Physics-Enhancing Machine Learning event (PEML2025)

1-3 October 2025
Institute of Physics, London, UK

iop.org

DATA-CENTRIC ENGINEERING

IOP Publishing

IOP Institute of Physics

Welcome to the Physics-Enhancing Machine Learning event 2025!

What are AI for Science, Scientific ML, and Physics-informed ML? Are they really accelerating discoveries and advancing mechanics, materials, physics, climate, and space?!

These are the questions that are addressed during this exciting 3-day PEML25 event (1-3 October 2025, Institute of Physics Headquarters and online) that is part of the activities organised by the Institute of Physics (IOP) and co-sponsored by the journal Data-Centric Engineering (DCE) and the Institute of Physics Publishing (IOPP). PEML25 is led by members of the IOP Applied Mechanics Group special interest group, and supported by the Computational Physics and Materials and Characterization groups.

Our ambitions are to widening participation and to facilitate exchange of knowledge in this exciting and fast-growing research area that is having fast impact also in real-world applications. Thanks to the support of the IOP, DCE and to the outstanding invited speakers (who agreed to contribute to the workshop without a refund for travel expenses), the PEML25 is organised with an incredibly low registration fee for in-person attendance and online participation. Moreover, we were able to offer 17 travel awards to early career researchers to facilitate their in-person participation!

PEML25 features speakers at various stages of their career and from all around the world and cover a broad range of applications. Each day there are outstanding keynote speakers driving the development of methods for enhancing machine learning by embedding physics-knowledge in mechanics and materials (day 1), computational physics (day 2), climate and space (day 3). Undoubtedly, they helped in attracting the overwhelming number of high-quality contributions. This event is building on the experience of the three previous PEML workshops, and features sessions dedicated to real-world applications providing invaluable inputs to current and future challenges on the application of PEML techniques. Moreover, it includes exciting special sessions on the UK strategy relating to AI for Science, the Town Hall for IOP's new AI & Machine Learning Special Interest Group, and Early Career Training sessions.

A big thanks to Fatima Kanu (IOP event manager) for making this event a reality. I would also like to personally thank Marta Varela and Eloisa Bentiveglia, from the Computational Physics IOP group, for their precious help in setting up day 2 of this event, and Andrew Hyde (from Data-Centric Engineering) for his immediate enthusiastic reaction in sponsoring for the 4th time the PEML event!

As of today, we know that about 60 people each day will participate in-person and more than 60 will join online the PEML25. These number show the importance of Physics for enhancing Machine Learning, and Machine Learning for enhancing Physics, both in academia and industry. On behalf of the IOP Applied Mechanics group, I would like to thank each person that has registered to the workshop and will join the exciting discussions in this rapidly evolving field where physics-knowledge is more than ever extremely important!

Dr Alice Cicirello
Assistant Professor in Applied Mechanics at the University of Cambridge
Chair of the workshop and member of the Institute of Physics Applied Mechanics group
25/10/2025

Mechanics and Materials	Wednesday, October 1, 2025
8:45 AM - 9:15 AM	Registration and Coffee
9:15 AM - 9:30 AM	Welcome on behalf of the IOP and structure of the day, **Dr Alice Cicirello, University of Cambridge**
9:30 AM - 10:00 AM	A physics-informed ML recipe book, **Dr Daniel Pitchforth, University of Sheffield**
	Lightning Talks – Session I
10:00 AM - 10:07 PM	System Parameter Identification with Partial Knowledge via Interpretable Hyperparameter Optimisation **Nathan Hinchliffe University Of Sheffield**
10:07 AM - 10:15 AM	A Physics-Informed GAN-ResNet Framework for Imbalanced Rotor Fault Diagnosis, **Yang Fang, Shanghai Jiao Tong University**
10:15 AM - 10:22 AM	Hybrid models for Active Noise Reduction **Matt Bryan, Cambridge University**
10:22 AM - 10:29 AM	Towards Fast and Interpretable Physics-Informed Learning: Second-Order Neurons and Mixed-Activation Networks **João Böger, Dtu**
10:29 AM - 10:36 AM	Vanishing Stacked-Residual PINN for State Reconstruction of Hyperbolic Systems, **Katayoun Eshkofti, Kth Royal Institute Of Technology**
10:36 AM - 10:43 AM	DNN–LSTM-Based Surrogate Modeling and Load Identification for Helicopter Rotor Blades **Ye Wei, Shanghai Jiao Tong University, Shanghai**
10:43 AM - 10:50 AM	Physics-Aware Interpolation of Microstructures with Descriptor-Conditioned Generative Models **Felix Mett**
10:50 AM - 11:20 AM	Coffee Break
	Lightning Talks – Session II
11:21 AM - 11:28 AM	On the use of VAE for ecomaterial selection with applications in structural and aircraft design optimization **Joseph Morlier, Isae-supaero**
11:28 AM - 11:35 AM	Modelling a MDOF Beam with Harmonically Coupled Modes using Lagrangian Neural Networks, **Alan Xavier, Imperial College London**

11:35 AM - 11:42 AM	Multi-Fidelity Neural Network for Predicting Frictional Nonlinear Dynamic System Performance **Peiyu Wang, University Of Southampton**
11:42 AM - 11:49 PM	Physics Guided Neural Surrogates for Microstructure Prediction in Directed Energy Deposition (DED) of Inconel 718 **Alireza Fadavi Boostani, Georgia Institute of Technology**
11:49 AM - 11:56 PM	Physics-Informed Neural Networks for Nonlinear Structural Dynamics **Hassaan Idrees, IMT Lucca**
11:56 AM - 12:03 PM	Self-Adaptive Physics-informed Neural Networks with Reduced Modeling and Basis Updating for Structural Damage Identification **Rui Zhang, ETH zurich**
12:03 PM - 1:00 PM	Lunch and networking
1:00 PM - 1:30 PM	Early career training session IOPP: Choosing the right journal and writing your paper
1:30 PM - 2:15 PM	Keynote, Olga Fink, EPFL – 30 mins + 15 mins for questions **Olga Fink, EPFL**
2:15 PM - 3:00 PM	Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): **Mattia Montanari, PhysicsX**
3:00 PM - 3:30 PM	Coffee & Tea
3:30 PM - 4:15 PM	Keynote, Maarten Schoukens, Eindhoven University of Technology – 30 mins + 15 mins for questions **Maarten Schoukens, Eindhoven University Of Technology**
4:15 PM - 5:00 PM	Real-world applications of PEML (30 minutes + 15 mins for question): Keith Butler, UCL **Dr Keith Butler, Ucl**
5:00 PM - 6:00 PM	Special Session 1

Computational Physics and Applications	Thursday, October 2, 2025
8:45 AM - 9:15 AM	Registration and Coffee 2
9:15 AM - 9:30 AM	Welcome on behalf of the IOP and structure of the day 2, **Dr Eloisa Bentivegna, IBM**
9:30 AM - 10:00 AM	A brief introduction to Physics Enhancing Machine Learning, **Dr Marta Varela, St George's University of London**
	Lightning Talks – Session I day 2
10:00 AM - 10:07 AM	A Comparative Evaluation of Physics-informed Neural Network Architectures for Glucose Modelling in Type 1 Diabetes **Harshith Yerraguntla**
10:07 AM - 10:14 AM	A physics-enabled graph neural network for characterising cardiac electrophysiology from electrode measurements **Annie Ching-En Chiu, Imperial College London**
10:14 AM - 10:21 AM	Bridging PINNs and KANs to Handle Noisy Partial Differential Equations **Siddhi Zanwar, PES University**
10:21 AM - 10:28 AM	Learning Exchange-Correlation Functionals via Differentiable Density Functional Theory **Antonius Freiherr Von Strachwitz, University Of Oxford**
10:28 AM - 10:35 AM	Machine-learning enhanced Density Functional Theory: Learning Exchange-Correlation from Data **Karim Kacper Alaa El-din, University Of Oxford**
10:35 AM - 10:42 AM	openCARP-PINNs: Towards Faster Prediction of Cardiac Signals Propagation **Balvinder Dhillon, Queen Mary University of London**
10:42 AM - 10:49 AM	Emulating CO Line Radiative Transfer with Deep Learning **Shiqi Su, University Of Leicester**
10:50 AM - 11:20 AM	Coffee Break 2
	Lightning Talks – Session II day 2
11:21 AM - 11:28 AM	AutoEmulate: Accelerating large-scale simulations with AI **Marjan Famili, The Alan Turing Institute**
11:28 AM - 11:35 AM	Sequential data assimilation using nudging particle filter

	Maneesh Kumar Singh, Imperial College London
11:35 AM - 11:42 AM	K-space Interpolation using Deep Koopman Autoencoders **Wassim Ben Salah, King's College London**
11:42 AM - 11:49 AM	Solving inverse problems of epoxy hardening with physics-informed neural networks. **Dr. Kateryna Morozovska, KTH Royal Institute of Technology**
11:49 AM - 11:56 AM	A Variational approach to Physics Informed Neural Network for Stochastic Partial Differential Equations **Robbie Slos, UGent**
11:56 AM - 12:03 PM	Transforming physics-informed machine learning to convex optimization (Online) **Letian Yi**, The Hong Kong University of Science and Technology (Guangzhou)
12:03 PM - 1:00 PM	Lunch and networking 2
1:00 PM - 1:30 PM	Early career training session IOPP: Peer-review and publication ethics
1:30 PM - 2:15 PM	Keynote, **Frederik De Ceuster, Leuven Gravity Institute**, KU Leuven – 30 mins + 15 mins for questions
2:15 PM - 3:00 PM	Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): **Carlos Peña Monferrer, SimZero**
3:00 PM - 3:30 PM	Coffee & Tea 2
3:30 PM - 4:15 PM	Keynote, **Payel Das, University of Surrey** – 30 mins + 15 mins for questions
4:15 PM - 5:00 PM	Keynote, **Miles Cranmer, University of Cambridge**, DAMTP & Institute of Astronomy – 30 mins + 15 mins for questions
5:00 PM - 6:00 PM	Special Session **- Daniel Smith, EPRSC**: unpacking the long-term AI for science opportunities

Space and Climate	Friday, October 3, 2025
8:45 AM - 9:15 AM	Registration and Coffee 3
9:15 AM - 9:30 AM	Welcome on behalf of the IOP and structure of the day 3, **Dr Alice Cicirello, University of Cambridge**
9:30 AM - 10:00 AM	A brief introduction to Physics Enhancing Machine Learning, **Dr Alice Cicirello, University of Cambridge**
	Lightning Talks – Session I day 3
10:00 AM - 10:07 AM	A perspective on fluid mechanical environments for challenges in reinforcement learning **Shruti Mishra**
10:07 AM - 10:14 AM	AI-driven Drifter Placement for Ocean Currents **Rui-Yang Zhang,** Lancaster University
10:14 AM - 10:21 AM	Identification of Time-Varying Modal Parameters in Offshore Wind Turbines **Melisa Bozaci University Of Cambridge**
10:21 AM - 10:28 AM	Multi-fidelity learning for physical system predictions **Paolo Conti, The Alan Turing Institute**
10:28 AM - 10:35 AM	Optimizing hp-Variational Physics Informed Neural Networks for real-world applications using the FastVPINNs framework **Divij Ghose,** Imperial College London
10:35 AM - 10:42 AM	Physics-Guided Graph Inference for District Heating Networks **Keivan Faghih Niresi, École Polytechnique Fédérale de Lausanne (EPFL)**
10:42 AM - 10:49 AM	HypeMARL: Multi-agent reinforcement learning for high-dimensional, parametric, and distributed systems **Nicolò Botteghi, Politecnico Di Milano**
10:50 AM - 11:20 AM	Coffee Break (30 minutes) day 3
	Lightning Talks – Session II day 3
11:21 AM - 11:28 AM	Neural Operators for Accelerating Flow Field Predictions **Alan Xavier, Imperial College London**
11:28 AM - 11:35 AM	A Surrogate Modelling Framework for the Correction of Structural Bias in SatelliteCloud Property Retrievals **Iarla Boyce, University Of Cambridge**

11:35 AM - 11:42 AM	Deep reinforcement learning for wall-bounded turbulent flows via wall measurements **Giorgio Maria Cavallazzi, City St. George's, University Of London**
11:42 AM - 11:49 AM	Design Optimisation of Locally Resonant Metamaterials under Uncertainty **Niccolo Klinger, University of Cambridge**
11:49 AM - 11:56 AM	Physics-enhanced Simulation-Based Inference: Likelihood-free MCMC via Normalizing Flows and Variational Autoencoders **Andrea Manzoni, Politecnico Di Milano**
11:56 AM - 12:03 PM	Reduced order modeling with shallow recurrent decoder networks **Matteo Tomasetto, Politecnico Di Milano**
12:03 PM - 1:00 PM	Lunch and networking 3
1:00 PM - 1:30 PM	Early career training session IOPP: maximising visibility and impact of your work
1:30 PM - 2:15 PM	Keynote, **Simon Driscoll, University of Cambridge** – 30 mins + 15 mins for questions
2:15 PM - 3:00 PM	Real-world applications of PEML (30 minutes + 15 mins for question), **Shiva Babu, Rolls-Royce**
3:00 PM - 3:30 PM	Coffee & Tea 3
3:30 PM - 4:15 PM	Keynote, **Anh Khoa Doan, TU Delft**, Faculty of Aerospace Engineering – 30 mins + 15 mins for questions
4:15 PM - 5:00 PM	Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): **Jean Kossaifi, NVIDIA**
5:00 PM - 6:00 PM	Town Hall for IOP's new AI & Machine Learning Special Interest Group

A physics-informed ML recipe book

Dr Daniel Pitchforth[1], Elizabeth Cross[1]
[1]The University of Sheffield, Sheffield, United Kingdom

A physics-informed ML recipe book, Dr Daniel Pitchforth, University of Sheffield, October 1, 2025, 09:30 - 10:00

Our physical understanding of how systems behave is a valuable resource. There are many ways in which we can utilise it to aid the capabilities of a machine learner, offering potential benefits in performance and interpretability. The method of physics-ML integration (recipe) will dictate the capabilities (deliciousness) of the final model and is an important decision to get right. Here, an overview of PIML model structures is presented in order of progression through a typical black-box machine learning pipeline, addressing how physics may be incorporated at each stage. The most appropriate meal to cook depends on many factors; the available ingredients, the palate of your guests and your available kitchen hardware.

System Parameter Identification with Partial Knowledge via Interpretable Hyperparameter Optimisation

Nathan Hinchliffe[1]

[1]University Of Sheffield, Sheffield, United Kingdom

System Parameter Identification with Partial Knowledge via Interpretable Hyperparameter Optimisation, October 1, 2025, 10:00 - 10:07

One of the many benefits of introducing physics into machine learning is the increased interpretability of the resulting models. In some cases, where the system form has been embedded into the model, the hyperparameter values can be interpreted directly as parameter estimates. However, should the assumed model form not accurately represent the true system, and this discrepancy is not accounted for, bias is likely to be introduced into the parameter estimates. In traditional system identification approaches, one common way to capture this discrepancy is with a machine learner in addition to the physical model, but this introduces significant complexity. Instead, it would be preferable to modify an existing model to handle these discrepancies internally.
Gaussian processes are a Bayesian, non-parametric approach to regression, characterised by a mean and a covariance function, and conditioned on data to form posterior predictive distributions. It has been demonstrated previously that a covariance function can be derived from the stochastic differential equation of a physical system, thereby encoding both knowledge of the model form and its parameters. However, any assumptions made during the derivation are also encoded, leading to a similar problem of model discrepancy. However, Gaussian process kernels are well known for their flexibility in kernel design, allowing for complex behaviours to be captured through the combination of simple individual covariance functions.
In this work, a covariance function derived from the single-degree-of-freedom oscillator under white noise loading, characterised by interpretable hyperparameters representing the system's natural frequency and damping ratio, in combination with existing black-box kernels, is optimised to obtain parameter estimates. The hyperparameter optimisation is conducted using signals generated from systems under various unknown load cases, demonstrating that, in some cases, the black-box components in the kernel are able to capture model discrepancies and remove bias from the parameter estimates.

A Physics-Informed GAN-ResNet Framework for Imbalanced Rotor Fault Diagnosis

Yang Fang[1], Ye Wei[1], Xinxing Ma[1], Zhenguo Zhang[1]
[1]Shanghai Jiao Tong University, Shanghai, China

A Physics-Informed GAN-ResNet Framework for Imbalanced Rotor Fault Diagnosis, October 1, 2025, 10:07 - 10:15

Rotor fault diagnosis faces significant challenges due to insufficient fault data, leading to severe data imbalance issues in intelligent fault diagnosis. The current data-level and algorithm-level imbalanced fault diagnosis methods have respective limitations, such as uneven data generation quality and excessive reliance on minority class information. To address this problem, this paper proposes a rotor fault diagnosis method driven by the fusion of mechanism-based data and feature enhancement for small-sample scenarios. Firstly, high-fidelity modeling of four typical rotor faults is performed based on rotor operating mechanisms to generate high-quality simulated fault data. Subsequently, using the simulated fault data as input, Generative Adversarial Networks are employed to enhance the feature quality of the simulated fault data. The enhanced dataset is then used to train a Residual Network (ResNet) model, enabling effective intelligent fault diagnosis. Experimental validation at rotational speeds of 2100 rpm and 2300 rpm demonstrated that the proposed approach significantly improved fault recognition accuracy, verifying its effectiveness and robustness in scenarios characterized by limited and imbalanced fault samples. This method provides a novel solution for small-sample intelligent fault diagnosis of modern equipment.

Hybrid models for Active Noise Reduction

Matt Bryan[1], Ole Mattis Nielsen[1,2], Tore Butlin
[1]Cambridge University Engineering Department, Cambridge, United Kingdom, [2]Bose Corporation , , United States

Hybrid models for Active Noise Reduction, October 1, 2025, 10:15 - 10:22

Active Noise Reduction (ANR) presents a challenging prediction task, especially in systems with nonlinearity; models must have both high accuracy and low latency for performance on embedded hardware. Purely physics-based modelling of the vibration pathway often proves insufficient, and data-driven methods are typically inefficient, and so hybrid models might show a performance benefit. The aim in particular is to separate the linear and nonlinear behaviour of a system: the former can be solved using a linear physics model (e.g. the governing ODE), and the latter can be the focus of a data-driven model, aiming to reduce the overall computation. However, separation is not always trivial, and is dependent on the coupling of the constituent elements.
To that extent, hybrid physics/machine-learning models are presented, taking inspiration from implicit numerical schemes. Then their performance on a nonlinear benchmark model (a linear system driven through a Duffing spring to represent an automotive suspension system) is analysed. The behaviour of models is characterised, and their performance is assessed using a scaling law argument, drawing comparison with purely data-driven models of the same system.

Towards Fast and Interpretable Physics-Informed Learning: Second-Order Neurons and Mixed-Activation Networks

João Böger[1]
[1]DTU, Copenhagen, Denmark

Towards Fast and Interpretable Physics-Informed Learning: Second-Order Neurons and Mixed-Activation Networks, October 1, 2025, 10:22 - 10:29

Complex simulators are central to scientific research, forecasting, and real-world applications. However, they often require intensive computational resources and suffer from scalability issues — challenges amplified in the big data era. The APEX project addresses these limitations by designing efficient architectures for scientific simulators, exploring inductive biases, causal relations, and architectural innovations that learn and generalize simulator dynamics in a fast, scalable, uncertainty-aware, and interpretable way.

Our goal is to build surrogate models (metamodels) that generalize beyond training data, even when conditions change—a common challenge where standard machine learning (ML) models often fail. Many simulators, like those in transport, climate, epidemiology, and hydrodynamics, are governed by ODEs/PDEs. Efficiently learning these dynamics is crucial. To improve over traditional ML, we introduce inductive biases based on known system constraints.

Quadratic neurons and domain-specific activation functions (MixFunn) [1], motivated by the analytic forms of differential equation solutions, integrate domain insights into network architectures. Combined with soft constraints in the loss function, as in Physics-Informed Neural Networks (PINNs), this design matches or surpasses PINNs with fewer parameters, enhancing efficiency and enabling closed-form solutions.

We hypothesize that polynomial neurons and tailored non-linearities are key to scalable, generalizable simulator models. We evaluate MixFunn, PINNs, neuralODEs, and Equation Learners across four ODE-based benchmarks: the SIR model, Lotka–Volterra system, Duffing oscillator, and Van der Pol oscillator. Two recent metrics [2] assess out-of-distribution (OOD) learning and robustness to noise.

This work highlights how architectural choices and domain-informed priors can close the gap between robustness, speed, and generalization, paving the way for compact, interpretable, and efficient scientific surrogates.

[1] arXiv:2503.22528, 2025

[2] arXiv:2402.18377, 2024

Vanishing Stacked-Residual PINN for State Reconstruction of Hyperbolic Systems

Katayoun Eshkofti[1], Matthieu Barreau[1]
[1]Division of Decision And Control Systems, Digital Futures, Kth Royal Institute of Technology, Stockholm, Sweden

Vanishing Stacked-Residual PINN for State Reconstruction of Hyperbolic Systems, October 1, 2025, 10:29 - 10:36

State reconstruction for systems governed by quasi-linear hyperbolic partial differential equations (PDEs) is significantly challenging due to discontinuities and shock formations, which violate the smoothness assumptions that typically enable physics-informed neural networks (PINNs) [1] to be effective. Moreover, to the best of our knowledge, no study has demonstrated that PINNs can effectively learn hyperbolic PDEs underlying traffic state models with acceptable accuracy. This provides a strong motivation to develop a more effective state estimation method for hyperbolic PDEs.

To address these issues, we propose the vanishing stacked-residual PINN [2], which incorporates prior knowledge in the form of a function series. This approach combines the vanishing viscosity method from applied mathematics, curriculum learning from machine learning, and stacked PINNs [3] in an effective manner. The methodology ensures convergence to the unique entropic solution of the hyperbolic PDE.

As illustrated in Fig. 1, training begins with a baseline PINN that solves a parabolic regularization of the main PDE using a sufficiently large viscosity coefficient. This guarantees that the approximated solution is entropic and that the PDE operator is sufficiently Lipschitz-continuous, allowing the PINN to learn the regularized solution effectively. The resulting low-fidelity solution is then fed into and refined by a sequence of residual blocks, each trained with a successively smaller viscosity coefficient. This forms a multi-stage correction process inspired by the vanishing viscosity method. Applied to traffic state reconstruction using the Lighthill-Whitham-Richards (LWR) model with Greenshield's flux, this method yields an order-of-magnitude improvement in accuracy over the vanilla PINN [1]. In addition to enhanced accuracy, the residual correction networks also stabilize learning, leading to lower errors and reduced variability in state estimations.

References

[1] M. Raissi, P. Perdikaris, and G.E. Karniadakis. "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: Journal of Computational Physics 378 (2019), pp. 686–707.

[2] Katayoun Eshkofti and Matthieu Barreau. "Vanishing Stacked-Residual PINN for State Reconstruction of Hyperbolic Systems". In: IEEE Control Systems Letters (2025).

[3] Amanda A. Howard, Sarah H. Murphy, and al. "Stacked networks improve physics-informed training: Applications to neural networks and deep operator networks". In: Foundations of Data Science (2025).

DNN–LSTM-Based Surrogate Modeling and Load Identification for Helicopter Rotor Blades

Ye Wei[1], Yang Fang[1], Xinxing Ma[1], Zhenguo Zhang[1]
[1]Shanghai Jiao Tong University, Shanghai, China, Shanghai, China

DNN–LSTM-Based Surrogate Modeling and Load Identification for Helicopter Rotor Blades, October 1, 2025, 10:36 - 10:43

A high-fidelity structural model, together with the reliable identification of distributed loads, is crucial for enhancing flight performance and ensuring operational safety. Although various techniques have been developed for model updating and load reconstruction, their effectiveness often depends on prior knowledge of the system and requires meticulous tuning of algorithmic parameters. In this study, a novel neural network framework is proposed to construct surrogate models for both forward and inverse problems of a helicopter blade. Specifically, the rotor blade model is first updated using a Deep Neural Network (DNN) based on modal information. Subsequently, the time-varying distributed loads are identified by a Long Short-Term Memory (LSTM) network trained on the measured bending moment responses. Application to the benchmark XH-59A rotor blade model showcases the method's effectiveness in refining structural models and accurately reconstructing distributed loads.

Physics-Aware Interpolation of Microstructures with Descriptor-Conditioned Generative Models

F. Mett[1,*)] and **M. Beer**[1,2,3)]

[1)]Institute for Risk and Reliability, Leibniz University Hannover, Callinstr. 34, 30167 Hannover, Germany.

[2)]Institute for Risk and Uncertainty, University of Liverpool, Liverpool L69 7ZL, 610101 Liverpool, United Kingdom.

[3)]International Joint Research Center for Engineering Reliability and Stochastic Mechanics, Tongji University, Siping road 1239, 200092 Shanghai, China.

[*)]Corresponding author, mett@irz.uni-hannover.de

Abstract

Material behaviour is strongly determined by its microstructure, which can change in response to processing parameters and evolution mechanisms. However, these dependencies are often only partially understood and expensive to capture through direct simulations or experiments. Generative artificial intelligence has recently shown great promise in producing synthetic microstructures that resemble training data, thereby reducing experimental effort. Previous work has demonstrated that generative models conditioned on processing parameters can interpolate between microstructure distributions when only discrete measurements at selected parameters are available [Tang, 2021], underscoring their potential to reduce experimental effort. Separately, conditioning on physical descriptors has been proposed as a means of achieving controlled microstructure generation [Kishimoto, 2023].

In this work, we combine these two perspectives and investigate physics-aware interpolation using Generative Adversarial Networks [Goodfellow, 2014] conditioned on the phase volume fraction, a low-order physical descriptor of the microstructure. To guide the generator, we introduce an additional loss term that explicitly communicates the physical meaning of this quantity. This setup allows the model to generate microstructures consistent with prescribed phase volume fraction values. Because these values are easily interpretable, they form a natural space in which interpolation can be carried out in a meaningful way. By supplying interpolated phase volume fractions to the trained generator, we obtain synthetic microstructures at unmeasured processing parameters or evolution steps, effectively letting the generator perform the interpolation while respecting expected physical trends.

We demonstrate this approach on a synthetically generated set of microstructures with prescribed and known nonlinear evolution. The results show how interpolation in the descriptor space can generate transitions between microstructure distributions, while also revealing the method's limitations. These findings highlight descriptor-conditioned generative models as a promising tool for physics-informed interpolation across microstructure evolution scenarios, potentially reducing the need for exhaustive experimental sampling.

References

J. Tang, X. Geng, D. Li, et al. Machine Learning-Based Microstructure Prediction during Laser Sintering of Alumina. *Scientific Reports*, 11(1):10724, 2021.

M. Kishimoto, Y. Matsui, H. Iwai. Conditional Generative Adversarial Network for Generation of Three-Dimensional Porous Structure of Solid Oxide Fuel Cell Anodes with Controlled Volume Fractions. *Journal of Power Sources*, 580:233411, 2023

I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, et al. Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 27, 2014

On the use of VAE for ecomaterial selection with applications in structural and aircraft design optimization

Joseph Morlier[1], Shantanu Sapre[1], Almudena Cobo-Urios[1], Alvaro Silva-Vilela-Caridade[1]
[1]Isae-supaero, Toulouse, France

On the use of VAE for ecomaterial selection with applications in structural and aircraft design optimization, October 1, 2025, 11:21 - 11:28

This work presents the use of variational autoencoders and mixed-variable solvers as a proposal for structural optimization and material selection. This approach allows for gradient-based, multi-objective optimization of mechanical and environmental trade-offs (cost, energy, waste, density, young modulus, and yield stress). Optimal material solutions are identified through the analysis of Pareto fronts, revealing hybrid, potentially non-existent materials that offer superior trade-offs across conflicting objectives. The proposed methodology integrates machine learning and design optimization to support sustainable aerospace innovation, and the selected material is later reintegrated into the Aircraft Design model to assess their impact on aircraft performance. This comprehensive framework bridges data-driven design with engineering constraints, offering a novel pathway for eco-efficient material selection in early-stage aircraft design. In the MOO context we adopt the Multi-Gradient Descent Algorithm (MGDA) to perform multi-objective optimization directly within the learned latent space

This work is an extension of this work (Yepes Llorente, L., Morlier, J., Sridhara, S., & Suresh, K. (2024). A hybrid machine learning and evolutionary approach to material selection and design optimization for eco-friendly structures. Structural and Multidisciplinary Optimization, 67(5), 69.) applied to aircraft design

Modelling a MDOF Beam with Harmonically Coupled Modes using Lagrangian Neural Networks

Alan Xavier[1], Ludovic Renson[1]
[1]Imperial College London, London, United Kingdom

Modelling a MDOF Beam with Harmonically Coupled Modes using Lagrangian Neural Networks,
October 1, 2025, 11:28 - 11:35

Accurately modelling the dynamic behaviour of nonlinear structures is challenging due to the wide range of potential nonlinearities and dynamic phenomena that they can exhibit. Physics-guided machine learning (PGML) has emerged as an attractive way to combine prior knowledge with data to solve a wide range of complex, nonlinear problems in science and engineering. Lagrangian Neural Networks (LNNs) are a particular PGML approach that models nonlinear systems' Lagrangian functions using artificial neural networks (NNs). The Euler-Lagrange equation is then reconstructed through automatic differentiation (AD) to derive the equations of motion, enforcing physical consistency during training.

Previously, we used LNNs to model the nonlinear vibrations of a single-degree-of-freedom (SDOF) Duffing oscillator. We extend our earlier works to address systems with multiple-degrees-of-freedoms (MDOF) and demonstrate the method on a beam structure exhibiting isolated responses. We analyse the physical consistency of the trained model and interpret the identified stiffness and damping nonlinearities from the partial derivatives of the potential energy and dissipation functions. The trained LNNs are then used to trace frequency response and bifurcation curves, which can be directly compared to curves measured with control-based continuation.

Multi-Fidelity Neural Network for Predicting Frictional Nonlinear Dynamic System Performance

Peiyu Wang[1], David Toal[1], Jie Yuan[1]
[1]Computational Engineering and Design Group, University of Southampton, Southampton, United Kingdom

Multi-Fidelity Neural Network for Predicting Frictional Nonlinear Dynamic System Performance, October 1, 2025, 11:35 - 11:42

Accurate analysis of amplitude-dependent nonlinear behaviour is essential for the design and optimisation of frictional nonlinear dynamical systems. However, repeated analyses during the design process often demand substantial computational resources. Surrogate modelling offers an effective alternative to accelerate this process. While high-fidelity numerical methods, such as the Harmonic Balance Method (HBM) enhanced by the Extended Periodic Motion Concept (EPMC) for nonlinear modal analysis, are effective in capturing amplitude-dependent responses, their computational cost escalates with the demand for higher accuracy, creating a trade-off between accuracy and efficiency. To address this challenge, a multi-fidelity neural network (MF-NN) is developed to integrate data from different fidelity levels, reducing the sampling cost while maintaining model accuracy. In this work, the MF-NN is applied to predict nonlinear modal characteristics, efficiently capturing the amplitude-dependent natural frequencies and damping ratios. Two case studies are conducted: first, a two-degree-of-freedom system with friction contact is examined to evaluate MF-NN performance under varying ratios of high- and low-fidelity samples, with results benchmarked against a Co-Kriging model. Second, a high-dimensional bladed-disk system with a ring damper is investigated to assess the approach in a practical engineering context. The results demonstrate that the MF-NN effectively integrates multi-fidelity data, achieving accurate predictions of dynamic performance with reduced computational effort compared to single-fidelity surrogate models. This research highlights the potential of multi-fidelity surrogate modelling as a robust and efficient tool for nonlinear dynamic analysis of rotating structures, enabling improved computational efficiency in practical applications.

Physics Guided Neural Surrogates for Microstructure Prediction in Directed Energy Deposition (DED) of Inconel 718

Alireza Fadavi Boostani, Hamid Garmestani
[1]George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, GA, USA., , United States, [2]School of Materials Science and Engineering, Georgia Institute of Technology, GA, USA., , United States

Physics Guided Neural Surrogates for Microstructure Prediction in Directed Energy Deposition (DED) of Inconel 718, October 1, 2025, 11:42 - 23:49

Additive manufacturing processes such as Directed Energy Deposition (DED) involve complex, coupled thermal–microstructural dynamics that challenge conventional numerical modeling in terms of accuracy and computational cost. Physics-Informed Neural Networks (PINNs) offer a promising alternative: they are mesh-free, enforce governing PDEs directly, and can learn from sparse experimental data, making them well-suited for in-process monitoring and control in DED.
This study introduces a Physics-Informed Neural Network (PINN) framework coupled with a simplified Kobayashi phase-field model to predict microstructural evolution during Directed Energy Deposition (DED) of Inconel 718, integrating experimentally sourced grain-size data (~104 µm average) for training. Governing thermal and grain-growth PDEs are embedded into the PINN loss function, alongside real measurement values. Importantly, all data within a 4 mm radius of the melt pool are held out for validation, enabling rigorous evaluation of model generalization.
Despite excluding melt-pool region data during training, the PINN accurately predicts temperature and grain-size distributions there, thanks to embedded physical laws—a capability supported by recent literature on PINN extrapolation into unseen domains. Benchmarking against finite element method (FEM) reveals that, for steady state 2D heat equations, FEM typically solves to 10^{-5} accuracy. However, once trained, the PINN delivers near-instantaneous, mesh-free inference—offering 10×–20× faster evaluation of new scenarios than traditional FEM interpolation.
Our optimized model, trained using the Adam algorithm with a composite loss (heat PDE residuals + phase-field residuals + experimental data loss), achieves $R^2 > 0.9$ on validation data. These results demonstrate that our PINN+phase-field framework provides a physically consistent, computationally efficient, and experimentally validated solution for microstructure modeling in DED processes. The method's rapid inference capability and flexibility position it as a powerful tool for in-situ microstructural control and process optimization in additive manufacturing of superalloys.

Physics-Informed Neural Networks for Nonlinear Structural Dynamics

Hassaan Idrees[1], Marco Paggi[1], Pattabhi Ramaiah Budarapu[2]
[1]IMT School of Advanced Studies, Lucca, Italy, [2]Indian Institute of Technology Bhubaneswar, , India
Physics-Informed Neural Networks for Nonlinear Structural Dynamics, October 1, 2025, 11:49 - 23:56

Physics Informed Neural Networks (PINNs) have emerged as a versatile framework for solving complex nonlinear dynamical systems. It offers a novel approach to modelling in scenarios where classical solvers struggle with strong nonlinearity, limited data, or the need for robust, interpretable predictions. In this work we investigate these challenges using PINNs over two nonlinear systems including Duffing oscillator and a belt-driven motor system with quadratic damping and cubic stiffness. To this end, we explored two key strategies: (1) ansatz-based trial solution, and (2) adaptive weighting of the loss function. The ansatz-based approach which embedded the initial condition analytically in the network output reduced the training cost by 5x in terms of simulation time cost and provided stable predictions across time windows which covered nearly the entire nonlinear transient regime. Our approach yielded comparative results as compared to Newmark-Beta benchmark with considerable error. Our study demonstrates that targeted architectural and training innovations can substantially mitigate spectral bias and improve the efficiency of PINNs for complex dynamic systems. Ongoing work explores dynamic loss weighting, while future directions will integrate residual-based adaptive sampling and curriculum learning to further automate and accelerate convergence for real-world, data-scarce nonlinear systems.

Self-Adaptive Physics-informed Neural Networks with Reduced Modeling and Basis Updating for Structural Damage Identification

Rui Zhang[1], Konstantinos Vlachas[1], Eleni Chatzi[1]
[1]ETH zurich, Zurich, Switzerland

Self-Adaptive Physics-informed Neural Networks with Reduced Modeling and Basis Updating for Structural Damage Identification, October 1, 2025, 11:56 - 12:03

Physics-Informed Neural Networks (PINNs) have emerged as a powerful tool for solving inverse problems in structural dynamics by embedding physical laws into neural network training. However, applying PINNs to real-world structural health monitoring remains challenging due to high computational demands and difficulties in balancing physics constraints with sparse, noisy measurements.

To overcome this challenge, this work proposes a hybrid framework that integrates PINNs with reduced-order modeling (ROM) to enable scalable and efficient damage identification. Structural dynamics are projected onto a low-dimensional subspace using a reduced basis (e.g., via Proper Orthogonal Decomposition), allowing predictions in the reduced coordinate space, thus improving efficiency prior to recovering full-field responses. A key contribution is the implementation of basis updating strategies that ensure consistency between the reduced basis and the evolving system stiffness affected by damage. This includes both online basis updating—using interpolation or clustering among precomputed bases during training—and sequential retraining with transfer learning, in which a PINN trained on the healthy-state basis is retrained after parameter updates using previously learned weights to accelerate convergence. Additionally, a self-adaptive loss weighting strategy based on Neural Tangent Kernel (NTK) theory is introduced to dynamically balance the influence of data and physics constraints during training. This improves convergence behavior and robustness under varying levels of noise and measurement sparsity.

The proposed framework is validated using the IASC-ASCE benchmark structure, which simulates realistic structural behavior of a lab-scale steel frame under various damage scenarios, sensor configurations, and modeling uncertainties. Results show accurate identification of damage location and severity, and high-fidelity reconstruction of structural response, even with incomplete and noisy data. By combining physics-informed learning with reduced modeling, basis adaptation, and adaptive training dynamics, the framework offers a scalable and interpretable solution to structural identification—advancing the real-world applicability of PEML techniques.

From Physics to Machine Learning and Back: Applications in Dynamical Systems

Olga Fink

Keynote, Olga Fink, EPFL – 30 mins + 15 mins for questions, October 1, 2025, 13:30 - 14:15

Deep learning has become an essential tool in many engineering applications. However, its effectiveness is often limited by its reliance on large, representative, and well-labeled datasets. In contrast, condition monitoring data from complex systems is typically sparse, unlabeled, and unrepresentative, making it difficult to apply purely data-driven methods effectively. Moreover, deep learning models often perform poorly in extrapolation scenarios—common in engineering systems with long service lives and evolving operational regimes.

To address these limitations, integrating physical laws and domain knowledge into deep learning frameworks has shown significant potential. This presentation will explore a range of approaches that integrate physics-based principles into machine learning models. Particular attention will be given to the use of structural inductive biases—such as those introduced by physics-informed graph neural networks—to improve model robustness, generalization and extrapolation.

Finally, the talk will examine emerging methods in symbolic regression that aim to close the loop between data-driven learning and physical understanding, enabling the discovery of interpretable, physics-consistent models from data.

When 99% Accuracy Means 100% Engineering Failure

Mattia Montanari[1]
[1]PhysicsX, , United Kingdom, [2]University of Oxford, , United Kingdom

Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): Mattia Montanari, PhysicsX, October 1, 2025, 14:15 - 15:00

In engineering, a 99% accurate AI model can mean 100% failure—not just mechanical disasters, but models gathering dust because engineers don't trust them, can't integrate them, or find them irrelevant to real-world constraints. Three decades after transitioning from physical prototypes to virtual simulations, we're entering a second revolutionary transformation with AI-driven Large Physics Models. At PhysicsX, we've learned that making physics-enhanced machine learning (PEML) truly useful requires more than impressive accuracy—it demands models that engineers actually use, trust, and seamlessly integrate into their workflows.

This talk explores real-world PEML applications in the aerospace sector, revealing how physics constraints transform theoretically useful models into practically useable tools. I'll share hard-won insights from deploying multi-physics surrogates that operate in real-time, enabling automated optimization over vast design spaces previously computationally intractable. The key to solving engineering's toughest challenges—from bridging the gap between ML researchers and domain engineers to quantifying uncertainties while ensuring physical consistency—lies in creating an interoperable ecosystem where specialized AI agents collaborate seamlessly. I'll demonstrate this through our agentic workflows powering foundation models: Large Language Models (LLMs) that understand engineering intent and constraints, Large Geometry Models (LGMs) that generate and manipulate complex 3D designs, and Large Physics Models (LPMs) that predict physical behavior across multiple scales and phenomena. By combining foundation models with physics on a modular platform, we are turning yesterday's 100% failure into tomorrow's engineering breakthroughs.

Improving Baseline Models through Model Augmentation: Model Structures and Learning Approaches

Maarten Schoukens[1]
[1]Eindhoven University Of Technology, Eindhoven , Netherlands

Keynote, Maarten Schoukens, Eindhoven University of Technology – 30 mins + 15 mins for questions,
October 1, 2025, 15:30 - 16:15

In most engineering applications, strong prior knowledge is present in the form of pre-existing models, provided by system designers and engineers, even though they do not capture all the nonlinear dynamics of the real-life system. These models are currently not accounted for during black-box system identification / data-driven modelling tasks.
We aim to develop a comprehensive data-driven modelling framework to obtain accurate and interpretable models of measured complex system dynamics by augmenting an approximate pre-existing model through black-box nonlinear system identification. During this talk we will explore new theory and algorithms to 1) provide model structures, algorithms and theory that flexibly interconnect the pre-existing model and the black-box completion 2) study the interpretability of augmented models.

Learning the Language of Crystal Chemistry: Using Concepts from Natural Language to Model Solid State Chemistry

Dr Keith Butler

Real-world applications of PEML (30 minutes + 15 mins for question): Keith Butler, UCL, October 1, 2025, 16:15 - 17:00

The discovery and design of new materials is critical for advancing carbon-emission reducing technologies such as renewable energy and electric vehicles. Experimental discovery of new materials is typically slow and costly, quantum mechanics (QM) calculations have brought computational materials design within reach. However, QM calculations are often limited to relatively small sets of materials, as their computational costs are too great for large-scale screening, this is the case for calculating properties required for new energy materials. In this talk I will present examples of how we have been adapting concepts from language models to help with building fast and efficient models for materials properties. I will show how we can learn distributed representations of atomic species directly from large databases of crystallographic structures [1, 2]. I will also show how a large language model trained on crystallographic data can help to solve one of the outstanding challenges of solid-state chemistry; the prediction of structure from chemical formula [3].

[1] npj Comput Mater 8, 44 (2022)

[2] APL Machine Learning 2, (2024)

A Comparative Evaluation of Physics-informed Neural Network Architectures for Glucose Modelling in Type 1 Diabetes

Harshith Yerraguntla[1], Thomas Plant[1], Lukas Schuster[2], Sufyan Hussain[3,4,5], Marta Varela[6,7]
[1]Department of Bioengineering, Imperial College London, , United Kingdom, [2]Syntactiq, , Austria, [3]Department of Diabetes and Endocrinology, Guy's and St Thomas' Hospital NHS Trust, London, , United Kingdom, [4]Department of Diabetes, King's College London, , United Kingdom, [5]Institute of Diabetes, Endocrinology and Obesity, King's Health Partners, London, , United Kingdom, [6]Department of Medicine, City St George's University of London, , United Kingdom, [7]National Heart & Lungs Institute, Imperial College London, , United Kingdom

A Comparative Evaluation of Physics-informed Neural Network Architectures for Glucose Modelling in Type 1 Diabetes, October 2, 2025, 10:00 - 10:07

Introduction

Type 1 diabetes (T1D) affects over 8.4 million people worldwide. As the pancreas of T1D patients cannot synthesise insulin, they require lifelong administration of this hormone to maintain healthy blood glucose levels. Recently, closed-loop artificial pancreas systems (APS), which integrate continuous glucose monitors (GCM), insulin pumps, and electronic control algorithms, have emerged as promising technologies to reduce the burden of insulin management and improve outcomes in T1D.

Despite their promise, fully automated insulin dosing remains elusive, partly due to difficulty forecasting blood glucose on a per-patient basis. Physics-enhanced machine learning offers a compelling solution, enabling both glucose prediction and estimation of physiological parameters related to insulin dynamics.

Methods

We propose a Physics-Informed Neural Network (PINN) approach that embeds the 2015 Magdelaine glucose-insulin-carbohydrate model into a deep learning framework via differentiable residual constraints. The PINN is trained exclusively on glucose trajectories, with insulin and carbohydrate inputs supplied but unsupervised. To evaluate trade-offs, we benchmark three PINN architectures: (a) a vanilla multilayer perceptron (MLP), (b) a modified MLP with encoder-mixing and gating, and (c) a gated recurrent unit (GRU)-based model. All models are assessed on 10 synthetic and 5 real-world patient CGM datasets using two criteria: blood glucose forecasting accuracy and, for synthetic data, insulin sensitivity estimation accuracy.

Results & Discussion

After individually optimising hyperparameters, we observe a trade-off between glucose prediction and physiological parameter estimation (see Figure). The MLP PINN (a) yields the most accurate parameter estimates but higher forecasting error. The GRU PINN (c) excels at prediction but fails to recover robust physiological states. The modified MLP (b) shows greater output variance and limited biological plausibility.

Future work will focus on identifying which PINN architectural trade-offs best support safe and effective insulin dosing under clinical constraints

A physics-enabled graph neural network for characterising cardiac electrophysiology from electrode measurements

Annie Ching-En Chiu[1], Danilo Mandic[1], Marta Varela[1,2]
[1]Imperial College London, , United Kingdom, [2]City St George's University of London, , United Kingdom

A physics-enabled graph neural network for characterising cardiac electrophysiology from electrode measurements, October 2, 2025, 10:07 - 10:14

Cardiac arrhythmias—irregular heart rhythms caused by abnormal generation or propagation of electrical signals—affect millions worldwide. Most often, treatment strategies rely on probing extracellular electrical potentials inside the heart and destroying cardiac regions adjacent to empirically-defined abnormal electrical signals. Targeting precise abnormal regions would improve the success of these procedures and make them shorter and safer.

This project aims to characterise the electrophysiological properties of cardiac tissue using physics-enabled graph neural networks (GNN).

We train a graph neural network (GNN) on physical simulations so that it can convert extracellular potentials measured by a clinical device to transmembrane potentials. These are not measurable clinically but enable the characterisation of local electrophysiological properties to allow better targeting of abnormal cardiac regions.

Our GNN model naturally mirrors the biophysics of cardiac extracellular potentials in the following sense. There are two types of nodes: 1) those representing the electrodes with extracellular potentials and 2) those representing points where transmembrane potentials are known. We connect nodes from one type to the other if their Euclidean separation falls below a predefined threshold and assign edge weights inversely related to that distance. This construction mimics how an electrode "sees" the tissue -- extracellular potentials are generated by a weighted summation of nearby electric fields, with closer sources contributing more strongly. By embedding the electrode-tissue coupling directly into the network's architecture, our approach paves way for physics informed guidance of arrhythmia diagnosis and treatments.

The transmembrane potentials retrieved achieved a root-mean-square error of 0.43 on test nodes. With the transmembrane potentials, we can then compute parameters such as local tissue conductivity and excitation thresholds using physics-informed neural networks or other inverse parameter estimation schemes.

Bridging PINNs and KANs to Handle Noisy Partial Differential Equations

Shusrith Srikanth[1], Siddhi Zanwar[1], Siddhi Zanwar[1], Bhaskarjyoti Das[1]
[1]PES University, Bangalore, India
Bridging PINNs and KANs to Handle Noisy Partial Differential Equations, October 2, 2025, 10:14 - 10:21

Physics-Informed Neural Networks solve partial differential equations by embedding physical laws into neural architectures but suffer significant noise sensitivity that leads to degraded performance and physically inconsistent predictions on real applications.
We introduce a hybrid physics-data framework resolving PINN's noise sensitivity through the incorporation of Kolmogorov-Arnold Networks (KANs) as preprocessing denoising modules. KANs utilize learned adaptive activation functions with inherent filtering of high-frequency noise and conservation of physical dynamics. Denoised inputs are subsequently fed to PINNs to solve forward PDEs, forming a resilient pipeline preserving physical consistency under noisy scenarios.

We assess our framework on Burgers', Heat, and Wave equations with skewed normal noise simulating real-world asymmetric error distributions. Results indicate impressive MSE reductions: Burgers' equation goes from 0.047874 (noisy) to 0.016759 (KAN-processed), nearly achieving clean performance (0.016251). Heat and Wave equations indicate comparable improvements, with Heat equation displaying spectacular recovery from extreme noise-induced degradation.

Our hybrid method illustrates Physics-Enhanced Machine Learning through the fusion of domain expertise (physical principles) and adaptive data processing (noise elimination), and the resulting models are physically sound and resilient to real-world imperfections. The framework solves core issues: prevailing over poor generalization and physically unrealistic forecasts while being interpretive through physics-inspired constraints.
Results show structure-dependent noise sensitivity among PDE types, offering physics-informed guidance for focused enhancements in scientific computing applications where resilience to noise and physical consistency are critical.

Learning Exchange-Correlation Functionals via Differentiable Density Functional Theory

Antonius Freiherr Von Strachwitz[1], Karim Alaa El-Din[1], Sam Vinko[1]
[1]University Of Oxford, Oxford, United Kingdom

Learning Exchange-Correlation Functionals via Differentiable Density Functional Theory, October 2, 2025, 10:21 - 10:28

Density Functional Theory (DFT) is a widely used method in quantum chemistry and materials science, where the ground-state energy of a many-electron system is determined from its electron density. At the core of DFT lies the exchange-correlation (XC) functional, whose exact form is unknown and typically approximated through empirical or semi-empirical models. Recent efforts in machine learning have aimed to improve these approximations by learning XC functionals from data, often relying on large datasets of spatially resolved densities.

In this talk, we present a physics-enhanced machine learning strategy to reconstruct the exchange-correlation functional using only sparse scalar energy data. Our approach is based on a differentiable implementation of Kohn-Sham DFT, where the self-consistent cycle is made fully differentiable and used to implicitly regularize the learning process. We train neural network functionals to recover the Perdew-Wang local density approximation (PW-LDA) solely from synthetic energy data, without accessing electron densities during training.

This method exploits the structure of the underlying physics to constrain and guide learning, resulting in accurate and physically consistent models from minimal data. By circumventing the need for high-dimensional density inputs, this framework opens new opportunities for learning from experimental observables and domain-specific data, where access to full electronic structure is limited. Our results illustrate the power of differentiable programming in embedding first-principles constraints into ML-driven model discovery.

Machine-learning enhanced Density Functional Theory: Learning Exchange-Correlation from Data

Karim Kacper Alaa El-din[1], Antonius von Strachwitz[1], Sam M. Vinko[1]
[1]University Of Oxford, Oxford, United Kingdom

Machine-learning enhanced Density Functional Theory: Learning Exchange-Correlation from Data,
October 2, 2025, 10:28 - 10:35

In density functional theory, simpler exchange-correlation (XC) approximations such as the local density approximation (LDA) are favoured for computational speed but entail a loss of information, leading to a trade-off between accuracy and generality. Here, we train a neural LDA on gold standard data using a differentiable Kohn–Sham solver, imparting system-specific expertise for water and sacrificing generality for accuracy. We demonstrate how this model achieves previously unseen accuracy on the training domain, and discuss how it generalizes to other domains. The low data requirements for this approach and high efficacy provide a promising avenue to enhance simulations going forward, and emphasize the utility of differentiable simulations in physics.

openCARP-PINNs: Towards Faster Prediction of Cardiac Signals Propagation

Alexander Zolotarev[1], Lucas Ip[1], Krishna Saveri[1], Balvinder Dhillon[1], Clara Herrero Martín[2], Semhar Misghina[1], Caroline Roney[1]
[1]Queen Mary University of London, London, United Kingdom, [2]Universidad Politécnica de Valencia, , Spain

openCARP-PINNs: Towards Faster Prediction of Cardiac Signals Propagation, October 2, 2025, 10:35 - 10:42

Investigation and prediction of cardiac electrophysiological (EP) properties are beneficial in developing personalised diagnostics and treatments for cardiac arrhythmias. Physics-Informed Neural Networks (PINNs) combine data-driven learning with underlying physical laws. In cardiac electrophysiology, PINNs have been used in conjunction with biophysical models to predict electrical propagation faster than conventional biophysical modelling. Here, we proposed a novel toolbox with PINNs using the Mitchell-Schaeffer and Aliev-Panfilov ionic models to show fast and precise prediction of electrical wavefront propagation through in-silico heart tissue. For the generation of training data, we utilised OpenCARP, an open-source cardiac electrophysiology simulator. We conducted the experiments on a 2D in-silico grid with 3 different scenarios of wavefront initiation (a centrifugal wave from single or double corners of a 2D grid and a planar wave). We showed that openCARP-PINNs can accurately (with a mean RMSE below 0.2 in all experiments) and quickly (execution time of 16 ms) predict the shape of the action potential for both ionic models. Moreover, openCARP-PINNs precisely define the wavefront of signal propagation (mean Dice score between masks of ground truth and predicted transmembrane potential for the nodes in the testing set is more than 0.9). The proposed framework has the potential to guide treatment based on faster digital twins of patient hearts.

Emulating CO Line Radiative Transfer with Deep Learning

Shiqi Su[1], Frederik De Ceuster, Jaehoon Cha, Mark Wilkinson, Jeyan Thiyagalingam, Jeremy Yates, Yihang Zhu, Jan Bolte
[1]Department of Physics and Astronomy, University Of Leicester, Leicester, United Kingdom, [2]Scientific Computing, Rutherford Appleton Laboratory, Science and Technology Facilities Council, , United Kingdom

Emulating CO Line Radiative Transfer with Deep Learning, October 2, 2025, 10:42 - 10:49

The adoption of AI-based techniques in theoretical research areas is often slower than in other fields due to a perception that AI-based methods lack rigorous validation against theoretical counterparts. In this talk, we introduce COEmuNet, a surrogate model designed to emulate carbon monoxide (CO) line radiation transport in stellar atmospheres.
COEmuNet is based on a three-dimensional residual neural network and is specifically trained to generate synthetic observations of evolved star atmospheres. The model is trained on data from hydrodynamic simulations of Asymptotic Giant Branch (AGB) stars perturbed by a companion. Given a set of input parameters, including velocity fields, kinetic temperature distribution, and CO molecular number densities, the COEmuNet model emulates spectral line observations with a median relative error of ~7% compared to a classical numerical solver of the radiative transfer equation, while being 1000 times faster.
This presentation will also include some of our preliminary results, demonstrating the improved performance achieved through Physics-Informed Machine Learning (PIML) applied to the same problem, highlighting its potential for accelerating radiative transfer modelling in AGB starts.

AutoEmulate: Accelerating large-scale simulations with AI.

Marjan Famili[1], Paolo Conti[1]
[1]The Alan Turing Institute, London, United Kingdom

AutoEmulate: Accelerating large-scale simulations with AI, October 2, 2025, 11:21 - 11:28

Physical systems are understood and developed through simulations. The complexity of these systems often results in costly and time-consuming simulations. Emulators (surrogate models) can approximate the solution of simulations at a fraction of the computational expense. However, building effective emulators for scientific simulations typically requires significant machine learning (ML) expertise, posing a barrier for domain experts. To overcome this limitation, we present AutoEmulate, an open-source Python package that automates the construction of the best emulator for a given simulation, employing state-of-the-art advancements in the field. AutoEmulate compares and optimises a wide range of emulator models as well as data preprocessing, calibration, and analysis methods. It allows the user to integrate their own emulator and benefit from instant benchmarking. Furthermore, AutoEmulate provides simulator-in-the-loop integration, allowing physical consistency to be maintained and leveraging active-learning strategies to iteratively improve the emulator with physical information. AutoEmulate's user-friendly interface enables domain experts to deploy high-performance emulators with minimal ML expertise, offering a reliable solution for simulation-driven exploration across various disciplines.

K-space Interpolation using Deep Koopman Autoencoders

Wassim Ben Salah[1,2], Sarah McElroy[1,2,3], Antoine Naegel[1,5], Jon Cleary[1,4], Sebastien Ourselin[2], Jonathan Shapey[2], Christos Bergeles[2], Radhouene Neji[1,2]
[1]Imaging physics and engineering research department, School of Biomedical Engineering and Imaging Sciences, King's College London, London, United Kingdom, [2]Surgical & Interventional Engineering research department, School of Biomedical Engineering and Imaging Sciences, King's College London, , United Kingdom, [3]MR Research Collaborations, Siemens Healthcare Limited, Camberley, , United Kingdom, [4]Guy's and St Thomas' NHS Foundation Trust, London, , United Kingdom, [5]Siemens Healthcare, Courbevoie, , France

K-space Interpolation using Deep Koopman Autoencoders, October 2, 2025, 11:35 - 11:42

Purpose: Fast and robust reconstruction of accelerated MRI acquisitions is crucial in interventional MRI, particularly for parallel imaging, which is still widely used in the clinic. Our approach aims to develop a scan-specific and an interpretable physics-driven neural network for the reconstruction of accelerated MRI data with reduced noise and improved image quality at higher acceleration rates.

Theory: GRAPPA suffers from a nonlinear noise-induced error in estimating the coefficients for linear interpolation. To capture these nonlinear relationships, a deep Koopman autoencoder is used on k-space neighbourhoods to identify a global nonlinear coordinate transformation of acquired data and then interpolate missing targets across all coil channels.

Methods: The Deep Koopman autoencoder consists of an encoder, a linear layer for interpolation and a decoder. The encoder and decoder are designed with a nonlinear layer, each consisting of a complex-valued fully connected layer with a ModReLU activation function. The network is implemented using mean square error loss function on k-space neighbourhoods and target points derived the autocalibration signals (ACS), using batch gradient descent algorithm. The approach was compared to GRAPPA and RAKI in 281 retrospectively undersampled brain FastMRI datasets, both qualitatively and quantitatively using reconstruction metrics evaluated with respect to the ground truth fully sampled reconstruction. The effect of varying the number of autocalibration signals, the kernel size, and that latent space dimension for the proposed technique was also studied.

Results: The proposed approach showed improved quantitative metrics compared to GRAPPA and RAKI. Visually, the method showed higher resilience to noise amplification and improved preservation of sharp details.

Conclusion: This work introduces an interpretable neural network for k-space interpolation, enabling good reconstruction quality and offering avenues for extensions to enable autoencoder-based scan-specific denoising, as well as dynamic real-time reconstruction.

Keywords:
Parallel imaging, accelerated imaging, Koopman Autoencoder, non-linear mapping, k-space interpolation, deep learning

Solving inverse problems of epoxy hardening with physics-informed neural networks. Authors:

Dr. Kateryna Morozovska

Solving inverse problems of epoxy hardening with physics-informed neural networks., October 2, 2025, 11:42 - 11:49

In this project, we apply recent advancement in using Physics-Informed Neural Networks for chemical kinetics and cellulose ageing modelling [1] to model a process of epoxy curing. Curing is a chemical, irreversible reaction where polymer chains crosslink to form a rigid, solid material. Curing process starts with a viscous resin and ends with a fully cured solid. During the process, the crosslinking between polymer chains increases and the degree of cure also increases. The main factor influencing the curing process is temperature.

In this project, we use Kamal PDE to model the change in the degree of cure over time. In this case, we use simplified form of the Kamal equation and further plan to extend the model to account for more parameters. The figure below shows the schematic of the inverse problem implementation using PINNs, where Kamal ODE is integrated in the loss function.

Figure 1. PINN model for finding parameters of the Kamal equation.
The resulting model was able to learn the progression of the degree of cure with the resulting MSE=0.0121. The results also showed the generalizability of the previous model [1] to other types of applications with lesser need for adjustments. The current findings will be further enhanced with implementation of more complex version of the Kamal equation.

[1] Bragone, F., Morozovska, K., Laneryd, T., Shukla, K. and Markidis, S., 2025. Discovering Partially Known Ordinary Differential Equations: a Case Study on the Chemical Kinetics of Cellulose Degradation. arXiv preprint arXiv:2504.03484.

A Variational approach to Physics Informed Neural Network for Stochastic Partial Differential Equations

Robbie Slos[1,2], Tom Lefebvre[1,2], Jolan Wauters[3], Guillaume Crevecoeur[1,2]
[1]Department of electromechanical, systems and metal engineering, Ghent University, , Belgium, [2]MIRO core lab Flanders Make @ UGent, , Belgium, [3]Department of Mechanical Engineering, KU Leuven, , Belgium

A Variational approach to Physics Informed Neural Network for Stochastic Partial Differential Equations, October 2, 2025, 11:49 - 11:56

Physics-Informed Neural Networks (PINNs) refer to a deep learning approach to represent the spatial and temporal characteristics of a distributed physical phenomenon, such as thermal fields, using Neural Networks (NNs). The loss function used to train PINNs relies on a term that penalises any violation of the Partial Differential Equation (PDE) that governs the distributed phenomena, in addition to a least-squares-error penalisation of any available observations in the relevant domain. PINNs have been shown to be successful in approximating the solutions to PDEs and have proven effective even in the low data regime. A critical shortcoming is the lack of a systematic treatment of the uncertainty of the approximation. State-of-the-art approaches rely on Bayesian NNs, though these are computationally heavy, both during training and at inference, nor does the associated training procedure derive from first principles. To remedy these limitations, we propose Variational Inference PINNs (VI-PINNs).
Our approach derives from first principles: the uncertainty intrinsic to the distributed phenomena is explained by adopting Stochastic PDEs, while relying on standard measurement uncertainty to explain the observational uncertainty. This leads to the formulation of a posterior probability for the distributed phenomenon. Drawing parallels with Bayesian inference in finite spaces and relying on VI techniques to circumvent the otherwise intractable posterior. We derive a training objective that allows us to train two NNs, representing the mean and covariance of the approximation, respectively. The solution may be interpreted as a Bayesian belief about the true distributed phenomenon. Importantly, in the limit, the original PINN framework is recovered. We compare our approach with Bayesian PINNs (B-PINNs). Our results suggest that VI-PINNs are easier to implement, have a lower training time, and yields results that better align with reality, especially when extrapolating outside the measurement range.

Transforming physics-informed machine learning to convex optimization

Letian Yi[1], Siyuan Yang[1], Ying Cui[1,2], Zhilu Lai[1,2]
[1]The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, [2]The Hong Kong University of Science and Technology, , China

Transforming physics-informed machine learning to convex optimization (Online), October 2, 2025, 11:56 - 12:03

Physics-Informed Machine Learning (PIML) offers a powerful paradigm of integrating data with physical laws to address important problems in engineering, such as parameter estimation, inferring hidden physics, equation discovery, and state prediction. However, PIML, such as Physics-Informed Neural Networks (PINNs), still lack the necessary accuracy, stability, and interpretability when applying in pratical engineering due to many serious optimization challenges including the spectral bias, non-convex optimization, multi-objective optimization, and non-smooth optimization. In this study, we propose the Convex-PIML based on convex optimization and basis functions widely used in well-established numerical solvers to overcome all these limitations. The linear combination of B-splines is utilized to approximate the data, promoting the convexity of the loss function. By dividing variables into blocks and replacing the non-convex loss terms with convex approximations, the problem is further converted into a sequence of successively refined approximated convex optimization problems. This conversion known as Block Successive Convex Approximation (BSCA) allows the use of well-established convex optimization algorithms, obtaining solutions effectively and efficiently. Furthermore, an adaptive knot optimization method is introduced to mitigate the spectral bias issue of PIML, further improving the performance. The proposed fully adaptive framework by combining the adaptive knot optimization and BSCA is tested in scenarios with distinct types of physical prior. The results indicate that optimization problems are effectively solved in these scenarios, highlighting the potential of the framework for broad applications. Note that the Convex-PIML is also flexible since many other basis functions can also be incorporated to handle different systems.

Physics Informed Machine Learning for Astronomy

Dr Frederik De Ceuster

Keynote, Frederik De Ceuster, Leuven Gravity Institute, KU Leuven – 30 mins + 15 mins for questions, October 2, 2025, 13:30 - 14:15

Astronomy has always been a domain where modelling must be deeply physics-informed. Observations often provide only partial or indirect glimpses of the objects we study, and without embedding fundamental physics into our methods, inference and prediction would be severely underconstrained. In this keynote, I will discuss recent advances in combining machine learning with physical modelling to address these challenges. I will highlight three examples: MACE, a machine-learning emulator for chemical kinetics, which accelerates complex astrochemical simulations while maintaining physical consistency; pomme, a physics-informed framework for the 3D reconstruction of stellar atmospheres from spectral line observations; and the use of probabilistic numerics in astrophysics, specifically in numerical relativity, enabling efficient and quantifiably accurate modelling of gravitational-wave sources for next-generation observatories such as the Einstein Telescope and LISA. Together, these case studies illustrate how physics-informed ML can both respect fundamental laws and unlock new possibilities in astronomical modelling.

Deploying CFD-ML Models: From Optimization to Immersive Telepresence

Carlos Peña-Monferrer[1], Carmen Diaz-Marin[1]
[1]SimZero, , Spain

Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): Carlos Peña Monferrer, SimZero, October 2, 2025, 14:15 - 15:00

Numerical methods such as computational fluid dynamics (CFD) have long been essential for modeling complex physical systems in engineering, infrastructure, and environmental science. While powerful, their adoption in real-time or interactive applications has traditionally been limited by high computational costs. Recent advances in machine learning (ML), particularly those informed by physical principles, are beginning to shift this paradigm, enabling applications such as real-time intelligent flow control, energy optimization, risk anticipation, and human-in-the-loop decision support.

To move beyond research prototypes, these models must be not only accurate but also accessible, interoperable, and easy to deploy. In this talk, we present SimZero's approach to the practical exploitation of CFD-ML models, focusing on a modular, cloud-native deployment pipeline. This architecture encapsulates trained models behind a standardized API, decoupling simulation, inference, and visualization while enabling integration into digital twins, IoT, robotics, web platforms, and extended reality (XR) environments.

We present real-world use cases demonstrating the deployment of CFD-ML models in operational settings. These applications highlight the system's ability to support interaction and physics-informed decision-making across diverse platforms. This user-centric approach minimizes friction, simplifies system complexity, and enables responsive, high-fidelity simulation at the point of need. Ultimately, this technology bridges the gap between advanced computational modeling and its real-time use in critical applications across domains such as Industry 5.0 and smart cities.

Physics-informed machine learning for galaxy evolution studies

Dr Payel Das

Keynote, Payel Das, University of Surrey – 30 mins + 15 mins for questions, October 2, 2025, 15:30 - 16:15

One of the central goals of galaxy evolution studies is to determine how galaxies assembled over time and how dark matter is distributed within them. One route to this is through the present-day distribution of stellar positions and velocities (i.e. the phase-space structure) in nearby galaxies, which encodes both their evolutionary histories and underlying dark matter potentials. The collisionless Boltzmann equation (CBE) provides the theoretical framework, linking the stellar distribution function (DF) to the gravitational potential under the assumption of dynamical equilibrium.

Classical approaches to dynamical modelling often impose restrictive parametric forms for the DF or use moment-based methods such as Jeans modelling. Recent work has demonstrated that normalizing flows can serve as flexible, non-parametric representations of the DF, enabling more detailed descriptions of the phase-space structure of galaxies. However, flexibility alone is insufficient: dynamical equilibrium must still be imposed. This has motivated the development of physics-informed approaches, in which the learned DF is required to satisfy the CBE or its steady-state moments, ensuring that the resulting models correspond to true equilibrium solutions.

I will review these developments, highlighting how they allow us to infer both the DF and the gravitational potential. I will also discuss possible future directions, including the use of symplectic flows, which preserve Hamiltonian structure by design and may provide a more natural representation of equilibrium dynamics. These methods promise to open new opportunities for constraining the distribution of dark matter and the dynamical history of galaxies.

Relativity Wasn't in the Training Set

Prof, Miles Cranmer
University Of Cambridge

Relativity Wasn't in the Training Set, October 1, 2025, 4:15PM-5:00PM

Machine learning is revolutionizing data analysis in the sciences, yet cannot handle out-of-distribution prediction. While neural networks excel at interpolation within their training domain, including when transformed to cover different distributions in physically reasonable ways, they cannot fundamentally extrapolate to novel regimes the way mathematical physics can. Think of how preposterous it is that physics could accurately predict black holes in 1916 from purely tangential evidence. Modern machine learning is nowhere near making such predictions.

So, what can we do? In this talk, I will present two complementary approaches to infuse physics, in a soft way, into machine learning to achieve improved generalization on new problems. First, I will discuss symbolic regression, which leverages the compositional geometric structure inherent to physical laws. This approach effectively embeds centuries of physics knowledge as an inductive bias, but in a remarkably different way than a typical machine learning pipeline. Second, I will discuss Polymathic AI and our efforts to build foundation models trained on diverse scientific data across disciplines. These approaches, despite the massive chasm between the size of models they produce, are actually quite similar in how they impose priors on learning. I will highlight the subtleties of this.

EPRSC: unpacking the long term AI for science opportunities

Daniel Smith[1]
[1]EPSRC-UKRI, Swindon, United Kingdom

Special Session - Daniel Smith, EPRSC: unpacking the long term AI for science opportunities, October 2, 2025, 17:00 - 18:00

Special Session

In this session EPSRC will provide a short overview of the latest UK strategy relating to AI for Science, and explore, via a facilitated session, the long term opportunities and how these can be best realised.

Specifically, we will explore what partnerships, capacity building and interventions are required within engineering and physical sciences related disciplines such that the UK can fully realise the long term benefits, particularly for breakthrough scientific discoveries.

This will both consider those in the community already working in this area and how we may best facilitate broader engagement with others.

A perspective on fluid mechanical environments for challenges in reinforcement learning

Shruti Mishra[1], Michael Chang[2], Vamsi Spandan[1], Shmuel Rubinstein[3]
[1]Sony AI, , , [2]Cohere Labs Community, , , [3]Hebrew University of Jerusalem, ,

A perspective on fluid mechanical environments for challenges in reinforcement learning, October 3, 2025, 10:00 - 10:07

We consider the challenge of developing agents that efficiently interact with high-dimensional, evolving environments, towards a view of practical reinforcement learning (RL) agents interacting with open worlds, of which they witness and affect only a small part. We argue that canonical fluid mechanics problems, and their simulations, present a compelling testbed for the development of such methods. These problems arise in nonlinear instabilities, where small disturbances can grow to transform the dynamics of a system. Nonlinear instabilities represent several open scientific challenges with industrial applications -- the droplet breakup of a liquid jet, mixing at an interface between two fluids, and the appearance of unusually tall rogue waves in the ocean. In these settings, agents may leverage preserved representations across the changing dynamics to learn efficiently.

We present two problem descriptions of agents interacting with a fluid mechanical environment, and describe the state and action spaces, and reward functions, for these agents. For these examples, we specify the aspects of the environment which are nonstationary and the preserved invariances. We note Dedalus and JAX-CFD as open-source simulators that can be used for the development of reinforcement learning methods (Burns et al., 2016, Kochov et al., 2021). We demonstrate the use of Dedalus for environment generation by creating RL agents that learn to navigate in a stationary environment that is simulated using Dedalus. This sets the stage for future development of reinforcement learning agents that learn to meaningfully interact with simulated environments that represent scientific challenges in the natural world.

AI-driven Drifter Placement for Ocean Currents

Rui-Yang Zhang, Henry Moss[1], Lachlan Astfalck[2], Edward Cripps[2], David Leslie[1]
[1]Lancaster University, , United Kingdom, [2]University of Western Australia, , Australia

AI-driven Drifter Placement for Ocean Currents, October 3, 2025, 10:07 - 10:14

We introduce a formal active learning methodology for guiding the placement of Lagrangian observers to infer time-dependent vector fields -- a key task in oceanography, marine science, and ocean engineering -- using a physics-informed spatio-temporal Gaussian process surrogate model. The majority of existing placement campaigns either follow standard `space-filling' designs or relatively ad-hoc expert opinions. A key challenge to applying principled active learning in this setting is that Lagrangian observers are continuously advected through the vector field and so they make measurements at different locations and times. It is, therefore, important to consider the likely future trajectories of placed observers to account for the utility of candidate placement locations. To this end, we present BALLAST: Bayesian Active Learning with Look-ahead Amendment for Sea-drifter Trajectories. We observe noticeable benefits and robustness to model misspecification of BALLAST-aided sequential observer placement strategies.

Identification of Time-Varying Modal Parameters in Offshore Wind Turbines

Melisa Bozaci[1], Alice Cicirello[1]
[1]University Of Cambridge, Cambridge, United Kingdom

Identification of Time-Varying Modal Parameters in Offshore Wind Turbines, October 3, 2025, 10:14 - 10:21

Offshore wind turbines (OWT) are becoming critical structures to reach net-zero targets and mitigate climate change. Nevertheless, they experience extreme environmental and operational conditions, which can reduce their life span. Therefore, it is crucial to monitor their health status to assess their reliability and performance. In order to monitor OWT, operational modal analysis (OMA) is employed as a vibration-based structural health monitoring method. It uses accelerometer measurements to monitor damage-sensitive parameters, namely natural frequency, damping ratio, and mode shapes. Natural frequency is often monitored as it is easier to identify compared to damping ratio and mode shapes. However, OMA assumes that the OWT is linear time invariant (LTI) and the ambient excitation is stationary and broadband. In practice, these assumptions are violated as the OWT shows time-varying behaviour due to environmental and operational conditions (EOC), as well as degradation; the excitation is nonstationary and narrowband, such as wind and wave loading. These violations hinder accurate natural frequency identification under real operating conditions. In order to overcome these assumptions, data-driven methods have been explored to monitor natural frequencies. Nevertheless, the accuracy of these methods is limited by insufficient and unlabelled data, unobserved operational conditions, and lack of physical interpretation.
To address this challenge, this study proposes a hybrid modal analysis that combines physics and domain knowledge with data-driven methods using Long-Short Term Memory (LSTM) network and Extended Kalman Filtering (EKF). In this approach, the LSTM network provide the data-driven method while EKF introduces physics through the implementation of governing equations. The approach is tested using synthetic data generated from a finite element model of a 2-blade monopile offshore wind turbine to explore the blade rotation-induced effects on the offshore wind turbine. However, blade rotation introduces significant changes in natural frequencies due to mode veering and coalescence which can be a challenge for accurate identification. Therefore, performance of the proposed method is evaluated by comparing its results with a state-of-the-art physics-enhanced machine learning approach using the same dataset.

Multi-fidelity learning for physical system predictions

Paolo Conti[1]
[1]The Alan Turing Institute, London, United Kingdom

Multi-fidelity learning for physical system predictions, October 3, 2025, 10:21 - 10:28

High-fidelity simulations of physical systems are often limited by computational resources, significantly restricting the potential of scientific computing. Conversely, low-fidelity data—abundant, inexpensive, and fast to compute—are increasingly available, though they exhibit reduced accuracy, potential biases, and heterogeneous modalities. Multi-fidelity surrogate modeling aims to leverage such low-fidelity data to improve predictive accuracy when high-fidelity data are scarce. However, integrating diverse low-fidelity sources remains challenging, especially when they vary in quality and bias, or when certain modalities are inconsistently available or too costly to assimilate within budget constraints.

We introduce a multi-fidelity transformer architecture that assimilates information from multiple low-fidelity sources, merges it via attention mechanisms, and predicts the evolution of complex, high-fidelity physical systems. By incorporating physically meaningful low-fidelity signals, the method enforces physical consistency, enhancing interpretability and robustness in scientific modeling while maintaining the non-intrusive nature of data-driven approaches.

We demonstrate the generality of our method on PDE-based benchmarks, where low-fidelity models arise from coarse spatial or temporal discretizations, misspecified physics, or incomplete modality availability. We also apply our approach to a real-world climate modeling task predicting sea ice levels using data from in situ measurements and satellite observations. Our approach improves performance over standard multi-fidelity and purely data-driven bas

Optimizing hp-Variational Physics Informed Neural Networks for real-world applications using the FastVPINNs framework

Divij Ghose[1,2]
[1]Imperial College London, London, United Kingdom, [2]Indian Institute of Science Bangalore, , India

Optimizing hp-Variational Physics Informed Neural Networks for real-world applications using the FastVPINNs framework, October 3, 2025, 10:28 - 10:35

In this work, we extend the FastVPINNs framework for hp-Variational Physics Informed Neural Networks (hp-VPINNs) to real-world applications in computational fluid dynamics and computational electromagnetics, and present a new method for training on low-memory edge devices. We present the first implementation of hp-VPINNs to model turbulent flows by employing Reynolds-Averaged Navier Stokes equations (RANS) as the governing equations. Built upon the FastVPINNs framework, which employs variational loss formulations and tensor-based optimization, our approach significantly improves computational efficiency while achieving similar solution accuracy. The proposed method is validated on two canonical turbulent flow problems: a zero pressure gradient (ZPG) boundary layer and an adverse pressure gradient (APG) boundary layer. Additionally, we solve an inverse problem of sparse data assimilation for turbulent flow over a periodic hill. For forward problems, we compare the accuracies and computational times with those reported in the literature. The relative l2 errors for the mean fields and turbulent components are less than 3% and 10%, respectively. Additionally, we achieve 1.7x and 1.4x speedups for ZPG and APG problems, respectively. In the inverse setting, our method successfully assimilates sparse upstream measurements to recover near-wall flow characteristics, highlighting its potential in experimental data assimilation for turbulent flows. In the second part of this work, we explore the use of FastVPINNs for electromagnetic simulations of permanent magnet moving coil (PMMC) motors. In particular, we conduct magnetostatic simulations on different stator-airgap-magnet configurations. In each case, we predict the magnetic vector potential and compare it with finite element method solutions. Lastly, we present a novel method to extend FastVPINNs for edge deployment by applying tensor decomposition methods such as SVD, Tucker, CP, and tensor-train decomposition (TTD) to both the network parameters and the test function tensor. These decompositions yield low-rank approximations that significantly reduce memory and computation requirements.

Physics-Guided Graph Inference for District Heating Networks

Keivan Faghih Niresi[1], Lucas Kuhn[1], Gaëtan Frusque[1], Olga Fink[1]
[1]École Polytechnique Fédérale de Lausanne (EPFL), Renens VD, Switzerland

Physics-Guided Graph Inference for District Heating Networks, October 3, 2025, 10:35 - 10:42

Many intelligent infrastructure systems rely on sensor networks to monitor and control complex physical processes. These sensor networks are naturally modeled as graphs, where nodes represent sensors and edges capture relationships between them. Graph-based machine learning methods have shown great promise for processing and analyzing such data, enabling tasks like anomaly detection, forecasting, and data imputation.

However, the performance of these methods strongly depends on the quality of the underlying graph structure. In many real-world scenarios, the true graph topology is unknown, incomplete, or noisy. Even when available, the graph may not fully capture the functional dependencies which are critical for various downstream tasks.

To address these challenges, we propose a novel physics-enhanced graph signal processing (GSP) framework that leverages domain knowledge to improve graph inference and signal reconstruction. Our method incorporates physical constraints and prior knowledge about the system into the graph learning process, resulting in a smooth graph signal representation that guided by the underlying physics.

We formulate this approach as a convex optimization problem, efficiently solved using a primal-dual algorithm, enabling scalable inference of meaningful graph structures from raw sensor data. We validate our method on a case study of district heating networks, where sensor data are corrupted by noise and missing values. Results show that our approach improves graph inference quality and signal recovery compared to previous graph learning techniques, highlighting the benefit of integrating domain knowledge.

This physics-informed GSP approach bridges the gap between data-driven methods and physical system understanding, opening new avenues for interpretable and reliable graph- based learning in intelligent infrastructure systems.

HypeMARL: Multi-agent reinforcement learning for high-dimensional, parametric, and distributed systems

Nicolò Botteghi[1], Matteo Tomasetto[1], Urban Fasel[2], Francesco Braghin[1], Andrea Manzoni[1]
[1]Politecnico Di Milano, , Italy, [2]Imperial College London, , United Kingdom

HypeMARL: Multi-agent reinforcement learning for high-dimensional, parametric, and distributed systems, October 3, 2025, 10:42 - 10:49

Deep reinforcement learning has recently emerged as a promising feedback control strategy for complex dynamical systems governed by partial differential equations (PDEs). When dealing with distributed, high-dimensional problems in state and control variables, multi-agent reinforcement learning (MARL) has been proposed as a scalable approach for breaking the curse of dimensionality. In particular, through decentralized training and execution, multiple agents cooperate to steer the system towards a target configuration, relying solely on local state and reward information. However, the principle of locality may become a limiting factor whenever a collective, nonlocal behavior of the agents is crucial to maximize the reward function, as typically happens in PDE-constrained optimization problems. In this work, we propose HypeMARL: a decentralized MARL algorithm tailored to the control of high-dimensional, parametric, and distributed systems. HypeMARL employs hypernetworks to effectively parametrize the agents' policies and value functions with respect to the system parameters and the agents' relative positions, encoded by sinusoidal positional encoding. Through the application on challenging control problems, such as density and flow control, we show that HypeMARL (i) can effectively control systems through a collective behavior of the agents, outperforming state-of-the-art decentralized MARL, (ii) can efficiently deal with parametric dependencies, (iii) requires minimal hyperparameter tuning and (iv) can reduce the amount of expensive
environment interactions thanks to its model-based extension, MB-HypeMARL, which relies on computationally efficient deep learning-based surrogate models approximating the dynamics locally, with minimal deterioration of the policy performance.

Neural Operators for Accelerating Flow Field Predictions

Alan Xavier[1], Ludovic Renson[1]
[1]Imperial College London, London, United Kingdom

Neural Operators for Accelerating Flow Field Predictions, October 3, 2025, 11:21 - 11:28

Scientific machine learning uses machine learning (ML) to analyse and understand the behaviour of dynamical systems, often described by ordinary and partial differential equations (ODEs and PDEs). These modern ML methods can then be used to provide fast approximations for all possible system's states. Previous research in approximating highly complex and nonlinear functions has shown that these models are able to capture a wide range of behaviours, especially in areas where the physical phenomena, such as nonlinearity and turbulence, are not fully understood. However, these techniques have primarily focused on using neural networks (NNs) to learn mappings between finite-dimensional vector spaces. However, these functions exist within continuous and infinite-dimensional domains.

In this research, we investigate the use of neural operators (NOs), which are designed to learn mappings between infinite-dimensional spaces. We use a popular NO architecture, the deep operator network (DeepONet), to solve the unsteady Reynolds-averaged Navier-Stokes (URANS) equations in turbomachinery applications. The aerodynamics in these systems are highly complex, three-dimensional, and unsteady due to the relative rotation of the components and can vary from subsonic to transonic along the span of a single blade. The flow is also turbulent, displaying large and small-scale vortices that must be resolved. Due to the time-dependent nature of the URANS equations, we investigate extensions with recurrent neural networks (RNNs) to better capture the temporal evolution of our systems.

We focus on predicting the flow properties for an entire steady-state vibration cycle within an empty duct. The outlet of the duct is subjected to a sinusoidally varying unsteady pressure wave at a fixed amplitude and at varying frequencies, representing multiple different operating conditions. These simulations are used for training our augmented NO, where the first vibration cycle of an unsteady simulation is used to map to the last vibration cycle. Our NO learns a mapping from transient to steady conditions. Learning this mapping reduces the computational expense of running an expensive CFD simulation with thousands of time steps to only the first 100, as we only use this portion to predict the last 100 steps – which can then be used for other aeroelastic analyses, such as flutter.

A Surrogate Modelling Framework for the Correction of Structural Bias in SatelliteCloud Property Retrievals

Iarla Boyce[1], Alice Cicirello[1], Edward Gryspeerdt[2]
[1]University Of Cambridge, London, United Kingdom, [2]Imperial College London, , United Kingdom

A Surrogate Modelling Framework for the Correction of Structural Bias in SatelliteCloud Property Retrievals, October 3, 2025, 11:28 - 11:35

Satellite-based cloud property retrievals are an important component of climate science, though their efficiency often relies on simplified physical assumptions that introduce structural bias into the data products. A key source of this bias is the assumption of a globally fixed value for the effective variance of the cloud droplet size distribution (veff), a simplification which introduces systematic errors into derived parameters such as Liquid Water Path (LWP) and Cloud Droplet Number Concentration(Nd). This study addresses this deficiency by developing a post-processing correction framework based on surrogate modelling. The biased operational algorithm is treated as a "physics-deficient" base model, for which a surrogate is developed to predict its residual error. To construct this surrogate, a dataset is generated from numerous high-fidelity radiative transfer simulations, characterising the retrieval bias across a wide parameter space. This dataset is then used to build the error corrector: a classical surrogate model implemented as a multi-dimensional lookup table . Application of the surrogate correction framework reduces the mean bias in LWP and Nd by over 90%, improving the accuracy of the retrieved products. The outcome demonstrates how classical data-driven techniques, when conceptualised as surrogate models, can correct for known structural deficiencies in physical algorithms. The study establishes a methodology that can guide future work where the framework could be advanced through the use of more sophisticated machine learning emulators.

Deep reinforcement learning for wall-bounded turbulent flows via wall measurements

Giorgio Maria Cavallazzi[1], Alfredo Pinelli[1]
[1]City St. George's, University Of London, , United Kingdom

Deep reinforcement learning for wall-bounded turbulent flows via wall measurements, October 3, 2025, 11:35 - 11:42

In the field of Deep Reinforcement Learning (DRL) for the control of turbulent wall-bounded flows, it is a common practice to use measurements of the velocity field, u(x,t), from locations far from the wall as input for the agent. This approach has led to significant successes, such as active suction and blowing strategies that, when guided by a DRL agent, have demonstrated drag reduction of up to 40%. The policy of such an agent, π, can be represented as a function that maps the observed state to a control action: a_t = π(u|y+ =15,t). However, the practical implementation of such control systems is hindered by the difficulty of obtaining these far-field velocity measurements non-intrusively, in real-time. To overcome this limitation, we propose a novel approach that leverages Generative Adversarial Networks (GANs) to act as a virtual sensor. The GAN, represented by a generator function G, is trained to infer the instantaneous velocity field far from the wall, u|y+ =15, by using more available, non-intrusive measurements of wall pressure, p_w , and wall-shear stress, τ_w ,as its input. Specifically, we aim to learn a mapping G: {p_w ,τ_w } → u|y+ =15, such that u|y+ =15 can serve as a suitable state for the DRL agent. This study investigates the feasibility of this methodology
and quantifies the tolerance for uncertainty in the inferred velocity fields, u|y+ =15 to achieve a DRL policy with a satisfactory level of performance. Our work thus aims to bridge the gap between effective but impractical control strategies and their real-world application, paving the way for more robust and deployable DRL-based flow control systems.

Design Optimisation of Locally Resonant Metamaterials under Uncertainty

Niccolo Klinger[1], Alice Cicirello[1]
[1]University of Cambridge, , United Kingdom

Design Optimisation of Locally Resonant Metamaterials under Uncertainty, October 3, 2025, 11:42 - 11:49

The past decade has witnessed a growing focus on vibro-acoustic metamaterials for noise absorption and vibration mitigation. This is possible thanks to their frequency bandgap properties, which can be tuned by design engineers based on the specific application of the material. This behaviour is exhibited alongside with low specific mass and makes their use particularly desirable in a wide range of applications, including buildings noise insulation and acoustic liners for jet engines.

In recent years, research has pivoted around the maximisation of the bandgap frequencies targeted by metamaterials, with the intent of enlarging the region of attenuation for better performance over traditional materials. The majority of design optimisation techniques for metamaterials rely exclusively on physics-based models and the results are then used to identify some configurations to be tested experimentally. However, these optimisation frameworks do not take into account uncertainties that arise in manufacturing, mounting and experimental data collection. Modelling these uncertainties at the digital design step is an ongoing challenge.
As a consequence, discrepancies between the models results and the experimental results arise. This makes experimental tests essential for identifying a true optimal design.
Since it is not feasible to experimentally test all the possible metamaterials configurations to identify the true optimal design, a combination of measurements and models is necessary. State-of-the-art machine learning based optimisation techniques cannot handle noisy data acquired from experimental tests as inputs.

This work focuses on highlighting the challenges in the development of a hybrid physics-data method for iterative design optimisation that actively guides the selection of a limited number of experiments, with focus on locally resonant metamaterials.
A benchmark case study is investigated both experimentally and numerically. Quantification of different sources of uncertainties is discussed. Finally, possible solutions for experimental data handling in design optimisation iterations are discussed.

Physics-enhanced Simulation-Based Inference: Likelihood-free MCMC via Normalizing Flows and Variational Autoencoders

Giacomo Bottacini[1], Matteo Torzoni[2], Andrea Manzoni[1]
[1]MOX - Department of Mathematics, Politecnico di Milano, , Italy, [2]Department of Civil and Environmental Engineering, Politecnico di Milano, , Italy

Physics-enhanced Simulation-Based Inference: Likelihood-free MCMC via Normalizing Flows and Variational Autoencoders, October 3, 2025, 11:49 - 11:56

We present a new methodology for solving inverse problems in a variational framework exploiting neural networks that not only yields predictive results, but also provides uncertainty quantification. We employ a Differential Evolution Metropolis sampling method introducing a Normalizing Flow structure, specifically using Real-valued Non-Volume Preserving transformations (RealNVP), to approximate the likelihood function when dealing with complex physics-driven problems. Observational data are processed by a Variational Autoencoder to reduce the dimensionality of the input to be handled by the RealNVP and to extract the most relevant features for parameter estimation. To further enhance the informativeness of the latent space, a supervised loss term allows the latent space to be more structured and informative for the downstream inference task. The proposed methodology is validated on two case studies: a conductivity estimation problem for a steady-state groundwater flow governed by Darcy's law and damage estimation for a railway bridge structure.

Reduced order modeling with shallow recurrent decoder networks

Matteo Tomasetto[1], Jan Williams[2], Andrea Manzoni[1], J. Nathan Kutz[2]
[1]Politecnico Di Milano, , Italy, [2]University of Washington, , USA

Reduced order modeling with shallow recurrent decoder networks, October 3, 2025, 11:56 - 12:03

Reduced Order Modeling is of paramount importance for efficiently inferring high-dimensional spatio-temporal fields in parametric contexts. However, conventional techniques are typically data-hungry, limited to known and constant parameters, and inefficient for nonlinear and chaotic dynamics. In this work, we exploit SHallow REcurrent Decoder networks to build Reduced Order Models (SHRED-ROM) capable of reconstructing high-dimensional state dynamics in multiple scenarios from the temporal history of limited sensor measurements. Through applications on chaotic and nonlinear fluid dynamics, we show that the proposed technique is a robust decoding-only strategies, capable of dealing with both fixed or mobile sensors, physical and geometrical (possibly time-dependent) parametric dependencies and different data sources, while being agnostic to sensor placement and parameter values.

AI for emulation, forecasting and extremes in weather and climate

Simon Driscoll[1]
[1]University Of Cambridge, Cambridge, United Kingdom

Keynote, Simon Driscoll, University of Cambridge – 30 mins + 15 mins for questions, October 3, 2025, 13:30 - 14:15

Whilst machine learning and AI are not entirely new in weather and climate research, recent years have seen a rapid increase in their use - aided by advances in ML techniques, GPUs and so on.

Ranging from emulators of subgrid-scale processes to fully data-driven models, AI is now ubiquitous. Success in weather forecasting has even seen the adoption of AI forecasting models by operational centres. Yet there remain questions around their physical behaviour that are also crucial for building trustworthiness. The importance of the physical consistency and soundness of these models is also integral when considering longer term climate emulation and hybrid modelling approaches.

Here we discuss AI models in this context, ranging from forecasting abilities to our research probing their underlying physics and behaviour. We furthermore show how emulators of subgrid-scale processes can be made from model observations as well as synthetic data. Here obeying physics is particularly important in the case of where these emulators might be used in context of future climate projections. We discuss the role of observations for physically sound emulators (particularly in the cryosphere) and other research.

Applications of AI centre not only around the physical soundness of emulators and forecasting models but also for use around African meteorology and extremes which we briefly discuss.

Generative-AI to support Preliminary Engineering Design

Dr Shiva Babu[1], Marco Nunez[1], Yashwanth Gurbani[1], Nima Ameri[1]
[1]Rolls-royce, Bristol, United Kingdom

Real-world applications of PEML (30 minutes + 15 mins for question), Shiva Babu, Rolls-Royce, October 3, 2025, 14:15 - 15:00

From literature, there is evidence suggesting that conditional Generative Adversarial Networks (cGAN) can provide a valuable means to support engineering design by accurately predicting the results of computationally expensive simulations and their underpinning physics mechanisms through the encoding of design information into 2D images. However, there is a need for further work to identify and address some of the roadblocks hindering a wider application of this technology. The first part of this presentation summarises recent developments in the investigation conducted to understand and address restrictions identified for the use of cGAN models on different preliminary design engineering use cases. Key considerations are also drawn, with regards to key tasks conducted as part of the pre-processing of training data. A summary of developments in the adoption of cGAN models for 3D use cases is subsequently presented, along with the techniques under investigation with the aim of providing more freedom and flexibility to the creation and manipulation of engineering geometries. Preliminary results outline the possibility of generative design capabilities through the provision of suggested design concepts based on a single model containing both a geometrical representation of the system under study, in conjunction with the information of its performance against design metrics of interest (e.g., stress field, flow field, displacements, temperatures, etc.). A key advantage from such approach is that it allows training the neural network model with legacy simulation data to then offer a semi-instant prediction for the performance associated with synthetically generated design candidates.

Scientific Machine Learning for Extreme Events Prediction in Turbulence

Anh Khoa Doan[1]
[1]Delft University Of Technology, Delft, Netherlands

Keynote, Anh Khoa Doan, TU Delft, Faculty of Aerospace Engineering – 30 mins + 15 mins for questions, October 3, 2025, 15:30 - 16:15

Extreme events appear in many fluids mechanical systems, such as in atmospheric flows, oceanography, or wind turbines. These extreme events are sudden, unsteady, transient large nonlinear deviation of the flow away from its mean state. All these events are generally accompanied by detrimental and potentially catastrophic consequences. Therefore, the ability to predict such events is of the utmost importance. However, such a task is extremely challenging because of the underlying complex chaotic dynamics, the high dimensionality of flows and the relatively rare occurrence of extreme events in any dataset.
In this talk, we will present our recent developments in scientific machine learning techniques to support the prediction of such extreme events. Specifically, we will tackle three different aspects. First, we will present a combined dimensionality reduction/clustering approach to identify pathway to extreme events in chaotic systems. Second, we will discuss a reduced-order modelling approach, based on convolutional autoencoder and echo state network, that can learn the dynamics of flow with extreme events. Finally, some aspects related to the possibility of using machine learning-based control will be discussed.

Neural operators: a framework for scalable scientific computing

Dr Jean Kossaifi

Invited Talk on Real-world applications of PEML (30 minutes + 15 mins for question): Jean Kossaifi, NVIDIA, October 3, 2025, 16:15 - 17:00

Traditional deep learning typically involves learning mappings between finite-dimensional vector spaces. Real-world applications such as weather forecasting and aerodynamics, by contrast, involve modeling complex spatiotemporal processes governed by partial differential equations (PDEs) defined on continuous domains and at multiple scales. In other words, they involve learning mappings between infinite-dimensional function spaces.

Neural operators enable this by generalizing deep learning to learn mappings directly between function spaces, while offering substantial speed improvements over traditional PDE solvers, often several orders of magnitude faster. In this talk, I will introduce the fundamental concepts behind neural operators, illustrate their effectiveness on practical problems such as weather forecasting. Finally, I will touch on computational efficiency and practical implementation aspects in Python, demonstrating how these concepts can be applied in practice using open-source software.

Physics-Enhancing Machine Learning event (PEML2025)

1-3 October 2025

Institute of Physics, London, UK

IOP Institute of Physics