Physics-Informed Machine Learning for Predictive Maintenance: Applied Use-Cases

Lilach Goren Huber
School of Engineering
Zurich University of Applied Sciences
Winterthur, Switzerland
lilach.gorenhuber@zhaw.ch

Thomas Palmé
School of Engineering
Zurich University of Applied Sciences
Winterthur, Switzerland
janthomas.palme@zhaw.ch

Manuel Arias Chao
School of Engineering
Zurich University of Applied Sciences
Winterthur, Switzerland
manuel.ariaschao@zhaw.ch

***Abstract*—The combination of physics and engineering information with data-driven methods like machine learning (ML) and deep learning is gaining attention in various research fields. One of the promising practical applications of such hybrid methods is for supporting maintenance decision making in the form of condition-based and predictive maintenance. In this paper we focus on the potential of physics-informed data augmentation for ML algorithms. We demonstrate possible implementations of the concept using three use cases, differing in their technical systems, their algorithms and their tasks ranging from anomaly detection, through fault diagnostics up to prognostics of the remaining useful life. We elaborate on the benefits and prerequisites of each technique and provide guidelines for future practical implementations in other systems.**

***Index Terms*—physics-informed Machine Learning, Condition-Based Maintenance, Predictive Maintenance, Anomaly Detection, Fault Diagnostics, Fault Prognostics, Deep Learning.**

I. Introduction

Decision making for optimal health management of industrial assets has been traditionally performed based on domain knowledge and physical models, whenever available. However, in recent years, with the abundance of machine data and the industry 4.0 revolution, there is an increasing trend towards data-driven solutions, focusing on condition-based and predictive maintenance algorithms [1]–[4]. A natural step forward is being made with the recent movement into hybrid methods, that combine the best of both worlds: exploiting the vast basis of domain knowledge and years of experience on one hand, and making use of data-driven innovations and resources on the other hand. This combination of physics with data-driven models is known as "physics-informed machine learning" (PIML) and can be applied in many fields and various ways, as summarized in a recent seminal review paper [5]. One of the approaches mentioned there is to use physics information in order to augment and supplement the training data for machine learning (ML) algorithms. As in other fields, some examples of PIML applications for equipment prognostics and health management (PHM) have been recently demonstrated [6]–[8].

In this paper we demonstrate different approaches to physics-informed (PI) data augmentation for ML algorithms applied to PHM problems. The approaches are demonstrated on 3 different industrial use cases. The use cases differ not only in their application fields but also in the task that the ML is aimed at, ranging from anomaly detection, through degradation trending and diagnostics and up to prognosis of the remaining useful life (RUL) of the machines.

The use cases concretely demonstrate the various benefits of PIML for practical applications. A central advantage over pure data-driven approaches is an enhanced prediction accuracy even when labeled data is scarce, which is a common challenge in PHM problems. The second obvious advantage is that physics allows for a high degree of interpretability of the model outputs compared to models based on data alone. This, in turn, has the added value of increasing the trust and acceptance of the local domain experts in the outcomes of the models. In the other direction, the possibility to supplement traditional knowledge-based approaches with modern data-driven ones offers higher fidelity of the models on individual units despite the heterogeneity of their operative conditions.

An important contribution of this paper is to point out universal concepts of PIML over diverse application fields, data types, and tasks, which are transferable to many other systems. At the same time, we elaborate on the differences between the use cases and provide guiding principles for practitioners that allow to select the most appropriate approach for their own use case.

The paper is organized as follows: Sections II, III and IV describe the 3 use cases. In Section V we compare the use cases and provide guidelines for future applications.

II. Use case I: PIML for fault detection in solar power plants

In this section we describe the application of PIML for automatic fault detection (FD) in utility-scale operational photovoltaic (PV) power plants. One of the common fault mechanisms in PV plants are tracker faults. Solar trackers are devices that orient the solar panels towards the sun, thereby maximizing the amount of energy produced from a fixed amount of installed power generating capacity [9]. A tracker fault usually occurs when the tracker gets stuck at a certain orientation instead of tracking the sun.

Tracker faults can lead to a significant reduction in the power produced by the PV strings that are mounted on the

L.G.H was supported by Innosuisse - Swiss Innovation Agency Grant No. 55018.1

faulty trackers. Early and automatic detection and localization of such faults can therefore prevent large production losses, thereby increasing the cost-effectiveness of solar energy and thus helping to accelerate the transition towards renewable energy sources. Despite the potential efficiency gains, the scientific literature addressing tracker faults is very scarce [10].

In our work [11] we suggest an algorithm for tracker FD based on operationally available power data and use a physical model to augment the data. We train a classification Convolutional Neural Network (CNN) on the augmented field data from an operational PV plant under various healthy (normal) conditions. The plant is monitored at string level, thus the input to the CNN is the measured produced power of the single PV strings (rather than from individual solar panels). The approach we take is somewhat uncommon: instead of developing a full system simulator, we rather collect field data from a healthy system and use a physical model of the tracker mechanism exclusively for the generation of synthetic faults out of the healthy field data. In this way we obtain an augmented training set, containing 50% healthy power profiles from the field data and 50% faulty profiles that were generated using a physical model to corrupt healthy profiles (see Figure 1). This data set is used to train a binary CNN classifier to distinguish faulty from healthy daily string profiles, thus detect the faulty strings every day. The trained CNN was tested on two data sets: i) a synthetic data set, where all faults are model-generated. ii) a field data set with real historical tracker faults. The test results are shown in Figure 2 (a) and (b) for test sets with synthetic and real faults respectively. The precision-recall curves (PRC) are compared with the ones of a purely data-driven anomaly detection convolutional autoencoder trained with the healthy data only. The performance of the PIML is better, both for synthetic and for real faults. However, on real faults, the data-driven model performs extremely poorly and suffers strong instability against randomness in the training process (observed through strong fluctuations of the PRC over multiple training repetitions). The PIML, on the other hand, generalizes very well from the synthetic to the real conditions and shows highly reproducible and robust results. For more details about the model we refer the reader to [11].

The motivations to take the described hybrid approach for this use case are multiple: i) field data under healthy conditions is typically abundant ii) the field data captures complex phenomena that are very hard to simulate. In our case, these are highly non linear weather dynamics and their effect on the power production. iii) the field data covers the realistic variability between strings within one power plant. iv) we were able to express the effect of the faults as function of the physical variables in the healthy component, i.e., to generate faulty power profiles by corrupting the healthy ones using a simplified approximate physical model. It is important to note that since the richness and complexity of the operative conditions is covered by the healthy field data, the level of accuracy of the fault simulator is allowed to be rather low.

Some key aspects of the approach are universal and may be found useful irrespective of the specific application:

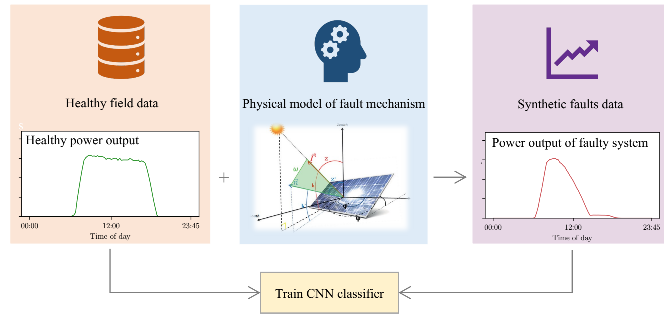

Fig. 1. Training scheme of the physics-informed deep learning model for fault detection in solar power plants. Healthy field data (left) is augmented using a physics-based fault simulator (middle) to generate synthetic faulty data (right). The augmented data is used to train a deep CNN classifier.

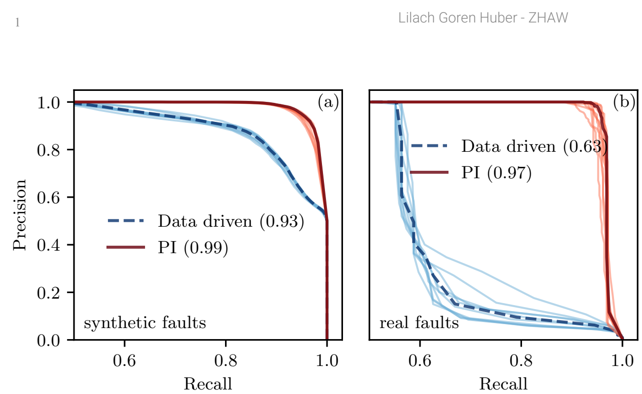

Fig. 2. physics-informed vs. a pure data-driven model. The fault classification performance of two models is evaluated in terms of their precision-recall curves. A purely data-driven convolutional autoencoder is compared with a physics-informed (PI) deep learning model on two test data sets: (a) with synthetic faults simulated using the physical model (b) with real annotated faults from field data. For each model, we show the results of 10 training runs and their median (thick curve). The AUC of the median curve is given in brackets.

- Using physical models to synthesize faulty data from field healthy data is a recommended approach when very little or even no faulty field data is available. The approach requires rich healthy data, which represents all operative conditions during normal operation of the machines.
- Another prerequisite for the approach is access to an approximate physical model of the effect of faults on the healthy data. this allows to ”corrupt” the healthy data in a physics-informed manner.
- If the field data covers the complexity and variability of the healthy states, the physical model need not be highly precise.
- Introducing noise to the synthetically augmented data can improve the performance and the robustness significantly. The noise should be physics-informed rather than completely random. In this way, potential sources of physical noise in the data can be simulated (for details, see [11]).

expect to see a residual as illustrated in figure 10 (by the hybrid model residuals). The residual is centred on 0, with a low variance, when the GT is performing as expected. However, several different mechanisms will affect the residuals, such as degradation, sensor faults and hardware failure. In many situations, time-based progressions of the deviations are of interest, especially when diagnosing any potential failures.

To illustrate the hybrid model response for a failing sensor, a sensor slow degradation rate is superimposed on unseen TAT data collected after the hybrid corrector development. The superimposed trend is shown in Figure 11. Without knowing the true value for the correct sensor reading, this sensor failure would be difficult to detect at an early stage.

The sensor degradation is imposed on the operational data is a linear trend with a degradation rate of:

$$\frac{dT}{dt} = 0.18°C/day$$

To illustrate the impact of this degradation rate on fault free measurement, figure 11 shows the both the operational fault free data as well as the data with superimposed degradation trend.

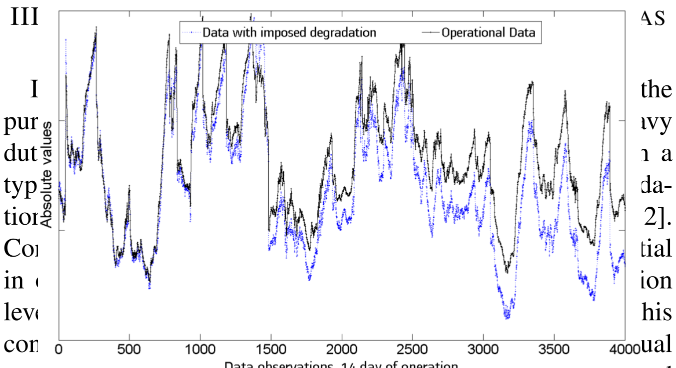

Figure 11: Sensor failure superimposed on fault free operational data

To provide an example of how this degradation trend is revealed and detected by the hybrid model, figure 12 shows the hybrid model residual as well as the superimposed degradation around the superimposed degradation trend; this is basically the high variance error as seen in figure 9. In a practical application, this would lead to a prolonged time for detection of the degraded sensor.

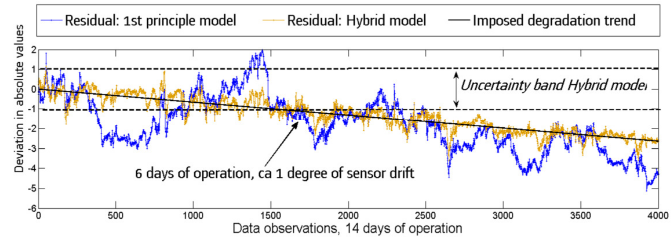

Figure 12: Residual evaluation of 1st principle model and hybrid model

CONCLUSIONS

Monitoring today is carried out in centralized M&D centres where experts analyse data from multiple plants. Efficient monitoring tools are needed to support the personnel in the evaluation of incoming data and flagging potential problems requiring detailed assessment. These tools need to be robust and efficient, both in their application as well as in their development. A hybrid model combines the in-house knowledge of GT operation as well as taking individual settings and characteristics to enable accurate monitoring of the unit's measured data from the unit. A pure data-driven modeling of the GT performance requires complex statistical models, while on the other hand day-to-day monitoring using complex 1st principle models of the GT is not efficient. By combining the two sources of information, we can have the "best of both worlds". The main reasons to generate a surrogate model of a 1st principle performance tool are:

- Faster than a 1st principle model, since there is no complex iterative solver once it is developed
- The GT was ... model equations
- Can be applied to new plants with limited operating history, with increased allowable residual thresholds
- Ease of integrating together different systems independently of the first principle model, which generated the surrogate. Modular construction

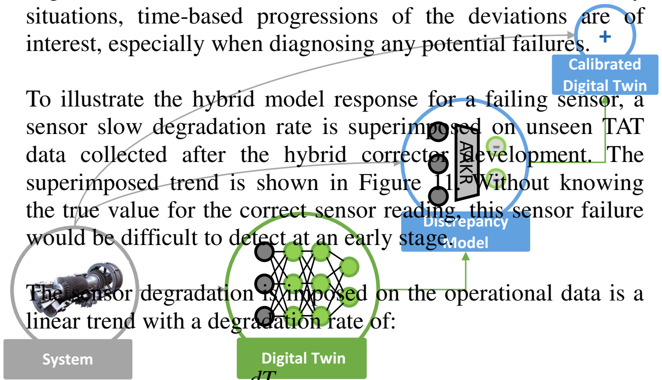

Fig. 3. Overview of the hybrid PIML framework for degradation trending and fault detection in gas turbines (GTs). A fleet-level digital twin is based on a physical model of the GT and is adapted to predict individual turbines using a transfer learning discrepancy model.

III ... AS

I... pur... dut... typ... tion... Co... in ... lev... con... measured variables (such as output power, temperature, and pressure) with the ones expected based on a physical model of the GT. Such physical models serve primarily for design purposes or performance evaluations during machine design and commissioning [13].

The main disadvantage of using a first-principle physical model is that it usually fails to capture the unit-to-unit variability arising from differences in sensor calibration, settings, and history of operation and maintenance [13]. In this way, discrepancies between the expected and actual behavior are not easily interpretable and thus cannot be clearly assigned to specific failure or degradation modes.

Here we describe a hybrid approach, combining a physical model of GTs with data from individual units in order to accurately monitor anomalies and degradation patterns and deduce their root cause. The approach can be divided into two steps as depicted in Figure 3:

1) A fleet-level digital twin (DT, green in Fig. 3). A fully connected regression neural network is trained to reproduce the outputs of a physical model. The inputs to the model are the operative conditions of the GT (inlet temperature, pressure, humidity, and guide vane). The DT is trained to output 12 performance variables of the GT (such as the power output, and various outlet values). It is calibrated to reproduce the mean fleet values. The advantages of the DT are its computational efficiency and numerical stability compared to the original physical model.
2) A fine-tuning unit-specific transfer model (blue). The output of the DT is used together with field condition monitoring (CM) data from a specific unit, in order to train a discrepancy transfer learning (TL) model based on Gaussian kernel regression (AAKR) [14]. The TL model is aimed at correcting the fleet-based physical output and adapting it to produce accurate predictions for the specific unit.

For a more detailed description of the hybrid method, we refer the reader to [15].

In Figure 4 we show an example of the application of the method for monitoring the condition of a GT. The data in the plot was generated as follows: field data from a healthy operational GT was recorded. Then, a degradation trend was mathematically added to performance data in order to emulate a failure and have a clear ground truth for the model validation. The Figure compares the prediction residuals of the physical model (blue) with those of the hybrid PIML approach (orange), and the ground truth degradation trend (black). The hybrid model is capable of detecting the degradation faster, more accurately and with a high degree of certainty compared to a purely physical model. A reliable degradation monitoring, in turn, is a key condition to diagnosing the root cause of the failure by monitoring all output variables simultaneously.

Fig. 4. Degradation trending of GT using the principle model residuals. The hybrid PIML approach (orange) captures the true degradation (black) much more accurately than the pure physical model (blue).

Some aspects of this approach are universal and may be relevant irrespective of the application:

- For many complex industrial systems there exists a rich domain knowledge in-house that can and should be taken advantage of, not only in the design stages, but also for condition monitoring purposes. A PIML approach allows to exploit this knowledge. This, in turn, is expected to increase the trust and acceptance of the local domain experts in the predictions of the model, which is familiar to them.
- A generic physical model is often unable to account for the unit-to-unit variability. Field data from the specific unit enhances the predictive power of the model for this very unit. The advantage of the hybridization of the physical model with the data-driven model is that even very little data can significantly improve the accuracy.

This allows for deployment at early stages, with t possibility to improve the accuracy with additional fie data.

- A data-driven TL method to adapt the physical moc for a specific unit increases the interpretability of t fault mechanism. In addition, it allows to understa the natural variability within the fleet, and thus suppo design improvements.

IV. Use case III: PIML for fault prognostics in turbofan engines

In this section, we describe the application of PIML f the robust prognostics of the remaining useful life (RUL) a small fleet of turbofan engines.

Some of the most common long-term deterioration mech nisms in aircraft engines are erosion, corrosion, fouling, a tip-clearance increases on rotating components [12]. The deterioration mechanisms have an impact on engine performance indicators like the thrust, the operating temperatures, and the fuel consumption, which in turn affect the lifetime, the operational cost, and even safety [7].

One of the main challenges for reliable diagnostics stems from the fact that the degradation mechanisms are only indirectly measurable or detectable through an in-depth engine inspection, with no possibility to monitor effects like fouling or erosion directly. Therefore, early detection and isolation of the degradation source (diagnostics) and the prediction of the failure time (prognostics) must be inferred from indirect measurements or, alternatively, from physical models of the engine [12].

In this work, we follow the design principle of the early work [8] and suggest a new algorithm combining physics-based and deep-learning models that enable accurate, robust, and simultaneous diagnostics and prognostics of turbofan engines. The proposed algorithm is depicted in Figure 5. Some of its building blocks are similar to the ones in the GT use case described above. The first building block is a digital twin (green) trained to emulate an existing physics-based performance model of the engine. The DT takes CM variables (i.e., throttle position, altitude, and ambient temperature) as inputs and provides measured variables as outputs, such as pressure and temperature in various locations on the engine. As a major difference to use-case II above but similarly to use-case I, the DT also takes as input health parameters that account for the performance impact of degradation at the sub-component level. In other words, the DT is aimed at modeling not only healthy conditions but also degraded ones. Similar to the GT use-case II, a discrepancy term (blue) is also computed to compensate for errors of the DT resulting from factors that are unaccounted for by the physical model (these can be unit-to-unit variations). But, contrary to use-case II, no modeling is involved for the discrepancy. The discrepancy-aware calibrated DT (blue) is able to provide accurate sensor readings estimations of degrading systems given the operative conditions and the discrepancy.

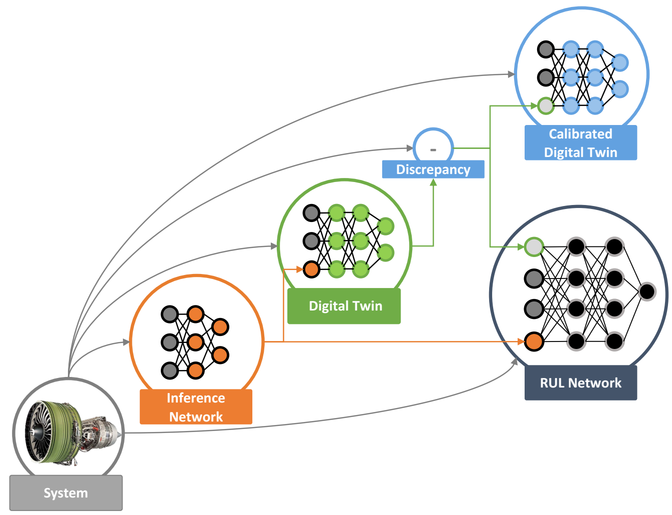

Fig. 5. Overview of the proposed PIML framework for fault prognostics and diagnostics in turbofan engines. Given a DT of the engine (green) and CM data, engine health parameters capturing the performance impact of the degradation on the engine components are obtained with an inference network (orange). To compensate for a possible incomplete physical representation of the DT, the discrepancy-aware model (blue) is incorporated to provide accurate sensor reading estimations of degrading systems. The DL-based RUL network (black) receives an augmented input space of CM features, model discrepancy, and known engine health parameters and provides a RUL prediction as its output.

In addition to the DT and the discrepancy blocks, here, there are two building blocks that allow for simultaneous RUL prediction and root-cause identification: the inference network (IN) (orange) and the RUL network (black). The IN is aimed at inferring the hidden internal variables (health parameters) that account for the performance impact of degradation on the engine components. In this way, degradation can be not only monitored but also accurately quantified and assigned to a specific root cause. The RUL network model receives as inputs the run-to-failure CM data, the inferred health (or degradation) parameters from the IN unit as well as the observed discrepancy of the DT to provide a prediction of the RUL.

The inference, the discrepancy, and the RUL networks are trained simultaneously in an end-to-end fashion using full run-to-failure degradation CM data from a small fleet of three turbofan engines [16]. For simplicity, fully connected layers were considered in this case study for all the networks.

The performance of the proposed method is compared with the one of a purely data-driven prognostic method. Figure 6 shows the error between the true and predicted RUL for a purely data-driven approach (left) and the proposed PI approach (right) on the test data, containing three turbofan engines sharing the same failure mode. The average root mean squared error (RMSE) in engine cycles is given in brackets. Under this metric, the proposed hybrid approach provides an average increase of 15% in RUL prediction accuracy. It is important to highlight that in addition to a RUL estimation, the PIML method provides early detection of the degradation

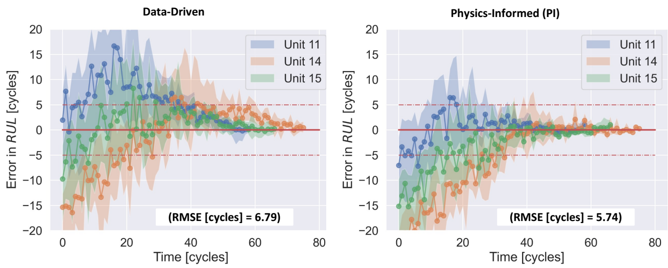

Fig. 6. physics-informed vs. a pure data-driven model. The RUL prediction performance of the two models is evaluated in terms of their prediction errors versus time along with their confidence bounds. The overall RMSE [cycles] is given in brackets.

along with the diagnostic information, which engine component is involved in the degradation.

There are several motivations to take the described hybrid approach for this use case: i) for safety-critical systems such as aircraft engines, a physics-based performance model or a surrogate model is generally available for control and performance evaluation from different stakeholders (e.g., OEM and O&M), ii) on the other hand, a lower-level representation of the physical process and, in particular, accurate damage models with a reasonable computation cost are not available, iii) the failure modes of interest result in a degraded performance which in turn produces discernible signatures on sensor readings, iv) the performance impact on the system components resulting from degradation of the system can be roughly represented by a set of known model parameters called 'health parameters.' v) the functional failure of interest has been observed and recorded for a small fleet of engines, i.e., run-to-failure trajectories in the form of time series sensor readings along with the corresponding time-to-failure labels are available.

Also here there are universal aspects of the approach, that may be relevant for other applications:

- The PIML approach provides accurate predictions of the RUL, even when available data sets to train data-driven models are sparse.
- Using physics-based performance models to enhance the interpretability of the degradation process in terms of physically meaningful features is recommended for the diagnosis and prognostics of safety-critical systems such as turbofan engines. This means that degradation can not only be interpreted in terms of its effect on the sensor readings but also in terms of known engine health parameters.
- The fidelity of the performance model in approximating the real process is not a prerequisite. It is assumed that complex physical processes cannot be modeled in full detail with reasonable computational cost. Therefore, the discrepancy term allows for the compensation of the unknown physics.

V. Guidelines for the Application of Physics-Informed Data Augmentation

Sections II to IV demonstrated three applications of PIML for the purpose of machine health monitoring and management. In all three examples, physical knowledge was exploited in order to augment the available data used as input to the model. However, each example took a different approach to data augmentation, selected according to the task, the data availability and the access to physics-based models. A summary of the comparison between the approaches is displayed in Table 1. Below we elucidate the advantages and prerequisites for the implementation of each approach.

A. *Synthetic Fault Generation from Healthy Field Data*

In use case I, the synthetic generation of faults requires rich and representative data from healthy regimes, but no labeled faults or anomalies. On the other hand, it requires a relatively basic physical model of the fault mechanism that serves to synthesize faults from field healthy data, thereby augmenting the healthy data and allowing to train a supervised classification model. The synthetic faults were found to be representative enough of the real faults, such that a DL model trained with the augmented data reached excellent performance on field data with real faults, outperforming pure data-driven models. This approach could be used ideally for anomaly detection or diagnostics, in cases of abundance of healthy data, when no labeled faults are available.

B. *Transfer Learning for Data-Driven Calibration of a Digital Twin*

In use case II, in contrast to the previous example, a high fidelity physical model was available and had been deployed in the past for design purposes. This is often the case for complex industrial machines, for which extensive domain knowledge has been accumulated over the years. The physical model is used for the development of a digital twin, which models the system behaviour under healthy conditions. In order to adapt the model to reproduce the behaviour of specific individual units, healthy data from these units must be collected and fed into a transfer learning discrepancy model. The resulting unit-specific model is used for monitoring degradation trends in the model residuals, thus allowing for interpretable fault detection and identification of anomalous sensor readings. This hybrid approach is suitable for degradation monitoring and diagnostics in case a well established physical model as well as healthy historical data from some fleet units is available. It can then be used in conjunction with relatively little field data from the individual units for which condition monitoring is needed.

C. *Performance Model Interpretation using Degradation Data*

Similarly to the previous case, in use case III, the availability of a physics-based performance model is a prerequisite. The main difference to use case II is that the training is done with full degradation trajectories and that machine health

System	**Solar power plants**	**Gas turbines**	**Aircraft engines**
Main idea	Synthetic fault generation from healthy field data	Transfer learning for data-driven calibration of a digital twin	Performance model interpretation using degradation data
Task	Anomaly detection	•Degradation trending •Fault detection and localization	RUL prediction with diagnostics
Data used for	•Healthy training data •Baseline for fault generation	•Calibration of fleet-based digital twin •Model transfer from fleet to individual units	•RUL Modelling •Identifying the root-cause of degradation
Physics used for	Synthetic fault generation from healthy data	A fleet–based digital twin	A digital twin, including a degradation model
Prerequisites of the approach	•Representative healthy field data •Simple approximate model of fault mechanism	•Physics based performance model •Small amount of healthy data from each unit	•Physics based performance model •Run-to-failure data of several units
Benefits of hybrid approach	•High accuracy with no fault data •Natural extension from detection to diagnostics •No need for high modelling accuracy of the physical model (complexity covered by the data)	•Early deployment with little field data •High acceptance of in-house domain experts •Interpretability •Support for design improvements	•Accurate RUL prediction with sparse training data •Interpretability of the degradation root-cause
Hybridization level	data physics	data physics	data physics

Table 1: comparison of PIML use-cases

parameters capturing the performance impact of the degradation on the machine components are inferred. The inferred parameters are used to predict the RUL of new units and, at the same time, provide insights into the root cause of degradation by identifying the faulty component. The hybridization of a physical performance model with a data-driven DL model allows for enhanced interpretability of the degradation process and requires significantly less training data than a purely data-driven approach. This approach is suitable for RUL prediction and diagnostics in setups for which a physics-based performance model, as well as run-to-failure data from some units, are available.

VI. Summary

In this paper we described three application cases of PIML for condition-based maintenance tasks, based on augmenting and supplementing the condition monitoring data. As shown in Table 1, the use cases differ not only in their application fields, but also in their tasks, their benefits and the level of data-physics hybridization (indicated by the small triangles along the data/physics scale). It thus becomes evident that physics-informed data augmentation for ML algorithms has a broad range of applications and can be used for one or more of the following purposes: (i) enhancing prediction accuracy despite data scarcity (ii) enhancing accuracy compared to incomplete physical models by using the data variability (iii) gaining interpretability compared to a pure data-driven approach (iv) enhancing the acceptance of the local domain experts (v) supporting improvements in the design and operation of machines.

We believe that the universal aspects of the described methods are relevant and applicable for condition-based and predictive maintenance purposes in various systems and application fields, for which both data and physical knowledge are available.

References

[1] S. R. Saufi, Z. A. B. Ahmad, M. S. Leong, M. H. Lim, Challenges and opportunities of deep learning models for machinery fault detection and diagnosis: A review, Ieee Access 7 (2019) 122644–122662.

[2] R. Liu, B. Yang, E. Zio, X. Chen, Artificial intelligence for fault diagnosis of rotating machinery: A review, Mechanical Systems and Signal Processing 108 (2018) 33–47.

[3] L. Zhang, J. Lin, B. Liu, Z. Zhang, X. Yan, M. Wei, A review on deep learning applications in prognostics and health management, Ieee Access 7 (2019) 162415–162438.

[4] O. Fink, Q. Wang, M. Svensen, P. Dersin, W.-J. Lee, M. Ducoffe, Potential, challenges and future directions for deep learning in prognostics and health management applications, Engineering Applications of Artificial Intelligence 92 (2020) 103678.

[5] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nature Reviews Physics 3 (6) (2021) 422–440.

[6] S. Frank, M. Heaney, X. Jin, J. Robertson, H. Cheung, R. Elmore, G. Henze, Hybrid model-based and data-driven fault detection and diagnostics for commercial buildings, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States) (2016).

[7] M. Arias Chao, Combining deep learning and physics-based performance models for diagnostics and prognostics, Doctoral thesis, ETH Zurich, Zurich (2021). doi:10.3929/ethz-b-000517153.

[8] M. A. Chao, C. Kulkarni, K. Goebel, O. Fink, Fusing physics-based and deep learning models for prognostics, Reliability Engineering & System Safety 217 (2022) 107961.

[9] S. Racharla, K. Rajan, Solar tracking system–a review, International journal of sustainable engineering 10 (2) (2017) 72–81.

[10] A. Triki-Lahiani, A. B.-B. Abdelghani, I. Slama-Belkhodja, Fault detection and monitoring systems for photovoltaic installations: A review, Renewable and Sustainable Energy Reviews 82 (2018) 2680–2692.

[11] J. Zgraggen, Y. Guo, A. Notaristefano, L. Goren Huber, Physics informed deep learning for tracker fault detection in photovoltaic power plants, in: 14th Annual Conference of the Prognostics and Health Management Society, Nashville, USA, 1-4 November 2022, Vol. 14, PHM Society, 2022.

[12] H. Saravanamuttoo, G. Rogers, H. Cohen, Gas Turbine Theory, Prentice Hall, 2001.

[13] L. A. . Urban, Gas path analysis applied to turbine engine condition monitoring.

[14] M. J. W. P. Wand, Kernel Smoothing, London, Chapman Hall, (1995).

[15] T. Palmé, F. Liard, D. Cameron, Hybrid modeling of heavy duty gas turbines for on-line performance monitoring, in: Turbo Expo: Power

for Land, Sea, and Air, Vol. 45752, American Society of Mechanical Engineers, 2014, p. V006T06A011.

[16] M. Arias Chao, C. Kulkarni, K. Goebel, O. Fink, Aircraft engine run-to-failure dataset under real flight conditions for prognostics and diagnostics, Data 6 (1) (2021) 5.