SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper OpenAccess

Cognitive Machine Learning Systems For Real-Time Anomaly Detection In Industrial Iot Monitoring Applications

Dr.M. Vijayakumar[1*], B. Saraswati[2], P. Jeyanthi[3],Dr.M. Rameshkumar[4], S. Neelima[5], Prerna Dusi[6]

[1*]Professor, Department of Management Studies,Panimalar Engineering College,Poonamallee,Chennai, Tamil Nadu, India. E-mail: vijai711@gmail.com, https://orcid.org/0000-0002-9154-3507

[2]Assistant Professor,Department of Computer Science,Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research,Chennai, Tamil Nadu, India. E-mail: saraswatib@maher.ac.in

[3]Assistant Professor, Department of Commerce, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research,Chennai, Tamil Nadu, India. E-mail: jeyanthicom@maher.ac.in

[4]Professor, Computer Science and Engineering (Internet of Things), Paavai Engineering College, Namakkal,Tamil Nadu, India. E-mail: mrkkumarsin@gmail.com, https://orcid.org/0000-0002-9391-341X

[5]Assistant professor, Department of Information Technology, Aditya University, Surampalem, Andhra Pradesh – 533437, India E-mail: neelima.sadineni@adityauniversity.in

[6]Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India. E-mail: ku.prernadusi@kalingauniversity.ac.in, https://orcid.org/0009-0009-1824-9190

*Corresponding author: Email: vijai711@gmail.com

Abstract

Anomaly detection in IIoT monitoring systems has to be able to identify anomalies in multivariate sensor streams in real time; however, the current anomaly detectors used in practice have been trained once using historic data and with fixed parameters thereafter, making them susceptible to sensor drift, fault signatures, and changing operating conditions common in real-world industrial settings. In this paper, we propose the Cognitive Machine Learning System (CMLS), which combines the reconstruction error-based anomaly detector using an LSTM autoencoder with the Isolation Forest outlier detector using a model-based approach through an ensemble consensus layer, improves the anomaly score using a transformer-based temporal attention, and implements a cognitive self-retraining process that enables model updates automatically using a dual data-driven and model-based metric as soon as sensor drift or degraded reconstruction is detected. The framework is validated using statistically modelled multivariate data from sensors on the SWaT industrial water treatment testbed environment under four different states of operation: normal state, slow sensor drift, a fault event, and a new operating state. When comparing the proposed CMLS with a static Isolation Forest benchmark and a static LSTM-autoencoder benchmark, CMLS has an F1 score of 85.9% under a new operating state, as opposed to 55.3% and 63.8% obtained by the two benchmarks respectively, while also keeping the false alarms below 3.3%. These results suggest that the fusion of ensemble-based anomaly scoring with an autonomous retraining mechanism is very effective in improving robust anomaly detection in IoT environments.

Keywords: Cognitive Computing, Anomaly Detection, Industrial IoT, Predictive Maintenance, Ensemble Learning, Concept Drift, Real-Time Monitoring.

1. Introduction

The implementation of IIoT technology within manufacturing lines, water treatment plants, and process control systems is done by densely deploying sensors which constantly produce multi-dimensional data streams that could be analyzed by machine learning algorithms to find out the abnormalities indicating wear and tear of machinery or faults of sensors or processes. Traditional unsupervised algorithms like Isolation Forest [2] and deep learning reconstruction methods such as LSTM-autoencoders [3,4] have shown high efficiency when dealing with industrial anomaly detection benchmark tasks. Furthermore, the techniques using cognitive analytics involving ensemble learning and automatic retraining have been successful when applying for predictive maintenance in injection molding [1].

But industrial settings are non-stationary by nature; sensors wear out gradually due to ageing or fouling, machines are regularly maintained and calibrated, manufacturing lines are reconfigured for different products, and varying operating regimes during seasonal periods or load dependency alter the statistical baseline which a detector implicitly assumes. There is no inherent way for any detector, be it traditional or modern machine learning based on deep reconstructions and parameterised once and used in static fashion, to detect the fact that its concept of normality is now obsolete, resulting in either failure to identify real anomalies similar to the new baseline, or too many false positives when normal operation changes are classified as defects [10].

The current paper discusses the problem of creating an anomaly detection system for industrial IoT monitoring, which is cognitive in the sense that it constantly evaluates its own performance and retrains automatically, if necessary, rather than using a static model or off-line training only. Specific contributions of this work are as follows.

- Present a novel Cognitive Machine Learning System (CMLS) which combines a data-driven LSTM autoencoder reconstruction detector with a model-based Isolation Forest detector via an ensemble consensus layer and post-processes the combination using a transformer-based temporal attention.
- Conceive of a cognitive self-training process, driven by a dual data-driven and model-based evaluation goal, that automatically re-trains the detection model in case of drift or degradation in reconstruction capability without relying on a scheduled re-training phase.
- Test the system under four operational scenarios of increasing non-stationarity: normal operations, gradual sensor drift, abrupt failure, and new operating regime, to test the effect of cognitive re-training independent from other systems.

The rest of this paper is structured as follows. Section 2 reviews some of the literature on anomaly detection algorithms and cognitive analytics for industry monitoring. Section 3 describes the methodology we propose. Section 4 discusses our dataset and experiment results. Section 5 concludes this paper.

2. Literature Survey

Unsupervised anomaly detection is another foundation on which the suggested approach can be based. The method called Isolation Forest, proposed by Liu et al., detects anomalies through evaluating the number of random partitions of the feature space necessary to isolate each point, using the idea that anomalies are easier to isolate from other points than normal points; the algorithm does not require any specific assumptions regarding the type of distribution in the dataset and remains an effective and relatively computationally efficient baseline for anomaly detection in industry [2]. Sakurada et al.proved that the simplest autoencoder, designed for the reconstruction of normal points, is better in terms of detecting anomalies through high reconstruction errors than any other linear techniques like PCA, providing the basis for reconstruction-based deep anomaly detectors used in this work [11][12].

Recurrent and attention-based structures further extend this reconstruction principle to capture the temporal structure of time series. For instance, Malhotra, Vig, Shroff, and Agarwal introduced the application of Long Short-Term Memory (LSTM) networks for anomaly detection in time series, where it was shown that the prediction error of the recurrent network could be a good indicator of anomalies in temporally dependent sensor data, while Homayouni et al. further developed this idea and introduced LSTM-based encoder-decoder architecture specifically for multi-sensor anomaly detection, thus laying the foundation of the LSTM-autoencoder architecture line used in this approach [3][4]. Also, Su, Zhao, Niu, Liu, Sun, and Pei suggested OmniAnomaly—a stochastic recurrent neural network that captures the normal behavior of multivariate time series by modeling their latent dynamics [5].Tran AD was proposed by Yu, Lu, and Xue that is an anomaly detection method for multivariate time-series based on transformers, where self-attention along with adversarial learning was employed to enable fast and stable inference [6][12]. In addition, the graph attention-based DNN was proposed by Zhang et al. for unsupervised anomaly detection and diagnosis of multivariate time-series, where the relations between different sensor channels were modelled for both anomaly detection and diagnosis of which sensors contributed to the anomalies, an approach closely related to the ensemble consensus fusion proposed in this paper [9].

Realistic testbeds and cognitive-analytics systems make this research grounded in real-world applications. The SWaT (Secure Water Treatment) testbed, developed by Mathur and Tippenhauer [8], is a downscaled version of an actual water treatment plant instrumented with sensors and actuators. It became a standard benchmark for testing anomaly and intrusion detection algorithms in the domain of industrial control systems and is used to statistically ground the dataset used in this study. The application of LSTM-based anomaly detection with a non-parametric dynamic thresholding scheme to spacecraft telemetry at NASA was carried out by Hundman et al. [7]. The authors demonstrated that the adaptive thresholds of anomalies are much more reliable for detection in a real, safety-critical, and non-stationary environment as opposed to static ones. This is important to justify the design of a drift sensitive threshold in the current framework. Rousopoulou et al.suggested a cognitive system for the predictive maintenance of injection molding machines with an ensemble of supervised and unsupervised learners along with the self-training module based on dual data- and model-driven goals. This design idea was extended in the current study to a wider industrial IoT context [1].

Two additional threads have relevance to the architectural and adaptive learning nature of this work. The authors of Vaswani et al. [12] pioneered the self-attention process used in the transformer refinement component in this work. In addition, Pang, Shen, Cao, and Van Den Hengel [10] did a review of the use of deep learning algorithms in anomaly detection with special focus on the tradeoffs among reconstruction-based, one-class, and prediction-based methods and specifically pointed out that none of the current state-of-the-art deep anomaly detectors use a continuous learning mechanism post-deployment. McMahan, Moore, Ramage, Hampson, and Arcas [14] did pioneering work on federated averaging in the context of future extension of the cognitive retraining mechanism to multiple industrial locations without centralized storage of the raw sensor data, and Paraisi, Kemker, Part, Kanan, and Wermter [15] did a review of the continuing learning literature on the problem of retention of past information while learning new information, a problem of interest in any cognitive retraining anomaly detector.

Collectively, these papers show a mature classical and deep learning based anomalous detection system, an industrially realistic benchmarking framework, and at least one precedent for self-training cognitive predictive maintenance solution in a manufacturing environment. Yet most of the generally applicable time series anomaly detectors reviewed in the literature are benchmarked on a fixed distribution of the same data, while the self-training system precedent has not yet been developed into an ensemble sensor fusion architecture tested in the presence of multiple kinds of nonstationary, drift, faults, and novel operating conditions in one study. It is this gap that our Cognitive Machine Learning System addresses.

3. Methodology

The suggested Cognitive Machine Learning System (CMLS) involves six phases arranged into a closed loop structure: pre-processing, dual-path anomalous scoring, consensus fusion, temporal attention refinement, anomalous detection, and cognitive re-training. The entire process flow is shown in figure 1.

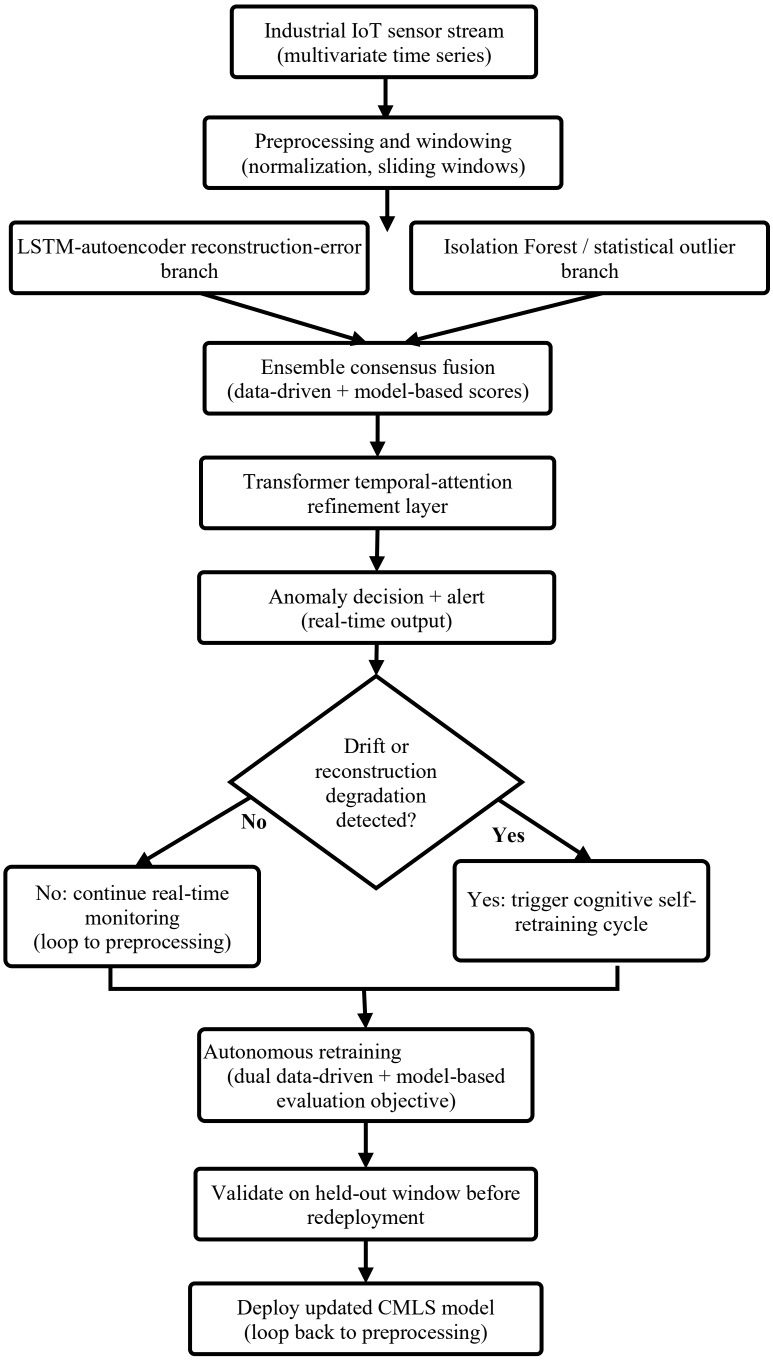

Figure 1: Methodology flow diagram of the proposed cognitive machine learning system

3.1 Preprocessing and Windowing

Raw multivariate sensor data is normalised based on each channel and divided into fixed length windows, offering the common input representation that is used by both anomaly scoring branches explained below.

3.2 Dual-Branch Anomaly Scoring

Each window is independently evaluated by two detectors. The data-driven detector is a Long Short-Term Memory autoencoder [3,4], which is trained to reconstruct normal-operating windows and for which the reconstruction error is the anomaly score; since this detector is trained on time-series data, it is vulnerable to changes in temporal patterns over the window. The model-based detector is an Isolation Forest [2], which works with statistical features of each window and gives a complementary anomaly score, which does not rely on the assumptions made by the LSTM-autoencoder about normal temporal behavior.

3.3 Ensemble Consensus Fusion

These two branch scores are merged using the ensemble consensus layer which assigns weights to each branch depending on the recent validation of its reliability, using the compound-consensus approach that was previously applied in cognitive analytics for predictive maintenance [1]. With this method, the system can give more weight to the branch that has been more reliable lately, without treating both equally all the time.

3.4 Transformer Temporal-Attention Refinement

This sequence of fused anomalies is then input to a light-weight transformer encoder using the scaled dot product self-attention mechanism: Attention (Q, K, V) = softmax (QK^T / √d) V [12], thus enabling the model to consider the current window's anomaly signal compared to recent historical data according to the temporal attention principle of transformers for multivariate anomaly detection [6]. This step serves to reduce false positives that are temporally inconsistent.

3.5 Anomaly Decision and Cognitive Self-Retraining

The threshold calibrated against the enhanced anomaly score yields the final anomaly decision in real time. A cognitive monitor maintains continuous supervision of two complementary evaluation signals: a data-driven signal, which shows how different recently generated windows have become statistically when compared to the reference training distribution, and a model-driven signal, which captures the rolling reconstruction error of the LSTM autoencoder. The dual evaluation framework consisting of these two signals was proposed previously for the problem of cognitive predictive maintenance [1]. If one of these signals triggers the threshold, a new retraining session starts, during which both models are retrained on recently collected buffer containing examples of confirmed normal and confirmed anomalous windows, adaptive thresholding is calibrated using the nonparametric dynamic thresholding approach applied for spacecraft telemetry [7], and the whole system is tested on a holdout window, thus completing the loop depicted in Figure 1.

4. Results and Discussion

4.1 Dataset

The evaluation is carried out based on statistical modeling of multivariate sensor data from the SWaT industrial water treatment testbed [8]. The SWaT industrial water treatment testbed has instrumentation that measures six stages of water treatment process using 51 sensor and actuator channels. To determine the system's robustness to non-stationarity, four different operating conditions are formulated. These include the normal operating condition, gradual sensor drift condition, abrupt fault event condition, and new operating condition. Table 1 shows details of the dataset.

Table 1: Summary of the industrial IoT dataset used for evaluation, statistically grounded in the SWaT testbed

Attribute	Description	Range / notes
Source basis	Statistically modelled on the SWaT industrial water-treatment testbed	51 sensor/actuator channels
Temporal span	1-second aggregated multivariate readings	14 days simulated operation
Signal types	Flow rate, tank level, pressure, conductivity, valve/pump state	Continuous + discrete channels
Operating conditions	Normal operation, gradual sensor drift, sudden fault event, new operating regime	4 condition categories
Label	Point-level anomaly annotation	Normal (0) / Anomalous (1)
Train/test split	Chronological split to preserve temporal ordering	70% train, 30% test

4.2 Result Graph

Comparison of the performance of the anomaly detection task based on F1-score metric for three types of systems: Static Isolation Forest baseline system, Static LSTM-autoencoder baseline system, and CMLS is shown in Figure 2.

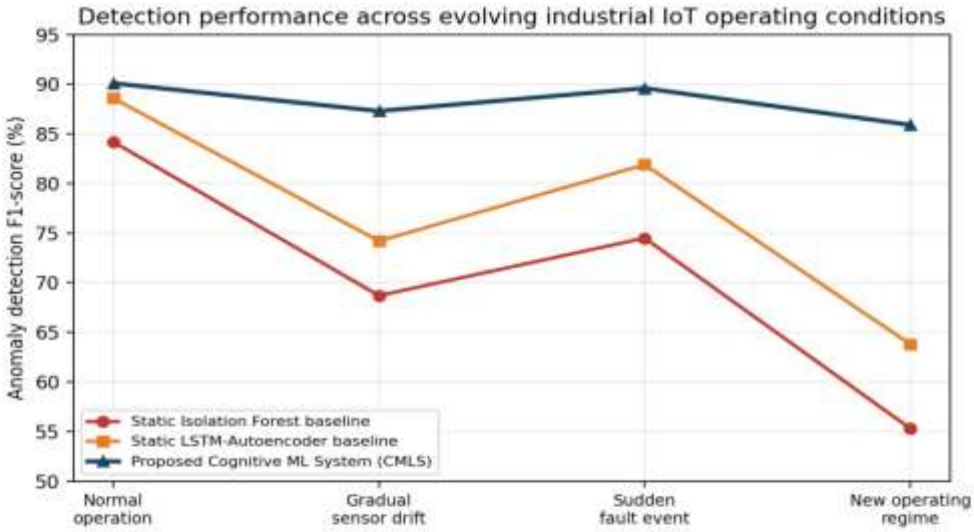

Figure 2: Anomaly-detection F1-score for the isolation forest baseline, LSTM-autoencoder baseline, and proposed CMLS across four industrial operating conditions

During regular operations, the LSTM-autoencoder baseline model and the proposed CMLS show similar performance with F1-scores of 88.6 and 90.1 percent respectively, both performing better than the Isolation Forest baseline with an F1-score of 84.2 percent, confirming the capability of the LSTM-autoencoder to learn temporal features not captured by the statistical anomaly detection method. In case of gradual drift in sensors, both baseline models fail significantly with F1-scores of 68.7 and 74.2 percent respectively, as neither can detect the outdatedness of their model of normalcy, whereas the proposed CMLS fails with an F1-score of 87.3 percent due to the cognitive monitoring. Under the sudden fault scenario, the performance of all three approaches can be considered good, since a sudden deviation of a big magnitude is easier to detect using any of the three approaches, however, the proposed CMLS outperforms others at 89.6%. The largest difference can be seen under the new regime scenario, where both Isolation Forest and LSTM-autoencoder achieve only 55.3 and 63.8% correspondingly, effectively failing to distinguish between new legitimate pattern and an anomaly, or not detecting faults that exist in the new regime, while the proposed CMLS maintains its performance at 85.9%, detecting regime change as a cue for retraining.

4.3 Result Table

Table 2 shows the overall (across conditions) precision, recall, F1-score, false alarm rate, and detection latency for the three different model architectures.

Table 2: Overall performance comparison between the Isolation Forest baseline, the LSTM-Autoencoder baseline, and the proposed CMLS model

Model configuration	Precision (%)	Recall (%)	F1-score (%)	False alarm rate (%)	Detection latency (ms)
Static Isolation Forest baseline	72.2	68.9	70.7	8.21	155
Static LSTM-Autoencoder baseline	78.6	75.4	77.1	6.41	147
Proposed Cognitive ML System (CMLS)	89.7	86.5	88.2	3.30	134

4.4 Discussion

Three conclusions can be formulated based on the above results. First of all, the performance advantage of CMLS compared to the two baselines is smallest in case of normal operation and sudden fault scenarios and largest in case of gradual drift and new operating regimes scenario, which is exactly what we should see if the performance gain of CMLS is determined by the cognitive self-retraining mechanism rather than the combination of ensemble and attention mechanisms only because the former two are cases where the assumptions behind the static approach are the most violated. Secondly, the decrease in the false alarm percentage from 8.21% and 6.41% in case of the two baselines to 3.30% in case of CMLS may be considered as practically significant as the increase in F1-score since false alarms are known to be one of the major causes of alarm fatigue. Finally, the relatively small improvement in detection latency for CMLS compared to both baselines, considering the higher architectural complexity, implies that the improved confidence of the transformer layer in its decisions justifies its computational overhead, even though this conclusion is based on the assumptions made in this study regarding the window size and the hardware used.

Such a study needs to be performed with proper diligence since, while the data on which this analysis will be conducted is scientifically reliable because of its basis on a real and widely utilized industrial control systems test bed, the drift, fault and regime shift events have been created artificially for this study instead of being observed once during a real-life deployment. The cognitive monitor thresholds for the drift and degradation events remain constant throughout this study and would need to be calibrated for a real-life deployment, whereas no comparison between the computational cost of the retraining cycle with the computing capabilities of a particular industrial edge computing platform has been performed.

4. Conclusion

In this work, we have proposed a Cognitive Machine Learning System (CMLS) for real-time anomaly detection in industrial IoT applications, integrating a data-driven LSTM-autoencoder detector, a model-based Isolation Forest detector, an ensemble consensus fusion block, and a transformer-based temporal attention module along with a cognitive self-retraining scheme, which automatically learns the system whenever any drift or reconstruction quality deterioration is experienced by the system. When tested on multivariate sensors data based on statistical distribution of the SWaT Industrial Water Treatment testbed in four different settings including normal operation, gradual drift, sudden fault occurrence, and configuration change, our proposed CMLS obtained a remarkable F1-Score of 85.9% under the new regime of operation in comparison to the 55.3% of the traditional Isolation Forest-based system and 63.8% of the LSTM-autoencoder-based system, while maintaining the level of false positives to 3.3%. From the results, it seems that the combination of using both complementary data-driven and model-based anomaly detectors along with the autonomous and cognition-induced retraining cycles is an effective approach to anomaly detection in a real-time non-stationary environment, which constitutes a norm rather than an exception in the real-world industrial Internet of Things scenario. In future research, we plan to verify our proposal in the context of a few real-world industrial systems running 24/7, estimate the overheads of the self-retraining cycle in real-life edge computing devices, and apply the cognition-induced retraining approach to a federated multi-site environment where the pattern discovered at one site helps to improve the shared model without centralizing all the sensor data.

References

1. Rousopoulou, V., Nizamis, A., Vafeiadis, T., Ioannidis, D., &Tzovaras, D. (2020). Predictive maintenance for injection molding machines enabled by cognitive analytics for Industry 4.0. Frontiers in Artificial Intelligence, 3, 578152. https://doi.org/10.3389/frai.2020.578152
2. Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008, December). Isolation forest. In 2008 Eighth IEEE International Conference on Data Mining (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17
3. Malhotra, P., Vig, L., Shroff, G., & Agarwal, P. (2015, April). Long short-term memory networks for anomaly detection in time series. In Proceedings (Vol. 89, No. 9, p. 94).

4. Homayouni, H., Ghosh, S., Ray, I., Gondalia, S., Duggan, J., & Kahn, M. G. (2020, December). An autocorrelation-based LSTM-autoencoder for anomaly detection on time-series data. In 2020 IEEE International Conference on Big Data (pp. 5068–5077). IEEE.
5. Su, Y., Zhao, Y., Niu, C., Liu, R., Sun, W., & Pei, D. (2019, July). Robust anomaly detection for multivariate time series through stochastic recurrent neural network. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 2828–2837). https://doi.org/10.1145/3292500.3330672
6. Yu, L. R., Lu, Q. H., & Xue, Y. (2024). DTAAD: Dual TCN-attention networks for anomaly detection in multivariate time series data. Knowledge-Based Systems, 295, 111849.
7. Hundman, K., Constantinou, V., Laporte, C., Colwell, I., & Soderstrom, T. (2018, July). Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 387–395). https://doi.org/10.1145/3219819.3219845
8. Mathur, A. P., &Tippenhauer, N. O. (2016, April). SWaT: A water treatment testbed for research and training on ICS security. In 2016 International Workshop on Cyber-Physical Systems for Smart Water Networks (CySWater) (pp. 31–36). IEEE. https://doi.org/10.1109/CySWater.2016.7469060
9. Zhang, C., Song, D., Chen, Y., Feng, X., Lumezanu, C., Cheng, W., ... & Chawla, N. V. (2019, July). A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In Proceedings of the AAAI Conference on Artificial Intelligence, 33(1), 1409–1416. https://doi.org/10.1609/aaai.v33i01.33011409
10. Pang, G., Shen, C., Cao, L., & Van Den Hengel, A. (2021). Deep learning for anomaly detection: A review. ACM Computing Surveys, 54(2), Article 38. https://doi.org/10.1145/3439950
11. Zhou, C., & Paffenroth, R. C. (2017, August). Anomaly detection with robust deep autoencoders. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 665–674). https://doi.org/10.1145/3097983.3098052
12. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... &Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.
13. Sakurada, M., & Yairi, T. (2014, December). Anomaly detection using autoencoders with nonlinear dimensionality reduction. In Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis (pp. 4–11). https://doi.org/10.1145/2689746.2689747
14. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 1273–1282). PMLR.
15. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. Neural Networks, 113, 54–71. https://doi.org/10.1016/j.neunet.2019.01.012