Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
 Published by Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P652-cd

Fully Unsupervised Fault Detection in Solar Power Plants using Physics-Informed Deep Learning

Jannik Zgraggen
School of Engineering, Zurich University of Applied Sciences, Switzerland. E-mail: jannik.zgraggen@zhaw.ch

Yuyan Guo
Fluence Energy, Switzerland. E-mail: yuyan.guo@fluenceenergy.com

Antonio Notaristefano
Fluence Energy, Switzerland. E-mail: antonio.notaristefano@fluenceenergy.com

Lilach Goren Huber
School of Engineering, Zurich University of Applied Sciences, Switzerland.
E-mail: lilach.gorenhuber@zhaw.ch

Machine learning algorithms for anomaly detection often assume training with historical data gathered under normal conditions, and detect anomalies based on large residuals at inference time. In real-world applications, labelled anomaly-free data is most often unavailable. In fact, a common situation is that the training data is contaminated with an unknown fraction of anomalies or faults of the same type we aim to detect. In this case, training residual-based models with the contaminated data often leads to increased missed detections and/or false alarms. While this challenge is rather common, in particular in technical fault detection setups, it is only rarely addressed in the scientific literature.
In this paper we address this problem by introducing a data refinement algorithm that is capable of cleaning the contaminated training data in a fully unsupervised manner, and apply the algorithm to a problem of fault detection in grid-scale solar power plants. The data refinement framework is based on an original physics informed deep learning classification algorithm that would require healthy data as its input, in order to generate from it synthetic faulty data and train a binary classifier. We show that in order to achieve high fault detection performance, it is essential to avoid contamination of the original healthy data with unlabelled faults. To this end, we introduce an algorithm that isolates the healthy data in a fully unsupervised manner prior to training the binary classifier. We test our algorithm with field data from an operational solar power plant which includes contamination of unlabelled faulty data and demonstrate its high performance. In addition, we demonstrate the robustness of the proposed refinement method against an increasing fraction of faults in the training data.

Keywords: Unsupervised Machine Learning, Deep Learning, Anomaly Detection, Data contamination, Data Refinement, Predictive Maintenance, Physics Informed Machine Learning, Solar Power.

1. Introduction

One of the central challenges in operational deployment of machine learning (ML) algorithms for fault detection and isolation is the scarcity of labelled data from faulty conditions Fink et al. (2020). This challenge is typically addressed by applying anomaly detection algorithms that rely on having anomaly-free training data. A model is trained to predict the system behaviour under normal (healthy) conditions, and anomalies or faults are detected by their dissimilarity to the predicted behaviour. Common examples for such residual-based models are various types of autoencoder neural networks (see for example Munir et al. (2018); Audibert et al. (2020); Provotar et al. (2019); Thill et al. (2021); Ferdousi and Maeda (2006); Zhang et al. (2019)). These approaches are often referred to as ”unsupervised anomaly detection” despite the fact that partial supervision is assumed in terms of healthy labels. Thus, a more appropriate term for this family of methods would be ”semi-supervised learning”.

Another approach, gaining attention recently, is to exploit physical knowledge in combination with ML models in order to augment the healthy data with synthetically generated faults Karniadakis et al. (2021); Frank et al. (2016). Generating an unlimited number of faults is attractive, as it allows the use of supervised classification methods.

Both of the aforementioned approaches to deal with the scarcity of real labelled faults require access to labelled healthy data. In one case, the healthy data is used for training normal state prediction models, and in the second case it is used for the synthetic generation of faults.

However, in an operational industrial environment there is usually no guarantee that labelled fault-free data with no (known or unknown) fault contamination is available for training. In fact, the most common situation is the one with no labels at all. In this case the only viable solutions are truly unsupervised anomaly detection algorithms. Despite the high relevance of this topic, very little research has been carried out in this direction until now Beggel et al. (2020); Zhou and Paffenroth (2017); Berg et al. (2019); Yoon et al. (2021); Qiu et al. (2022) .

In this paper we suggest a truly unsupervised fault detection algorithm based on a physics-informed deep learning (PIDL) approach. The suggested framework includes a two step procedure. In the first step we apply the PIDL on all available data, which is assumed to contain both healthy and faulty samples. We develop a refinement score that allows a clear separation between the healthy and faulty samples of the original data. After the data refinement, we retrain the same PIDL algorithm using only the refined fault-free data. We demonstrate the approach on field data from an operational solar power plant, and focus on the detection of faults in the solar tracking system.

Tracker faults are common in solar power plants. Solar trackers are rotating devices that are used to keep the solar panels mounted on them at an optimal angle considering the sun position at any given moment Racharla and Rajan (2017). The most common fault mechanism of the tracking system is that a tracker gets stuck at a certain rotation angle instead of tracking the sun. In this case, all solar panels that are mounted on a faulty tracker suffer from power production losses compared to the optimal production of the healthy system. The effect of a tracker fault on the daily power production profile of a solar string (typically containing a multitude of panels) can strongly vary not only with the solar position but also with varying weather conditions.

In a previous publication we showed that PIDL is a powerful method to detect tracker anomalies. The approach uses healthy power profiles from field data and corrupts them based on a physical model in order to generate power profiles corresponding to tracker faults. The augmented data, containing field healthy samples and synthesized faults, is used to train a binary classifier, assigning anomaly scores to the daily power profile of each solar string.

In this paper we report the extension of this case study to a truly unsupervised setup where fault-free field data is not necessarily available. Instead, the starting point of the present study is a data set which may or may not contain a considerable fraction of contamination of faults of the same type we wish to detect.

We evaluate the performance of the algorithm on contaminated field test data, and show that it significantly outperforms the approach of ”blind” training on the contaminated data. In addition we perform tests with various contamination ratios, introduced by synthetic generation of faults, in order to evaluate the robustness of the method towards an increasing number of faults in the training data.

The contribution of this paper is the following:

- We tackle the problem of truly unsupervised anomaly detection, that is, not relying on access to anomaly-free training data. This topic is rarely addressed in the scientific literature, despite its high relevance for most anomaly detection setups and in particular, across all prognostics and health management (PHM) applications.

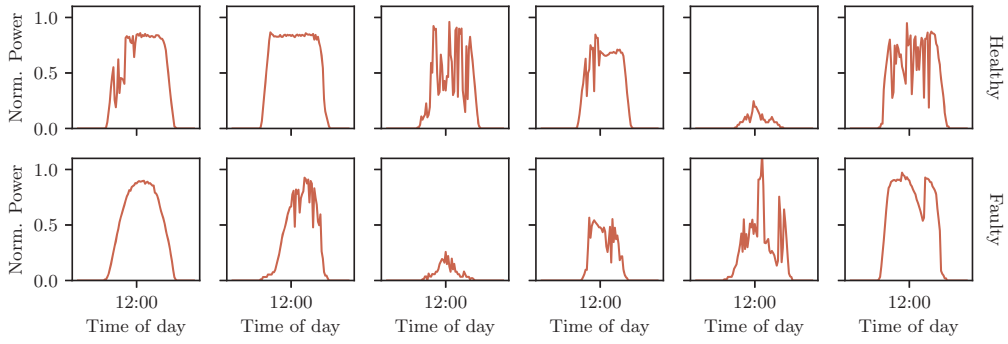

Fig. 1. Daily power profiles from an operational solar power plant. The normalized power is plotted against the time of the day. Upper row: Healthy profiles (no tracker fault). Lower row: Faulty profiles (with known tracker faults).

- We develop an end-to-end framework for solar-tracker fault detection which is applicable to mixed normal/abnormal data with no labels at all. This framework is inherently a physics-informed machine learning approach. It can be easily extended to other application fields in which abundant healthy data is expected to be available, and the fault mechanism can be modeled using a simple physics-based transformation on the healthy data.

2. Description of the use case

The proposed unsupervised data refinement method is demonstrated here on a real field data set from a solar power plant, with the aim of detecting faults in the solar trackers. The data set contains a single variable, the output power, measured at each solar string in a grid-scale power plant. In this power plant, each solar tracker carries 5 solar strings, whose power production is affected in case the respective tracker is faulty.

The initial data set $\mathcal{D}_0$ contains daily power profiles measured from healthy solar strings as well as from strings that are mounted on faulty trackers. Some examples of such (normalized) power profiles are shown in Figure 1. The upper row displays healthy profiles and the lower row faulty profiles. Due to the strong heterogeneity in weather conditions as well as individual degradation history of the solar strings, the profile shapes are very heterogeneous and it is not a trivial task to distinguish between the healthy and the faulty ones.

In the following we describe the details of an unsupervised data refinement (USDR) approach that enables the detection of faulty daily power profiles at the single string level at the end of each day.

3. Method

We assume an initial field data set $\mathcal{D}_0 = \{x_{0i}\}_{i=1}^{N}$ containing N unlabelled normal as well as abnormal samples. The goal of our method is to train a physics informed deep learning (PIDL) binary classifier, that will be able to detect abnormal samples in an unseen test data $\mathcal{D}_f$.

Given the initial unlabelled field data set $\mathcal{D}_0$, the USDR method includes the following steps, as depicted in Figure 2:

(1) Data refinement (left frame in Figure 2).
(2) Retraining the classifier on the refined data for anomaly detection (right frame in Figure 2.

Both steps are executed using a PIDL algorithm which is itself comprised of (i) physics informed (PI) data augmentation and (ii) training a deep learning (DL) classifier of the augmented data.

In the refinement step (1), the classifier is used to obtain a refinement score for each sample of the initial data set. Based on the refinement score,

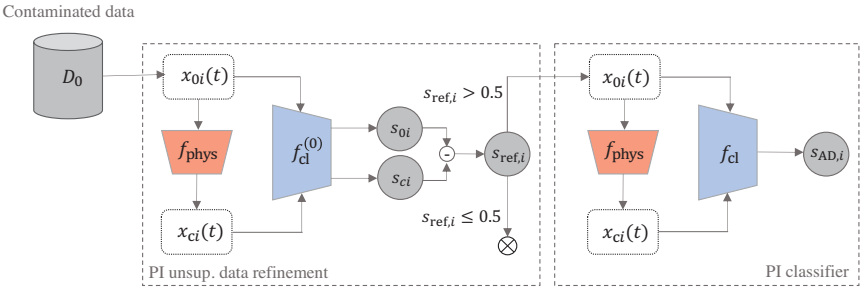

Fig. 2. A framework for unsupervised fault detection using physics informed deep learning (PIDL). Each time series sample $x_{0i}(t)$ out of an unlabelled training set is assigned a refinement score $s_{\text{ref},i}$ (left panel). The refinement score for sample i is calculated based on the difference in classification outputs between the original sample $x_{0i}(t)$ and its corrupted counterpart $x_{ci}(t)$, obtained by applying the physical model f_{phys}. All samples with $s_{\text{ref},i} > 0.5$ are assigned a "healthy" label and are used to train a PIDL classifier (right panel).

part of the samples are selected as normal and are propagated to the next step (2) of training the anomaly classifier.

The next subsections describe all of the above steps.

3.1. *Physics Informed Data Augmentation*

In our previous work Zgraggen et al. (2022) we described an algorithm that detects tracker faults by classifying the daily power profiles of the individual solar strings. The starting point of the algorithm is a data set with labelled healthy power profiles from an operational power plant. The (presumably healthy) profiles of the field data $\{x_{0i}(t)\}_{i=1}^{N}$ are then corrupted by a physical model f_{phys} in order to simulate the effect of various possible tracker faults and to generate a set of synthetic faulty profiles $\{x_{ci}\}_{i=1}^{N}$ (subscript c stands for "corrupted"). A faulty profile $x_c(t)$ is generated from a healthy profile $x_0(t)$ using the equations

$$\begin{aligned} x_c(t) &= c_p\left[(1-\gamma)g(\theta_0,\theta_{\text{i}}^*(t))+\gamma\right]x_0(t) \\ g(\theta_0,\theta_{\text{i}}^*(t)) &= \frac{\cos\theta_0\cdot f_{\text{IAM}}(\theta_0)}{\cos\theta_{\text{i}}^*(t)\cdot f_{\text{IAM}}(\theta_{\text{i}}^*(t))}. \end{aligned} \tag{1}$$

with $f_{\text{IAM}}(\theta_{\text{i}}) = 1 - b_0(1/\cos\theta_{\text{i}} - 1)$ and where $\theta_{\text{i}}^*(t)$ is the optimal tilt angle of the tracker at time t, θ_0 is the stuck angle of the faulty tracker, b_0 and γ are model parameters estimated empirically using the data, and c_p is a degradation loss coefficient. By sampling the values of θ_0, γ, b_0 and c_p from uniform distributions within realistic ranges, we generate a rich collection of possible faulty profiles. For details of the physical model we refer the reader to Zgraggen et al. (2022).

The empirical-physical model of the fault mechanism is used to augment the normal data set, such that it now contains both healthy and faulty power profiles. Each daily profile is a time-series of size 96 (due to a 15 minute resolution). At a next step, the augmented data set, containing balanced healthy and faulty samples is pre-processed by subtracting from each power profile the 0.9 quantile over the entire plant at any given moment in time. The normalized profiles are used to train a 1d-CNN classifier f_{cl} that assigns an anomaly score s_{AD} to each daily profile. This allows to detect faulty strings at the end of each day (which is the relevant time resolution for decision making in practice). The CNN contains three one-dimensional convolutional layers followed by two fully-connected layers, with a total of around 30'000 trainable parameters. The network architecture was optimized using a grid search to tune the number of layers and filters and the learning rate.

In our previous work the PIDL algorithm was shown to outperform a pure data driven approach based on a convolutional autoencoder, both in terms of accuracy and in terms of robustness of the outcomes Zgraggen et al. (2022).

Here we extend the previous framework to a fully unsupervised anomaly detection scenario. The crucial difference is that we do not assume any prior knowledge about the initial data set $\mathcal{D}_0$, which can now include a non-negligible ratio of faulty power profiles in addition to the healthy ones, also known as contamination ratio r_c. Since this is the common scenario in real applications, we use a field data set of one year from a grid-scale solar power plant as our starting point for training the fault detection algorithm.

3.2. *Unsupervised Data Refinement*

The use of the original PIDL classification algorithm described above assumes a data set of healthy (normal) samples, from which faulty (abnormal) samples can be synthesized. Here, on the other hand, we assume an initial data set with contaminated data, containing healthy and faulty samples. In order to be able to use the PIDL classifier we apply a step of data refinement prior to the classification. The data refinement step itself utilizes the same PIDL algorithm and applies it "blindly" on the initial contaminated data, using the resulting score to extract meaningful information about each sample. This physics informed (PI) unsupervised data refinement (USDR) step is shown as the left frame in Figure 2.

Note that we use the term "unsupervised" even though we apply a standard classification framework. The reason is that we use the classifier assuming that the training data is partially mislabelled due to fault contamination, such that its true labels are unknown. The only underlying assumption is that the majority of the original data belongs to the normal (healthy) class, which is a standard assumption in any anomaly detection task.

The essential step of the data refinement is aimed at assigning a refinement score $s_{\mathrm{ref},i}$ to each sample $x_{0i}(t)$ of the unlabelled data set $\mathcal{D}_0$. This score is an indicator of the likelihood of this sample to be healthy. At the next step only samples with a high likelihood to be healthy are used to retrain the PIDL classifier. In the following we explain the individual steps of this approach.

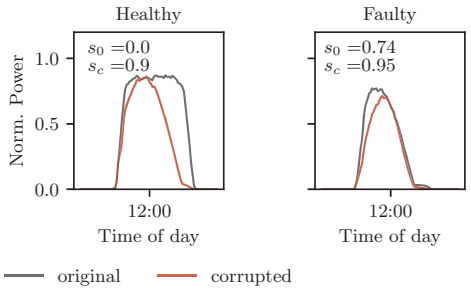

Fig. 3. Examples for physics-informed data corruption. Left: a daily power profile of a healthy solar string before (black) and after (red) the synthetic fault generation. Right: the same for a faulty string. s_0 and s_c are the classification scores before and after corruption respectively.

3.2.1. *Refinement scores*

We denote an arbitrary original sample (drawn from the field data $\mathcal{D}_0$) as $x_{0i}(t)$, where $t = 1...96$ is the discretized time of the day. Using the physical model f_{phys} we can now corrupt this sample to generate the corrupted power profile $x_{ci}(t)$. We note that if the original sample $x_{0i}(t)$ is healthy, applying the physical corruption model to it generates a sample $x_{ci}(t)$ with a characteristic tracker fault. An example for such a sample can be seen in the left panel of Figure 3, where the original (black) and the corrupted (red) profiles are shown. However, since the data $\mathcal{D}_0$ is contaminated with unlabelled faults, the original sample may also be faulty. In this case the physical corruption model applied to it generates a power profile $x_{ci}(t)$ which is no longer characteristic to a faulty string, nor to a healthy string. The right panel in Figure 3 displays an example of the original (black) and corrupted (red) samples in a case where the original sample is already faulty.

Both original and corrupted samples are used to train the refinement classifier $f_{\mathrm{cl}}^{(0)}$. This binary classifier is trained to predict 0 for $x_{0i}(t)$ and 1 for $x_{ci}(t)$ (even though some of the samples may be mislabelled). After training the classifier, we evaluate the predicted classification scores of all samples in the training data itself, both before and after corruption. For each sample, we denote the

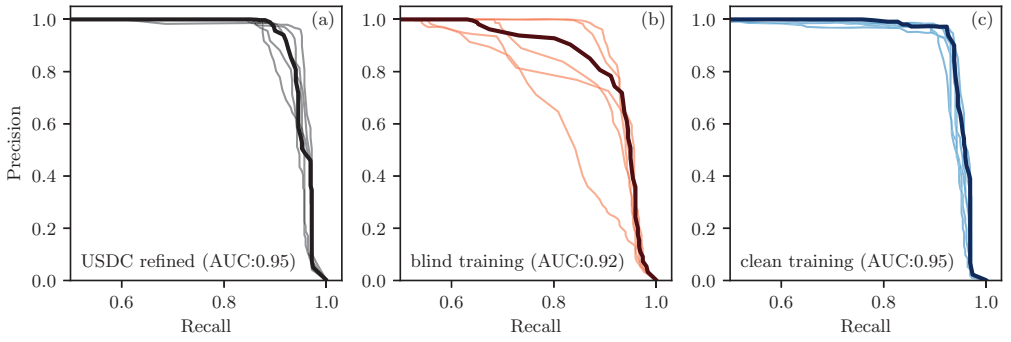

Fig. 4. Fault detection performance. Precision recall plots for field test data of (a) the proposed PI USDR method (b) a PI classifier trained blindly on the contaminated data (c) a PI classifier trained on labelled clean data for reference. For each model we show the results of 5 training runs and their mean (thick curve). The AUC of the mean curve is given in brackets.

score of the original (uncorrupted) sample $x_{0i}(t)$ as s_{0i}, and the score of its corrupted counterpart $x_{ci}(t)$ as s_{ci}. We then define a new *refinement anomaly score* for each sample based on the difference between the two scores:

$$s_{\mathrm{ref},i} = s_{ci} - s_{0i} \tag{2}$$

A healthy sample i is likely to obtain a high refinement score (close to 1) whereas a faulty sample obtains a low refinement score (0.5 or lower). The reason behind it is the following: if the original sample $x_{0i}(t)$ is healthy, it should get a low classification score $s_{0i} \sim 0$. When corrupting it, we generate a synthetic faulty profile $x_{ci}(t)$ which should obtain a high classification score $s_{ci} \sim 1$ (see left panel of Figure 3). The difference between them is expected to be $s_{\mathrm{ref},i} \sim 1$. However, if the original sample $x_{0i}(t)$ is faulty, it gets a high classification score $s_{0i} \sim 1$. Corrupting it will generate a profile $x_{ci}(t)$, that is expected to be far from the typical healthy samples, and get a classification score $0.5 < s_{ci} < 1$ (see right panel of Figure 3). The score difference in this case is expected to be $s_{\mathrm{ref},i} \sim 0$, and may also be negative.

3.2.2. *Data refinement*

We clean the contaminated data set $\mathcal{D}_0$ by setting a threshold on the refinement score at 0.5, and removing any original sample $x_{0i}(t)$ with $s_{\mathrm{ref},i} <$ 0.5 from the data set, assuming that this sample is likely to be faulty. The refined data set, denoted by $\mathcal{D}$ contains now only samples which are likely to represent healthy strings.

3.2.3. *Retraining with refined data*

The selected power profiles $x_{0i}(t)$ are now fed into the physical model f_{phys} in order to generate faulty profiles $x_{ci}(t)$. As a result we obtain an augmented data set with samples $x_{0i}(t)$ and $x_{ci}(t)$ that are likely to be healthy and faulty respectively. The resulting augmented data set is fed into the CNN classifier f_{cl} as training data. The trained model can now be used to infer the anomaly scores $s_{\mathrm{AD},i}$ for new unseen data. The anomaly scores $s_{\mathrm{AD},i}$ range between 0 (healthy) and 1 (faulty). Depending on the selected threshold, we obtain different fault detection outputs. To evaluate the outcomes and compare models we typically look at precision-recall curves, which are the most informative for highly unbalanced data, and calculate the area under them (AUC) as a performance metric.

4. Results

To evaluate the performance of the unsupervised data refinement (USDR) method we use test data of 18 Months from the same operational solar plant used for training. During this year we manually labelled 417 daily profiles as clearly suffer-

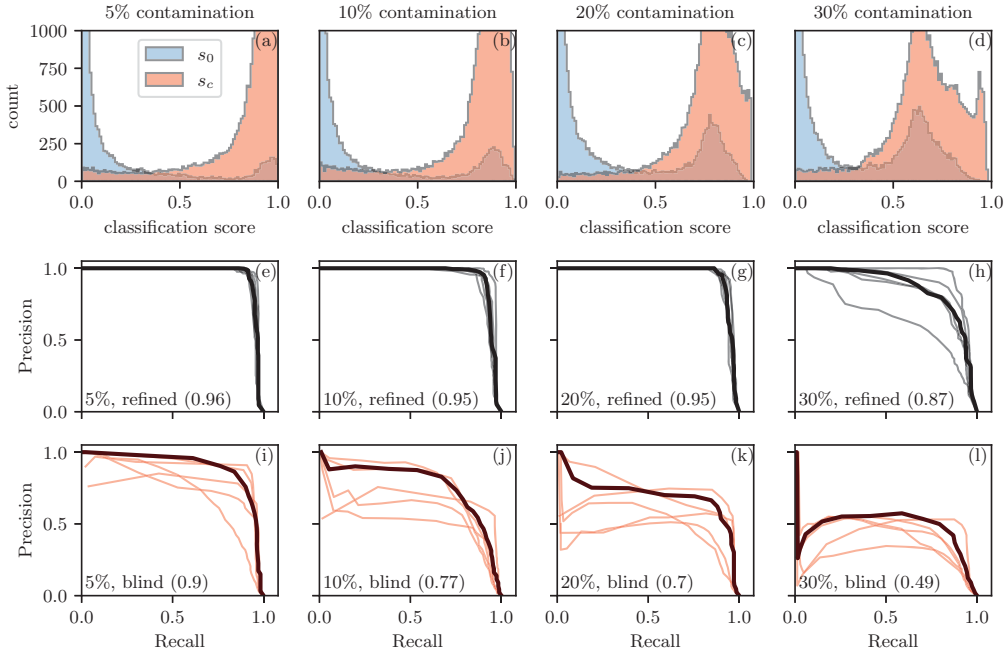

Fig. 5. The effect of the contamination ratio (CR) of the training data on the performance of the USDR algorithm. Panels (a)-(d): distributions of the classification scores s_{0i} and s_{ci} of the original and corrupted training data samples respectively for CR of $5\%, 10\%, 20\%$ and 30%. Panels (e)-(h): precision-recall curves (PRCs) trained with the USDR algorithm for the same CR values. Panels (i)-(l): baseline PRCs for blind training with the contaminated data for the same CR values. Each PRC panel contains 5 training repetitions and their mean (thick curve), with AUC of the mean in brackets. All PRCs are evaluated on the same field test set.

ing from tracker faults, and 334'176 profiles as healthy. Other known fault types or unclear cases were filtered out from the training and the test data.

Figure 4(a) shows the precision-recall (PRC) plot achieved by the USDR framework. In order to evaluate the reproducibility of the algorithm we show the results of 5 training repetitions and their mean. The result is compared with two different baselines. In panel (b) the "blind" baseline, training the classifier blindly without refining the training data. In panel (c) we show the results of the ideal case, in which the data is *manually* cleaned prior to training, such that it contains exclusively fault-free samples, with no contamination.

We observe that the USDR framework is clearly superior to blind training with contaminated data, with $AUC = 0.95$ compared to $AUC = 0.92$ of the mean of all blind baselines. More importantly we notice that the variability between training repetitions of the same model is very large in case of blind training as opposed to the robust outcomes of the refinement method USDR. Some of the blind training repetitions are severely impaired by the fault contamination of the training data whereas other are less affected. Both the fault detection performance and the reproducibility of the USDR framework are close to the ideal model trained with fault-free data ($AUC = 0.95$). For practical applications reproducibility is a central performance metric that indicates the robustness of the model towards random changes in the training data or stochastic properties of the training process.

4.1. *The effect of the contamination ratio.*

Similar to other unsupervised anomaly detection methods, the performance of the USDR method

strongly depends on the initial contamination ratio of the data, that is, the fraction of faulty samples in the initial data set. In order to test the sensitivity of the suggested method towards the contamination ratio, we introduced a growing fraction of faults into the initial data set by randomly selecting healthy samples and synthetically corrupting them. In this way we could control the effect of the contamination ratio without reducing the number and thus the representativeness level of the healthy samples. For each contamination ratio we applied the USDR framwork and evaluated the PRC performance for the same test data set, containing 18 months of field data of an operational solar power plant.

Figure 5 displays the results of the comparison for 4 different contamination ratios between 5% and 30%. Panels (a)-(d) in the upper row of the figure display the histograms of the anomaly scores s_{0i} (blue) and s_{ci} (red) of the samples before and after corruption respectively. As expected, we observe a majority of $s_{0i} \sim 0$ with a smaller peak at $0.5 < s_{0i} < 1$ which grows with the growing contamination ratio. These are the faulty samples of the initial data. After corruption, the majority earns the score $s_{ci} \sim 1$, with a tail towards $s_{ci} = 0$ that becomes more significant as the contamination ratio grows.

Panels (e)-(h) in the middle row display the PRC curves achieved by the USDR framework for the different contamination ratios. Using the data refinement framework, we observe almost no performance deterioration up to contamination ratio of 30%. The AUC of the 5-run mean is given in brackets in each panel.

Comparing the results to the blind training baseline in panels (i)-(l), it is evident that the refinement step leads to a significant improvement of the fault detection performance. Moreover, the suggested method obtains remarkably reproducible results also at high contamination ratios. We note that a contamination ratio of 20% anomalies in the initial data is considerably higher than any realistic scenario for the use case we consider.

5. Conclusions

Anomaly or fault detection tasks are typically addressed with semi-supervised residual-based algorithms, which assume the availability of fault-free training data. In real applications this assumption is rarely valid. Thus, there is a need for truly unsupervised anomaly detection methods, that are applicable to contaminated training data. Despite their high relevance, such methods are only rarely addressed in the scientific literature.

In this paper we presented a framework for fully unsupervised anomaly detection using physics-informed deep learning, introducing a data refinement step that cleans the contaminated data prior to the anomaly classification step. We demonstrated the performance of the framework on field data from an operational solar power plant with tracker faults. We showed that whereas training blindly with the contaminated data severely impairs the performance, the suggested framework is able to reach high detection performance and robustness, which are comparable to the ideal (but rather unrealistic) case in which fault-free training data is available.

Our future research will validate the results against other approaches for anomaly detection with contaminated data. In addition we will extend the applicability of the framework for further fault types, as well as for data sets with mixed fault types in order to allow fault diagnostics.

Acknowledgment

This research was funded by Innosuisse - Swiss Innovation Agency under grant No. 55018.1 IP-ICT.

References

Audibert, J., P. Michiardi, F. Guyard, S. Marti, and M. A. Zuluaga (2020). Usad: Unsupervised anomaly detection on multivariate time series. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 3395–3404.

Beggel, L., M. Pfeiffer, and B. Bischl (2020). Robust anomaly detection in images using adversarial autoencoders. In *Machine Learning and Knowledge Discovery in Databases: European*

Conference, ECML PKDD 2019, Würzburg, Germany, September 16–20, 2019, Proceedings, Part I, pp. 206–222. Springer.

Berg, A., J. Ahlberg, and M. Felsberg (2019). Unsupervised learning of anomaly detection from contaminated image data using simultaneous encoder training. *arXiv preprint arXiv:1905.11034*.

Ferdousi, Z. and A. Maeda (2006). Unsupervised outlier detection in time series data. In *22nd International Conference on Data Engineering Workshops (ICDEW'06)*, pp. x121–x121. IEEE.

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

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

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

Munir, M., S. A. Siddiqui, A. Dengel, and S. Ahmed (2018). Deepant: A deep learning approach for unsupervised anomaly detection in time series. *Ieee Access 7*, 1991–2005.

Provotar, O. I., Y. M. Linder, and M. M. Veres (2019). Unsupervised anomaly detection in time series using lstm-based autoencoders. In *2019 IEEE International Conference on Advanced Trends in Information Theory (ATIT)*, pp. 513–517. IEEE.

Qiu, C., A. Li, M. Kloft, M. Rudolph, and S. Mandt (2022). Latent outlier exposure for anomaly detection with contaminated data. In *International Conference on Machine Learning*, pp. 18153–18167. PMLR.

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

Thill, M., W. Konen, H. Wang, and T. Bäck (2021). Temporal convolutional autoencoder for unsupervised anomaly detection in time series. *Applied Soft Computing 112*, 107751.

Yoon, J., K. Sohn, C.-L. Li, S. O. Arik, C.-Y. Lee, and T. Pfister (2021). Self-trained one-class classification for unsupervised anomaly detection. *arXiv e-prints*, arXiv–2106.

Zgraggen, J., Y. Guo, A. Notaristefano, and L. Goren Huber (2022). 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*, Volume 14. PHM Society.

Zhang, C., D. Song, Y. Chen, X. Feng, C. Lumezanu, W. Cheng, J. Ni, B. Zong, H. Chen, and N. V. Chawla (2019). A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In *Proceedings of the AAAI conference on artificial intelligence*, Volume 33, pp. 1409–1416.

Zhou, C. and R. C. Paffenroth (2017). Anomaly detection with robust deep autoencoders. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 665–674.