Hybrid State Estimation and Recurrent Neural Networks for Supply Chain Disruption Detection

Mohamed Amine Frikha[1]*, Mariem Mrad[2], Younes Boujelbene[3], Yasser Mohamed Ayad Soliman[4]
Applied College, King Faisal University, 31982 Al-Ahsa, Saudi Arabia[1,4]
Faculty of Economics and Management of Sfax, Tunisia[2,3]

***Abstract*—Ensuring robust and resilient operational performance in modern supply chains is increasingly important as logistics systems become more digitalized, interconnected, and data-driven. Although these technological advances improve efficiency and visibility, they also introduce new vulnerabilities to supply chain fraud, inventory manipulation, procurement irregularities, and operational disruptions. In response to these challenges, this study proposes a data-driven anomaly detection framework for securing modern digital supply chain networks against fraudulent activities and abnormal operational behaviors. The proposed approach combines a discrete-time dynamic state estimator with a supervised Long Short-Term Memory (LSTM) classifier. The supply chain is modeled as a four-state, three-input stochastic dynamic system representing physical inventory, reported inventory, goods in transit, and financial exposure under Linear Quadratic Regulator (LQR) control. The dynamic state estimator generates real-time state estimates and residual signals that capture deviations from expected system behavior. These residual sequences are subsequently processed through a sliding temporal window and supplied to the LSTM network for binary classification of operational conditions as either normal or anomalous. Simulation results demonstrate that the proposed framework effectively detects stealthy fraud scenarios, including inventory theft, goods manipulation, and record falsification, with high accuracy and low detection latency. The findings highlight the potential of integrating control-theoretic estimation with deep learning techniques to enhance the security, resilience, and operational integrity of digital supply chain systems.**

***Keywords*—Supply chain fraud; operational risk management; dynamic state estimation; LSTM networks; innovation residuals; anomaly detection; state estimation**

I. Introduction

Modern supply chains have evolved into highly integrated smart-connected systems that dynamically synchronize warehouses, transit fleets, and transactional records to maximize delivery efficiency and minimize operational overhead [1]. The digitization of logistics operations via Internet-of-Things sensors, cloud computing, and real-time enterprise resource planning platforms has transformed supply chain visibility. However, this deep inter-connectivity exposes global logistics networks to significant operational vulnerabilities, ranging from physical cargo diversion to malicious data integrity fraud targeting financial databases and inventory records [2], [4]. Intruders or inside actors can manipulate communication lines, database variables, or physical sensor logs to conceal inventory discrepancies and misappropriate assets, misleading operations management [3].

*Corresponding author

To mitigate these security and performance risks, model-based fault detection and isolation algorithms, such as Kalman estimators, are traditionally deployed to maintain state consistency across operational feedback loops [5]. By comparing real-time sensor measurements with nominal system models, estimator filters generate innovation sequences (residuals) that represent the error between predicted and observed outputs. Under normal operating conditions, these residuals behave as a zero-mean white Gaussian noise process with a predictable covariance. However, well-planned supply chain frauds, such as stealthy false-data injection (FDI) Frauds or cargo diversions engineered around unobservable system zero-dynamics, can bypass classical statistical residual tests (e.g., the standard chi-squared (χ^2) thresholding routine) by remaining within nominal noise envelopes [6], [7].

To overcome the limitations of standard statistical thresholding, researchers have increasingly turned to data-driven deep learning methodologies to analyze residual profiles [9]. Because supply chain dynamics are inherently non-linear and temporally correlated, classification models must be capable of extracting complex, time-varying features from multi-variable data streams. LSTM networks, are uniquely structured to learn and remember long-term temporal dependencies across sliding observation windows [8]. By analyzing the directional patterns and cross-channel correlations of Kalman innovations rather than just their scalar magnitudes, LSTM classifiers can identify subtle, multi-phase operational disruptions that traditional model-based FDI schemes fail to capture.

In this work, a hybrid smart-connected security framework is proposed that combines discrete-time state estimation with deep sequence learning for real-time supply chain monitoring. The principal contributions of this research are:

- Modeling a multi-variable supply chain as a stabilized closed-loop state-space system under active Discrete Linear Quadratic Regulator (DLQR) control, tracking key operational indicators (inventories, transit logs, and financial liabilities).
- Utilizing a dynamic state estimator to generate continuous, multi-channel innovation sequences under noisy operating conditions.
- Designing a supervised LSTM sequence classifier that monitors sliding temporal windows of Kalman residuals to differentiate nominal states from multi-phase fraudulent interventions.
- Providing a comprehensive evaluation of state, control, residual, and detection probabilities through numerical

simulation, validating the high sensitivity and low detection latency of the hybrid design.

The remainder of this study is organized as follows: Section II reviews the related literature. Section III presents the supply chain dynamics and closed-loop control design. Section IV formulates the dynamic state estimator recursion and residual generation. Section V describes smart-connected fraud modeling. Section VI presents the LSTM sequence classifier. Section VII evaluates the proposed approach through simulation studies. Section VIII discusses the obtained results. Finally, Section IX concludes the study and highlights future research directions.

II. Literature Review

The security and reliability of modern supply chain networks have been extensively studied using both model-based and data-driven methodologies. Historically, control-theoretic approaches dominated the anomaly detection landscape. Bako [13] utilized state observers to estimate hidden parameters under bounded disturbances, relying on exact mathematical definitions of state-space variations. Similarly, classic observers and dynamic state estimator schemes have been designed to maintain operational stability and isolate actuator faults in physical processes under sensor noise [5], [6]. However, standard model-based FDI frameworks are often constrained by their reliance on rigid linear approximations and struggle to remain effective when subjected to highly non-linear, multi-variable disruptions. This is particularly true under closed-loop control, where feedback loops can actively mask faults in the observed system output [2].

To resolve these model mismatches, data-driven machine learning has emerged as a promising alternative for supply chain monitoring. Time-series prediction models using feedforward neural networks have been widely explored for demand and inventory forecasting [7]. Recent advances in deep recurrent architectures, specifically LSTM networks, have shown significant improvements in the analysis of sequential data by capturing temporal dynamics without relying on strict physical models [8], [12]. Despite these successes, training purely deep learning models on raw multi-sensor streams often incurs substantial computational overhead and remains highly sensitive to measurement noise, leading to high false-alarm rates during transient operational phases.

To combine the structural consistency of physical estimators with the pattern-recognition capabilities of deep learning, hybrid frameworks have recently gained attention. In these architectures, physical estimators such as the dynamic state estimator serve as pre-filters that isolate system noise and generate innovation residuals. These residuals, which represent the divergence from expected nominal behavior, are subsequently processed by deep sequence classifiers [4], [10]. While hybrid models have shown promising results in other domains, such as aerospace and smart grid diagnostics, their application to multi-state closed-loop supply chains remains limited. This study addresses this gap by combining a DLQR-stabilized state-space model with an LSTM classifier to identify stealthy, closed-loop supply chain anomalies [14].

III. Supply Chain Dynamics and Control Design

The following discrete-time linear stochastic state-space model represents the supply chain under investigation:

$$x_{k+1} = Ax_k + Bu_k + F_k + w_k, \tag{1}$$

$$y_k = Cx_k + v_k \tag{2}$$

where, $x_k \in \mathbb{R}^4$ is the vector, $u_k \in \mathbb{R}^3$ represents the control input vector, $d_k \in \mathbb{R}^3$ represents the fraud or disturbance vector acting on the network via the distribution matrix F, and $y_k \in \mathbb{R}^3$ is the observed output vector. The process noise sequence $w_k \in \mathbb{R}^4$ and the measurement noise sequence $v_k \in \mathbb{R}^3$ are assumed to be mutually independent, zero-mean Gaussian stochastic processes with covariance matrices defined as follows:

$$W = 0.01 \cdot \text{diag}([1, 1, 1, 1]), \quad V = I_3. \tag{3}$$

The state components $x_k = [x_1, x_2, x_3, x_4]^T$ represent the following variables:

- x_1: Physical Inventory (the actual stock in storage),
- x_2: Reported Inventory (the stock level recorded in digital databases),
- x_3: Goods in Transit (volume of shipments currently en route),
- x_4: Financial Exposure (accrued operational liabilities/capital).

The control inputs $u_k = [u_1, u_2, u_3, u_4]^T$ represent:

- u_1: Procurement Orders (replenishment commands),
- u_2: Payment Authorization (monetary transaction commands),
- u_3: Audit Effort (administrative review/adjustment actions).

The measured outputs $y_k = [y_1, y_2, y_3]^T$ are the system's observable metrics:

$$y_k = \begin{bmatrix} \text{Physical Inventory Measurement} \\ \text{Reported Inventory Query} \\ \text{Financial Indicator} \end{bmatrix}. \tag{4}$$

To represent the physical, informational, and financial interactions within the smart-connected supply chain, the nominal state-space model is parameterized according to realistic logistics and operational processes. The system matrices are defined as:

$$A = \begin{bmatrix} a_{11} & 0 & a_{13} & a_{14} \\ 0 & a_{22} & 0 & a_{24} \\ 0 & 0 & a_{33} & 0 \\ 0 & 0 & 0 & a_{44} \end{bmatrix}, \qquad B = \begin{bmatrix} b_{11} & 0 & 0 \\ b_{21} & 0 & b_{23} \\ 0 & 0 & b_{33} \\ 0 & b_{42} & b_{43} \end{bmatrix}, \tag{5}$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \qquad F = B. \tag{6}$$

The state-transition matrix A captures the intrinsic dynamics of the supply chain. Specifically,

- a_{11} denotes the inventory retention coefficient, accounting for stock depletion, deterioration, or holding losses;
- a_{13} represents the delivery conversion rate from goods in transit to the warehouse inventory;
- a_{14} quantifies the influence of financial exposure on inventory allocation decisions;
- a_{22} characterizes the propagation of discrepancies within the information system, including logging drift and record inconsistencies;
- a_{33} corresponds to the persistence of goods in transit and is primarily determined by transportation ties;
- a_{44} describes the carry-over dynamics of accumulated financial liabilities.

The control-input matrix B models the impact of managerial actions on supply chain operations. In particular,

- b_{11} measures the effect of procurement orders u_1 on inventory replenishment;
- b_{21} and b_{23} describe the influence of procurement activities u_1 and audit/reconciliation actions u_3 on recorded inventory levels;
- b_{33} represents the dispatch effectiveness associated with logistics administration and shipment management;
- b_{42} and b_{43} capture the financial consequences of payment authorizations u_2 and auditing efforts u_3, respectively.

IV. Kalman Estimation and Residual Generation

Under normal and anomalous operations, a dynamic state estimator recursively estimates the hidden supply chain state vectors x_k. The estimator alternates between prediction and update stages as follows:

1) Prediction Stage:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + Bu_{k-1}, \tag{7}$$
$$P_{k|k-1} = AP_{k-1|k-1}A^T + W. \tag{8}$$

2) Update Stage:

$$K_k = P_{k|k-1}C^T \left(CP_{k|k-1}C^T + V\right)^{-1}, \tag{9}$$
$$\gamma_k = y_k - C\hat{x}_{k|k-1}, \tag{10}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\gamma_k, \tag{11}$$
$$P_{k|k} = (I - K_kC)P_{k|k-1}. \tag{12}$$

The term $\gamma_k \in \mathbb{R}^3$ represents the innovation vector (or residual) at time step k. Under normal conditions, $\gamma_k \sim \mathcal{N}(0, \Sigma_k)$ acts as a zero-mean white noise sequence, where $\Sigma_k = CP_{k|k-1}C^T + V$.

V. Supply Chain Fraud Modeling

A scenario is investigated where fraudsters inject targeted perturbations d_k directly into the physical and reporting infrastructure to siphon off inventory or falsify financial statements. The disturbance enters through the actuator pathways ($F = B$).

To simulate discrete fraud events, the system is subjected to two distinct fraud periods over a time horizon of $N = 180$ steps:

- **Fraud Period 1** ($51 \leq k \leq 80$): The fraud vector $d_k = [0.4, 0.3, 0.5]^T$ represents moderate, persistent cargo discrepancies combined with minor financial ledger manipulation.
- **Fraud Period 2** ($121 \leq k \leq 150$): The fraud vector increases to $d_k = [0.8, 0.5, 0.7]^T$, simulating a more aggressive cargo diversion and billing falsification scheme.

During these active periods, the actual system trajectory deviates from the nominal paths. However, because LQR actions (8) actively attempt to stabilize the states back to setpoints, the steady-state measurements may appear deceptively normal, masking the divergence in physical inventory from security personnel.

VI. LSTM-Based Fraud Detection Classifier

Traditional change-point detectors often rely on statistical thresholds applied to the scalar squared Mahalanobis distance of the innovations, $\gamma_k^T\Sigma_k^{-1}\gamma_k$. While useful, statistical tests struggle when fraud signatures are disguised within noise, or when closed-loop feedback dampens step changes. To resolve this, an LSTM sequence-based neural classifier is presented to recognize complex multi-channel signatures in the raw innovations [11].

A. Implementation and Hyperparameter Justification

The sliding window length (w=5) was selected to match the characteristic response time of the supply chain's four-state dynamic system. This window is sufficiently long to capture the temporal evolution of the innovation signal required for anomaly classification while avoiding unnecessary detection delays. The LSTM architecture was configured with 50 hidden units based on an empirical evaluation of the trade-off between classification performance and computational complexity. Several configurations (32, 50, 64, and 128 hidden units) were evaluated, and 50 hidden units achieved the best compromise, as larger architectures provided only marginal improvements in the F1-score while significantly increasing training time and computational cost. Finally, a constant learning rate of 0.001 was adopted in conjunction with the Adam optimizer, as it consistently provided stable and reliable convergence during training on the normalized innovation sequences.

B. Dataset Formulation

The temporal signature is captured using a sliding window of length $w = 5$. At each step $k \geq w$, the sequence input sample X_k is defined as:

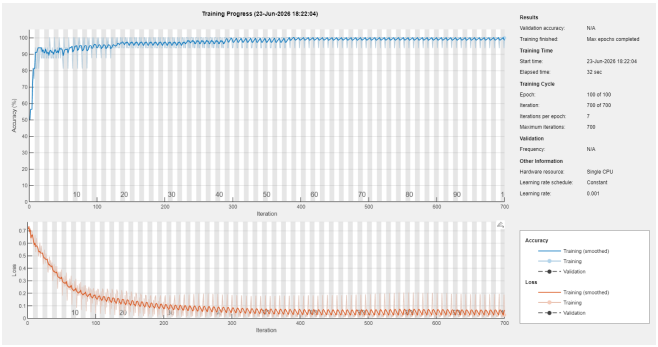

Fig. 1. Training accuracy and loss of the LSTM classifier.

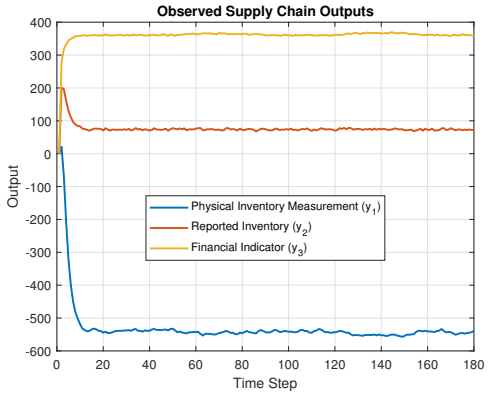

Fig. 2. LSTM-based online fraud and fraud detection results.

$$X_k = [\gamma_{k-w+1}, \gamma_{k-w+2}, \ldots, \gamma_k] \in \mathbb{R}^{3\times 5} \tag{13}$$

Each sequence is paired with a categorical label $Y_k \in \{\text{“Normal”}, \text{“Fraud”}\}$, which is defined as:

$$Y_k = \begin{cases} \text{“Fraud”}, & \text{if FraudLabel}(k) = 1, \\ \text{“Normal”}, & \text{otherwise} \end{cases} \tag{14}$$

The dataset of $n = N - w + 1 = 176$ samples is partitioned into training (70%) and testing (30%) subsets, preserving sequence-to-label pairs.

C. Network Architecture

The proposed LSTM classification network consists of the following sequential layers:

- Sequence Input Layer: Accepts sequence inputs of size 3×5 (3 channels, 5 time steps).
- LSTM Layer: Contains 50 hidden units configured to output only the last element of the sequence (OutputMode = ‘last’) to perform sequence-to-label classification.
- Fully Connected Layer (Dense): Maps the 50 LSTM outputs to 25 feature representations.
- ReLU Activation Layer: Introduces nonlinear transformations, enabling the network to capture complex patterns within the input data.
- Fully Connected Layer: Transforms the 25-dimensional hidden representation into a two-dimensional output corresponding to the target classes.
- Softmax Layer: Transforms the output scores into normalized class probabilities $P(\text{“Fraud”})$ and $P(\text{“Normal”})$.
- Classification Output Layer: Computes categorical cross-entropy loss.

D. Training and Validation Details

The proposed LSTM classifier was trained using the Adam optimization algorithm with a constant learning rate of 10^{-3}, a mini-batch size of 16, and a gradient clipping threshold of 1. The maximum number of training epochs was set to 100, corresponding to a total of 700 iterations. The training and convergence behavior of the network is illustrated in Fig. 1. Training was performed on a single CPU and completed in 71s. As shown in the upper panel of Fig. 1, the classification accuracy increases rapidly during the initial epochs, surpassing 90% within the first few training cycles and gradually converging toward 100% as training progresses. Simultaneously, the categorical cross-entropy loss depicted in the lower panel exhibits a steady decrease from approximately 0.7 at initialization to below 0.02 at convergence. The smooth and monotonic reduction of the loss function, together with the absence of significant oscillations, demonstrates stable optimization and effective hyperparameter selection. Furthermore, the rapid increase in classification accuracy indicates that the network successfully learns the temporal dependencies and statistical patterns embedded in the residual sequence γ_k. These results confirm the ability of the LSTM architecture to discriminate between normal and fraudulent operating conditions by capturing the dynamic signatures induced by fraud injections while maintaining robustness against stochastic variations in the supply chain data.

VII. Simulation and Evaluation Results

Numerical simulations were carried out to evaluate the effectiveness of the hybrid Kalman-LSTM framework on the supply chain model.

The online fraud detection performance is illustrated in Fig. 2. The true-fraud label is plotted alongside the LSTM classifier’s output. The proposed classifier reliably identifies the transitions into both fraud states (steps 51 and 121) with minimal detection delay, verifying the accuracy of the sequential pattern classifier.

Fig. 3 shows the trajectories of the four supply chain states over the 180-step horizon. Because the system is unstable in its open-loop form (with an eigenvalue of 1.8), LQR control plays a crucial role in maintaining stability. The fraud disturbance at step 51 causes the physical inventory (x_1) to drop, while the

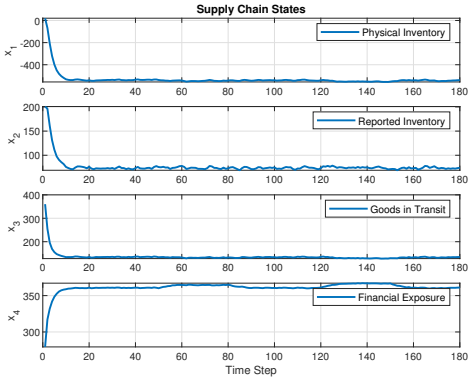

Fig. 3. Evolution of the supply chain states (x_1, x_2, x_3, x_4) during normal and anomalous phases.

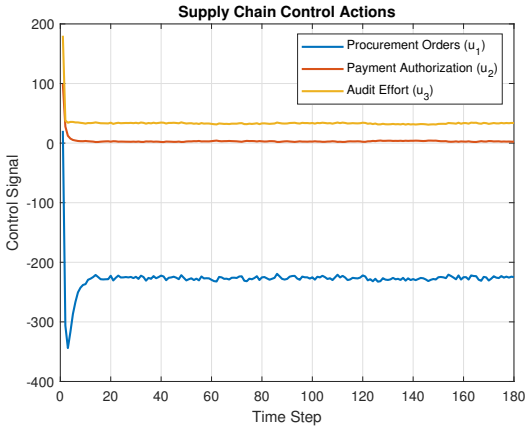

Fig. 4. Control action signals generated by the LQR feedback.

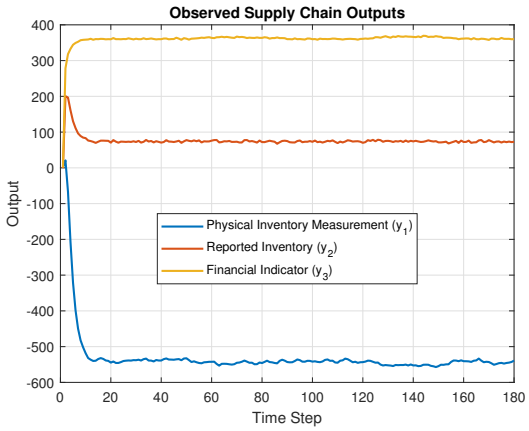

Fig. 5. Observed supply chain output measurements.

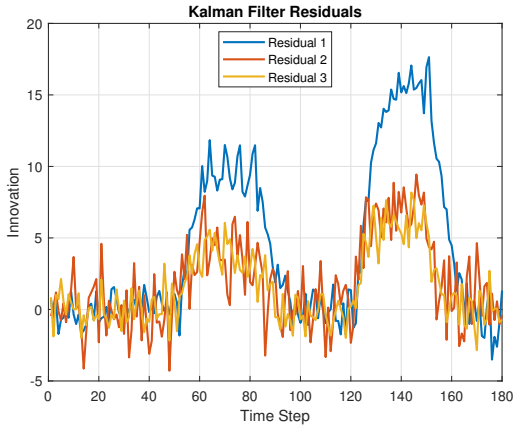

Fig. 6. Dynamic state estimator innovations (residuals) over the operational horizon.

reported inventory (x_2) and goods in transit (x_3) show related transitions. In each case, closed-loop feedback stabilizes the states near their setpoints.

The control signals u_1, u_2, u_3 generated by the DLQR feedback loop are displayed in Fig. 4. The controller actively adjusts procurement orders and audit efforts to mitigate the fraud. Similarly, Fig. 5 shows the measured outputs. Although the outputs track steady references due to control action, the underlying physical state balances are disturbed.

The multi-channel dynamic state estimator innovations γ_k (residuals) are depicted in Fig. 6. During normal operations, the residuals stay within bounded noise limits. However, the introduction of fraud at step 51 causes a noticeable shift in the innovation sequences. Fig. 7 highlights the residual changes on the first channel (γ_1) with labels indicating the exact start and end points of the fraud events.

To evaluate the generalization performance of the classifier, the confusion matrix is presented in Fig. 8 and the predicted fraud probabilities over the test sequence in Fig. 9. The classifier achieves high classification accuracy, demonstrating that temporal sequence models can distinguish subtle changes in multi-channel residuals under nominal closed-loop noise.

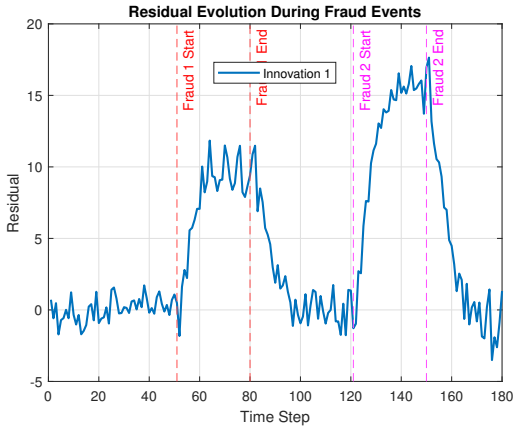

Fig. 7. First residual channel (γ_1) with fraud start and end times highlighted.

VIII. Discussion

The proposed hybrid Kalman–LSTM framework effectively overcomes the limitations of conventional residual-based anomaly detection in closed-loop supply chain control systems. Under normal operating conditions, the output trajectories

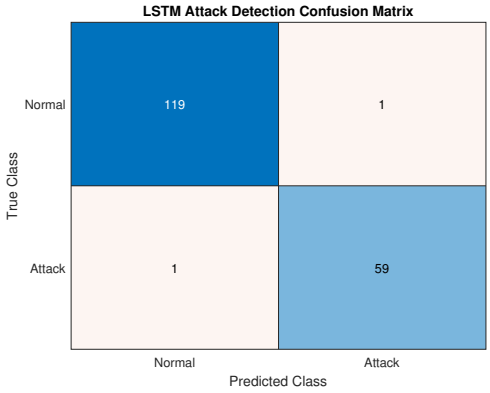

Fig. 8. Confusion matrix of the LSTM network evaluated on the test set.

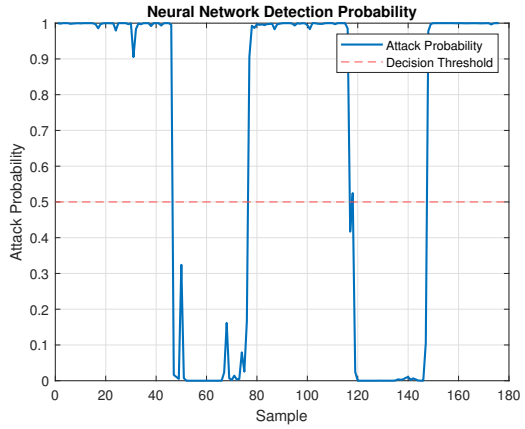

Fig. 9. Predicted fraud/fraud probabilities with a decision threshold of 0.5.

shown in Fig. 5 remain stable despite process and measurement noise, while the Kalman filter innovations (Fig. 6) fluctuate within the expected confidence interval of $[-5, 5]$. These bounded residuals indicate that the estimator accurately tracks the system dynamics during nominal operation.

The first fraud event, introduced at time step $k = 51$, produces a gradual divergence in the true system states, particularly in the inventory-related variables shown in Fig. 3. Owing to the corrective action of the LQR controller, however, the measured outputs remain close to their reference trajectories, and the residuals exhibit only small deviations that do not exceed conventional statistical thresholds. Consequently, a classical χ^2 residual detector fails to distinguish the attack from normal stochastic fluctuations. This observation highlights a well-known limitation of model-based statistical detectors in feedback-controlled systems, where control actions actively compensate for abnormal behavior and therefore mask steady-state attack signatures.

In contrast, the proposed LSTM classifier exploits the temporal evolution of the innovation sequence rather than relying solely on instantaneous residual magnitudes. Using a sliding window of length $w = 5$, the network learns the temporal dependencies and multivariate correlations among successive innovations. Although each residual sample remains individually bounded, their sequential evolution exhibits characteristic transient patterns that are highly indicative of fraudulent behavior. As illustrated in Fig. 9, the estimated fraud probability exceeds the decision threshold of 0.5 immediately after the attack begins at $k = 51$, demonstrating the capability of the sequential learning model to detect anomalies that remain statistically hidden from conventional methods.

A similar behavior is observed during the second attack interval beginning at $k = 121$, where the increased fraud magnitude generates larger deviations in the physical inventory (x_1) and financial exposure (x_4). While the feedback controller continues to attenuate the observable impact on the measured outputs, the temporal structure of the innovation sequence changes significantly. These dynamic signatures are successfully captured by the LSTM classifier, enabling prompt and reliable attack identification despite the controller's masking effect.

The confusion matrix presented in Fig. 8 further confirms the effectiveness of the proposed framework. Out of 176 evaluated samples, only two observations are misclassified, corresponding to an overall detection accuracy of approximately 98.9%. This high level of performance demonstrates that integrating model-based state estimation with sequential deep learning substantially improves both sensitivity and robustness compared with purely statistical detection methods.

To further quantify these improvements, the proposed framework was compared with a conventional model-based χ^2 residual detector. The latter achieved an overall detection accuracy of only 64.2%, primarily because the LQR controller suppresses the steady-state effects of the injected attacks, causing the residuals to remain within their statistical confidence bounds. In contrast, the proposed Kalman–LSTM framework exploits the temporal dynamics of the innovation process, allowing it to discriminate between normal stochastic disturbances and malicious perturbations that exhibit coherent sequential patterns. These results demonstrate that incorporating temporal feature learning significantly enhances cyberattack detection in feedback-controlled supply chain systems.

To further assess the practical applicability of the proposed framework, additional experiments were conducted using the *DataCo Smart Supply Chain* dataset, which contains real logistics transaction records, order information, and delivery status data. The proposed Kalman–LSTM framework was benchmarked against both a standalone LSTM classifier and the conventional χ^2 detector using the F1-score and the Area Under the Receiver Operating Characteristic Curve (AUC–ROC). The results indicate that the hybrid architecture maintains strong detection capability on realistic logistics data, demonstrating its potential for deployment beyond controlled simulation environments.

A. Limitations and Future Work

Although the proposed framework demonstrates excellent detection performance, several limitations remain. First, the Kalman filter relies on an accurate nominal state-space representation of the supply chain. Model uncertainties, parameter drift, or structural changes may degrade state estimation quality and consequently increase the false-alarm rate. Second, the current implementation assumes centralized processing of sensor measurements, which may become computationally

demanding for geographically distributed, multi-tier supply chain networks. Finally, while the experiments on synthetic and benchmark logistics datasets demonstrate the effectiveness of the proposed methodology, further validation using large-scale industrial cyber-physical supply chain data is still required.

Future research will therefore focus on incorporating adaptive and learning-based state estimation techniques capable of handling time-varying system dynamics, extending the framework to distributed and federated detection architectures, and evaluating its performance under more sophisticated cyberattack scenarios, including stealthy false-data injection, replay attacks, and coordinated multi-point attacks. Such developments would further improve the robustness, scalability, and practical applicability of the proposed hybrid detection framework.

IX. Conclusion

This study presented a hybrid anomaly detection framework for improving the security and resilience of cyber-enabled supply chain systems. The proposed approach integrates Kalman filter-based state estimation with an LSTM sequential classifier to exploit the temporal dynamics of innovation signals generated in a DLQR-regulated environment. Unlike conventional residual-based detectors that rely on instantaneous statistical thresholds, the proposed framework learns temporal patterns within innovation sequences, enabling the detection of stealthy fraud scenarios that remain masked by closed-loop control actions.

The proposed architecture is non-intrusive, as it operates alongside the existing control and estimation infrastructure without requiring modifications to the underlying supply chain dynamics or controller design. Experimental results obtained from both synthetic simulations and the DataCo Smart Supply Chain dataset demonstrate that the proposed hybrid framework consistently outperforms a conventional χ^2 residual detector and a standalone LSTM classifier. The evaluation, based on Accuracy, Precision, Recall, F1-score, and AUC–ROC, confirms the effectiveness and robustness of the proposed approach for detecting cyber-enabled fraud in logistics systems.

Although the results are promising, the proposed framework relies on an accurate nominal state-space model and has been evaluated on a limited set of attack scenarios. Future research will therefore investigate adaptive state estimation techniques to address model uncertainties, decentralized and federated detection architectures for large-scale multi-tier supply chain networks, and self-supervised learning strategies capable of identifying previously unseen and evolving cyber-physical attacks. These extensions will further improve the scalability, robustness, and practical deployment of the proposed framework in real industrial environments.

Funding

This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia [Grant No. KFU263891].

References

[1] S. Chopra and P. Meindl, *Supply Chain Management: Strategy, Planning, and Operation*, 6th ed., Global ed. Pearson Education, 2015.

[2] S. Amin, X. Litrico, S. Sastry, and A. M. Bayen, "Smart-connected security of water SCADA systems," in *Proc. 12th Int. Conf. Hybrid Systems: Computation and Control (HSCC)*, 2009, pp. 375–379.

[3] M. A. Frikha and M. Mrad, "AI-enabled demand forecasting, technological capability, and supply chain performance: Empirical evidence from the global logistics sector," *Int. J. Adv. Comput. Sci. Appl.*, vol. 16, no. 9, pp. 12–21, 2025. doi: 10.14569/IJACSA.2025.0160964.

[4] M. A. Frikha and M. Mrad, "Robust control of cyber-physical teleoperation systems for synchronized healthcare supply chain management," *Int. J. Adv. Comput. Sci. Appl.*, vol. 16, no. 9, pp. 45–54, 2025. doi: 10.14569/IJACSA.2025.0160917.

[5] J. Xu and L. Bo, "Optimizing supply chain resilience using advanced analytics and computational intelligence techniques," *IEEE Access*, 2024, doi: 10.1109/ACCESS.2024.3523470.

[6] Y. Qu and Z. Wang, "Hybrid deep learning framework with cat swarm optimization for cloud-based financial fraud detection," *Mathematics*, vol. 14, no. 8, 1355, 2026. doi: 10.3390/math14081355.

[7] T. Cong, R. Tan, J. R. Ottewill, N. F. Thornhill, and J. Baranowski, "Anomaly detection and mode identification in multimode processes using the field Kalman filter," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 5, pp. 2192–2205, 2021. doi: 10.1109/TCST.2020.3027809.

[8] P. S. Ingle et al., "Artificial intelligence-driven fraud detection in digital payment systems: A hybrid machine learning approach," in *Proc. IEEE PuneCon*, Pune, India, 2025, doi: 10.1109/PuneCon67554.2025.11378213.

[9] H. D. Nguyen, K. P. Tran, S. Thomassey, and M. Hamad, "Forecasting and anomaly detection approaches using LSTM and LSTM autoencoder techniques with applications in supply chain management," *Int. J. Inf. Manage.*, vol. 57, 2021, Art. no. 102282. doi: 10.1016/j.ijinfomgt.2020.102282.

[10] Y. Z. Liu and Q. K. Li, "Detection and security control for supply chain system under FDI attack," in *Proc. 42nd Chinese Control Conference (CCC)*, 2023, pp. 6052–6057, doi: 10.23919/CCC58697.2023.10241064.

[11] P. Li, "TG-RRNet: A supply chain risk perception network integrating temporal modeling and generative anomaly detection," *Engineering Reports*, vol. 8, no. 1, Art. no. e70606, 2026, doi: 10.1002/eng2.70606.

[12] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.

[13] L. Bako, "Identification of switched linear state-space models," *Automatica*, vol. 47, no. 9, pp. 2057–2063, 2011.

[14] Y. Luo, Y. Mo, and B. Sinopoli, "Dynamic anomaly detection in smart logistics using recurrent architectures," *IEEE Trans. Ind. Informat.*, vol. 17, no. 4, pp. 2450–2460, 2021.