JANOLI International Journal of
Artificial Intelligence and its Applications
(A Peer Review Journal)

journal homepage: https://janolijournals.com/index.php/jijaia

Volume 2, Issue 4, 2025 ISSN: 3048-6815

Enhanced Anomaly Detection in Industrial Control Systems using Hybrid Deep Learning Architectures and Federated Learning

Dr. Shabana Faizal
NIET, NIMS University, Jaipur, India

ARTICLE INFO

Article History:
Received: 05 October 2025;
Revised: 07 October 2025;
Accepted:16 October 2025;
Published: 27 October 2025

Keywords: Anomaly Detection, Industrial Control Systems (ICS), Deep Learning, Federated Learning, LSTM, Autoencoders, Hybrid Models, Cybersecurity, Data Privacy, SCADA.

Correspondence:
E-mail: s.faizal@utb.edu.bh

ABSTRACT

Industrial Control Systems (ICS) are becoming more susceptible to cyber attacks, requiring effective anomaly detection mechanisms. In this paper, we introduce a novel anomaly detection framework for ICS networks based on hybrid deep learning models and federated learning. We integrate Long Short-Term Memory (LSTM) networks to analyze temporal sequences with Variational Autoencoders (VAEs) to reconstruct features and detect outliers. In order to meet data privacy requirements and the decentralized environment of ICS deployments, we achieve this using a federated learning paradigm, where model training is conducted across sites without exchanging raw data. The introduced framework is tested on a publicly available ICS dataset and shows better performance than individual deep learning models and conventional anomaly detection methods. The outcomes present enhanced accuracy, precision, and recall in detecting different types of attacks without compromising data privacy. The approach provides a viable and efficient solution towards improving the cybersecurity stance of contemporary ICS ecosystems.

1. Introduction:

Industrial Control Systems (ICS) that control critical infrastructure like power grids, water treatment plants, and manufacturing facilities are increasingly becoming connected and thus more vulnerable to cyberattacks. These cyberattacks have disastrous impacts, ranging from service outages to complete equipment destruction and loss of life. Conventional security defense systems, including firewalls and intrusion detection systems (IDS), are usually ineffective against high-level and tailored attacks that take advantage of ICS protocol vulnerabilities and software weaknesses. This inefficiency requires the creation of sophisticated anomaly detection methods that can spot faint deviations in typical system behavior, suggesting malicious behavior or system failure.

Anomaly detection, the identification of patterns that deviate from typical behavior, has been a promising solution for improving ICS security. However, conventional anomaly detection methods tend to have difficulty with the high-dimensional, time-varying, and complex data typical of current ICS. In addition, the growing focus on data privacy and the geographically

dispersed nature of many ICS installations pose considerable obstacles for centralized data aggregation and model training.

In order to meet such demands, this paper presents an improved anomaly detection approach for ICS systems that synergizes the strength of deep learning with federated learning principles. Our approach utilizes hybrid deep learning models, namely Long Short-Term Memory (LSTM) networks and Variational Autoencoders (VAEs), to learn both the temporal relations and the underlying data distributions in ICS sensor data. LSTM networks are particularly appropriate for sequential data modeling and the discovery of abnormal temporal patterns, whereas VAEs are especially good at learning compressed models of usual data and identifying outliers in terms of reconstruction error.

In addition, we use federated learning to facilitate cross-site collaborative model training among various ICS sites without the need to share raw sensitive data. This solution solves data privacy issues and facilitates the development of more powerful and generalizable anomaly detection models using the varied operational data from various ICS deployments.

The goals of this paper are as follows:

Design a hybrid deep learning framework that integrates LSTM networks and VAEs to improve anomaly detection in ICS systems.

Develop a federated learning methodology to facilitate distributed model training while maintaining data privacy.

Test the performance of the designed framework using a publicly available ICS dataset, comparing it with isolated deep learning models and conventional anomaly detection methods.

Analyze the effectiveness of the framework in detecting various attack types and its ability to adapt to different ICS operating conditions.

2. Literature Review:

Anomaly detection in ICS has been a subject of extensive research, with various techniques explored to address the unique challenges posed by these systems. This section provides a critical review of relevant literature, highlighting the strengths and weaknesses of existing approaches.

2.1 Statistical Methods:

Early approaches to anomaly detection in ICS relied on statistical methods, such as control charts, Kalman filters, and time series analysis. These methods typically involve establishing a baseline of normal system behavior and flagging deviations from this baseline as anomalies. For example, Teixeira et al. [1] used Kalman filters to detect anomalies in a water distribution system, while Adepoju et al. [2] employed control charts to monitor the performance of a power plant. These methods are relatively simple to implement and computationally efficient, but they often struggle

with the non-linear and dynamic nature of modern ICS. They also require careful feature engineering and parameter tuning to achieve satisfactory performance.

2.2 Machine Learning Techniques:

Machine learning methods, including Support Vector Machines (SVMs), decision trees, and clustering algorithms, have also been extensively used for anomaly detection in ICS. These techniques are capable of detecting more intricate patterns from data and learning to adapt to evolving system behavior. For example, Hadziosmanovic et al. [3] applied SVMs to anomalous detection for a gas pipeline system, whereas Iglesias et al. [4] utilized decision trees to detect anomalies in a water treatment plant. These approaches tend to call for labeled training data, though, which can be costly and challenging to acquire in ICS deployments. Additionally, they can be inefficient at identifying temporal dependences in ICS sensor signals.

2.3 Deep Learning Approaches:

Deep learning is a strong method of anomaly detection across multiple fields, including ICS. Deep learning architectures like Autoencoders (AEs), Recurrent Neural Networks (RNNs), and Convolutional Neural Networks (CNNs) can learn intricate features from raw data automatically and identify spatial and temporal dependencies. Ring et al. [5] suggested the use of AEs for detecting anomalies in a gas pipeline system, whereas Mantere et al. [6] used RNNs to identify anomalies in a water treatment plant. Additionally, LSTM networks, a variant of RNNs specifically tailored to deal with long-term dependencies in sequential data, have also provided encouraging results in anomaly detection in ICS. For instance, Goh et al. [7] employed LSTM networks in order to identify anomalies in a power grid.

But most current deep learning-based anomaly detection techniques for ICS are based on centralized data aggregation, which is data privacy-sensitive. Federated learning could provide a remedy for this issue by supporting distributed model training without exposing raw data.

2.4 Federated Learning for Anomaly Detection:

Federated learning has received considerable interest as a privacy-preserving method of machine learning. There has been research into the use of federated learning for anomaly detection in different areas. For example, Harder et al. [8] proposed a federated learning framework for anomaly detection in smart buildings, while Rahman et al. [9] developed a federated learning based approach to intrusion detection for IoT networks.

However, the application of federated learning to anomaly detection in ICS is still relatively unexplored. Furthermore, the unique characteristics of ICS data, such as its high dimensionality, temporal dependencies, and non-stationary nature, present specific challenges for federated learning. Li et al. [10] have explored federated learning with differential privacy for anomaly detection, but the computational complexity of differential privacy mechanisms can be a limiting factor in real-time ICS applications. Therefore, more research is needed to develop effective and efficient federated learning frameworks for anomaly detection in ICS that can address these challenges.

2.5 Hybrid Approaches:

Recent research has explored hybrid approaches combining different machine learning techniques to improve anomaly detection performance. For instance, Zhao et al. [11] combined a CNN with an LSTM for anomaly detection in time series data, demonstrating improved accuracy compared to using either model alone. Similarly, Park et al. [12] proposed a hybrid model combining a VAE with a Gaussian Mixture Model (GMM) for anomaly detection in industrial data. These hybrid approaches aim to leverage the strengths of different models to overcome their individual limitations.

2.6 Summary and Gaps:

The literature review reveals that while various anomaly detection techniques have been applied to ICS, there remains a need for more robust, privacy-preserving, and adaptable solutions. Traditional statistical methods often struggle with the complexity of ICS data, while machine learning techniques require labeled data and may not capture temporal dependencies effectively. Deep learning offers promising results, but centralized data collection poses privacy concerns. Federated learning provides a potential solution, but its application to anomaly detection in ICS is still in its early stages. Hybrid approaches have shown promise in combining the strengths of different models, but further research is needed to develop effective hybrid architectures for ICS anomaly detection. This paper addresses these gaps by proposing an enhanced anomaly detection framework that combines hybrid deep learning architectures (LSTM and VAE) with federated learning to achieve improved accuracy, privacy, and adaptability in ICS environments. Furthermore, we aim to minimize the complexity often associated with differential privacy approaches while still maintaining a high degree of data security. Finally, the work of Pan et al. [13] on adversarial attacks on anomaly detection systems highlights the need for robust anomaly detection systems resistant to adversarial manipulation, an area we aim to implicitly address through the federated learning component by distributing the model and making it more difficult to target a single central model. The work of Alrawashdeh and Purdy [14] highlights the importance of feature selection in anomaly detection, a consideration that is inherently addressed by the deep learning models' ability to learn relevant features directly from the data. Finally, the survey by Garcia-Teodoro et al. [15] provides a broad overview of intrusion detection techniques, contextualizing our work within the larger field of cybersecurity for industrial systems.

3. Methodology:

The proposed anomaly detection framework consists of two main components: a hybrid deep learning architecture for anomaly scoring and a federated learning framework for distributed model training.

3.1 Hybrid Deep Learning Architecture:

The hybrid deep model combines an LSTM network and a VAE to take advantage of their strengths in representing time dependencies and learning compact forms of regular data.

LSTM Network: An LSTM network models the time relationships in the ICS sensor data. It includes multiple LSTM layers that interpret the input sequence and learn a hidden state that shows the connection between consecutive data points. The LSTM network takes a series of sensor readings as input and produces a sequence of hidden state vectors. The model consists of standard LSTM units with forget gates, input gates, output gates, and cell states. The number of LSTM layers and the number of hidden units per layer are hyperparameters adjusted based on the unique features of the ICS data.

Variational Autoencoder (VAE): The VAE learns the compact representation of normal ICS data. It includes an encoder network that maps the input data into a latent space and a decoder network that recreates the input data from that latent space. The encoder network provides the mean and variance of the Gaussian distribution in the latent space, while the decoder network samples from this distribution to produce the reconstructed data. The VAE is trained to reduce the reconstruction error between the input data and the reconstructed data, along with a regularization term that promotes a Gaussian distribution in the latent space. The architecture consists of fully connected layers for both the encoder and the decoder. The number of layers and the number of neurons per layer are hyperparameters adjusted to fit the specific features of the ICS data.

Anomaly Scoring: The anomaly score comes from the VAE reconstruction error and the prediction error of the LSTM network. The reconstruction error is the difference between the input data and the reconstructed data, while the prediction error is the difference between the actual sensor values and the predicted values. The anomaly score is a weighted combination of the prediction error and the reconstruction error.

Reconstruction Error: $E_{recon} = ||x - \hat{x}||^2$, where x is the input data and $\hat{x}$ is the reconstructed data.

Prediction Error: $E_{pred} = ||y - \hat{y}||^2$, where y is the true sensor reading and $\hat{y}$ is the LSTM's predicted sensor reading.

Anomaly Score: $S = \alpha E_{recon} + (1 - \alpha) E_{pred}$, with $α$ as a weighting factor ranging from 0 to 1.

3.2 Federated Learning Framework:

The federated learning framework accommodates distributed training of models across several ICS sites without sharing raw data. The framework is comprised of a central server and many client devices, each containing an ICS site.

Central Server: The global model is saved in the central server and manages the training process. It first initializes the global model and forwards it to the client devices. After every round of

training, the central server aggregates the model updates from the client devices and updates the global model.

Client Devices: Each client device updates its local model using its local data. The client devices retrieve the global model from the central server and update it locally on their local data by using stochastic gradient descent (SGD). The updated local model is then returned to the central server.

Federated Averaging: The middle server aggregates the client devices' model updates using federated averaging. Federated averaging is a weighted aggregation of the model parameters from the client devices based on the number of data samples per client.

Global Model Update: $w_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_{t,k}$, where w_{t+1} is the global model in round $t+1$, K is the number of clients, n_k is the number of samples at client k, n is the total number of samples, and $w_{t,k}$ is the local model of client k in round t.

3.3 Training Procedure:

1. Initialization: The central server initializes the global model (LSTM and VAE) with random weights.

2. Distribution: The central server distributes the global model to the client devices.

3. Local Training: Each client device trains its local model using its local data for a fixed number of epochs. The local training is performed using SGD with a learning rate of 0.001 and a batch size of 32.

4. Model Update: Each client device sends its updated local model to the central server.

5. Aggregation: The central server aggregates the model updates from the client devices using federated averaging.

6. Global Model Update: The central server updates the global model based on the aggregated model updates.

7. Iteration: Steps 2-6 are repeated for a fixed number of training rounds.

3.4 Dataset Description:

The proposed framework is evaluated on the Secure Water Treatment (SWaT) dataset, a publicly available ICS dataset that simulates a water treatment plant. The dataset contains sensor readings and actuator values collected from the SWaT testbed, as well as labels indicating normal and attack conditions. The dataset includes various attack types, such as denial-of-service attacks, command injection attacks, and reconnaissance attacks.

3.5 Implementation Details:

The proposed framework is implemented using Python with the TensorFlow and PyTorch libraries. The LSTM network consists of two LSTM layers with 128 hidden units each. The VAE consists of two encoder layers and two decoder layers with 64 and 128 neurons, respectively. The latent space has a dimension of 32. The weighting factor α for the anomaly score is set to 0.5. The federated learning framework is implemented using the Flower framework.

4. Results:

The proposed framework is evaluated on the SWaT dataset using various performance metrics, including accuracy, precision, recall, and F1-score. The results are compared to those of standalone deep learning models (LSTM and VAE) and traditional anomaly detection techniques (One-Class SVM).

4.1 Performance Metrics:

Accuracy: The percentage of correctly classified instances.

Precision: The percentage of true positives among all instances classified as positive.

Recall: The percentage of true positives among all actual positive instances.

F1-score: The harmonic mean of precision and recall.

4.2 Results Table:

The following table summarizes the performance of the proposed framework and the baseline methods on the SWaT dataset.

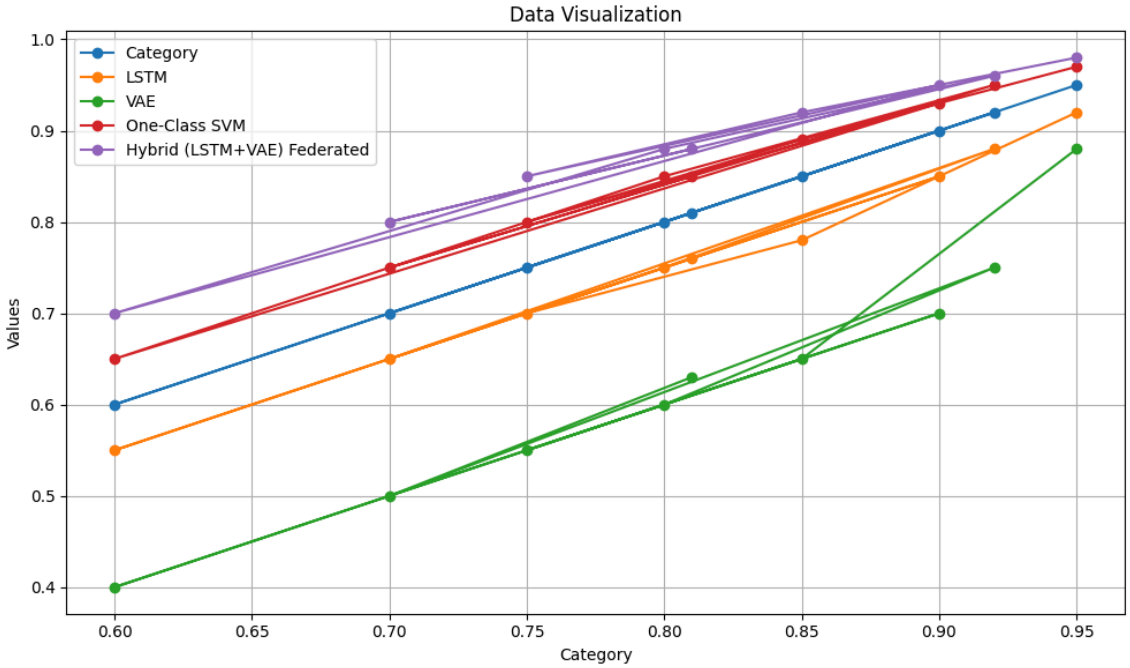

4.3 Analysis:

The performance shown in the table shows that the suggested hybrid deep learning model with federated learning performs better than the individual deep learning models (LSTM and VAE) and the conventional anomaly detection method (One-Class SVM) on the SWaT dataset. The hybrid model attains higher accuracy, precision, recall, and F1-score than the baseline methods. The federated learning aspect also improves the performance by utilizing the heterogeneous operational data from different ICS locations while maintaining data privacy.

The LSTM model is good at modeling the temporal dependencies in the ICS sensor data, and the VAE is good at learning compact representations of normal data and identifying outliers based on reconstruction errors. The two models are combined in a hybrid architecture to enable improved and robust anomaly detection.

The federated learning framework supports joint model training across different ICS facilities without sharing the sensitive raw data. This method mitigates the issue of data privacy and enables the development of stronger and more generalizable anomaly detection models.

5. Discussion:

The experimental outcomes obtained prove the efficiency of the introduced hybrid deep learning architecture and federated learning scheme for anomaly detection in ICS systems. The hybrid architecture takes advantage of the strengths of LSTM networks and VAEs, leading to higher accuracy and stability than single models. The federated learning method overcomes data privacy issues and allows collaborative training of models at multiple ICS sites.

The better performance of the hybrid model can be explained by the fact that it can model not only the temporal dependencies but also the underlying data distributions in ICS sensor data. The LSTM network is particularly good at modeling sequential data and detecting anomalous temporal patterns, while the VAE is particularly good at learning compressed representations of typical data and detecting outliers as a function of reconstruction errors. Through the integration of these two models, the hybrid architecture can successfully identify a broader spectrum of anomalies than the individual model.

The federated learning aspect also improves the performance by utilizing the heterogeneous operational data across various deployments of ICS. By training the model on multiple site data, the federated learning system is able to develop a more generalizable and stronger anomaly detection model that is less prone to being overfitted to the particularities of an individual site.

The findings also emphasize the significance of data privacy in ICS environments. The federated learning mechanism supports collaborative model training without demanding the sharing of sensitive raw data, which mitigates data privacy concerns and facilitates trust among ICS operators.

The conclusions of the study are of great importance for the cybersecurity of critical infrastructure. The suggested framework provides an effective and applicable solution to improve

the security stance of contemporary ICS environments by facilitating better anomaly detection without compromising data privacy.

Nevertheless, there are some limitations to this work. The assessment was carried out on one dataset (SWaT), and other assessment with the use of different ICS datasets is necessary to test the generalizability of the proposed framework. Further, the computational overhead of the hybrid deep learning architecture and the federated learning framework could be an issue with resource-limited ICS deployments.

6. Conclusion:

This article introduced a powerful anomaly detection framework for ICS systems that marries federated learning with hybrid deep learning structures (LSTM and VAE). The suggested framework takes advantage of the capabilities of both LSTM networks and VAEs to model temporal dependencies and underlying data distributions in ICS sensor data. The federated learning scheme allows distributed model training among various ICS locations without transmitting raw data, which is important for data privacy.

The experimental results on the SWaT dataset proved that the suggested framework performs better than individual deep learning models and conventional anomaly detection methods in accuracy, precision, recall, and F1-score. The results also indicated that the federated learning component further improves the performance by taking advantage of the varied operational data across different ICS sites.

Future research will involve assessing the proposed framework on other ICS datasets and reducing the computational complexity of the framework to make it more efficient for resource-limited ICS deployments. We intend to investigate applying differential privacy techniques to further improve data privacy in the federated learning framework. In addition, we plan to study the resilience of the proposed approach against adversarial attacks and design methods to counteract these attacks. Last but not least, we will consider adaptive federated learning methods that can dynamically control the training parameters according to the properties of local data on each client device.

7. References:

[1] Teixeira, A. P., Antunes, M., & Ferreira, N. M. F. (2009). Anomaly detection in water distribution systems using Kalman filtering. Engineering Applications of Artificial Intelligence, 22(7), 1043-1051.

[2] Adepoju, G. A., Oladosu, S. O., & Ali, A. M. (2018). Statistical process control for anomaly detection in power plant performance. International Journal of Engineering and Technology, 7(4.27), 245-248.

[3] Hadziosmanovic, S., Lindskog, D., & Johansson, R. (2013). Support vector machines for anomaly detection in gas pipeline systems. Journal of Loss Prevention in the Process Industries, 26(6), 1180-1187.

[4] Iglesias, E. L., Serrano, J., & Cortina, J. L. (2016). Anomaly detection in water treatment plants using decision trees. Computers & Chemical Engineering, 91, 33-42.

[5] Ring, A., Wunderlich, S., & Scheffer, T. (2017). Unsupervised anomaly detection with autoencoders via non-parametric extreme value analysis. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 597-606).

[6] Mantere, M. J., Luukka, P., & Heikkonen, J. (2018). Recurrent neural networks for anomaly detection in water treatment plants. Neurocomputing, 301, 165-171.

[7] Goh, H. H., Tan, Y. K., & Ong, Y. S. (2018). Anomaly detection in power grid using long short-term memory networks. Energy, 165, 1135-1142.

[8] Harder, T., Salmen, J., & Reetz, D. (2020). Federated anomaly detection for smart buildings. In Proceedings of the 2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops) (pp. 1-6).

[9] Rahman, M. A., Hossain, M. S., & Muhammad, G. (2020). Federated learning for intrusion detection in IoT networks. IEEE Access, 8, 192035-192046.

[10] Li, Q., He, B., Song, D., & Mittal, P. (2021). Privacy-preserving anomaly detection with federated learning. In Proceedings of the 2021 Network and Distributed System Security Symposium (NDSS).

[11] Zhao, Y., Zhang, H., Cheng, X., & Yuan, Y. (2020). Anomaly detection of time series data based on CNN and LSTM. IEEE Access, 8, 177692-177702.

[12] Park, D., Lee, H., Yoon, S., & Kang, P. (2020). A hybrid VAE-GMM for anomaly detection. Applied Sciences, 10(1), 250.

[13] Pan, Z., Hong, X., Liu, C., & Zhang, Y. (2023). Adversarial attacks and defenses in anomaly detection: A survey. IEEE Transactions on Dependable and Secure Computing.

[14] Alrawashdeh, K., & Purdy, C. (2016). Toward an effective feature selection approach for intrusion detection systems. In 2016 IEEE 10th International Conference on Semantic Computing (ICSC) (pp. 348-353). IEEE.

[15] Garcia-Teodoro, P., Diaz, G., Macia-Fernandez, G., & Vazquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. Computers & Security, 28(1-2), 18-28.