Physics Informed Machine Learning Part I: Different Strategies to Incorporate Physics into Engineering Problems

Eleonora Maria Tronci[1,2], and Austin R.J. Downey[3,4], Azin Mehrjoo[2], Puja Chowdhury[3], Daniel Coble[3]

[1]Department of Civil and Environmental Engineering, Northeastern University, Boston, MA
[2]Department of Civil and Environmental Engineering, Tufts University, Medford, MA
[3]Department of Mechanical Engineering, University of South Carolina, Columbia, SC 29208
[4]Department of Civil and Environmental Engineering, University of South Carolina, Columbia, SC 29208

ABSTRACT
Physics-informed machine learning (PIML) is a methodology that combines principles from physics with machine learning (ML) techniques to enhance the accuracy and interpretability of predictive models. By incorporating physical laws and constraints into the learning process, physics-informed machine learning enables more robust predictions and reduces the need for large amounts of training data. PIML has a wide range of applications in science and engineering, such as modeling physical systems, solving partial differential equations, and performing inverse analysis and optimization.

In part I of this two-part series, the authors will provide attendees with an overview of the main concepts, methods, applications, and challenges of PIML. According to the way that a first-principle model is integrated with a data-driven ML model, it is possible to classify physics-informed strategies. In this overview, seven strategies will be covered: physics-constrained ML; physics-guided ML; physics-encoded ML; data-augmentation via physics principles; transfer learning from physics-based synthetic data to experimental data; delta-learning physics correction to improve physics generalization and delta-learning unknown physics to represent unmodeled physical phenomena. The benefits of these approaches including better generalization, explainability, and efficiency of the ML models will be addressed. This work will present related challenges and limitations of each approach. Finally, the authors will discuss some open research questions and future directions for PIML. By the end of this tutorial, the participants will have a comprehensive understanding of the principles and potential of PIML, as well as the ability to critically evaluate PIML models.

Keywords: physics-informed, first-principle, machine learning

INTRODUCTION

The task of modeling and predicting the behavior of complex systems that encompass multiple scales and disciplines remains a stimulating challenge within the scientific community.

Physics-based strategies are widespread in many engineering fields and they are certainly the most popular approaches used to characterize complex phenomena focusing on the use of mathematical models to describe the physical laws and principles governing the behavior of a dynamic system [1]. These methods have the advantages of being highly interpretable and generalizable, enabling the understanding of the underlying logic of the model and the applicability to different systems with similar characteristics. Additionally, these strategies allow the flexibility to incorporate prior knowledge and constraints into the model structure and parameters, which further enhances their accuracy. However, they often have large and time-varying modeling errors and are characterized by a heavy computational burden, due to the sometimes complex dynamics characterizing the systems of interest.

The field of machine learning and data-driven strategies has seen remarkable advancements due to the abundance of data and computing resources. These methods utilize data to learn the system dynamics and control without the need for an explicit model [1]. They are flexible, adaptive, and scalable, making them ideal for complex and nonlinear systems. Additionally, they can handle uncertainties and disturbances in the data and learn from new data without requiring the entire model to be retrained. However, due to their intrinsic structure, these models can limitedly represent only the datasets they were trained to learn, without any flexibility or inference capability towards unseen conditions and a low level of interpretability and explainability. The accuracy and performance of these strategies depend on the availability of large and high-quality datasets and since the majority of the scientific fields are not big and comprehensive data-oriented domains, these models are often trained on sparse datasets and fail to generalize well.

A new paradigm, Physics informed machine learning (PIML), was born to overcome these challenges: leverage the wealth of prior knowledge, the underlying physics, and domain expertise (differential equations, conservation laws, or boundary conditions) incorporating it into the data-driven models, making it possible to build physically consistent predictive models which are faster to train, more generalizable, interpretable, and trustworthy. In the past decade, PIML raised the interest of several researchers in the engineering field and has shown its potential to revolutionize scientific discovery and technological innovation. The ongoing research in this field continues to push the boundaries of what is possible and it is possible to find several detailed review papers covering and keep researchers updated on the evolution of this topic [1–5].

INTEGRATING FIRST-PRINCIPLE AND DATA-DRIVEN MACHINE LEARNING MODELS

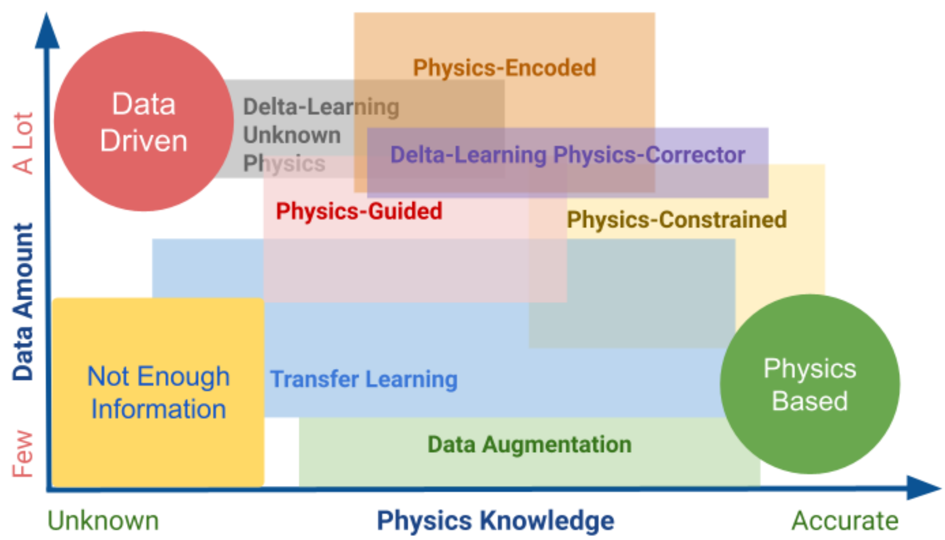

Figure 1: Physics informed machine learning categories according to the amount of data and physics knowledge.

In the realm of PIML, there are a variety of strategies that can be employed. These methods are typically classified according to three major characteristics:

1. the amount and quality of data that is utilized to describe a given model
2. the strategy chosen to incorporate the physics into the problem
3. the level of physical knowledge and understanding representing the phenomena of interest

Figure 1 addresses the classification of the major PIML categories according to the data availability and trustworthiness/restrictiveness of the physical principle.

In this overview, seven strategies will be covered: 1) Data augmentation; 2) Delta Learning - Physics-Corrector; 3) Physics Constrained; 4) Physics-Guided; 5) Transfer learning; 6) Delta Learning - Unknown Physics, and; 7) Physics Encoded. Figure 2 gives a short visual description of these strategies.

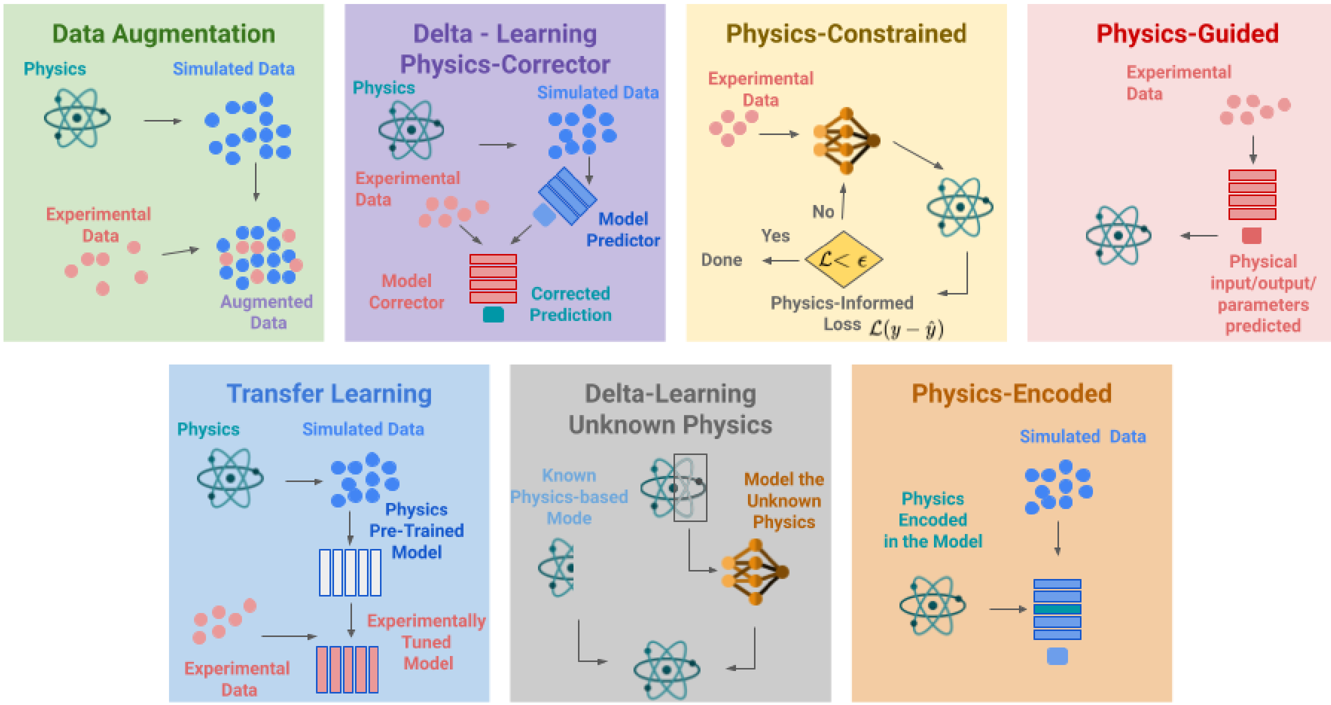

Figure 2: Seven approaches to constructing hybrid PIML models.

1) Data augmentation is a set of techniques to artificially increase the amount of data by generating new data points from existing data [6]. This includes making small changes to data or using deep learning models to generate new data points. Data augmentation can help to improve the performance and generalization of machine learning models, especially when the original data is insufficient or noisy. When an accurate and robust knowledge of the first principle is available, it is possible to leverage this information and run first-principle simulations to generate data at various states and operating conditions of a physical system. The synthetic data can be later combined with the experimental data if available to create an augmented dataset. A straightforward application of this strategy can be found in the work from Ritto et al. [7] that proposes a strategy to integrate a physics-based model with a machine learning classifier to create a fast and accurate digital twin. The physics-based model is a discrete computational model that simulates the dynamics of a damaged structure under different scenarios. The machine learning classifier is trained with data generated from the physics-based model and serves as the digital twin that can identify the damaged state of the structure in real-time.

2) Delta Learning - Physics-Corrector: is an approach that uses a physics-based model that is representative of the first principle model to generate training data for an ML model. The predictions of this physics-trained model will be used to inform a second ML model together with experimental data and have this second model learn the residuals on top of initial predictions. The final predictions are the sum of the initial predictions and residuals. The first principle model is used as baseline physics to produce initial, population-based predictions that may not be accurate for individual system units but can generalize well to different operating conditions adding the data-driven residuals to compensate for missing physics and unit-to-unit variability. The resulting performance gains are improved generalization over pure data-driven ML and improved prediction accuracy over purely physics-based modeling.

3) Physics Constrained: This approach is centered around the strict enforcement of physical laws on machine learning models. By imposing constraints on the outputs or parameters, such as mass conservation, energy conservation, or entropy production, Physics Constrained ensures that the model's predictions are always physically consistent. This is especially crucial when data is incomplete or uncertain.

Physics-constrained or physics-informed neural networks are the most popular implementation of this strategy. They rely on the first principle being embedded in the loss function of the network to be minimized. Integrating both data-driven insights and physical principles ensures that the model's outputs or parameters always adhere to established physical laws. Among the most known publications in this area; the work from Raissi et al. (2019) [8] stands out. The authors introduce physics-informed neural networks to be used to solve supervised learning tasks while respecting any given laws of physics described by general nonlinear partial differential equations. The effectiveness of the proposed framework is demonstrated through a collection of classical problems in fluids, quantum mechanics, reaction–diffusion systems, and the propagation of nonlinear shallow-water waves. An example of In part II of this two-part series, the authors present structural response forecasting using a physics-constrained methodology to solve the homogeneous and non-homogeneous second-order differential equations that constitute the equation of motion of a linear structural system.

Another group of constrained strategies is represented by physics-constrained Gaussian processes are a type of machine learning model that use Gaussian processes (GPs) to learn from data while respecting any given laws of physics described by partial differential equations or other constraints. Physics is usually incorporated in GPs by using physical constraints or prior knowledge to modify the kernel function, the likelihood function, or the posterior distribution of the GP model [9, 10]. Cross et al. (2023) [10] propose to use physics-informed Gaussian processes to learn the acoustic emission maps from a small number of measurements, while incorporating the physical knowledge of the structure's geometry and boundary conditions.

4) Physics-Guided Machine Learning: integrates domain-specific physical knowledge into the machine learning process, but rather than enforcing strict constraints, it uses this knowledge as a guide. The aim is to enhance the model's generalization, interpretability, and robustness, particularly in scenarios with limited or noisy data. While the model is informed by physical principles, it isn't strictly bound by them. An example would be a neural network that incorporates physical equations into its architecture or loss function, ensuring that its predictions are both data-driven and aligned with known physical behaviors, but without the strict enforcement seen in Physics Constrained. Andersen et al. (1990) [11] explore the use of artificial neural networks for modeling and controlling the arc welding process. The authors discuss how the neural network model can be used to model the weld bead geometry in terms of the welding parameters, such as the current, the voltage, the wire feed rate, and the travel speed. The proposed strategy can accurately capture the nonlinear and complex relationship between the input and output variables, and can also handle noisy and incomplete data.

5) Transfer Learning: consists of those strategies focused on using a model already trained on one problem to help solve another problem that is similar but not the same. This strategy can save time and resources by using existing models instead of training new ones and can also improve the performance of models when there is not enough data for the new problem. In the context of PIML, it is possible to build a transfer learning strategy by pre-training an ML model on physics-based synthetic data and then fine-tuning it on experimental data is a learning paradigm that leverages both empirical and theoretical knowledge to improve the performance of the model on a specific task. Physics-based synthetic data are generated by a generative model that is trained on real data and incorporates physical laws and principles that govern the system of interest. Experimental data are collected from real-world observations or measurements of the system. By pre-training the ML model on the synthetic data, the model can learn general and robust features and representations that capture the underlying physical mechanisms. By finetuning the model on the experimental data, the model can adapt to the specific characteristics and variations of the real data and correct the bias or errors introduced by the synthetic data. However, this approach requires a lot of physics knowledge and expertise to design and implement the generative model, to select and encode the physical prior, and to evaluate and interpret the results. Wang et al (2021) [12] propose a new method for fault diagnosis of rolling bearings under multiple working conditions using a deformable convolutional neural network, a deep long short-term memory network, and transfer learning strategies addressing the challenges of insufficient and dynamic data for bearings working under different conditions, such as speed, load, and temperature.

6) Delta Learning - Unknown Physics: represents the set of strategies where a data-driven ML model is used as a surrogate to learn and recover the unmodeled physics. The data-driven model is then combined with a first-principle simulation model to provide a comprehensive representation of the physical phenomenon of interest. This hybrid modeling approach is commonly used in multi-fidelity modeling and quantification of simulation model discrepancy. Jiang et al. (2022) [13] aims to improve the accuracy of failure prognostics for miter gates that are subject to degradation and damage over time, and require timely maintenance and repair to avoid costly and disruptive failures. The author proposes a dynamic model correction framework that uses strain measurements to update a simplified physics-based degradation model of the miter gate. The paper uses a polynomial chaos expansion (PCE) model to compensate for the missing physics in the simplified model, and a maximum likelihood estimation method to estimate the uncertain parameters of both models.

7) Physics Encoded: is a framework that embeds physical knowledge into the architecture or design of machine learning models, such as convolutional layers and recurrent layers. Physics-encoded aims to enhance the representation and learning capabilities of machine learning models, especially when data is complex or high-dimensional. An example of physics-encoded is physics-encoded neural networks, which use neural networks to learn the solution of forward or inverse problems by using physics-inspired layers or modules. Innes et al. [14] encoded the ordinary differential equation of motion, as the transformation function, into a neural network to simulate the trebuchet's inverse dynamics. The network with classical layers takes the target location and wind speed as input and estimates the weight and angle of the projectile to hit the target. These outputs are fed into the ODE solver to calculate the achieved distance. The model compares the predicted value with the target location and backpropagates the error through the entire chain to adjust the weights of the network.

CHALLENGES AND OPEN RESEARCH QUESTIONS

Physics-informed machine learning is a learning paradigm that leverages empirical data and available physical prior knowledge to improve performance on a set of tasks that involve a physical mechanism. It is an interdisciplinary research area that has many potential applications and challenges in science and engineering domains. Despite, its widespread application and development, there are some of the areas and domains still open for future research interest [4, 15]. Figure 3 provides a summary of the most relevant open research questions on the topic.

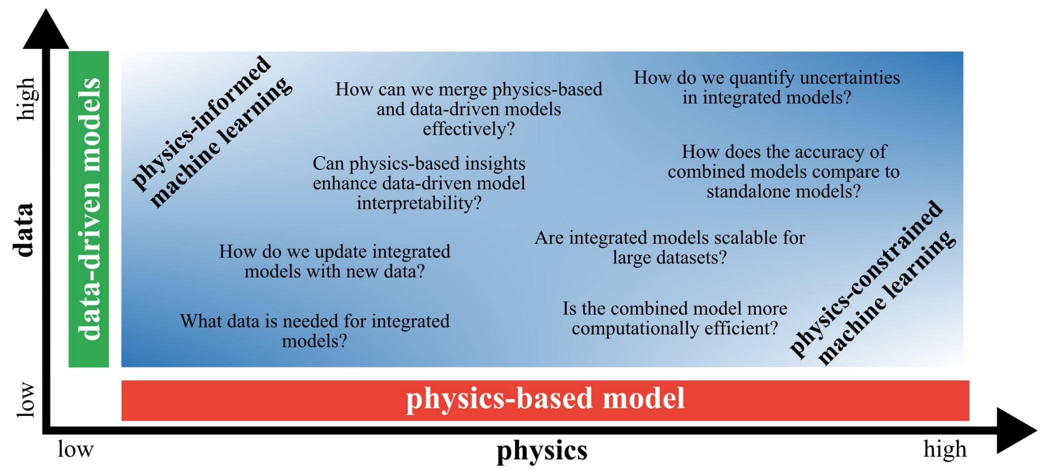

Figure 3: Open research questions on PIML

The majority of the publications related to PIML focus on simple problem validation with rare applications on full-scale dynamic systems. It is fundamental to develop applications scaled up to realistic, high-dimensional, and complex systems in different engineering domains. In this context, it is crucial to balance the trade-off between model complexity, interpretability, and generalization.

It is important, for any physics-based, data-driven, or hybrid strategy to identify and categorize the different sources of uncertainty. For a hybrid strategy, this task becomes even more complicated, distinguishing between model uncertainty and noise from uncertainty arising from missing data and unrepresented physical phenomena. An open research field concerns the investigation of uncertainty quantification and propagation of these sources of uncertainty from data, model, and physical prior and to eventually design robust and reliable PIML models.

Given the flexibility and several possibilities of incorporating physics into a data-driven modeling strategy, it is challenging to evaluate, validate, and compare hybrid models in different domains and scenarios. An interesting field of study focuses on the definition of appropriate metrics and bench-marking PIML methodologies against other PIML, data-driven, and physics-based methods.

Future research or lay in the expansion of PIML methods to novel and challenging problems, such as inverse engineering design, optimal control, active learning, and reinforcement learning, integrating and updating at the same time PIML methods with newborn machine learning techniques such as generative models.

REFERENCES

[1] Adam Thelen, Xiaoge Zhang, Olga Fink, Yan Lu, Sayan Ghosh, Byeng D Youn, Michael D Todd, Sankaran Mahadevan, Chao Hu, and Zhen Hu. “A comprehensive review of digital twin—part 1: modeling and twinning enabling technologies”. *Structural and Multidisciplinary Optimization*, 65(12):354, 2022.

[2] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. “Integrating physics-based modeling with machine learning: A survey”. *arXiv preprint arXiv:2003.04919*, 1(1):1–34, 2020.

[3] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. “Physics-informed machine learning”. *Nature Reviews Physics*, 3(6):422–440, 2021.

[4] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. “Scientific machine learning through physics–informed neural networks: Where we are and what’s next”. *Journal of Scientific Computing*, 92(3):88, 2022.

[5] Salah A Faroughi, Nikhil Pawar, Celio Fernandes, Subasish Das, Nima K Kalantari, and Seyed Kourosh Mahjour. “Physics-guided, physics-informed, and physics-encoded neural networks in scientific computing”. *arXiv preprint arXiv:2211.07377*, 2022.

[6] Alhassan Mumuni and Fuseini Mumuni. “Data augmentation: A comprehensive survey of modern approaches”. *Array*, page 100258, 2022.

[7] TG Ritto and FA Rochinha. “Digital twin, physics-based model, and machine learning applied to damage detection in structures”. *Mechanical Systems and Signal Processing*, 155:107614, 2021.

[8] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. *Journal of Computational physics*, 378:686–707, 2019.

[9] Guofei Pang and George Em Karniadakis. “Physics-informed learning machines for partial differential equations: Gaussian processes versus neural networks”. *Emerging Frontiers in Nonlinear Science*, pages 323–343, 2020.

[10] Elizabeth J Cross, Timothy J Rogers, Daniel J Pitchforth, Samuel J Gibson, and Matthew R Jones. “A spectrum of physics-informed gaussian processes for regression in engineering”, 2023.

[11] Kristinn Andersen, George E Cook, Gabor Karsai, and Kumar Ramaswamy. “Artificial neural networks applied to arc welding process modeling and control”. *IEEE Transactions on industry applications*, 26(5):824–830, 1990.

[12] Zheng Wang, Qingxiu Liu, Hansi Chen, and Xuening Chu. “A deformable cnn-dlstm based transfer learning method for fault diagnosis of rolling bearing under multiple working conditions”. *International Journal of Production Research*, 59(16):4811–4825, 2021.

[13] Chen Jiang, Manuel A Vega, Michael D Todd, and Zhen Hu. “Model correction and updating of a stochastic degradation model for failure prognostics of miter gates”. *Reliability Engineering & System Safety*, 218:108203, 2022.

[14] Mike Innes, Alan Edelman, Keno Fischer, Chris Rackauckas, Elliot Saba, Viral B Shah, and Will Tebbutt. “A differentiable programming system to bridge machine learning and scientific computing”. *arXiv preprint arXiv:1907.07587*, 2019.

[15] Chayan Banerjee, Kien Nguyen, Clinton Fookes, and Maziar Raissi. “A survey on physics informed reinforcement learning: Review and open problems”, 2023.