains. In the particle physics community, CFD studies have examined cooling schemes for silicon detectors, including forced air convection, liquid cooling with evaporative refrigerants, and dry gas flushing to prevent condensation. The ATLAS Inner Tracker upgrade at CERN required detailed CFD analysis of dry nitrogen flushing to ensure that the detector volume remained below the dew point specification during operation.

These studies typically employ Reynolds-averaged Navier-Stokes equations with turbulence models such as k-ε or k-ω SST, solved on unstructured meshes that conform to complex detector geometries. The simulations resolve conjugate heat transfer between solid detector components and the surrounding gas, often requiring millions of computational cells to capture the range of length scales from millimetre-scale electronic components to metre-scale enclosure dimensions.

The computational cost of these simulations is substantial. A typical steady-state simulation of the ATLAS ITk volume required approximately 8.8 million cells and several hours of wall-clock time on a high-performance computing cluster. Transient simulations, which are necessary to capture thermal cycling and flow oscillations, require one to two orders of magnitude more time. This cost structure motivates the development of reduced-order models that can provide rapid estimates of thermal behaviour for control and monitoring applications.

2.2 Surrogate Modelling for Thermal-Fluid Systems

Surrogate modelling, also known as metamodelling or response surface methodology, seeks to construct computationally cheap approximations of expensive simulation models. Classical approaches include polynomial regression, radial basis functions, and Kriging (Gaussian process regression). These methods have been applied successfully to thermal-fluid problems ranging from heat exchanger design to building energy simulation.

More recently, machine learning approaches have gained popularity for surrogate construction. Neural networks, support vector machines, and ensemble methods such as random forests have been used to predict performance metrics such as Nusselt number, pressure drop, and thermal resistance from geometric and operational parameters. These methods can capture nonlinear relationships that are difficult to represent with classical response surfaces.

However, most surrogate modelling studies focus on predicting scalar quantities of interest rather than full field reconstruction. For thermal monitoring applications, the ability to estimate the complete temperature and velocity fields from sparse sensor data is essential. This requires a different approach that can interpolate spatially distributed measurements while respecting physical constraints.

2.3 Physics-Informed Machine Learning

Physics-informed machine learning encompasses a family of methods that incorporate physical laws into the learning process. The most prominent example is the physics-informed neural network, which adds residual terms from the governing differential equations to the loss function during training. This approach has been applied successfully to forward problems (solving PDEs given boundary conditions) and inverse problems (estimating parameters or boundary conditions from observations).

The computational cost of physics-informed neural networks remains substantial because the loss function must be evaluated at many collocation points, and automatic differentiation is required to compute the derivatives appearing in the physical residuals. For thermal-fluid problems involving turbulence, the cost is further increased by the need to resolve multiple scales.

Alternative approaches have explored the use of simpler machine learning algorithms combined with physical constraints. For example, physics-informed Kriging incorporates known trends or conservation laws into the mean function of the Gaussian process. Physics-informed KNN, the focus of the present study, modifies the distance metric and prediction step to ensure that interpolated values satisfy local conservation constraints.

2.4 K-Nearest Neighbours for Field Reconstruction

The K-nearest neighbours algorithm is a non-parametric method that predicts the value at a query point by averaging the values at the K closest training points. Despite its simplicity, KNN has proven effective for a wide range of regression and classification tasks, particularly when the underlying function is smooth and the training data are dense enough to capture local variations.

For field reconstruction, KNN offers several advantages over parametric methods. It makes no assumptions about the functional form of the field, allowing it to adapt to complex spatial patterns. It is inherently local, meaning that predictions at a given location depend only on nearby data, which is consistent with the local nature of many physical transport processes. And it is computationally efficient for prediction, requiring only distance calculations and averaging, with no need for iterative solving or matrix inversion.

The application of KNN to thermal-fluid field reconstruction has been limited, however. Most studies use KNN for scalar prediction or time-series forecasting rather than spatial field reconstruction. The incorporation of physical constraints into KNN has received little attention, despite the potential for improving accuracy and physical consistency.

3. Methodology

3.1 Case Study Description

The case study for this investigation is a simplified enclosure representative of a particle detector thermal management system. The enclosure is a rectangular box with dimensions of 2.0 m × 1.5 m × 1.0 m, with heated walls on two opposite sides representing detector surfaces at 30°C and 40°C, respectively. The remaining walls are adiabatic. The enclosure is filled with dry nitrogen at 1 atm pressure, and buoyancy-driven circulation develops due to the temperature differences.

This configuration captures the essential physics of the ATLAS ITk dry nitrogen flushing scheme, where temperature gradients between detector surfaces and the surrounding gas drive natural circulation patterns. While the actual detector geometry is more complex, involving cylindrical layers and support structures, the simplified enclosure provides a well-defined test case for developing and validating the physics-informed KNN framework before applying it to more complex configurations.

The CFD simulation of this enclosure was performed using ANSYS Fluent 2023 R1 with the realizable k-ε turbulence model and full buoyancy effects. The mesh consisted of 1.2 million hexahedral cells with refinement near the heated walls to resolve the thermal boundary layers. The SIMPLE algorithm was used for pressure-velocity coupling, with second-order upwind discretization for momentum and energy equations. Convergence was achieved after approximately 2,500 iterations, with residuals below 10^{-5} for all equations.

The simulation produced three-dimensional fields of temperature, velocity magnitude, and pressure. These fields served as the ground truth for training and validating the physics-informed KNN model. A total of 50,000 sample points were extracted from the CFD solution, with coordinates (x, y, z) and values of temperature and velocity components.

3.2 Physics-Informed KNN Framework

The physics-informed KNN framework modifies the standard KNN algorithm in three ways to incorporate physical constraints.

First, the distance metric is augmented with a physics-based component. In addition to the Euclidean distance in physical space, the algorithm considers the similarity of local flow characteristics. Specifically, the distance between a query point and a training point is computed as:

$$d_{total} = d_{spatial} + \lambda \times d_{physics}$$

where $d_{spatial}$ is the Euclidean distance in the x, y, z coordinate system, $d_{physics}$ is a measure of dissimilarity in local Reynolds number and Rayleigh number, and λ is a weighting coefficient that balances the two contributions. This modification ensures that training points in similar flow regimes are given greater weight in the prediction, improving accuracy in regions where the flow behaviour differs significantly from the average.

Second, the prediction step incorporates a conservation constraint. Standard KNN predicts the value at a query point as the average of the K nearest neighbours. The physics-informed version modifies this average to ensure that the predicted temperature satisfies the local energy balance. Specifically, if the query point is near a heated wall, the predicted temperature is adjusted to be consistent with the wall heat flux and the local thermal conductivity. This prevents the unphysical situation where interpolated temperatures violate the boundary conditions.

Third, the algorithm includes an uncertainty estimate based on the local density of training points and the variability among neighbours. This uncertainty estimate is used to flag regions where the prediction is less reliable and additional sensor data would be beneficial.

3.3 Training and Validation Procedure

The training set was constructed by randomly selecting 5,000 points from the CFD solution, representing a sparse sampling of the full field. These points were treated as "sensor measurements" for the purpose of evaluating the reconstruction accuracy. The remaining 45,000 points served as validation data.

The physics-informed KNN model was trained by storing the training points and their associated field values. No iterative optimization was required, as KNN is a lazy learning algorithm that performs all computation at prediction time. The weighting coefficient λ and the number of neighbours K were treated as hyperparameters and optimized using a separate validation set.

Prediction accuracy was evaluated using three metrics: mean absolute error (MAE), root mean square error (RMSE), and the coefficient of determination (R^2). These metrics were computed for temperature, velocity magnitude, and each velocity component. The computational cost was measured as the wall-clock time required to predict the full field at all validation points.

4. Results and Discussion

4.1 CFD Solution Characteristics

The CFD simulation revealed a buoyancy-driven circulation pattern within the enclosure. Hot fluid rises along the heated wall at 40°C, travels across the top of the enclosure, and descends along the opposite wall, which is at 30°C. This creates a clockwise circulation cell in the vertical plane, with secondary vortices in the corners.

The temperature field exhibits stratification, with warmer fluid accumulating in the upper portion of the enclosure. The temperature difference between the top and bottom of the enclosure is approximately 8 K, with a maximum temperature of 40°C near the hot wall and a minimum of 30°C near the cold wall. The thermal boundary layers on the heated walls have a thickness of approximately 15 mm, as indicated by the steep temperature gradients in these regions.

The velocity field shows maximum velocities of approximately 0.35 m/s in the core of the circulation cell, with lower velocities near the walls due to viscous effects. The flow is turbulent, with Reynolds numbers based on enclosure height and maximum velocity of approximately 12,000. The turbulent kinetic energy is highest in the shear layers between the rising and descending plumes.

These flow characteristics are representative of buoyancy-driven convection in enclosed spaces, which is the dominant heat transfer mechanism in the ATLAS ITk dry nitrogen environment. The simplified enclosure thus provides a relevant test case for the physics-informed KNN framework.

4.2 Reconstruction Accuracy

The physics-informed KNN model achieved high accuracy in reconstructing the temperature and velocity fields from sparse training data. Table 1 summarizes the performance metrics for different values of the hyperparameter K.

Table 1: Reconstruction Accuracy as a Function of K

K	Temperature MAE (K)	Temperature RMSE (K)	Temperature R^2	Velocity MAE (m/s)	Velocity RMSE (m/s)	Velocity R^2
1	4.12	5.87	0.891	0.142	0.198	0.823
2	3.28	4.52	0.934	0.118	0.162	0.881
3	2.56	3.41	0.962	0.094	0.128	0.925
5	2.34	3.12	0.968	0.087	0.119	0.935
7	2.41	3.28	0.965	0.091	0.124	0.929
10	2.78	3.89	0.951	0.108	0.147	0.901
15	3.45	4.87	0.923	0.135	0.184	0.845

The optimal value of K is 5, which minimizes both MAE and RMSE for temperature and velocity. At this value, the temperature MAE is 2.34 K, representing approximately 29% of the maximum temperature difference in the enclosure. The velocity MAE is 0.087 m/s, representing approximately 25% of the maximum velocity.

The coefficient of determination exceeds 0.96 for temperature and 0.93 for velocity at K = 5, indicating that the model captures more than 93% of the variance in the CFD data. This level of accuracy is sufficient for many monitoring and control applications, where the goal is to detect trends and identify anomalous conditions rather than to resolve every detail of the flow field.

The comparison between the physics-informed KNN and standard KNN (without physical constraints) shows a consistent improvement. At K = 5, the physics-informed version reduces the temperature MAE by 18% and the velocity MAE by 22% compared to standard KNN. The improvement is most pronounced near the heated walls, where the conservation constraint prevents the unphysical temperature values that standard KNN occasionally produces.

4.3 Spatial Distribution of Errors

The reconstruction error is not uniform across the enclosure. Figure 1 presents contour plots of the temperature error magnitude in three orthogonal planes through the enclosure.

The largest errors occur in the core of the circulation cell, where the temperature gradient is relatively small and the flow is well-mixed. In these regions, the temperature is close to the mean value, and the absolute error is consequently small. The relative error, however, may be larger because the signal is small.

The smallest errors occur near the walls, particularly in the thermal boundary layers where the temperature gradient is steep. This counterintuitive result reflects the effectiveness of the conservation constraint in the physics-informed KNN. By enforcing consistency with the wall heat flux, the algorithm accurately reproduces the near-wall temperature profiles even when the training data are sparse in these regions.

The velocity error follows a similar pattern, with the largest errors in the low-velocity regions near the corners and the smallest errors in the high-velocity core of the circulation cell. This is consistent with the physics of the flow, as the velocity field is smoother and more predictable in the core than in the complex corner regions.

4.4 Effect of Training Data Density

The performance of the physics-informed KNN model depends on the density of training data. Figure 2 shows the temperature MAE as a function of the number of training points for both the physics-informed and standard KNN algorithms.

For both methods, the error decreases as the number of training points increases, following a power-law relationship with an exponent of approximately -0.6. This is consistent with the theoretical convergence rate for non-parametric regression in three dimensions.

The physics-informed KNN consistently outperforms the standard version across all training densities. The improvement is most significant at low training densities, where the physical constraints compensate for the lack of data. At 1,000 training points, the physics-informed version achieves a temperature MAE of 3.87 K compared to 5.42 K for standard KNN, a 29% improvement. At 5,000 training points, the improvement is 18%, and at 20,000 points, it is 12%.

This result has important implications for sensor placement in practical monitoring systems. The physics-informed approach allows accurate reconstruction with fewer sensors than would be required for purely data-driven methods, reducing installation and maintenance costs.

4.5 Computational Performance

The computational cost of the physics-informed KNN prediction is dominated by the distance calculations between query points and training points. For a training set of N points and a query set of M points, the cost scales as $O(N \times M)$ for the distance computation and $O(N \times M \times K)$ for the neighbour selection.

For the present case study with 5,000 training points and 45,000 validation points, the prediction of the full field required approximately 2.3 seconds on a standard desktop workstation. This compares to approximately 32 minutes for the CFD simulation, representing a speedup of approximately 835 times. The physics-informed version required 15% more computation time than standard KNN due to the additional physics-based distance calculations, but this overhead is negligible compared to the overall speedup.

For real-time monitoring applications, the prediction time can be reduced further by using spatial data structures such as k-d trees to accelerate neighbour searches. With such optimizations, predictions can be performed in milliseconds, enabling control decisions at rates compatible with typical sensor sampling frequencies.

4.6 Comparison with Alternative Methods

To place the performance of physics-informed KNN in context, Table 2 compares it with two alternative approaches: radial basis function (RBF) interpolation and proper orthogonal decomposition (POD) with Kriging.

Table 2: Comparison of Reconstruction Methods

Method	Temperature MAE (K)	Velocity MAE (m/s)	Training Time (s)	Prediction Time (s)
Physics-Informed KNN	2.34	0.087	0.1	2.3
Standard KNN	2.85	0.112	0.1	2.0
RBF Interpolation	2.67	0.098	1.2	3.1
POD-Kriging	2.18	0.081	45.3	0.8

The POD-Kriging method achieves the highest accuracy, with a temperature MAE of 2.18 K and velocity MAE of 0.081 m/s. However, this comes at the cost of a training time of 45 seconds, which is required to compute the POD modes and fit the Kriging models. The physics-

informed KNN is competitive in accuracy while requiring essentially no training time.

The RBF interpolation method falls between the two extremes, with intermediate accuracy and training time. Its prediction time is slightly longer than KNN because of the need to evaluate the radial basis functions at all training points.

The choice of method depends on the specific application requirements. For applications where the training data are fixed and many predictions will be made, POD-Kriging may be preferred for its higher accuracy. For applications where the model must be updated frequently or where training time is constrained, physics-informed KNN offers a favorable trade-off between accuracy and computational cost.

4.7 Sensitivity to Sensor Placement

The spatial distribution of training points significantly affects the reconstruction accuracy. To assess this effect, three sensor placement strategies were compared: uniform random sampling, Latin hypercube sampling, and physics-informed sampling that places sensors in regions of high gradient.

Figure 3 shows the temperature MAE for each strategy as a function of the number of sensors. Latin hypercube sampling outperforms uniform random sampling by approximately 15% for the same number of sensors, reflecting its better coverage of the parameter space. Physics-informed sampling achieves a further 12% improvement by concentrating sensors in the thermal boundary layers where accurate reconstruction is most challenging.

These results suggest that the effectiveness of the physics-informed KNN framework can be enhanced through strategic sensor placement. In practical applications, the optimal placement will depend on the specific geometry and flow conditions, but the general principle of prioritizing regions of high gradient and physical importance applies broadly.

5. Implications for Thermal Management of Particle Detectors

The physics-informed KNN framework developed in this study has direct implications for thermal management of particle detector systems such as the ATLAS Inner Tracker.

During operation, the ITk must be maintained at a dew point below -60°C to prevent condensation on detector surfaces. Humidity sensors provide point measurements at discrete locations, but the dew point can vary significantly across the detector volume due to temperature gradients and local moisture sources. The ability to reconstruct the full dew point field from sparse sensor data would enable more precise control of the dry nitrogen flushing system and earlier detection of leak events.

The simplified enclosure studied here captures the essential physics of the ITk thermal environment, including buoyancy-driven circulation and thermal stratification. The accuracy of the physics-informed KNN reconstruction (temperature MAE of 2.34 K) is sufficient to identify regions where the local dew point approaches the specification limit. The computational speedup (835 times faster than CFD) enables real-time monitoring at rates compatible with sensor sampling.

The framework could be extended to include humidity as an additional scalar field, with the transport equation providing the physical constraint. The extension to the full ITk geometry would require training data from a detailed CFD simulation, but the methodology would remain the same.

6. Conclusions

This investigation has developed and validated a physics-informed K-nearest neighbours framework for thermal-fluid field reconstruction in complex enclosures. The principal findings are as follows:

The physics-informed KNN model achieves a temperature MAE of 2.34 K and velocity MAE of 0.087 m/s when reconstructing fields from 5,000 training points in a buoyancy-driven enclosure. This represents an 18% improvement over standard KNN for temperature and a 22% improvement for velocity.

The optimal number of neighbours is 5 for the present case study, balancing the need for local adaptivity against the desire for noise reduction. Lower values of K perform better in regions of high gradient, while higher values are preferable in well-mixed regions.

The physics-informed approach is most advantageous when training data are sparse. At 1,000 training points, the improvement over standard KNN is 29% for temperature, compared to 12% at 20,000 training points. This makes the method particularly attractive for applications where sensor installation is constrained.

The computational cost of the physics-informed KNN prediction is approximately 2.3 seconds for a field of 45,000 points on a desktop workstation, representing a speedup of 835 times compared to the full CFD simulation. With spatial data structure optimizations, this can be reduced to milliseconds for real-time applications.

The framework has direct applicability to thermal management of particle detector systems, where sparse sensor measurements must be used to monitor temperature and dew point fields that vary significantly across the detector volume. The method provides a practical tool for real-time monitoring and control that preserves physical consistency while achieving the computational efficiency required for online applications.

Future work should extend the framework to include additional physical fields such as humidity and to validate it against experimental data from operating

detector systems. The integration of the physics-informed KNN with model predictive control algorithms for active thermal management represents a promising direction for further research.

References

Bhamjee, M., Connell, M., Connell, S., Igumbor, E., Leeuw, L., Mafa, P., Oriunno, M., & Vreeswijk, M. (2024). Characterisation of the thermoflow due to the dry nitrogen flushing scheme in the ATLAS Inner Tracker using computational fluid dynamics. *Journal of Instrumentation*, *19*(5), P05012. https://doi.org/10.1088/1748-0221/19/05/P05012

Fluent, A. N. S. Y. S. (2023). *ANSYS Fluent theory guide* (Release 2023 R1). ANSYS Inc.

Haywood-Alexander, M., Arcieri, G., Kamariotis, A., & Chatzi, E. (2025). Response estimation and system identification of dynamical systems via physics-informed neural networks. *Advanced Modeling and Simulation in Engineering Sciences*, *12*, 8. https://doi.org/10.1186/s40323-025-00291-9

Kalesh, D., Merembayev, T., Omirbekov, S., & Amanbek, Y. (2025). Application of physics-informed neural networks for two-phase flow model with variable diffusion and experimental validation. *Results in Engineering*, *26*, 105439. https://doi.org/10.1016/j.rineng.2025.105439

Kontolati, K., Goswami, S., Karniadakis, G. E., & Shields, M. D. (2024). Learning nonlinear operators in latent spaces for real-time predictions of complex dynamics in physical systems. *Nature Communications*, *15*, 5101. https://doi.org/10.1038/s41467-024-49411-w

Li, Y. (2024). Physics-informed neural network-based nonlinear model predictive control for automated guided vehicle trajectory tracking. *World Electric Vehicle Journal*, *15*(10), 460. https://doi.org/10.3390/wevj15100460

Moon, H., Lee, S., Demeke, W., Ryu, B., & Ryu, S. (2025). Physics-informed neural operators for generalizable and label-free inference of temperature-dependent thermoelectric properties. *Npj Computational Materials*, *11*, 272. https://doi.org/10.1038/s41524-025-01769-1

Ramadhan, M. F., Taufik, M., Ramandha, M. E. P., Safitri, E., & Anggara, A. (2026). Towards intelligent physics experiments: A physics-informed KNN approach for projectile motion reconstruction using Tracker video analysis in Python. *Jurnal Pendidikan dan Ilmu Fisika*, *6*(1), 1–16. https://doi.org/10.52434/jpif.v6i1.44010

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, *378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Ryck, T. D. S. M. (2024). Numerical analysis of physics-informed neural networks and related models in physics-informed machine learning. *Acta Numerica*, *33*, 633–713. https://doi.org/10.1017/S0962492923000089

Siddiqui, M. S., Durrani, N., & Akhtar, I. (2015). Quantification of the effects of geometric approximations on the performance of a vertical axis wind turbine. *Renewable Energy*, *74*, 661–670. https://doi.org/10.1016/j.renene.2014.08.068

Wu, Y., Sicard, B., & Gadsden, S. A. (2024). Physics-informed machine learning: A comprehensive review on applications in anomaly detection and condition monitoring. *Expert Systems with Applications*, *255*, 124678. https://doi.org/10.1016/j.eswa.2024.124678

Xie, T., Chen, L., Yi, B., Li, S., Leng, Z., Gan, X., & Mei, Z. (2024). Application of the improved K-nearest neighbor-based multi-model ensemble method for runoff prediction. *Water*, *16*(1), 69. https://doi.org/10.3390/w16010069

Yokota, K., Kurahashi, T., & Abe, M. (2024). Physics-informed neural network for acoustic resonance analysis in a one-dimensional acoustic tube. *Journal of the Acoustical Society of America*, *156*(1), 1–15. https://doi.org/10.1121/10.0026459

Zapf, B., Haubner, J., Kuchta, M., Ringstad, G., Eide, P. K., & Mardal, K. A. (2022). Investigating molecular transport in the human brain from MRI with physics-informed neural networks. *Scientific Reports*, *12*, 15475. https://doi.org/10.1038/s41598-022-19157-w

Zhang, W. (2024). A general method for solving differential equations of motion using physics-informed neural networks. *Applied Sciences*, *14*(17), 7694. https://doi.org/10.3390/app14177694

Zou, Z., Meng, X., & Karniadakis, G. E. (2024). Correcting model misspecification in physics-informed neural networks (PINNs). *Journal of Computational Physics*, *505*, 112918. https://doi.org/10.1016/j.jcp.2024.112918