Automated Malfunction Detection for Robotic Arms in Panel Manufacturing Using Deep Latent State Space Model

Qisen Cheng, Kaushik Balakrishnan*, Zhihong Pan*
Samsung Display America Lab, San Jose, CA, USA

Yunyoung Kyung, Younggil Jin, Jaewon Kim
Samsung Display Co., Yongin-si, Gyeonggi-do, South Korea

Abstract
The paper describes a novel deep-learning-based system for malfunctions detection for the robotic arms used in display panel manufacturing process. Our method automates the preventive inspection of the manufacturing robotic arms by using a deep learning model that replaces the laborious human inspection. More importantly, our method achieves continuous monitoring of robotic arm condition, and alarms the maintenance team ≥2 days prior to the eventual failure, thus eliminating wastes of production materials and time consumed in the unexpected machine downtime. To achieve this automated and continuous malfunction detection, we leverage the deep latent state space model to learn the normal behavior of robotic arms from multivariate time-series sensor data collected in daily production, and then identify various types of malfunctions by looking at the difference between the expected normal signals and the measured signals. The deep latent state space model incorporates specialized latent embeddings formulation that is tailored to capture the full dynamics in time-series data. On an internal dataset that is directly sampled from production line, our approach achieves 100% malfunction detection (i.e. recall rate) while maintaining low false alarm rate (~2% measured in the span of a year). Our method can significantly advance the efficiency of preventive maintenance and reduce the cost in display panel manufacturing line.

Keywords
Machine Learning; Deep Learning; Preventive Maintenance; Timeseries Anomaly Detection; Deep State Space Model; Manufacturing Robots.

1. Introduction
Detection of equipment malfunctions is a common yet crucial task in the preventive maintenance of modern manufacturing [1-3]. In modern display panel manufacturing lines particularly, a large group of robotic arms are usually employed. Typical tasks performed by the robotic arms include picking up components for product assembly, wiring or soldering electronics, and constructing functional layers through welding or even additive fabrication [4-6]. During the daily operation, maintenance of those robotic arms is inevitable and occurs quite frequently as exemplified in [13]. A typical workflow requires a team of engineers to inspect the robotic arms regularly, and perform maintenance (e.g. change of parts, re-calibration, repairing, etc.) based on a certain schedule. The effectiveness of such workflow heavily relies on the occurrences of inspections performed, which consequently leads to exponentially increase of labor costs consumed in the frequent machine checkups. What makes things worse is that failures could actually occur to any robotic arm unexpectedly, demanding urgent trouble-shooting and causing extra downtime and wastes of materials. A more desirable

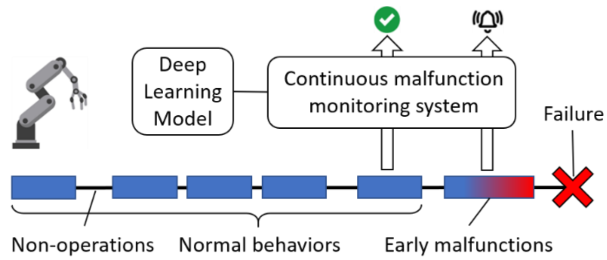

Figure 1. During daily operations of robotic arms, their condition gradually decays and eventually runs into catastrophic failure. In this paper, we design an automated malfunction monitoring system that is enabled by a deep timeseries model. By continuously monitoring the sensor measurements extracted from operations, our system can detect early malfunctions and alarm for preventive maintenance.

approach would be automated monitoring of robotic arm condition, which could significantly reduce the human inspection costs and achieve continuous detection of robotic arm malfunctions, minimizing the downtime and the waste.

In this paper, we present a deep-learning-based approach to achieve the desired automated and continuous monitoring. At the core of our solution, we develop a novel timeseries generative model to perform unsupervised robotic arm malfunction detection. Specifically, it models the normal behaviors of different robotic arms by learning to generate expected sensor measurements of normal operations. In the test time, the expected signals of normal operations are compared against the real measurements, and the deviation is used as anomaly score to identify malfunctions. We use the deep latent state space model as the main generative model. By modeling the sensor measurements as continuous signals instead of discrete sequences, it provides superior robustness to handle variations in sampling rates and signal waveforms observed from different robotic arms. We further equip the model with a latent embedding formulation that is tailored to encode the temporal transition between time points, capturing the entire dynamics of timeseries. In practice, sensor measurements directly from the real production environment are often noisy due to involvement of non-operational periods (i.e. robotic arm initialization, testing, manual stop, etc.). To handle the noise, we design a multi-level data pipeline to perform data cleansing and extraction at the raw signal level, and aggregate the predictions on useful signal segments using sliding windows (Fig. 1).

On an internal dataset, our approach achieves highly accurate early malfunction detection (i.e. 100% recall rate and 2 days prior to the actual failure), while maintaining very low false alarm rate (~2% measured in the span of a year).

* Equal contribution

ISSN 0097-996X/25/5601-0274-$1.00 © 2025

2. Methodology

Our method includes 3 main parts as shown in Fig. 1: 1) data preprocessing of the raw sensor signal, 2) latent state space model, and 3) anomaly score generation.

2.1 Data Preprocessing

The sensor signal directly from the production line is often very noisy. One cause is that normal operational operation periods are often interleaved with many non-operational periods (i.e. robotic arm initialization, testing, manual stops, etc.) [7]. The latter do not represent the normal robotic arm behaviors, and hence could act as outliers deteriorating the model training and inference. Another issue roots in the variations of hardware or mechanical components between different robotic arms. Even for the same movements, sensor signal magnitudes could vary from robotic arm to robotic arm.

Operational periods extraction: To clean the data, we first perform rule-based extraction of operational periods from the raw signal. Specifically, we first segment the signal timeseries by thresholding the delta time between adjacent time points. And then those short segments lasting less than a certain period of time are filtered out. The remaining segments are regarded as continuous operation periods and used in model training and inference.

Dynamic window-wise timeseries standardization: Conventional standardization techniques require statistics of the timeseries (e.g. min and max for 0-1 standardization). However, getting those statistics for all robotic arms could be practically challenging. Alternatively, we follow a more convenient approach that dynamically standardizes each window of signal being fed into the model with statistics computed for the window [8], eliminating the need of pre-calibration for different robotic arms. This approach is effective under the assumption that different measurements of the same motion vary in magnitudes but are similar in terms of waveforms, which is aligned with our observations.

2.2 Latent State Space Model

The latent state space model is a generative model that learns the normal behavior of robotic arms through a reconstruction task. As in Fig. 2, it first consists of an encoder that projects sensor measurements (i.e. timeseries signal) into latent embeddings. The latent embeddings are learned by using a latent embedding module with tailored designs to fully capture the temporal dynamics of input signal. Then a decoder is used to reconstruct the inputs again from the embeddings.

Latent embedding for timeseries. Given a timeseries ($x_{\le T}$) of length T, we want to learn some latent embeddings (z) to encode the signal patterns. Following [9], this posterior mapping – conditional probability distribution of latent embeddings given the input signal – can be factorized as product of marginals:

$$p_\varphi(z_{\le T}|x_{\le T}) = \prod_{i=0}^{L-1} p_\varphi\left(z_{t_i}|x_{\le t_i}\right) \quad (1)$$

This mapping is realized with a sequence model (i.e. an encoder neural network that is parameterized by φ). Particularly, we tailor the design of the model such that it employs a sequence of latent embeddings ($z_{\le T}$) of the same length as the input signal, each corresponding to one time point in the input. This design maximizes the latent embedding resolution, hence facilitating downstream detection of subtle anomalous patterns.

Learning of latent embeddings. To obtain the supervision signal for learning the latent embeddings, we need to transform from the embedding space back to the signal space. This can be achieved by taking a two-step procedure: to generate a new signal value at time t, 1) we first generate the corresponding latent embedding at the same time point, based on all historical latent embeddings up to time $t-1$, and then 2) generate the signal value (x_t) based on historical latent embeddings plus the latest latent embedding generated in the first step. The historical signal values ($x_{<t}$) can also be incorporated in the second step to account for the autoregressive structure of the signal, as in Equation (2). Formally, the two subsequent generation steps can be expressed using two conditional distributions, product of whom represents the full probabilistic relation to be learned:

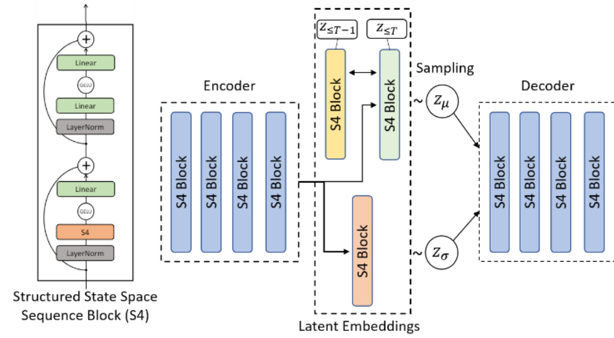

Figure 2. Illustration of architecture of the deep latent state space model, which consists of an encoder network, a latent embedding module, and a decoder network. All three networks are constructed by stacking multiple structured state space sequence (S4) blocks.

$$p_{\theta,\lambda}(x_{\le T}, z_{\le T}) = \prod_{i=0}^{L-1} p_\theta\left(x_{t_i}|x_{<t_i}, z_{\le t_i}\right) p_\lambda\left(z_{t_i}|z_{<t_i}\right) \quad (2)$$

Here, the conditional distribution for the first step is commonly referred to as prior distribution. We deliberately formulated it as a conditional distribution to model the dependency among latent embeddings, in turn learning the full temporal dynamics of the input signal. Similarly, the conditional distribution for the second step models the likelihood of signal given latent embeddings. In this work, we assume both distributions are conditional Gaussian distributions, which have been demonstrated to be effective in reconstruction tasks [3]. The two distributions are realized by two individual sequence models, namely a latent embedding module that is parameterized by λ and the decoder that is parameterized by θ. We train our model by minimizing the following variational lower bound, i.e. training loss:

$$L = \mathbb{E}_{p_\varphi(z_{\le T}|x_{\le T})}\left[\sum D_{KL}(p_\varphi \parallel p_\lambda) - log p_\theta\right] \quad (3)$$

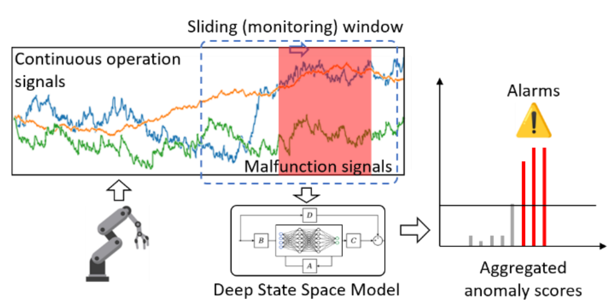

Figure 3. Our malfunction detection system continuous examine the latest sensor signals in a sliding window fashion, and outputs anomaly scores for each unit of time.

where KL stands for the calculation of Kullback–Leibler divergence between the prior and posterior. To ensure proper learning of temporal dependency of latent embeddings, we specify the operands in KL as follows:

$$D_{KL}(p_\varphi(z_t|x_{\le t}) \parallel p_\lambda(z_t|z_{<t})) \quad (4)$$

In this case, the posterior converts each time point of input signal to corresponding embeddings whereas the prior has to predict the current latent embedding from the history and matches the output from the posterior. This setup enforces the prior network to properly learn temporal dynamics in the latent space.

Structured state space sequence (S4) block. The encoder, decoder, and latent embedding module share the same backbone which comprises multiple stacked S4 blocks. S4 gets increasingly popular for sequence learning, largely due to its competitive performance compared to the State-of-the-art transformer counterparts, yet with a much lower time complexity (O(n) vs. O(n^2)). This benefits the learning on long sequences such as raw signals, which is evidenced by its leading rank on the Long Range Arena, a widely used benchmark for timeseries data [14].

Another advantage of S4 is that it models the sequence as an underlying continuous-time waveform, which demonstrates excellent robustness against frequency shift that is common in sensor readouts from the factory. The waveform is assumed to be governed by the following recurrent differential equations:

$$\frac{d}{dt}\boldsymbol{h}_t = \boldsymbol{A}\boldsymbol{h}_t + \boldsymbol{B}x_t \quad (5)$$

$$y_t = \boldsymbol{C}\boldsymbol{h}_t \quad (6)$$

where $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$ stands for parameter matrices to be learned through gradient descent. The recurrent system can be re-formulated into the following equivalent convolutional form after discretization:

$$x_k = \overline{\boldsymbol{A}}x_{k-1} + \overline{\boldsymbol{B}}u_k \quad \overline{\boldsymbol{A}} = (\boldsymbol{I} - \Delta/2\cdot\boldsymbol{A})^{-1}(\boldsymbol{I} + \Delta/2\cdot\boldsymbol{A}) \quad (7)$$

$$y_k = \overline{\boldsymbol{C}}x_k \quad \overline{\boldsymbol{B}} = (\boldsymbol{I} - \Delta/2\cdot\boldsymbol{A})^{-1}\Delta\boldsymbol{B}) \quad \overline{\boldsymbol{C}} = \boldsymbol{C} \quad (8)$$

This form allows parallel processing of all time points simultaneously, while preserving the state transitions in the recurrent form. Furthermore, S4 imposes the following HIPPO-structure on the coefficient matrices, offering a principled modeling of long-range dependencies with much reduced complexity for optimization:

$$\boldsymbol{A}_{nk} = -\begin{cases} (2n+1)^{1/2}(2k+1)^{1/2} & if\ n > k \\ n+1 & if\ n = k \\ 0 & if\ n < k \end{cases} \quad (9)$$

2.3 Anomaly Scoring

To detect malfunctions, we design a system to monitor the input signals continuously. Specifically, it scans the latest input signals that reside in a sliding window of length T. For each window, the continuous operation periods are first extracted for examination. The trained latent state space model transforms the extracted signal to latent embeddings, and then attempts to reconstruct the signal through the aforementioned generation process. Since the model has only been trained to generate signals of normal behaviors, the quality of reconstruction of anomalous signals should be relatively poorer than the reconstruction of normal state signals. The quality gap in turn gives signals of abnormality. In this work, we quantify the abnormality using two different scores: 1) KL divergence (i.e. the first term of training objective in Equation 3), and 2) Mean-Absolute-Error (MAE) between the reconstructed and input signals. The two scores are complementary because the former captures the deviation in temporal transition in latent space, while the later reflect the divergence of waveform in signal space.

We aggregate the two scores by summing them up and taking the average as the window-wise anomaly score. By advancing the sliding window one point a time, we obtain the consecutive anomaly scores. We finally take the maximum among the window-wise scores within a certain reporting cycle (e.g. each hour) as the final anomaly score. The threshold for detection is determined through calibration on robotic arms that operate in normal state. Specifically, we select 99 percentiles of the anomaly scores measured in normal state as the threshold, and our system makes verdict of malfunction as anomaly score is found to be higher than the threshold.

3. Experiment

We tested our method on an internal dataset which was collected directly from the production line. A total of 12 one-month cases were collected for training. Another 12 one-month cases was reserved for testing: 7 cases contain failures caused by malfunctions (Cases F1-F7 in Fig. 4), and 5 cases represent operations in normal conditions (Case N1-N5 in Fig. 4). To validate the effectiveness of our method, we compared it with 3 other popular baselines – Graph Deviation Network [10],

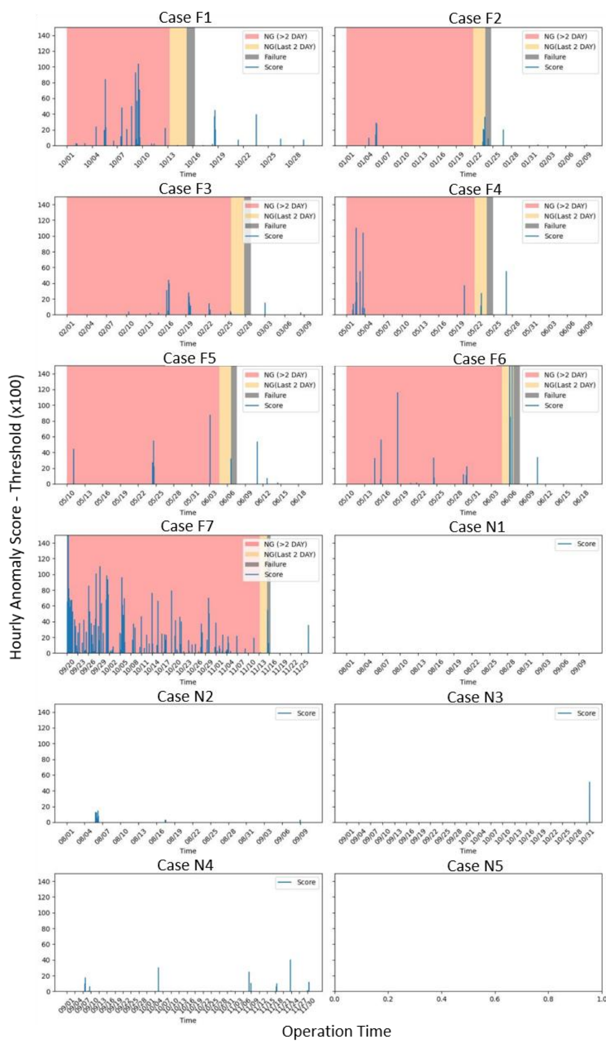

Figure 4. Anomaly scores (scaled by 100x for visualization) that are aggregated by each hour for the 12 testing cases, each contains about 1 month of data.

Anomaly Transformer [11], and LSTM-DVAE [12]. The metrics we employ are Recall and False Alarm Rate (FAR). Note the Recall rate was defined on case level because we don't have labelling of malfunctions down to each hour, whereas the FAR was found on hour level.

As shown in Table 1, our method was able to detect malfunctions in all 7 cases, and made few wrong hourly detections (~2%), surpassing all baseline methods.

Our method was also capable of detection of early malfunctions. As shown in Fig 3, a number of hours with high anomaly scores were reported in the red zone, which represented at least 2 days prior to the occurrence of actual failure (yellow zone). This further demonstrate the usefulness of our method in preventative maintenance.

Table 1. Malfunction detection performance of different deep learning models.

Model	Recall (case)	FAR (hour)
Graph Deviation Network [10]	71.4%	11.6%
Anomaly Transformer [11]	85.7%	4.3%
LSTM-DVAE [12]	85.7%	6.8%
Latent State Space Model (ours)	**100%**	**2.0%**

4. Conclusions

In this paper, we present a malfunction detection system for continuous monitoring of robotic arms employed in display panel manufacturing lines. The system is enabled by a specially designed deep latent state space model that is able to robustly detect anomalous patterns in sensor measurements.

5. References

1. Pan Z, Shenoy R, Balakrishnan K, Cheng Q, Lee J, Jeon Y, Jeong D, Kim J. 32 - 1: Improving QD Backplane Defect Image Generation Using Automatic Masking in Diffusion Models. In SID Symposium Digest of Technical Papers 2024 Jun (Vol. 55, No. 1, pp. 409-412).
2. Balakrishnan K, Cheng Q, Lee J, Jeong D, Kim E, Kim J. 6-4: Deep Learning for Classification of Repairable Defects in Display Panels Using Multi-Modal Data. In SID Symposium Digest of Technical Papers 2023 Jun (Vol. 54, No. 1, pp. 58-61).
3. Cheng Q, Qu S, Lee J. 72 - 3: Deep Learning Based Visual Defect Detection in Noisy and Imbalanced Data. InSID Symposium Digest of Technical Papers 2022 Jun (Vol. 53, No. 1, pp. 971-974).
4. Bogue R. The role of robots in the electronics industry. Industrial Robot: the international journal of robotics research and application. 2023 Aug 9;50(5):717-21.
5. Cheng Q, Zhang C, Shen X. Estimation of Energy and Time Usage in 3D Printing With Multimodal Neural Network. In2022 4th International Conference on Frontiers Technology of Information and Computer (ICFTIC) 2022 Dec 2 (pp. 900-903). IEEE.
6. Zhang C, Cheng Q. Predicting Melt-Crystal Interface Position and Shape during the Manufacturing Process of Single Crystal via Explainable Machine Learning Models. InIOP Conference Series: Materials Science and Engineering 2022 Oct 1 (Vol. 1258, No. 1, p. 012029). IOP Publishing.
7. Brogårdh T. Present and future robot control development—An industrial perspective. Annual Reviews in Control. 2007 Jan 1;31(1):69-79.
8. Kollovieh M, Ansari AF, Bohlke-Schneider M, Zschiegner J, Wang H, Wang YB. Predict, refine, synthesize: Self-guiding diffusion models for probabilistic time series forecasting. Advances in Neural Information Processing Systems. 2024 Feb 13;36.
9. Zhou L, Poli M, Xu W, Massaroli S, Ermon S. Deep latent state space models for time-series generation. InInternational Conference on Machine Learning 2023 Jul 3 (pp. 42625-42643). PMLR.
10. Deng A, Hooi B. Graph neural network-based anomaly detection in multivariate time series. InProceedings of the AAAI conference on artificial intelligence 2021 (Vol. 35, No. 5, pp. 4027-4035).
11. Xu J, Wu H, Wang J, Long M. Anomaly Transformer: Time Series Anomaly Detection with Association Discrepancy. InInternational Conference on Learning Representations.
12. Park D, Hoshi Y, Kemp CC. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. IEEE Robotics and Automation Letters. 2018 Feb 2;3(3):1544-51.
13. Wyatt H, Wu A, Thomas R, Yang Y. Life cycle analysis of double-arm type robotic tools for LCD panel handling. Machines. 2017 Mar 4;5(1):8.
14. Tay Y, Dehghani M, Abnar S, Shen Y, Bahri D, Pham P, Rao J, Yang L, Ruder S, Metzler D. Long Range Arena: A Benchmark for Efficient Transformers. InInternational Conference on Learning Representations.
15. Xiong W, Lee J, Qu S, Jang W. 81-1: Invited Paper: Data Augmentation for Applying Deep Learning to Display Manufacturing Defect Detection. In SID Symposium Digest of Technical Papers 2020 Aug (Vol. 51, No. 1, pp. 1210-1213).
16. Yu J, Liu J. Multiple granularities generative adversarial network for recognition of wafer map defects. IEEE Transactions on Industrial Informatics. 2021 Jun 25;18(3):1674-83.
17. Liu J, Zhang F, Yang B, Zhang F, Gao Y, Wang H. Focal auxiliary classifier generative adversarial network for defective wafer pattern recognition with imbalanced data. In 2021 5th IEEE Electron Devices Technology & Manufacturing Conference (EDTM) 2021 Apr 8 (pp. 1-3). IEEE.