



Physics-informed neural networks as virtual sensors for domestic heat pumps

Ellen Louise Nicholson^{a,b*}, Yang Song^b, Adam Buruzs^a, Bernd Windholz^a, Davide Ronaldo^b, Michael Lauermann^b, Hatef Madani^{a,b}

^aUnit of Applied Thermodynamics and Refrigeration, Dept. of Energy Technology, KTH Royal Institute of Technology, Brinellvägen 68, Stockholm 10044, Sweden

^bAIT Austrian Institute of Technology GmbH, Center of Energy, Giefinggasse 2, Vienna 1210, Austria

Abstract

To ensure efficient operation of heat pump systems, comprehensive monitoring of the operational status is needed. However, due to cost constraints in the domestic setting, key parameters such as refrigerant mass flow rate are not monitored. This has led to growing interest in virtual sensors which use machine learning along with commonly available measurements to predict these missing parameters. To enhance the accuracy as well as interpretability of machine learning models, research has shifted to physics-informed machine learning. In this paper a short description of physics-informed machine learning is provided, followed by a case study detailing a physics-informed virtual sensor for the refrigerant mass flow rate of a domestic heat pump. The proposed model achieves a prediction accuracy of 93% with greater interpretability than that of a standard machine learning model.

© HPC2026.

Selection and/or peer-review under the responsibility of the organizers of the 15th IEA Heat Pump Conference 2026.

Keywords: Heat pumps; physics-informed machine learning; virtual-sensors

1. Introduction

In the EU, residential heating accounts for more than 15% of the final energy consumption across all sectors. Unfortunately, only a quarter of this heat is produced using renewable resources [1]. This is due to the large dependency on natural gas as fuel for residential boilers [2]. However, the heat pump (HP) has emerged as a very attractive alternative to gas boilers. With both its higher efficiency and ability to run on green electricity, the HP is set to decarbonise residential heating. To reach the EU's 2030 Net Zero target the European Heat Pump Association predicts that an additional 35 million HP units need to be installed. Despite its promising outlook, if recent sales trends continue, this target will not be met [3].

Alarming, studies have found that more than 50% of installed residential HPs do not operate as efficiently as they could due to faults in the system [4]. This leads to severe degradation of system performance, resulting in both a decrease in energy efficiency and thermal comfort. Therefore, to promote a greater HP roll out and reach the EU targets, it is necessary to streamline their installation and maintenance.

To monitor the operational status of a HP, and determine if any faults are present, it is imperative to study key parameters of HP operation. Among others, these include the electricity consumption, heating/cooling capacity, mass flow rates, temperatures and pressures. Some of these parameters, such as temperatures, are relatively easy and cheap to monitor. However, measurements like mass flow rates, which require more specialised equipment, would noticeably increase the price of a domestic HP system [5]. Therefore, these measurements are not taken, limiting the ability for accurate status monitoring.

* Corresponding author. Tel.: +421917135307
E-mail address: elnic@kth.se

One solution to incomplete empirical measurements comes in the form of virtual sensor. Here easy-to-measure variables are mathematically connected to the not measured variables of interest. For example, [6] developed three separate mass flow rate sensors based on a compressor map, the energy balance and an empirical correlation of an electronic expansion valve. All three of the virtual sensors performed with a relative maximum error of less than 5%. In [7], a virtual flow meter was developed based on energy conservation. The reported accuracy was around 6% within the tested temperature range.

With the rise in popularity of machine learning (ML), as well as the increasing amount of data available, research has shifted towards using data-driven models to make up for these missing measurements. In [8], four different ML models were developed as a virtual sensor for refrigerant charge. The four tested models were a multiple linear regression model, support vector regression, extreme random tree and gradient boosting tree and they achieved testing R^2 scores between 0.79 and 0.815. By introducing a ML pipeline to an existing virtual refrigerant charge sensor, authors in [9] managed to reduce the mean average percentage error of the sensor by about 16%. Despite the demonstrated success of ML in these applications, many issues remain unsolved. This includes the lack of generalisation capability outside of the specific training scope of individual solutions [10]. Because of the complex thermodynamical relations between different components in a HP, using the data-driven models for extrapolation purposes produces inaccurate results. This means that it is hard to rely on these models in real systems.

It is possible to improve a data-driven model's accuracy by incorporating physical knowledge of the system into the model. This field of research has been named 'Physics-informed Machine Learning' (PIML). The following paper will present the underlying principles of PIML and present a case study of a PIML model.

The rest of this paper is structured as follows. In Section 2, an overview of what PIML is provided. In Section 3, a case study of a virtual sensor for the refrigerant mass flow rate is presented, providing the methodology developed. Section 4 includes the results and discussions of the case study, with Section 5 give the overall conclusions of this paper.

2. Physics-informed machine learning

Following a literature review of PIML with regards to HP systems, it was commonly agreed that there was no specific criterion that made a data-driven model physics-informed. In the words of Karniadakis et. al. (2021), PIML is "the process by which prior knowledge stemming from our observational, empirical, physical or mathematical understanding of the world can be leveraged to improve the performance of a learning algorithm" [10]. In the 2021 research paper, three categories of prior knowledge were introduced as biases. These include observational bias, using sensor data to weakly embed physics principles; inductive bias, using ML model architectures to implicitly embed prior knowledge; and learning bias, embedding prior knowledge through soft constraints in the loss function of ML models [10]. These categories can be classified into four prevailing methodological groups of applying prior knowledge to a data-driven model.

These four groups are:

- Physics-informed inputs,
- Physics-informed loss functions,
- Physics-informed architectures and
- Physics-informed ensemble models.

These groups are described in more detail in the following subsections.

2.1. Physics-informed inputs

There are numerous ways in which we can make inputs physically informed. The first group of methods can be classified as data-preprocessing. Today, many of these methods are considered standard practice in ML in general.

The main idea behind data-preprocessing to make a model physics-informed is by ensuring that the inputs have a physical meaning attached to them. Generally, if the model is using physically relevant inputs, it is able to learn the physical correlations between inputs [11]. This can be seen in [12] where a hybrid deep neural network (NN) was designed using measurement data together with selected physics variables calculated using a thermodynamic model. The benefit of having these physics variables is that the model has the needed components to learn physical correlations between the inputs. [12] reported an increase in predictive performance from 95% to 99% when using the hybrid deep NN.

The second group of methods attempt to introduce physics into the feature space of the problem. This includes using simulation data to enhance empirical measurement data [11, 13]. The idea is that the simulation

data includes the known mathematical relationship between the input features allowing the data-driven model to learn these relations. As seen in [14] this method served to increase prediction accuracy of electricity performance across multiple buildings compared to the ASHRAE guidelines.

Following similar reasoning, transfer learning can be used to learn the underlying physical relationships of a dataset, with the fine-tuned model adjusting the model parameters to the empirical data. This can be seen in [15]. Here, simulation data of a wind turbine is used to pretrain a NN and monitoring data is used to fine-tune the model. This process resulted in a 1.64% decrease of the normalised root mean square error in the transfer learning model compared to the model solely trained on the monitoring data.

2.2. Physics-informed loss functions

Generally, loss functions in machine learning models include some form of a weighted regularisation term [16]. A physics-informed loss function is then simply a loss function where the regularisation term, or part of the regularisation term, comes from some prior physical knowledge of the system. Here, the prior knowledge of the system is being added a soft constraint during the learning phase of a data-driven model [11]. The idea is that the physics-informed term in the loss function guides the problem towards the realm of physically possible solutions. Because of this it acts as an additional source of information, enabling greater extrapolation and generalisation capabilities as well as ensuring a smaller dependence on data.

One of the largest applications of physics-informed loss functions is in physics-informed neural networks (PINN) [17]. Conceptualised by [17] in 2019, PINNs incorporate partial differential equations (PDEs) into the loss function of a NN as seen in Eq. (1).

$$L_{Total} = \lambda_1 L_{Data} + \lambda_2 L_{PDE} \quad (1)$$

More generally, the physics-informed loss terms can also include residuals from initial and boundary conditions [18]. In [19], a PINN was developed to predict the thermal profile of a heat exchanger. In addition to using the mean squared error (MSE) of the residuals of the governing equations, the MSE of the boundary conditions, gradient conditions and interface equality constraints are also used. The PINN consistently achieved a lower mean error in predicting temperatures compared to Naïve inference [19].

2.3. Physics-informed architectures

It is possible to add the prior physical knowledge as a hard constraint to a data-driven model [12]. This happens during the design phase of the model. Similarly to the physics-informed loss function, these hard constraints serve to constrain the predicted solution to adhere to our physical understanding of the system.

One way to add these hard constraints is to add specific constraint layers to the network architecture which will constrain outputs to physically feasible values. This method is demonstrated in [20], where, to predict the outlet temperature of a solution heat exchanger, the pinch point temperature and first and second laws of thermodynamics are implemented as temperature constraints on the prediction. [20] reported that the multi-layer perceptron (MLP) using hard constraints outperformed, in terms of thermal consistency, both a soft-constrained and unconstrained MLP. The constrained MLP also demonstrated the second-best coefficient of determination among the six studied models.

Another way of utilising a physics-informed architecture is to implement custom activation functions which assign physical meaning to the connections between layers [18]. This can be seen in [21] where neurons are assigned a physical model in a NN designed to predict fatigue life of materials. Here the authors showed that when only a small amount of data was available, the PINN obtained more accurate results than standard ML models.

Finally, studies have used specific types of architectures to reflect the physics of the problem. For example, [22] used NNs to represent the state space model of a room heated by a HP to predict the future temperatures in the system. The authors showed that their proposed architecture outperformed other NN models, especially when extrapolating for unknown conditions. [23] used a residual NN to predict the heat transfer rate in a finned tubed evaporator, as the residual blocks can learn the three separate thermodynamic behaviours governing the heat transfer at different points in the evaporation cycle. The proposed residual network outperformed the other tested networks such as a multilayer perceptron and a support vector machine.

2.4. Physics-informed ensemble models

Finally, a PIML model can consist of a mixture of a physics-based model and a data-driven model. Here the strengths of both the physics-based and data-driven approaches can be utilised. The majority of these models rely on problem decomposition where the well understood physics-driven components are separated from the stochastic ones. This is shown in [24] where the ensemble model provided better accuracy in forecasting air-conditioning power consumption than a purely data-driven models.

Another approach is to use a data-driven model to learn the residuals between the output of a physics-based model and the experimental data available. This allows a general physics-based model to be applied to a wider range of real systems, providing greater robustness of the solution. This has been demonstrated in [25] where the ensemble model obtained lower errors than the purely physics-based model.

3. Case Study

A key challenge identified when working with domestic HPs is the missing of key measurements related to system operation, such as mass flow rates. Based on a study conducted by [5], the typically collected measurements are presented in Table 1. In the following section a case study will be presented which describes how a physics-informed architecture model can be applied to a domestic HP system to accurately predict the refrigerant mass flow rate from collected measurements.

Table 1. Common measurements collected for a domestic HP installation

Measurement type	Measurement label
Compressor speed	F_{com}
Sink inlet and outlet temperature	$T_{w,in}$ and $T_{w,out}$
Source inlet and outlet temperature	$T_{b,in}$ and $T_{b,out}$
Refrigerant temperature at compressor inlet	$T_{evap,out}$
Refrigerant temperature at compressor outlet	$T_{con,in}$
Refrigerant at condenser outlet	$T_{con,out}$

3.1. Data augmentation

The dataset used for the case study consists of operational data from a field test installation. The data contains of 16946 steady state operational points taken from one ground source heat pump in 30 second intervals over 8 days in October of 2020. For this ground source HP test installation, in addition to the measurements in Table 1, the mass flow of the sink as well as the condensing and evaporating temperatures are available. For the distribution of the operating conditions of the dataset, see Fig. 1.

With the sink mass flow measurement, it is possible to calculate the heating capacity of the ground source HP and so estimate the refrigerant mass flow rate. This allows for supervised training and validation of the developed model. The heating capacity is calculated following Eq. (2) while the refrigerant mass flow is then derived through Eq. (3), where the enthalpies are calculated using CoolProp and heat losses are neglected on account of good thermal insulation.

$$\dot{Q}_{heating} = \dot{m}_{sink} c_{p,water} (T_{w,in} - T_{w,out}) \quad (2)$$

$$\dot{m}_{ref} = \frac{\dot{Q}_{heating}}{h_{con,in} - h_{con,out}} \quad (3)$$

As the data being used is operational data, there are inherent uncertainties present in the dataset. This includes both random uncertainties, such as statistical artefacts and systematic uncertainties, such as equipment precision. However, due to the dataset coming from a test installation these uncertainties are minimised. Another source of uncertainty in the data is that the refrigerant mass flow rate is calculated rather than empirically measured. This means that we do not have a ground truth value to compare the model prediction to. However, Eq. (2) and (3) provide thermodynamically accepted correlations between the available measured variables and the variable of interest, ensuring minimal uncertainty.




Figure 1. Range of operating conditions included in the dataset

3.2. Physics-informed architecture model

A common feature in physics-informed architectures has been the separation of the problem into smaller blocks. Then instead of representing the whole problem as a single ML model, each block becomes a separate ML model and the results of each block are combined to provide the solution to the overarching problem.

This approach can be seen in [5] where a HP is broken down into subcomponent models connected in a fixed point iteration. In this paper we modify the methodology presented in [5] to produce a PIML model.

In our PIML model, NN are used to generate the subcomponent models. All the NN models are composed of 2 hidden layers with 64 and 32 nodes respectively, using Rectified Linear Unit (ReLU) activations. The output layer is a single node with a linear activation. The network is trained using MSE loss with the ADAM optimiser at a learning rate of 0.001. Batch updates and early stopping with a patience of 50 epochs are used to prevent overfitting.

The NN subcomponent models are trained with normalised data obtained from a numerical model of the respective component. These numerical models are the dedicated manufacturer software provided for each component. Because of the way these software are designed, they are not directly compatible with the available measurements from domestic HP systems seen in Table 1. However, they can be used to create a database that

is utilised to train the NN models. Given the definition of physics-informed inputs above, this makes the subcomponent models physics-informed. To determine the inputs to the NN subcomponent models, it was necessary to operate within the constraints of the commonly measured parameters given in Table 1. However, the configurations chosen for the fixed-point iteration were also validated using physical correlations [5].

The fixed-point iteration is illustrated in Fig. 2. The iteration loop starts with an initial guess of the refrigerant mass flow rate. This guess is combined with $T_{con,in}$, $T_{con,out}$, $T_{w,in}$ and $T_{w,out}$ which serve as inputs to the condenser NN that predicts the condensation temperature. The condensation temperature is then used to calculate the condensation pressure using CoolProp, and combining the two values provides the enthalpy of the refrigerant as it enters the evaporator. This is under the assumption that the expansion valve is isenthalpic. This enthalpy, $T_{evap,out}$, $T_{b,in}$ and $T_{b,out}$ are then used as inputs to the NN evaporator subcomponent model. The output of this NN is the evaporating temperature of the refrigerant. The evaporating temperature, along with $T_{evap,out}$ and F_{com} are then inputs to a third NN. This NN is the compressor subcomponent model and the output corresponds to the mass flow rate of the refrigerant. For the iteration loop to converge, the output mass flow rate from the compressor model needs to coincide with the input mass flow rate to the condenser model. The converged refrigerant mass flow rate is then used to predict the evaporator and condenser states (temperature and pressure). Now, these predicted values correspond to the true thermodynamic values describing the refrigerant cycle of the HP.




Figure 2. Figure taken from [5] explaining the fixed-point iteration.

3.3. Accuracy metrics

To quantify the overall accuracy of the developed models, both relative root mean square error (RRMSE) and the coefficient of determination (R^2) are used. For temperature predictions the root mean square error (RMSE) rather than the RRMSE is presented due to the scale of the results.

4. Results and Discussion

The fixed-point iteration was initialised using an initial refrigerant mass flow rate of 0.018 kg/s and the threshold of mass flow rate convergence error was 0.0005 kg/s. The rest of the inputs into the PIML model

come from measurement of the field test ground source HP so that the outputs of the PIML can be compared to the true value. This comparison between the true and predicted outputs can be seen in Fig. 3.




Figure 3. Comparison between real measurements and model results.

The performance between the two models appears similar in Fig. 3. So, to validate the benefit of using the PIML model, the regression performance is compared to those of the integrated model developed in [5]. These results are presented in Tables 2 and 3.

Table 2. RMSE of the two prediction models

Model Type	RMSE (°C)		RRMSE (%)		
	T_{con}	T_{evap}	P_{con}	P_{evap}	$\dot{m}_{ref}$
Integrated Model [5]	0.18	0.53	0.51	1.65	9.25
PIML Model	0.26	0.31	0.87	1.03	6.11

Table 3. R^2 of the two prediction models

Model Type	R^2				
	T_{con}	T_{evap}	P_{con}	P_{evap}	$\dot{m}_{ref}$
Integrated Model [5]	0.998	0.667	0.997	0.844	0.845
PIML Model	0.996	0.891	0.991	0.870	0.933

From Tables 2 and 3 it is possible to see that for the prediction of the refrigerant mass flow rate, as well as the evaporating conditions, the PIML model outperforms the Integrated model. However, for the condensing conditions, the integrated model shows slightly better predictions.

Comparing the obtained results to other virtual sensors for the refrigerant mass flow rate, the PIML's RRMSE of 6.11% is in a similar range to the previously published work. This RRMSE corresponds to a R2 of 93%. This shows that this method achieves competitive performance in comparison to other methods. Similarly [26] states that models with excellent regression accuracy have RRMSE values of under 10%. This means that not only is the PIML model developed competitive, it also has excellent regression accuracy making it suitable for practical implementation.

A possible limitation of the model is its generalisation capabilities. The data used to train the NN subcomponent models represents steady state data under different operational conditions, limiting the application of the PIML model to steady state predictions. However, both the simulation data for the NN subcomponent models and the operational data used for the fixed-point iteration covers large ranges of operational conditions. This means that the PIML model has shown its ability to accurately predict the refrigerant mass flow rate over a range of operational conditions. These results therefore validate the usefulness of using a physics-informed methodology for the development of a virtual sensor for the refrigerant mass flow rate of a HP. Future studies could further validate the PIML methodology by expanding the dataset to include more operational conditions and different weather conditions.

A further limitation of the PIML is that its transferability between different HP installations is not known. Because the PIML model was trained with data corresponding to a specific HP installation, it is likely that it will perform worse when used with data of a different HP installation. However, to overcome this limitation, simply retraining the NN subcomponent models with data corresponding to the desired HP installation is required. Therefore, while the trained PIML may not be transferable, the methodology behind the PIML is.

5. Conclusion

In conclusion, there are multiple possibilities for introducing physics informed elements in classical machine learning models. Though introducing more computational complexity to the models, the integration of prior knowledge acts to improve the accuracy of the model as it imposes physically derived constraints on the outputs of the models. The physical priors also serve to enhance the generalisation capabilities of the models, improving their trustworthiness in real world applications.

In this paper, four categories of methods by which physics priors are integrated into ML models are presented. The first category consists of physics-informed inputs, where the inputs contain the physics for the model to learn. A second category involves the use of physics-informed loss functions, where soft constraints are used to constrain the model to predict physically plausible results. A third category takes advantage of a model's architecture to reflect that of the physical problem being solved. This generates models with physics-informed architectures. A final category combines traditional numerical models with ML models in different ways to generate physics-informed ensemble models.

Following the definition of a physics-informed architecture, a PIML model is developed to act as a virtual sensor for the refrigerant mass flow rate of a domestic HP. The model also includes aspects of physics-informed inputs. This is because the refrigerant cycle of a HP was separated into three underlying subcomponent NN models and each model was trained using simulated data. It is shown that this PIML can predict the refrigerant mass flow rate of a HP with a 93% accuracy. Because of the subcomponent structure of the PIML model, virtual sensors for the condensation temperature and pressure are also included. This provides a greater interpretability of the model, with possible future applications in fault detection and diagnosis strategies. The results presented validate the PIML methodology as a virtual sensor for HPs. In future work, the generalisation and transferability of the methodology can be shown by testing the PIML with new datasets corresponding to different operating conditions or different HP installations.

References

- [1] Eurostat. Energy consumption in households. Date published: June 2025. Data accessed: November 2025. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households

- [2] European Commission. Heating and Cooling. Date accessed: November 2025. https://energy.ec.europa.eu/topics/energy-efficiency/heating-and-cooling_en
- [3] EHPA, (2024). European Heat Pump Market and Statistics Report 2024. https://www.ehpa.org/wp-content/uploads/2024/08/Executive-summary_EHPA-heat-pump-market-and-statistic-report-2024-2.pdf
- [4] Belen Llopis Mangual (2025): Development of Methodologies for Fault Detection and Diagnosis in Residential Air Conditioning and Heat Pumps. Doctoral Thesis. Universitat Politècnica De Valencia, Valencia.
- [5] Song, Yang; Rolando, Davide; Marchante Avellaneda, Javier; Zucker, Gerhard; Madani, Hatf (2023): Data-driven soft sensors targeting heat pump systems. In *Energy Conversion and Management* 279, p. 116769. DOI: 10.1016/j.enconman.2023.116769.
- [6] Kim, Woohyun; Braun, James E. (2015): Development and evaluation of virtual refrigerant mass flow sensors for fault detection and diagnostics. *International Journal of Refrigeration* 63, pp. 184–198. DOI: 10.1016/j.ijrefrig.2015.11.005.
- [7] Han, Lichao; Yang, Qiang; Shao, Shuangquan; Wang, Zhichao; Xu, Zhaowei; Li, Hao; Li, Jiandong; Gou Xiaoxi; (2024): A virtual flow meter approach for measuring heating capacity of air source heat pumps under field condition. In *Flow Measurement and Instrumentation* 102, p. 102769. DOI:10.1016/j.flowmeasinst.2024.102769.
- [8] Shao, Wencong; Yang, Tianyang; Zhao, Lijin; Zou, Huiming; Smierciew, Kamil; Zhang, Shuo; Huang, Mingxuan (2025): Online monitoring and estimating of refrigerant charge in automotive transcritical CO₂ heat pump systems. In: *Applied Thermal Engineering* 278, p. 127225. DOI:10.1016/j.applthermaleng.2025.127225.
- [9] Zheng, Yue; Han, Hua; Xiong, Jun; Zhang, Hua; Cao, Xun; Gu, Bo; Dai, Wenjie; Gao, Xu; Yi, Weiqi (2025): Improvements and intelligence integration of virtual refrigerant charge (VRC) sensor. In: *International Journal of Refrigeration* 182, pp. 456–468. DOI:10.1016/j.ijrefrig.2025.12.017.
- [10] Karniadakis, George Em; Kevrekidis, Ioannis G.; Lu, Lu; Perdikaris, Paris; Wang, Sifan; Yang, Liu (2021): Physics-informed machine learning. In *Nat Rev Phys* 3 (6), pp. 422–440. DOI: 10.1038/s42254-021-00314-5.
- [11] Farea, Amer; Yli-Harja, Olli; Emmert-Streib, Frank (2024): Understanding Physics-Informed Neural Networks: Techniques, Applications, Trends, and Challenges. In *AI* 5 (3), pp. 1534–1557. DOI: 10.3390/ai5030074.
- [12] Hwang, Jun Kwon; Duhirwe, Patrick Nzivugira; Yun, Geun Young; Lee, Sukho; Seo, Hyeongjoon; Kim, Inhan; Santamouris, Mat (2020): A Novel Hybrid Deep Neural Network Model to Predict the Refrigerant Charge Amount of Heat Pumps. In *Sustainability* 12 (7), p. 2914. DOI: 10.3390/su12072914.
- [13] Maria Tronci, Eleonora; Downey, Austin R.J.; Azin, Mehrjoo; Chowdhury, Puja; Coble, Daniel (2025): Physics-Informed Machine Learning Part I: Different Strategies to Incorporate Physics into Engineering Problems. In *Data Science in Engineering Vol. 10*, Conference Proceedings of the Society for Experimental Mechanics Series. DOI: 10.1007/978-3-031-68142-4_1
- [14] Nutkiewicz, Alex; Yang, Zheng; Jain, Rishie K. (2018): Data-driven Urban Energy Simulation (DUE-S): A framework for integrating engineering simulation and machine learning methods in a multi-scale urban energy modeling workflow. In *Applied Energy* 225, pp. 1176–1189. DOI: 10.1016/j.apenergy.2018.05.023
- [15] Schroder, Laura; Dimitrov, Nikolay Krasimiro; Verelst, David Robert; Sorensen, John Aasted (2022): Using Transfer Learning to Build Physics-Informed Machine Learning Models for Improved Wind Farm Monitoring. In *Energies* 15(2), p. 558. DOI: 10.3390/en15020558
- [16] Kukacka, Jan; Golkov, Vladimir; Cremers, Daniel (2017): Regularization for Deep Learning: A Taxonomy. In ArXiv. DOI: 10.48550/arXiv.1710.10686
- [17] Raissi, M.; Perdikaris, P.; Karniadakis, G. E. (2019): Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. In *Journal of Computational Physics* 378, pp. 686–707. DOI: 10.1016/j.jcp.2018.10.045.
- [18] Wu, Yuandi; Sicard, Brett; Gadsden, Stephen Andrew (2024): Physics-informed machine learning: A comprehensive review on applications in anomaly detection and condition monitoring. In *Expert Systems with Applications* 255, p. 124678. DOI: 10.1016/j.eswa.2024.124678
- [19] Majumdar, Ritam; Jadhav, Vishal; Deodhar, Anirudh; Karande, Shirish; Vig, Lovekesh; Runkana, Venkataramana (2024): HxPINN: A hypernetwork-based physics-informed neural network for real-time monitoring of an industrial heat exchanger. In *Numerical Heat Transfer, Part B: Fundamentals*, pp. 1–22. DOI: 10.1080/10407790.2024.2325648.
- [20] Chen, Xin; Zhang, Lin; Huang, JiangBo; Jin, Lei; Song, YongShi; Zheng, XianHua; Zou, ZhiXiong (2025): A thermodynamics-consistent machine learning approach for ammonia-water thermal cycles 315, p. 134443. DOI: 10.1016/j.energy.2025.134443.

- [21] Chen, Dong; Li, Yazhi; Ke, liu; Li, Yi (2023): A physics-informed neural network approach to fatigue life prediction using small quantity of samples. In *International Journal of Fatigue* 166, p 107270. DOI: 10.1016/j.ijfatigue.2022.107270
- [22] Mun, Jeeye; Jo, Hyeong-Gon; Park, Cheol Soo (2025): Toward scalable prediction of indoor thermal dynamics: Neural-network-implanted state-space (NNiSS) model. In *Energy and Buildings* 331, p. 115359. DOI: 10.1016/j.enbuild.2025.115359.
- [23] Liang, Xing-Yu; Zhang, Bo; Zhang, Chun-Lu (2024): Physics-informed deep residual neural network for finned-tube evaporator performance prediction. In *Energy* 302, p. 131822. DOI: 10.1016/j.energy.2024.131822.
- [24] Dong, Bing; Li, Zhaoxuan; Rahman, S.M. Mahbobur; Vega, Rolando (2016): A hybrid model approach for forecasting future residential electricity consumption. In *Energy and Buildings* 117, pp. 341-351. DOI: 10.1016/j.enbuild.2015.09.033
- [25] Ma, Zhihao; Jiang, Gang; Chen, Jianli (2024): Physics-informed ensemble learning with residual modeling for enhanced building energy prediction. In *Energy and Buildings* 323, p. 114853. DOI: 10.1016/j.enbuild.2024.114853
- [26] Li, Mao-Fen; Tang, Xiao-Ping; Wu, Wei; Liu, Hong-Bin (2013): General models for estimating daily global solar radiation for different solar radiation zones in mainland China. In: *Energy Conversion and Management* 70, pp. 139–148. DOI:10.1016/j.enconman.2013.03.004.