

# Anomaly detection in irregular image sequences for concentrated solar power plants

Sukanya Patra, Le Thi Khanh Hien and Souhaib Ben Taieb \*

Faculty of Science, University of Mons, Belgium

**Abstract.** Operations at extremely high temperatures can lead to various malfunctions in Concentrated Solar Power (CSP) plants, emphasizing the need for predictive maintenance (PdM). We study PdM as an anomaly detection (AD) problem from irregular image sequences, which represent the minute-by-minute solar receiver's surface temperature from a CSP plant. Contrary to standard benchmark image datasets in AD research, our data shows distinct characteristics such as non-stationarity, temporal dependence, and irregular sampling, which are unaddressed by current image-based AD techniques. Therefore, we introduce a forecast-based AD method to address these characteristics, drawing inspiration from irregular sequence modelling. The results show that the proposed method outperforms classical image-based AD methods on our dataset.

## 1 Introduction

There has been a growing focus on renewable energy sources to counteract the devastating effects of climate change. A key challenge in adapting renewable energy is generation on demand. With the help of CSP plants, power sectors are utilizing Thermal Energy Storage (TES) as a solution to this challenge. CSP plants accumulate energy by using arrays of mirrors to concentrate sunlight onto *solar receivers* placed on top of solar towers, and TES temporarily stores the energy by heating or cooling a storage medium such as water or molten salt. However, due to operation in extreme temperatures, the solar receivers are impacted in multiple ways, such as deformation, corrosion of the heat exchanger tubes, and freezing of the medium blocking the heat exchanger tubes. Therefore, CSP plants require close monitoring to ensure uninterrupted power generation.

PdM refers to techniques that estimate the condition of a machine in-service to predict necessary repairs. We reduce PdM in CSP plants to an AD problem on the solar receiver dataset (see Section 2). AD is defined as the process of detecting rare data that are different from the normal behaviour of the data set [1]. Through exploratory data analysis, we identified some key characteristics of the data from the CSP plants, namely *non-stationarity*, *temporal dependence* and *irregular sampling*. These characteristics distinguish our data from the classical image benchmark datasets used in the AD literature and need to be addressed for an effective PdM.

---

\*This work is a part of the FLARACC research project, a MecaTech Cluster program funded by the Walloon Region, Belgium.

To address the characteristics mentioned above of the data, we propose a forecast-based AD [2] method addressing the challenges in the context of PdM of CSP plants drawing inspiration from research on irregular sequence modelling. Furthermore, we compare the proposed method with classical image-based anomaly detection methods.

## 2 Solar receiver dataset and challenges

The solar receiver consists of multiple solar panels. Each panel contains vertical heat exchanger tubes through which the heat transfer medium flows, absorbing the heat from the concentrated sun rays. The data<sup>1</sup> is collected using Infrared (IR) cameras placed at various locations of the solar receiver. IR cameras recorded the solar receiver’s surface temperature, creating a *heat profile* represented by an “image” roughly at an interval of one minute. The unlabelled dataset contains a heat profile which is a two-dimensional (2D) matrix of size  $184 \times 608$ , and is accompanied by the timestamp when it was recorded. For the sake of simplicity, we will use the terms heat profile and image interchangeably. Under normal operating conditions, as the heat transfer medium flows from one end of the panel to the other through the vertical tubes, its temperature rise due to the heat from the concentrated sun rays. Thus, the surface temperature recorded by the IR cameras is expected to have a smooth gradient in the same direction as the flow of the medium. Based on this knowledge and following discussion with domain experts, a small subset of the dataset is manually labelled for testing purposes. Examples of two normal and two anomalous samples, as provided by the domain experts, are shown on the left and right of Figure 1a, respectively. For example, in the anomalous sample on the top-right of the figure, we can observe a streak of high temperature, which is unexpected under normal operating conditions.

**Data characteristics.** Through extensive data analysis, we identified the following characteristics:

- C1. Non-stationarity.** The distribution of image temperature changes over time. Figure 1c shows the average of *heat profiles* recorded at various points in time over a particular week. It can be observed that the data exhibits daily seasonality.
- C2. Temporal dependence.** The heat profiles are highly dependent on recent weather conditions, leading to interdependence on sequential images and daily seasonal trends.
- C3. Irregular sampling.** The images are recorded in irregular time intervals. Due to different factors such as equipment failure, and non-operational periods of the CSP plant, there can be no recorded data for certain days. Figure 1b shows the distribution of the inter-arrival times.

---

<sup>1</sup>The dataset used for this study is confidential and thus cannot be released publicly. Please reach out to the authors for any clarification regarding the reported results.

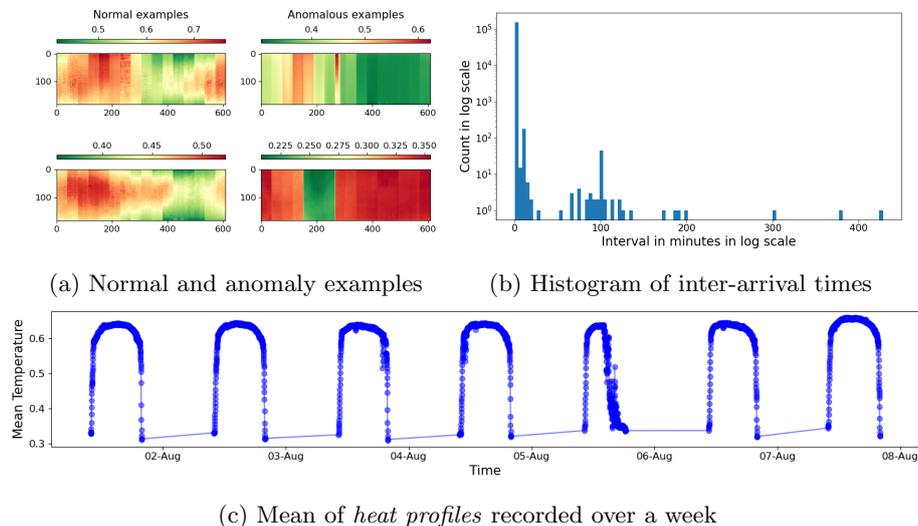


Fig. 1: Visualisation of different properties of the data

### 3 Anomaly detection methods

Consider an unlabelled dataset  $\mathcal{D}_u := \{(x_i^u, t_i^u)\}_{i=1}^{n_u}$  consisting of  $n_u$  samples. Each sample consists of an input  $x_i^u \in \mathcal{X} = \mathbb{R}_+^d$  where  $d = 184 \times 608$  is the dimension of the input space. The corresponding timestamp  $t_i^u$  is drawn from  $\mathcal{T} = \mathbb{R}_+$ , denoting the time when the sample  $x_i^u$  was recorded. Given the unlabelled dataset  $\mathcal{D}_u$ , our problem can be formulated as a unsupervised AD problem. For evaluating the performance of all methods, an annotated dataset  $\mathcal{D}_v := \{(x_i^v, t_i^v, y_i^v)\}_{i=1}^{n_v}$  is kept aside consisting of  $n_v$  samples. Each sample in  $\mathcal{D}_v$  also contains the ground truth  $y_i^v \in \mathcal{Y} := \{0, 1\}$  corresponding to the input  $x_i^v$ . The samples are labelled in consultation with the domain experts. Here anomalous samples are labelled 1, and normal samples are labelled 0.

#### 3.1 Image-based anomaly detection

**Global models.** We start by adopting methods from the extensive literature on image AD, disregarding the temporal aspects of the dataset. Primarily, we focus on deep learning-based methods, namely autoencoder and Deep Support Vector Data Description (DeepSVDD) [3]. Deep methods have proven more effective than shallow ones for image anomaly detection [1], leveraging deep neural networks' ability to extract representative features through multiple abstraction layers. This study considers both SoftBoundary and OneClass variants of DeepSVDD, as proposed by Ruff et al. [3]. The objective is to learn a mapping function  $\Phi(\cdot; W) : \mathcal{X} \rightarrow \mathcal{F}$ , using a deep neural network parameterized by  $W$ , to map data samples  $x \in \mathcal{X}$  to an output space  $\mathcal{F} = \mathbb{R}^p$ . In both DeepSVDD vari-

ants,  $\mathcal{F}$  has a significantly lower dimension than the input, i.e.,  $p \ll d$ . Whereas an autoencoder is trained with a reconstruction objective, and thus the input and output space is the same, i.e.  $\mathcal{F} = \mathcal{X}$ . An *anomaly score*  $s^{(a)}$  is computed as the reconstruction error when using an autoencoder and as the distance from the estimated hypersphere centre  $c \in \mathcal{F}$  for both variants of DeepSVDD.

**Hour-wise models.** To address non-stationarity (C1), we further train individual models  $\Phi_h(\cdot; W_h)$  corresponding to each hour  $h \in \{0, \dots, 23\}$  of the day with associated parameters  $W_h$ . For training  $\Phi_h(\cdot; W_h)$ , a subset  $\mathcal{D}_h^u := \{(x_k^u, t_k^u)\}_{k=1}^{n_h}$  of the dataset  $\mathcal{D}_u$  is used such that  $t_k^u$  lies in the interval  $[h-1, h+1)$ . The number of samples in  $\mathcal{D}_h^u$  is denoted by  $n_h$ . We consider both autoencoder and DeepSVDD as the individual models in this setting. Given a new test sample  $(x_k^v, t_k^v)$  we assign an anomaly score  $s_k^{(a)}$  using the model  $\Phi_h$  corresponding to the hour  $h$  closest to timestamp  $t_k^v$ .

### 3.2 Forecasting-based anomaly detection

To consider non-stationarity (C1) and temporal dependence (C2), we adopt a forecasting-based AD method [4]. In this approach, a model is trained to predict the next image in a sequence, and anomalies can be detected when the trained model fails to predict the next image. In deciding whether an image at time  $t$  is anomalous or normal, we consider the ‘‘context’’ (or history)  $\mathcal{H}_{t_i} = \{(x_{i-k+1}^u, t_{i-k+1}^u), \dots, (x_{i-1}^u, t_{i-1}^u)\}$  which is a sequence of  $k$  samples prior to  $t$ . A key constraint in directly applying a sequence model is the irregularity in the inter-arrival time (C3) between the  $i$ -th and  $(i-1)$ -th sample  $\tau_i = t_i - t_{i-1}$ . To overcome this challenge, we draw inspiration from Neural Temporal Point Processes. Firstly, inter-arrival time embedding  $\psi_i = f_{sin}(\tau_i)$  is obtained using a sinusoidal encoding function  $f_{sin}(\cdot)$  [5, Equation 65]. Secondly, we embed each image  $x_i^u$  as  $z_i = \Phi_e(x_i^u; W_e)$ , where  $\Phi_e(\cdot; W_e)$  is a deep neural network with learnable parameter  $W_e$ . Then,  $z_i$  is concatenated with the time embedding  $\psi_i$  to obtain  $\mathcal{S}_{t_i} = \{[z_{i-k+1}, \psi_{i-k+1}], \dots, [z_{i-1}, \psi_{i-1}]\}$  where  $[\cdot, \cdot]$  denotes concatenation.  $\mathcal{S}_{t_i}$  is passed through the sequence model  $\varphi(\cdot; W_c)$  for generating the context  $c_i = \varphi(\mathcal{S}_{t_i}; W_c)$ , where  $W_c$  is the parameter of the sequence model. To predict the next image  $\hat{x}_i^u = \Phi_d((c_i, \psi_i); W_d)$  after an interval  $\tau_i$ , a decoder network is used with learnable parameters  $W_d$ .

After predicting an image  $\hat{x}_i^u$ , the prediction error  $\xi_i^u = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N (\hat{x}_i^u(m, n) - x_i^u(m, n))$  is averaged over the image pixels at coordinates  $(m, n)$ , where  $d = MN$ . Given  $\{\xi_1^u, \dots, \xi_{n_u}^u\}$ , we model the distribution  $f_\xi(\cdot)$  of mean prediction errors by fitting a univariate Gaussian distribution and Kernel Density Estimation (KDE). The anomaly scores  $s_i^{(a)}$  for the image  $x_i^v$  is then computed as the negative likelihood of observing the associated error  $\xi_i^v$ , given the estimated distribution  $\hat{f}_\xi$ , i.e.  $s_i^{(a)} = -\hat{f}_\xi(\xi_i^v)$ .

## 4 Experiments

**Training setup**<sup>2</sup>. The solar receiver dataset contains images obtained over a year starting from October. The unlabelled images from the first two months are used for training, consisting of  $n_u = 29,186$  samples. Additionally,  $n_v = 1,312$  samples are randomly selected from the rest of the data and labelled manually for testing purposes. We use the network architecture introduced in [3] for all our models. We further modify the architectures by using an Average-Pooling layer with a rectangular kernel of size  $23 \times 8$  as the first layer for dimensionality reduction. For the forecast-based AD, we use an LSTM as the sequence model with a context length  $k = 30$ . All the methods are trained for 200 epochs. We use Adam with a learning rate 0.001 and weight decay  $10^{-6}$ . We compare the Area under the Receiver Operating Characteristic Curve (AUC) of the methods.

Table 1: Experimental results

Method	AUC
Autoencoder	0.03
DeepSVDD OneClass	0.51
DeepSVDD SoftBoundary	0.41
Hour-wise autoencoder	$0.53 \pm 0.24$
Hour-wise DeepSVDD OneClass	$0.48 \pm 0.24$
Hour-wise DeepSVDD SoftBoundary	$0.49 \pm 0.26$
Forecasting model with gaussian	<b>0.84</b>
Forecasting model with KDE	<b>0.86</b>

**Comparison of anomaly detection methods.** The results from the above experiments are summarized in Table 1. For hour-wise models, as we have different models for each hour of the day, the mean of the AUC of all the models  $\Phi_h(\cdot, W_h)$  and their standard deviations are reported. The standard deviations are high as the available training sample is significantly smaller for some  $\Phi_h(\cdot, W_h)$ . We can observe that the hour-wise models perform better than their counterparts that disregard the non-stationarity. By further considering the temporal dependence along with non-stationarity in the data, the forecasting-based AD demonstrates the best performance. The results thus support our reasoning and validate the importance of incorporating the temporal dependence to learn effective representations for anomaly detection in the solar receiver dataset.

**Effect of context length  $k$ .** We report the AD performance of the forecast-based method with varying context lengths  $k$  in Figure 3. We observe a considerable increase in AD performance with increasing context length. Additionally, the performance is better when using KDE to estimate the distribution of mean prediction errors for assigning anomaly scores.

<sup>2</sup>The year is not disclosed due to data confidentiality

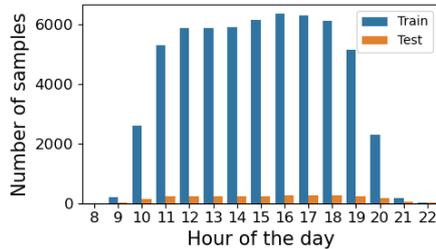


Fig. 2: Samples per hour-wise model

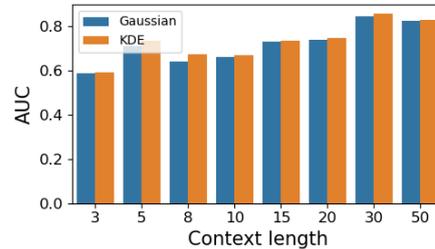


Fig. 3: AUC vs context length

## 5 Conclusion

We addressed AD in irregular image sequences collected from IR cameras in a CSP plant. Our data exhibits distinct traits differing from AD image benchmark datasets such as MVTEC [6] as evidenced by the underperformance of existing image-based AD methods. To better capture the characteristics of our dataset, we proposed a forecast-based AD model inspired by irregular sequence modelling. Experimental results demonstrate that the proposed model significantly improves AD performance as measured by the AUC. The results also showed the significance of capturing temporal information, including inter-image temporal relationships, as it holds valuable information for the AD task.

## 6 Acknowledgements

We thank Adrien Farinelle from John Cockerill for helping us understand the solar receiver dataset and the associated anomalous behaviours.

## References

- [1] Lukas Ruff, Jacob R. Kauffmann, Robert A. Vandermeulen, Gregoire Montavon, Wojciech Samek, Marius Kloft, Thomas G. Dietterich, and Klaus Robert Muller. A Unifying Review of Deep and Shallow Anomaly Detection. *Proceedings of the IEEE*, 109(5):756–795, 5 2021.
- [2] Manish Gupta, Jing Gao, Charu C. Aggarwal, and Jiawei Han. Outlier Detection for Temporal Data: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 26(9):2250–2267, 9 2014.
- [3] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep One-Class Classification. In *International Conference on Machine Learning*, pages 4393–4402. PMLR, 2018.
- [4] Pankaj Malhotra, L. Vig, Gautam M. Shroff, and P. Agarwal. Long Short Term Memory Networks for Anomaly Detection in Time Series. *The European Symposium on Artificial Neural Networks*, 2015.
- [5] Joseph Enguehard, Babylon Health, Dan Busbridge, Adam Bozson, Claire Woodcock, and Nils Hammerla. Neural Temporal Point Processes For Modelling Electronic Health Records. *Proceedings of Machine Learning Research LEAVE, UNSET:2020*, 7 2020.
- [6] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTEC ad-A comprehensive real-world dataset for unsupervised anomaly detection. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 9584–9592, 2019.