

Anomaly Detection in Industrial Internet of Things Systems Using Unsupervised Learning Techniques

Dr. Daniel Cruz

Assistant Professor, Polytechnic University of the Philippines, Philippines

Abstract

The Industrial Internet of Things (IIoT) has introduced new opportunities for connecting industrial equipment, monitoring operational conditions, and analyzing information generated through distributed sensors and communication systems. Industrial networks continuously collect measurements associated with machine temperature, vibration, electrical current, pressure, and other operating characteristics. These observations can support equipment monitoring and operational assessment. However, irregular sensor behavior, changing operating conditions, communication disturbances, and equipment deterioration may introduce anomalous patterns that are difficult to identify through conventional monitoring procedures. Detecting such irregularities is particularly challenging when verified fault labels are limited or unavailable. The present paper develops a hypothetical methodological framework for investigating unsupervised learning techniques for anomaly detection in IIoT environments. Four computational approaches are considered: One-Class Support Vector Machine, Local Outlier Factor, Isolation Forest, and Autoencoder-based anomaly detection. The proposed framework incorporates multivariate sensor preprocessing, temporal window construction, unsupervised representation learning, anomaly scoring, and comparative performance evaluation.

A fictional industrial monitoring dataset containing 12,000 sensor observation windows is introduced to demonstrate the intended research design. Two hypothetical analytical tables illustrate the distribution of normal and anomalous observations and the comparative detection performance of the selected algorithms. Precision, recall, and F1-score are considered principal evaluation measures. All numerical results are manually constructed and do not represent executed experiments, collected industrial measurements, or verified machine-learning outcomes. The methodological discussion examines sensor heterogeneity, data contamination, operational variability, computational limitations, and model interpretability. Particular attention is given to the distinction between statistically unusual observations and confirmed industrial faults. The paper concludes that unsupervised learning provides a relevant methodological direction for examining irregular patterns in industrial sensor data. However, reliable implementation requires representative observations, appropriate threshold selection, independent validation, and careful interpretation of operational context.

Keywords: Industrial Internet of Things, Anomaly Detection, Unsupervised Learning, Industrial Sensor Data, Isolation Forest, Local Outlier Factor, One-Class SVM, Autoencoder, Predictive Maintenance, Industrial Data Analytics.

1. Introduction

The Industrial Internet of Things has become an important technological development within modern manufacturing and automated industrial operations. Industrial equipment increasingly incorporates interconnected sensors, controllers, communication devices, and computational monitoring systems. These technologies enable the collection and exchange of operational information across interconnected industrial environments. Sensor measurements provide insights into machine behavior, production conditions, and variations in equipment performance. As the quantity of available industrial information expands, organizations require computational methods capable of identifying meaningful operational changes. Automated anomaly detection has consequently emerged as a relevant area of investigation within industrial informatics and intelligent system monitoring.

Industrial operations produce complex streams of measurements associated with mechanical, electrical, thermal, and process-related conditions. These measurements may vary according to equipment characteristics, production schedules, environmental conditions, and normal changes in operating load. Consequently, an unusual sensor reading does not necessarily indicate equipment failure or an abnormal process condition. Conventional monitoring procedures commonly use predefined thresholds to identify measurements outside expected ranges. Although such procedures remain useful, they may provide limited information when abnormalities involve relationships among multiple sensors or develop gradually without exceeding individual measurement limits.

Machine learning offers alternative analytical approaches for investigating patterns within industrial sensor information. Supervised learning methods rely on examples with known classifications, whereas unsupervised methods examine data characteristics without requiring labels during model fitting. This distinction is important because industrial datasets may contain relatively few examples of verified equipment faults or abnormal process events. Certain operating conditions may occur infrequently, making comprehensive labeled datasets difficult to obtain. Unsupervised anomaly detection addresses this limitation by identifying observations that differ statistically from learned patterns, although these differences require further examination before being interpreted as actual faults.

Several unsupervised learning techniques have been developed to identify unusual observations within complex datasets. One-Class Support Vector Machines estimate boundaries around typical observations, while Local Outlier Factor evaluates differences in neighborhood density. Isolation Forest identifies unusual observations through randomized partitioning, and Autoencoder-based methods examine reconstruction differences between observed and learned representations. These approaches rely on different mathematical principles and may respond differently to sensor correlations, measurement noise, and operational variation. Their comparative evaluation is therefore necessary when considering anomaly detection within heterogeneous industrial environments.

The present paper examines unsupervised anomaly detection techniques through a hypothetical IIoT monitoring framework. It proposes a structured comparison of four algorithms using multivariate industrial sensor observations organized into temporal windows. The research design considers data preparation, anomaly scoring, threshold selection, and independent evaluation. A fictional dataset provides the basis for demonstrating classification outcomes and analytical interpretation. The investigation does not involve actual industrial machinery, recorded sensor measurements, or executed computational experiments. Instead, it establishes a methodological foundation for future empirical research on unsupervised anomaly detection in industrial information systems.

2. Objectives of the Study

The principal objective of this study is to examine the methodological suitability of unsupervised machine learning techniques for identifying anomalous patterns in Industrial Internet of Things systems. The investigation concentrates on industrial sensor characteristics, comparative anomaly detection performance, and evaluation reliability. Three specific objectives establish the analytical direction of the proposed research framework.

2.1 Examination of Anomalous Sensor Behavior

The first objective is to examine how irregular patterns may be represented within multivariate industrial sensor observations. Temperature, vibration, pressure, electrical current, and equipment operating indicators can exhibit relationships that reflect normal machine behavior. Deviations from these relationships may provide useful information for anomaly detection. The proposed study considers how individual measurements and their temporal combinations could be organized into analytical features. Particular attention is given to distinguishing isolated measurement fluctuations from patterns that differ meaningfully from established operating conditions.

The investigation further aims to examine the influence of operational context on the interpretation of unusual sensor behavior. Industrial equipment may exhibit substantial variation during startup, shutdown, maintenance, and changing production loads. Such conditions may generate observations that differ from typical operating patterns without representing equipment deterioration.

2.2 Comparative Assessment of Unsupervised Detection Models

The second objective focuses on establishing a comparative framework for four unsupervised anomaly detection approaches. One-Class Support Vector Machine, Local Outlier Factor, Isolation Forest, and Autoencoder-based detection are selected to represent boundary-based, density-based, isolation-based, and reconstruction-based methods. The proposed comparison examines their ability to identify unusual sensor observations under consistent evaluation conditions. Precision, recall, and F1-score are considered complementary measures for investigating detection reliability when retrospective reference labels are available.

A further objective involves examining the relationship between anomaly detection performance and computational characteristics. Different algorithms may respond differently to feature dimensionality, sensor correlation, and dataset size. Some approaches may require extensive parameter optimization, while others may offer comparatively efficient scoring procedures. The proposed framework therefore considers the methodological importance of comparable preprocessing, consistent threshold selection, and independent evaluation. This objective seeks to establish a balanced basis for model assessment rather than assuming that any particular unsupervised technique is universally superior.

2.3 Identification of IIoT Monitoring Limitations

The third objective is to identify methodological difficulties that may affect unsupervised anomaly detection within industrial environments. Missing sensor measurements, inconsistent sampling intervals, measurement noise, incomplete operational histories, and changing equipment conditions can complicate model development. These characteristics may produce unusual observations that are unrelated to genuine system faults. The proposed study examines how data cleaning, temporal feature construction,

and appropriate evaluation procedures could address selected limitations while preserving information relevant to industrial monitoring.

The investigation also considers the practical requirements associated with implementing anomaly detection models in interconnected industrial systems. Computational capacity, communication latency, model maintenance, and information security can influence monitoring reliability. Furthermore, anomaly scores require appropriate interpretation because statistical deviations may not correspond directly to operational failures. This objective emphasizes transparent modeling procedures, responsible handling of industrial data, and appropriate verification mechanisms. These considerations support the evaluation of unsupervised learning as part of a broader industrial monitoring process.

3. Review of Literature

3.1 Industrial IoT Architecture and Process Monitoring

Gubbi et al. (2013) examined the architectural foundations of the Internet of Things and discussed interconnected sensing, information processing, and computational infrastructure. Xu et al. (2014) investigated industrial IoT technologies and their application across different operational environments. Lee et al. (2015) subsequently presented a cyber-physical systems architecture for manufacturing applications. These contributions established important foundations for understanding how connected industrial equipment generates and exchanges operational information. Their methodological relevance lies in connecting sensor-based observation with data-driven industrial system monitoring.

Qin (2012) examined data-driven industrial process monitoring and diagnosis, emphasizing the detection of abnormal operational changes through statistical analysis. Venkatasubramanian et al. (2003) reviewed process-history-based fault detection techniques and discussed their methodological characteristics. These studies established a foundation for using historical operational measurements to examine process variation. Their contributions also indicate that fault detection and fault diagnosis represent related but distinct analytical activities. Identifying abnormal behavior does not necessarily establish the underlying cause of an industrial equipment problem.

3.2 Development of Unsupervised Anomaly Detection Algorithms

Breunig et al. (2000) introduced Local Outlier Factor as a density-based approach for identifying observations that differ from their neighboring data points. Schölkopf et al. (2001) investigated support estimation for high-dimensional distributions, providing foundations relevant to one-class classification. Tax and Duin (2004) developed Support Vector Data Description for identifying target-class boundaries. These contributions established important computational strategies for anomaly identification without requiring conventional multiclass supervision. Their relevance extends to industrial environments where comprehensive labels for unusual operational events may be unavailable.

Liu et al. (2008) introduced Isolation Forest as an approach based on identifying observations through randomized partitioning. Liu et al. (2012) subsequently expanded the theoretical and experimental treatment of isolation-based anomaly detection. Chandola et al. (2009) provided a comprehensive survey of anomaly detection research, while Aggarwal (2013) examined broader methods for outlier analysis. These studies established important distinctions among density-based, boundary-based, and isolation-based approaches. They also demonstrated that method suitability depends on the underlying data distribution and the definition of an anomalous observation.

3.3 Reconstruction Learning and Industrial Data Variability

Hinton and Salakhutdinov (2006) investigated dimensionality reduction through neural networks, contributing to the development of representation-learning approaches. Sakurada and Yairi (2014) examined Autoencoder-based anomaly detection using nonlinear dimensionality reduction. Malhotra et al. (2015) investigated recurrent neural networks for identifying unusual patterns within time-series information. These contributions established a methodological foundation for learning representations of complex observations and using deviations from expected patterns as potential indicators of abnormal behavior.

Zimek et al. (2012) reviewed methodological difficulties associated with unsupervised outlier detection in high-dimensional data. Gama et al. (2014) examined concept drift adaptation, while Saito and Rehmsmeier (2015) investigated evaluation methods for imbalanced classification problems. These studies emphasize that anomaly detection depends on both model structure and evaluation design. Their findings support careful attention to changing data distributions, irrelevant attributes, and suitable performance measures when examining unusual observations within complex industrial sensor systems.

4. Research Methodology

4.1 Research Design and Unsupervised Learning Framework

The present paper adopts a hypothetical comparative quantitative design for examining unsupervised anomaly detection techniques in an industrial monitoring environment. Four algorithms are considered: One-Class Support Vector Machine, Local Outlier Factor, Isolation Forest, and an Autoencoder. The proposed framework evaluates multivariate sensor observations represented through fixed-length temporal windows. Each method would produce an anomaly score indicating the extent to which an observation differs from learned patterns. No computational implementation or actual industrial experiment was conducted.

The proposed analytical procedure consists of sensor data preparation, temporal window construction, unsupervised model fitting, anomaly scoring, and independent performance evaluation. Model fitting would use unlabeled training observations, while retrospective reference labels would be reserved for reporting detection performance. Thresholds would be determined using training-score distributions rather than test labels. Consistent evaluation conditions would support comparisons among algorithms. This methodological design is intended to distinguish unsupervised anomaly learning from the subsequent supervised assessment of detection outcomes.

4.2 Hypothetical Industrial Sensor Dataset

For illustrative purposes, the study assumes a fictional monitoring dataset containing 12,000 temporal observation windows from four interconnected industrial machines. Each window is assumed to contain derived measurements associated with temperature, vibration, electrical current, pressure, and operating conditions. The dataset includes 11,400 observations designated as normal and 600 designated as anomalous for retrospective evaluation. These values are manually constructed and do not describe actual industrial sensor measurements, recorded equipment faults, or executed simulation procedures.

The proposed preprocessing procedure would examine missing measurements, inconsistent sampling intervals, duplicate records, and sensor-specific scaling requirements. Temporal windows would be constructed to preserve selected relationships among consecutive observations. Feature normalization parameters would be estimated using development information only. The dataset would

be divided chronologically to reduce inappropriate access to future observations. Although fictional reference labels are included for evaluation, the selected anomaly detection algorithms would not use these labels during unsupervised model fitting.

4.3 Model Development and Performance Evaluation

The One-Class Support Vector Machine would estimate a boundary around typical sensor observations, while Local Outlier Factor would examine differences in neighborhood density. Isolation Forest would assign anomaly scores through randomized feature partitioning. The Autoencoder would learn a compressed representation and estimate reconstruction error. Each model would be configured using predefined unsupervised criteria. A nominal contamination assumption of 5% would guide illustrative anomaly threshold selection without using retrospective test labels.

The proposed evaluation would report precision, recall, and F1-score based on independently reserved reference labels. Precision would describe the proportion of anomaly alerts corresponding to labeled abnormal observations, while recall would represent the proportion of labeled anomalies identified. F1-score would summarize the balance between these measures. A completed investigation should additionally examine false-alert rates, detection delay, computational latency, and uncertainty estimates. The present manuscript uses fictional classification counts exclusively to demonstrate evaluation procedures.

5. Analysis

5.1 Illustrative Distribution of Industrial Sensor Observations

The distribution of normal and anomalous observations represents an important consideration in industrial monitoring research. Abnormal operating conditions may occur relatively infrequently, creating a substantial imbalance between routine observations and unusual events. Table 1 presents a hypothetical distribution of sensor observation windows across three temporally ordered partitions. The fictional example assigns 5% of the observations to the anomalous category. These values are intended to demonstrate a possible evaluation structure and do not represent measured industrial event frequencies.

The illustrative distribution emphasizes why anomaly detection performance should not be assessed through overall accuracy alone. A model predicting every observation as normal would achieve 95% accuracy within this constructed dataset while identifying none of the labeled anomalies.

Table 1. Hypothetical Distribution of Industrial Sensor Observation Windows

Observation Category	Training Set	Validation Set	Test Set	Total
Normal Observations	7,980	1,710	1,710	11,400
Anomalous Observations	420	90	90	600
Total Observations	8,400	1,800	1,800	12,000

The hypothetical dataset contains an overall anomaly proportion of 5%. This figure is used only as an analytical assumption and should not be interpreted as an established anomaly rate for industrial IoT systems.

5.2 Illustrative Comparison of Unsupervised Detection Techniques

Unsupervised anomaly detection algorithms identify unusual observations through different statistical and computational mechanisms. Boundary-based, density-based, isolation-based, and reconstruction-based approaches may therefore produce different anomaly scores and detection outcomes. Table 2 presents hypothetical true-positive and false-positive counts for four selected methods. The fictional test partition contains 90 anomalous and 1,710 normal observations. Precision, recall, and F1-score are calculated from these assumed counts to demonstrate a consistent comparative reporting procedure.

The illustrative results assign the highest F1-score to the Autoencoder, followed by Isolation Forest, Local Outlier Factor, and One-Class Support Vector Machine. These differences do not represent actual experimental findings or establish algorithmic superiority. A genuine investigation would require implemented models, independent testing, careful threshold selection, and repeated evaluation across different industrial conditions. The practical importance of performance differences would also depend on the consequences of missed anomalies and unnecessary alerts during industrial operations.

Table 2: Hypothetical Comparative Performance of Unsupervised Anomaly Detection Models

Unsupervised Technique	True Positives	False Positives	Precision	Recall	F1-Score
One-Class SVM	55	49	0.529	0.611	0.567
Local Outlier Factor	60	41	0.594	0.667	0.628
Isolation Forest	69	31	0.690	0.767	0.726
Autoencoder	74	28	0.725	0.822	0.771

The example demonstrates that higher anomaly recall does not necessarily eliminate false alerts. A suitable industrial monitoring system would need to consider both missed abnormalities and unnecessary intervention requirements.

6. Discussion

6.1 Unsupervised Learning and Industrial Anomaly Identification

Unsupervised learning provides a methodological approach for examining industrial sensor information when comprehensive fault labels are unavailable. Rather than relying exclusively on predefined abnormal categories, these algorithms assess whether observations differ from the patterns represented in their training information. This characteristic may support the identification of unusual operational behavior. However, statistical difference does not establish the existence of equipment failure. The proposed framework therefore distinguishes anomaly scores from confirmed operational faults and emphasizes the importance of contextual verification.

The hypothetical comparative framework demonstrates how alternative unsupervised methods could be evaluated under consistent monitoring conditions. Each algorithm applies a different definition of unusual behavior, which may influence its sensitivity to sensor relationships and data distributions. One-Class Support Vector Machine emphasizes decision boundaries, while Local Outlier Factor examines neighborhood density. Isolation Forest uses partitioning characteristics, and Autoencoders

analyze reconstruction behavior. These differences justify comparative investigation but do not establish that any single approach is appropriate for every industrial monitoring problem.

6.2 Multivariate Sensor Relationships and Detection Reliability

Industrial monitoring frequently involves multiple sensor measurements that reflect interconnected aspects of equipment operation. Temperature, vibration, current, and pressure may change together during normal production activities. An observation can therefore appear unusual because of an unexpected relationship among variables even when individual measurements remain within conventional operating limits. Multivariate anomaly detection provides an analytical approach for examining these relationships. Its effectiveness, however, depends on suitable feature construction and representative observations of normal operating variation.

The proposed Autoencoder configuration illustrates how reconstruction-based learning could be applied to multivariate sensor information. A network trained to represent common measurement patterns may assign larger reconstruction errors to observations that differ substantially from learned relationships. Nevertheless, complex models can sometimes reconstruct unusual observations effectively or respond strongly to benign operating changes. Consequently, reconstruction error alone cannot establish a fault. Reliable evaluation requires appropriate thresholds, independent reference information, and examination of the operational meaning of identified deviations.

6.3 Comparative Performance and Operational Interpretation

The hypothetical results present different balances between correctly identified anomalies and false-positive alerts. These differences demonstrate the importance of evaluating precision and recall together rather than selecting a model exclusively according to one measure. A highly sensitive detector may identify more abnormal observations while also producing additional unnecessary alerts. Conversely, a conservative detector may reduce false alarms but fail to identify some relevant events. The preferred balance depends on the consequences associated with different types of detection errors.

The proposed methodology also distinguishes offline classification assessment from practical industrial monitoring performance. Retrospective F1-score does not fully capture detection delay, operational downtime, alert processing requirements, or the consequences of equipment interruptions. A model achieving favorable evaluation metrics may still require modification before deployment within a resource-constrained industrial environment. Consequently, future research should examine temporal stability and operational suitability alongside classification performance. The illustrative findings cannot establish these properties because no real monitoring system was implemented.

7. Challenges

7.1 Sensor Noise, Missing Measurements, and Data Quality

Industrial sensor systems may produce inconsistent measurements because of calibration differences, communication interruptions, environmental disturbances, or temporary equipment conditions. Such variations can introduce statistical irregularities unrelated to actual process faults. Missing values and inconsistent sampling intervals may further complicate the construction of reliable temporal features. Unsupervised algorithms can respond strongly to these characteristics because they depend on the statistical structure of the available observations. Consequently, data quality assessment is essential before interpreting anomaly scores.

Preprocessing introduces additional methodological decisions concerning missing-value handling, scaling, outlier treatment, and feature aggregation. Excessive smoothing may remove short-duration abnormalities, while insufficient noise reduction may increase false alerts. Similarly, inappropriate normalization can distort relationships among sensors operating at different measurement scales.

7.2 Changing Operating Conditions and Model Adaptability

Industrial equipment may operate under different loads, production schedules, maintenance conditions, and environmental temperatures. These variations can alter normal sensor distributions without indicating a system fault. An unsupervised model trained under limited conditions may therefore classify legitimate operating changes as anomalous. This difficulty becomes particularly relevant when equipment performance evolves gradually over time. Effective monitoring requires careful evaluation of distributional changes and an understanding of the operating regimes represented within the training observations.

Model adaptability presents a further challenge because updating anomaly detectors with new observations can introduce contaminated information. If anomalous measurements are incorporated into the normal reference distribution without adequate screening, subsequent models may become less sensitive to relevant deviations. Conversely, retaining an outdated reference model may increase unnecessary alerts. The proposed framework therefore emphasizes controlled model maintenance, appropriate evaluation periods, and transparent updating procedures. These safeguards are important for preserving detection reliability under changing industrial operating conditions.

7.3 Computational Limitations and Anomaly Interpretability

Industrial IoT systems may operate under limitations involving processing capacity, communication bandwidth, memory availability, and response time. Complex anomaly detection algorithms can introduce computational requirements that affect their suitability for continuous monitoring. Autoencoder models, for example, may require additional training resources compared with some isolation-based approaches. However, practical computational efficiency depends on implementation details and available hardware. Future evaluation should therefore document inference time, resource consumption, and monitoring frequency alongside conventional detection metrics.

Interpretability also presents difficulties because an anomaly score may not clearly identify the source or operational significance of unusual behavior. Industrial personnel may need to determine which sensor relationships contributed to a detection and whether further investigation is justified. Complex representations can make this explanation challenging. Consequently, anomaly detection systems should support appropriate diagnostic review rather than automatically equating statistical irregularities with equipment faults. Transparent reporting and suitable human oversight remain important considerations in practical industrial monitoring.

8. Conclusion

Unsupervised machine learning provides a relevant methodological foundation for identifying irregular patterns within Industrial Internet of Things systems. The present paper developed a hypothetical comparative framework involving One-Class Support Vector Machine, Local Outlier Factor, Isolation Forest, and Autoencoder-based detection. Its primary emphasis was the organization of multivariate sensor information, anomaly scoring, and independent performance evaluation. The proposed design

demonstrates how different computational approaches could be compared when verified fault labels are limited during model development.

The illustrative tables demonstrate possible procedures for presenting sensor observation distributions and anomaly detection outcomes. However, all numerical values were manually constructed and do not represent recorded industrial measurements or executed computational experiments. Meaningful comparative findings would require authentic or reproducibly simulated sensor data, implemented algorithms, carefully selected detection thresholds, and independently evaluated reference labels. The hypothetical performance differences therefore cannot establish the effectiveness or operational superiority of any particular unsupervised learning technique.

Overall, the study identifies data quality, changing operating conditions, computational efficiency, and interpretability as central considerations in industrial anomaly detection. Future empirical research should investigate these factors through representative sensor datasets and carefully controlled experimental procedures. Unsupervised learning may support industrial monitoring by identifying observations requiring further examination, but its practical value depends on reliable validation and appropriate operational interpretation. The proposed framework therefore serves as a methodological basis for subsequent research rather than a validated industrial monitoring solution.

References

1. Adhikari, D., Jiang, W., Zhan, J., Rawat, D. B., & Bhattarai, A. (2024). Recent advances in anomaly detection in Internet of Things: Status, challenges, and perspectives. *Computer Science Review*, 54, Article 100665. <https://doi.org/10.1016/j.cosrev.2024.100665>
2. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15. <https://doi.org/10.1145/1541880.1541882>
3. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422). IEEE. <https://doi.org/10.1109/ICDM.2008.17>
4. Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., & Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7), 1443–1471. <https://doi.org/10.1162/089976601750264965>
5. Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., & Kloft, M. (2018). Deep one-class classification. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 4393–4402). PMLR.
6. Sakurada, M., & Yairi, T. (2014). 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). Association for Computing Machinery. <https://doi.org/10.1145/2689746.2689747>
7. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. In *Proceedings of the Network and Distributed System Security Symposium (NDSS 2018)*. Internet Society. <https://doi.org/10.14722/ndss.2018.23204>
8. Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Breitenbacher, D., Shabtai, A., & Elovici, Y. (2018). N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17(3), 12–22. <https://doi.org/10.1109/MPRV.2018.03367731>

9. Rafique, S. H., Abdallah, A., Musa, N. S., & Murugan, T. (2024). Machine learning and deep learning techniques for Internet of Things network anomaly detection—Current research trends. *Sensors*, 24(6), Article 1968. <https://doi.org/10.3390/s24061968>
10. Rodríguez, M., Tobón, D. P., & Múnera, D. (2023). Anomaly classification in industrial Internet of things: A review. *Intelligent Systems with Applications*, 18, Article 200232. <https://doi.org/10.1016/j.iswa.2023.200232>
11. Ni, C., & Li, S. C. (2024). Machine learning enabled Industrial IoT security: Challenges, trends and solutions. *Journal of Industrial Information Integration*, 38, Article 100549. <https://doi.org/10.1016/j.jii.2023.100549>
12. Anomaly detection based on Artificial Intelligence of Things: A systematic literature mapping. (2024). *Internet of Things*, 25, Article 101063. <https://doi.org/10.1016/j.iot.2024.101063>
13. Anomaly detection in industrial machinery using IoT devices and machine learning: A systematic mapping. (2023). *IEEE Access*, 11, 128288–128305. <https://doi.org/10.1109/ACCESS.2023.3333242>
14. Liu, J., Xie, G., Wang, J., Li, S., Wang, C., & Zheng, F. (2024). Deep industrial image anomaly detection: A survey. *Machine Intelligence Research*, 21, 104–135. <https://doi.org/10.1007/s11633-023-1459-z>
15. A review of machine learning methods for IoT network-centric anomaly detection. (2024). In 2024 47th International Conference on Telecommunications and Signal Processing (TSP). IEEE. <https://doi.org/10.1109/TSP63128.2024.10605928>
16. Moustafa, N., Keshky, E., Debiez, E., & Hameed, S. (2021). TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems. *IEEE Access*, 9, 165130–165150. <https://doi.org/10.1109/ACCESS.2021.3130875>
17. Mathur, A. P., & Tippenhauer, N. O. (2016). 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>
18. Goh, J., Adepu, S., Tan, M., & Lee, Z. S. (2017). Anomaly detection in cyber physical systems using recurrent neural networks. In 2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE) (pp. 140–145). IEEE. <https://doi.org/10.1109/HASE.2017.36>
19. Ring, M., Wunderlich, S., Scheuring, D., Landes, D., & Hotho, A. (2019). A survey of network-based intrusion detection data sets. *Computers & Security*, 86, 147–167. <https://doi.org/10.1016/j.cose.2019.06.005>
20. Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. *arXiv*. <https://doi.org/10.48550/arXiv.1901.03407>