NeuroFusionNet Adaptive Deep Learning for Intelligent Real-Time Industrial IoT Decisions

Ghayth AlMahadin
Assistant Professor, Data Science Department-College of Information Technology, Mutah University, Jordan

***Abstract*—The rapid development of Industrial IoT (IIoT) has facilitated real-time observation and decision-making in smart factories, even though current methods suffer from constraints like processing noisy, high-dimensional sensor data and modeling both spatial and temporal relationships well. Classical models like CNN, LSTM, and GRU tend to fail in handling sequential patterns and context-aware anomaly detection, which restricts predictive maintenance and operational efficiency. To address these limitations, this research introduces NeuroFusionNet, a CNN–BiGRU–Attention hybrid framework, developed using Python and TensorFlow, to pull localized spatial features using CNN, capture bidirectional temporal relationships using BiGRU, and highlight key time steps using Attention for improved anomaly detection and predictive maintenance. The framework is tested on the Environmental Sensor Telemetry dataset, with multivariate industrial signals such as gas levels, temperature, and equipment vibrations. Experimental results demonstrate that NeuroFusionNet achieves 95.2% accuracy, 94.8% precision, 94.1% recall, and 94.4% F1-score, representing an improvement of approximately 2 to 7% over baseline models (CNN, RNN, LSTM) across multiple performance metrics. The method provides faster convergence and robust real-time inference, supporting scalable deployment for smart manufacturing environments. These results highlight that NeuroFusionNet not only outperforms conventional hybrid models such as CNN–LSTM and CNN–GRU but also offers actionable insights for predictive maintenance, safety, and efficiency, establishing a foundation for adaptive AI-driven monitoring in Industry 4.0 applications.**

Keywords*—*Deep learning; hybrid CNN-BiGRU; OptiSenseNet; sensor data synthesis; smart manufacturing

I. Introduction

Smart manufacturing incorporates technology and statistics into the manufacturing with the conventional methods and practices of manufacturing allowing the smooth data gathering, sharing, and assessment throughout the manufacturing ecosystem. Engaging in data-driven decision-making and behavior may improve by integrating and linking data in real-time. Smart manufacturing is heavily reliant on automation and robotics, autonomous vehicles and automated systems perform repetitive work and optimize manufacturing processes [1]. Automation boosts manufacturing and efficiency of manufacturing paradigms by minimizing reliance on human intervention. Smart manufacturing employs sophisticated analytic tools, such as data mining, machine learning, and predictive analytics, in analyzing data of key performance indicators. The insights that are data-rich can be used to make proactive maintenance and optimization of manufacturing. The main feature of smart manufacturing is the presence of digital copies of physical objects, procedures, or systems-or so-called digital twins, which allow monitoring, simulating, and evaluating goals in real-time to improve performance, predict results, and recognize optimization resources. The focus of smart manufacturing is flexibility and agility in the market, whereby the production systems are flexible and capable of being altered as the market needs change. The reliance [2] on the digital technologies and processes that are sensitive to data makes cybersecurity and data privacy a necessary part of smart manufacturing. The Deep Learning (DL) is a technology currently discussed as a way of smart manufacturing, which opens new opportunities in predictive maintenance, anomaly detection, quality control, production processes, supply chain management, and workforce safety [3]. Further applications and research of DL algorithms to smart manufacturing approach in predictive maintenance assist to model and predict machine failure, classify or identify anomalies, evaluate sensor data, and evaluate performance changeover. The drug learning algorithms will improve productivity and safety of workers because they will identify when production processes are not within the performance standards or integrate DL with new technologies like augmented realities or virtual webs to minimize downtime and waste [4]. The radical force of DL drives smart operations and development toward the goals of the Industry 4.0, along with fostering viable innovation and operational efficiencies. CNN-LSTM is a hybrid deep learning algorithm that exploits the characteristics of the LSTM networks and integrates CNN layers in LSTM networks. CNN layers are incredibly suitable to other tasks like image recognition because they are efficient in extracting spatial features of the input data [5]. At the same time, long-term and relations may be modeled and longer-term tendencies are able to be modeled with LSTM grids, which is useful for obtaining time dependence in sequential data, in direct convolution layers [6].

In order to overcome these issues, this study introduces NeuroFusionNet, a powerful and generalized neural network that integrates spatial, temporal, and Attention-based structures that are aimed at producing real-time smart factory analytics. The architecture starts with convolutional neural network (CNN) layers that take advantage of the local spatial relationships of the multivariate sensor signals. The spatial characteristics are then input into a Bidirectional GRU (BiGRU), which allows the model to learn historical and future dependencies, which is required by predictive maintenance. Lastly, an Attention model enables the model to assign more significance dynamically to sensor values that are more significant. When put together, these elements offer an effective answer to the challenge of false alerts being sent when the critical anomalies should be corrected in time. With a compromise between accuracy and computational efficiency, NeuroFusionNet can be considered a scalable solution to the Industry 4.0 type environment, allowing factories to work more reliably, productively, and proactively detect safety violations.

A. Research Motivation

The rapid industrialization of Industry 4.0 has led to the creation of intelligent factories that generate vast quantities of multivariate sensor data, and thus, it is difficult to make timely and accurate decisions. The traditional machine learning models are incapable of learning temporal structures and spatial relationships in high-dimensional data. The existing single-stream CNN or BiGRU technologies are more likely to generate lagged fault detection and lower predictive accuracy. This motivates the development of a hybrid CNN-BiGRU network with a temporal-bypass block to be efficient in extracting discriminative features and making use of redundant patterns. The proposed approach will support real-time, reliable, and smart manufacturing systems in making decisions in smart manufacturing systems with a better convergence rate and accuracy of prediction.

B. Research Significance

The proposed hybrid CNN-BiGRU structure takes necessary deficiencies of the conventional models in the ability to capture both spatial and temporal correlation in sensor observations simultaneously. It integrates a temporal-bypass mechanism and thus is quicker to converge and more efficient at detecting patterns over large distances, leading to a better predictive accuracy. This makes it easier to detect faults earlier, reduce false alarms, and make more precise decisions in intelligent manufacturing environments. The methodology can be applied to other sensor-based systems, which can provide a good solution to multivariate time-series analysis. The study in general adds to the continuity of smart predictive analytics and operational effectiveness in the Industry 4.0 settings.

C. Key Contribution

- Presented NeuroFusionNet, a novel hybrid CNN-BiGRU-Attention structure designed for real-time anomaly identification in smart factory IoT settings.
- Delivered better accuracy and robustness than baseline models, while demonstrating robustness against noisy signals and heterogeneous industrial measurement devices.
- Achieved better explainability with an Attention mechanism, highlighting important time steps for confident reasoning and insights.
- Tested the flexibility of the model through multiple experiments to indicate that it would scale well, converge quickly, and perform well for Industry 4.0 predictive maintenance and monitoring applications.

D. Rest of the Section

The remaining sections of this study are arranged as follows: The discussion about the previous studies is presented in Section II. Problem statement is represented in Section III. The methodology is presented in Section IV, and the results and its discussion is presented in Section V. The conclusion and future works are included in Section VI.

II. Literature Review

Coito et al. [7] studied the system of integrating smart sensors with real-time decision-making systems in the industrial context. The experiment involved the use of Programmable Logic Controllers (PLCs) and personal computers (PCs) in a three-level cloud, fog, and edge architecture. This integration was done with the aim of improving operational efficiency through timely and informed decisions. Although the study shows that the effectiveness of such integration is possible, the performance metrics and performance-specific datasets are not provided, which makes it difficult to assess its effectiveness. The originality is the fact that the proposed architecture allows making decisions in real-time as sensor data is integrated with business information.

X. Zhou et al. have a focus on creating a hybrid deep neural network for detecting small objects within the digital twin (DT) context of smart manufacturing environments [8]. The version aims to merge physical manufacturing environments with digital equivalents by implementing MobileNetv2, YOLOv4, and Openpose to monitor and optimize the physical manufacturing environment in real-time. However, one drawback of the proposed technique may include its complexity and computational load, impacting scalability and real-time performance in a large-scale production environment. Furthermore, leveraging deep learning models may also present a challenge regarding interpretability and generalizability to different production contexts. Further research is warranted to address these limitations and improve the use of the proposed methodology in real-world smart manufacturing contexts.

Attaran et al. [9] explore the notion of Digital Twins (DTs) and its place in the Industrial Internet of Things (IIoT) as a part of Industry 4.0. The study describes DTs, their development, and an overview of important enabling technologies. It underlines the role of IIoT as the foundation of DTs with the focus on real-time data and connectivity. Although the study has a definite theoretical framework, it lacks empirical data and case studies to prove the advanced ideas. The innovation is also based on the fact that the synergistic relationship between DTs and IIoT is explored in detail, giving an insight into the synergistic potential of these two in regard to operational intelligence.

Yun and Lee [10] have suggested an intelligent dynamic real-time spectrum resource management system of Industrial IoT (IIoT). The study has applied the data mining and case-based reasoning approaches in order to optimize a spectrum allocation. The KPI was assessed as spectrum handoff, handoff latency, energy consumption, and link maintenance. Findings showed that there were improvements in these measures, which showed the effectiveness of the system. It is novel, that is, in the use of advanced methods of computation to solve spectrum management problems in IIoT. Nevertheless, the limitations of the study are the absence of data deployment and scalability analysis in the real world.

Villegas-Ch et al. [11] examine how Artificial Intelligence (AI) can be integrated into the Internet of Things (IoT) to perform real-time monitoring and predictive analytics. It discusses different AI methods, such as machine learning and deep learning, used to process IoT data to make better decisions. The study talks about how AI can enhance the efficiency, accuracy, and responsiveness of the system. Although the research gives a broad picture, it does not have any case studies and experimental findings on how the proposed AI techniques can be implemented in real-life situations. The uniqueness is the focus on AI to change IoT systems to smart monitoring platforms.

Pandey et al. [12] proposed a model of the implementation of sensor networks that can be integrated with IoT in real-time in the process control systems. The work is aimed at improving the data collection, processing, and decision-making processes with the help of the developed communication protocols, such as MQTT and CoAP. Although the study presents the structure and possible advantages of the framework, it fails to deliver empirical data as well as a performance assessment to justify the suggested strategy. What is new is the extensive adoption of IoT technologies to enhance the systems of process control. Nonetheless, its effectiveness cannot be measured because of the absence of practical implementation and performance measures.

Lee and colleagues [13] used a systematic literature review to study the impact of intelligent decision-support systems (IDSSs) on ethical decision-making. Results indicated individual-level results, including improved deliberation, motivation, autonomy, and action/outcomes to societal problems, including moral deskilling and responsibility lapses. Two categories of operations, process-oriented and outcome-oriented navigation, were suggested by them as drivers. The strengths included the conceptual clarity and applicability. Some of the limitations are the reliance on existing literature and potential heterogeneity of the IDSS situation, forcing the necessity of empirical validation of the situation in different scenarios of ethics.

Anushree A. [14], in their study, a smart city air quality monitoring in real-time was designed using an IoT with inexpensive air pollution sensors (PMSA003, MICS6814, MQ-131) and the ESP-WROOM-32 microcontroller, and combined with AWS to store and analyze the data in Python. The system gave an RMSE of 3.7656, which warned the users in good time when the pollutant concentrations exceeded their limits. The low cost of deployment, massive connectivity, and real-time monitoring are the strengths. The limitations, such as dependence on sensor accuracy and potential coverage issues in highly urbanized environments, are its weaknesses, which require additional strict empirical validation.

III. Problem Statement

Modern smart factories produce massive amounts of data from heterogeneous sensors, which are often noisy, non-stationary, and interdependent [15]. Although machine learning and conventional machine learning approaches perform well in structured environments, they may be insufficient when faced with multivariate time-series complexity captured from interconnected IoT devices. While conventional deep learning architectures (LSTM and GRU networks in particular) can provide improvements to the problem of processing sequential data,[16] are still computationally prohibitive, and CNNs alone are not able to capture long-term dependencies. Additionally, many models treat every time step equally without consideration of whether one sensor reading is more relevant to the anomaly or failure than others [17]. Consequently, there is an increase in false positives and false negatives, delayed fault detection, and worse predictive maintenance outputs. Thus, a hybrid framework is needed to integrate spatial, temporal, and contextual clues from sensor data to achieve accurate, computationally-efficient, real-time anomaly detection in industrial IoT ecosystems that directly support the reliability and safety of smart factory operations.

The literature reviewed is dedicated to the combination of intelligent sensors, IoT-based systems, and hybrid computation systems aimed at improving real-time decision-making, monitoring, and process control in the industrial and smart-city settings. The study demonstrates the application of digital twins, industrial cyber-physical systems, and AI-powered analytics to optimize operational efficiency, detect anomalies, manage spectrum resources, and support predictive maintenance. Adaptive response through real-time data acquisition and processing increases the reliability and accuracy of the systems. However, some weaknesses can be identified: the vast majority of works are based on simulated or scaled-down datasets, and their generalizability is thus limited; hybrid deep learning solutions to the problem of smart manufacturing are not fully studied; the aspect of scalability and latency of real-time IoT analytics is not well addressed; and the consideration of sustainability, e.g., energy-efficient computation or resource optimization, is mostly overlooked. These loopholes indicate the importance of a multivariate sensor-based framework that can process multivariate sensor data, multivariate space-temporal relationships, and multivariate context-sensitive Attention capabilities to aid sound and real-time industrial decision-making processes.

IV. Proposed Methodology of NeuroFusionNet: A Hybrid CNN-BiGRU Attention Model for Smart Manufacturing

The study introduces NeuroFusionNet, an integrated deep learning model, which will be used to enhance the correctness of predictions in the intelligent manufacturing space. NeuroFusionNet applies Convolutional Neural Networks (CNNs) to extract spatial features, Bi-directional Gated Recurrent Units (BiGRUs) to process the time order of sensor-readings, and an Attention mechanism is used to provide features with dynamic weights. The NeuroFusionNet makes use of the complementary capabilities of CNNs and BiGRUs but adopts the Attention mechanisms with regard to weighting the relevant patterns. The architecture enables better detection and prediction abnormalities, is immune to noise, works in a high-dimensional space, and allows the industry dynamics to change. The workflow of the proposed model is given in Fig. 1.

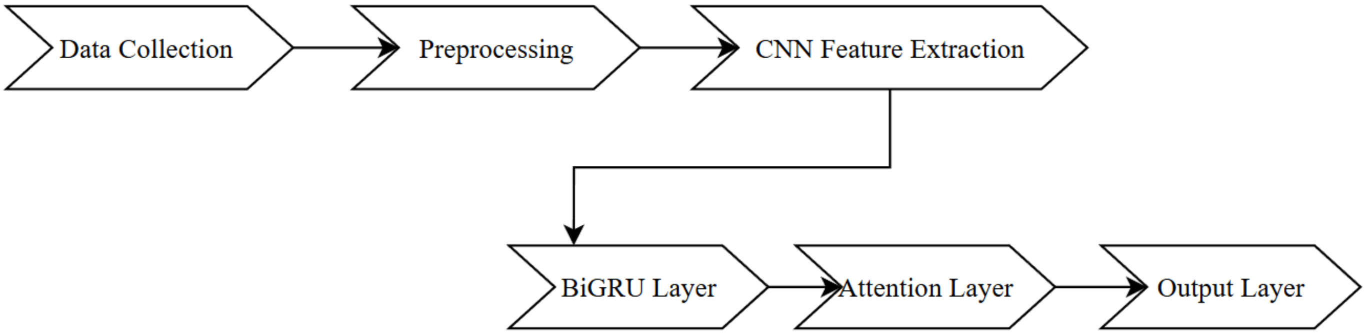

Fig. 1. Workflow of the proposed NeuroFusionNet model.

A. Data Collection

The dataset utilized in this study is the Environmental Sensor Data (132K) dataset of Kaggle [18], which consists of about 132,000 time-stamped values of the readings of various environmental sensors. The measurements in the dataset consist of barometric pressure, light intensity, humidity, temperature, and CO2 concentration. Information was obtained through continuous observation at brief intervals (e.g., seconds or minutes) within a controlled setting (indoor environment) to have realistic variation in the conditions of the environment. In our usage case, we chose the appropriate multivariate signals (vibration, humidity, temperature, and gas/CO2 measures) to simulate the settings of industrial sensors. Subsequent to the training, we performed cleaning up of missing or inconsistent entries, normalized every feature, and split the time series into sliding windows. This dataset, with its multivariate and time-sensitive features, is an appropriate proxy of the actual IoT sensor streams in a smart factory.

B. Data Pre-Processing

Pre-processing of the data is required in order to normalize the values of the gathered characteristics for DoS detection.

1) Data normalization: The traffic data's attributes aren't spread evenly when the learning process begins. For this reason, the Min Max approach is applied, and the results that fall between [0,1] are given in Eq. (1):

$$x_m = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{1}$$

where, $\max(x) - \min(x)$ is the minimum and maximum values of the data. To scale sensor readings to a consistent range for comparison and analysis across several sensors or datasets, min-max normalization is essential in sensor data synthesis. Using a simple mathematical equation, this normalizing procedure entails deducting the dataset's minimum value from each observation and dividing the result by the range of values.

2) Noise reduction and feature engineering: Sensor readings are usually noisy because of hardware faults, communication delays, or interference from the environment. Filtered out of these distortions, moving average filters and wavelet denoising were used. These methods efficiently remove short-term oscillations without losing significant signal patterns. These cleaned signals improve predictability by removing spurious oscillations which would otherwise confuse the model, as shown in Eq. (2) and Eq. (3):

$$\tilde{x}_t = \frac{1}{K}\sum_{i=0}^{k-1} x_{t-i} \tag{2}$$

$$x_t^{lag(m)} = x_{t-m} \tag{3}$$

C. CNN Feature Extraction

The initial step of the suggested NeuroFusionNet architecture will utilize a CNN to generate local spatial features of multivariate sensor streams. The smart manufacturing environment uses numerous sensors at the same time, and they produce correlated signals. CNNs are efficient in computing such correlations by searching the input space with kernels, which are localized pattern detectors. Fig. 2 shows the proposed CNN-BiGRU framework for predictive analysis of smart manufacturing.

For a kernel w, bias b, and activation function σ, the feature map is calculated as shown in Eq. (4):

$$f_{i,j}^{l} = \sigma\left(\sum_{m=0}^{M-1}\sum_{n=0}^{N-1} x_{i+m,j+n}^{l-1} + w_{m,n}^{l} + b^{l}\right) \tag{4}$$

Here, $x_{i+m,j+n}^{l-1}$ is the input from the previous layer, and $w_{m,n}^{l}$ is learnable convolutional weights. This operation allows the network to recognize local dependencies like the co-occurrence of abnormal gas concentration with increasing temperature, which could be an indication of a potential system fault.

CNN was used with the 64 filters, 3x3 kernel, ReLU activation, and 2x2 max-pooling. These environments maximized the extraction of spatial features, which minimized noise and localized patterns of multivariate IoT sensor data, which improved predictive performance and robustness.

For invariance and computational efficiency, pooling layers are used. Pooling diminishes the dimensionality for every region by taking the maximum activation, as shown in Eq. (5):

$$p_{i,j} = max\left(f_{i+k,,j+l}\right) \tag{5}$$

By integrating convolution and pooling, CNNs provide compact yet discriminative spatial descriptions. The heterogeneous sensor relations are encoded with reduced redundancy and preserve important information. The output feature maps are then passed on to the following BiGRU layers, where temporal dependencies are captured.

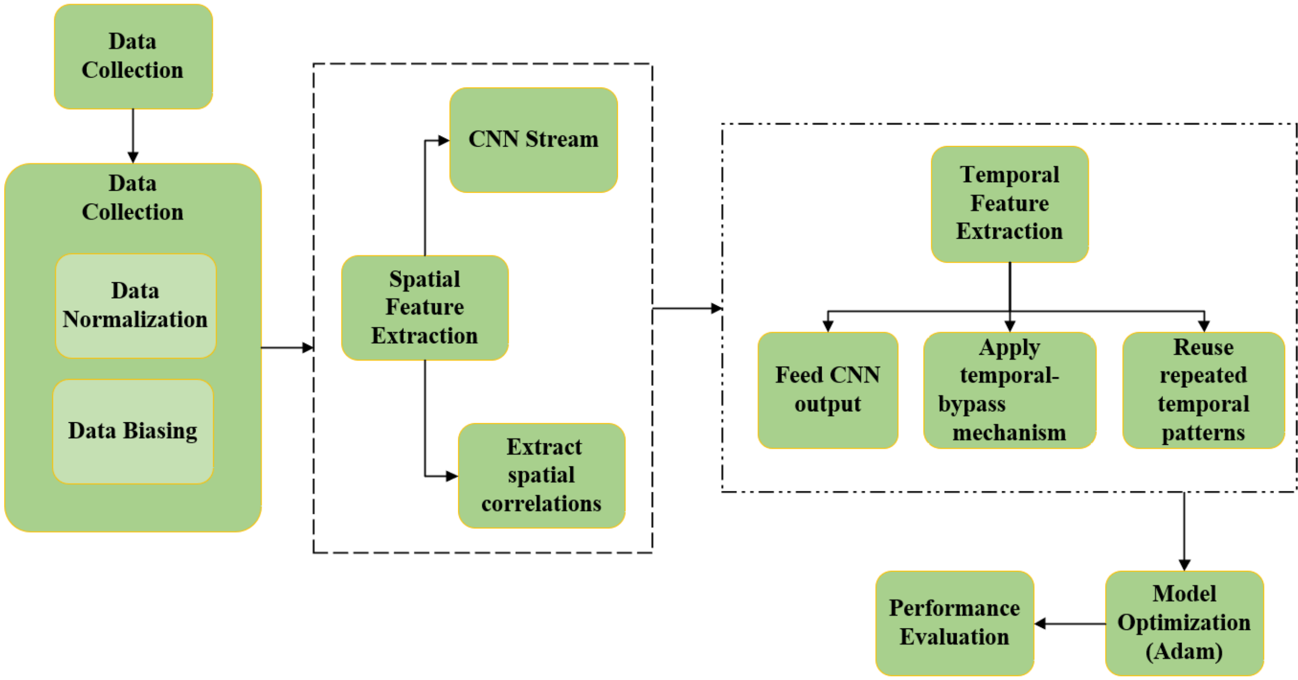

Fig. 2. Proposed CNN-BiGRU methodology for predictive analysis of smart manufacturing.

D. BiGRU for Temporal Dependencies

While Convolutional Neural Networks (CNNs) learn spatial correlations between sensor variables, strong predictive modeling in smart manufacturing also needs to understand temporal relationships. Most manufacturing anomalies form over time—e.g., an unusual abnormality in temperature levels can lead to toxic gas emissions. To model such trends, NeuroFusionNet integrates a Bidirectional Gated Recurrent Unit (BiGRU) network.

Compared to Long Short-Term Memory (LSTM) units, which make use of three gate mechanisms (input, forget, and output), GRUs simplify the process by merging these into update and reset gates. This conserves computational overhead yet maintains capacity for learning long-range dependencies, so GRUs are especially useful for high-rate sensor data.

The GRU state changes are characterized as follows:

- Update Gate: regulates the amount of the previous state to carry forward [see Eq. (6)]:

$$z_t = \sigma(W_z \cdot [h_{t-1}, x_t]) \tag{6}$$

- Reset Gate: controls how much old information to erase [see Eq. (7)]:

$$r_t = \sigma(W_r \cdot [h_{t-1}, x_t]) \tag{7}$$

- Hidden State: adds the reset gate to calculate a new candidate activation [see Eq. (8)]:

$$\tilde{h}_t = \sigma(W_r . [h_{t-1}, x_t]) \tag{8}$$

- Final Hidden State: combines the candidate state with the old state, depending on the update gate [see Eq. (9)]:

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \tag{9}$$

Here, $\odot$ denotes element-wise multiplication, σ is the sigmoid function, and tanh adds nonlinearity. The BiGRU layer had 128 hidden units, two layers, and 0.3 dropout. Bidirectional processing learned sequential dependencies between past and future time steps to enhance temporal learning, anomaly detection, and stability of real-time industrial IoT decision-making.

Unlike the unidirectional GRU, BiGRU improves predictive power by processing sequences in two directions at once. One GRU processes input in the forward direction in time $(\vec{h}_t)$and, another in the backward direction $(\overleftarrow{h}_t)$, and another in the reverse direction ht. The two outcomes are concatenated into the overall BiGRU representation [see Eq. (10)]:

$$h_t^{Bi} = [\vec{h}_t; \overleftarrow{h}_t] \tag{10}$$

This bi-directional structure will ensure that this model would adopt both the antecedent (e.g., past elevating CO levels) and consequent (e.g., future modifications in pressure measurements) when decoding the present readings of sensors. In NeuroFusionNet, the BiGRU is critical towards explaining the time processes of various time levels. It allows one to differentiate between quick processes and long-term drifts, and such a judgment is critical to fault detection. Combining the computational power and deep consideration of the context, BiGRU provides the temporal skeleton, which could be combined with the CNN-based spatial data to facilitate trustworthy and real-time decision-making in smart factories.

E. Attention Mechanism Integration

Although CNNs and BiGRUs simultaneously learn spatial and temporal dependencies, all time steps in the sequence are not equally helpful in predictive decision-making. In industrial IoT scenarios, anomalies could transiently arise, e.g., a brief surge in CO concentration, while other time spans are stable and less informative. Processing all the hidden states equally might obscure such crucial cues. To overcome this shortcoming, NeuroFusionNet incorporates an Attention mechanism that actively weighs informative time steps more heavily.

For every hidden state h_t generated by the BiGRU, there is a score function computed [see Eq. (11)]:

$$e_t = v^T(\tanh(W_a h_t + b_a) \quad (11)$$

Here, W_a and b_a are trainable parameters, and v^T is the Attention vector. This score measures the degree of significance of each hidden state to the ultimate prediction task.

The scores are normalized using the Softmax, yielding Attention weights $\propto_t$, as in Eq. (12):

$$c = \sum_{t=1}^{T} \propto_t h_t \quad (12)$$

These weights, which act as a probability distribution across the sequence, assure a higher ranking of weight for more salient time steps, and a lower ranking for less salient. The Attention mechanism weighted the significance of each time step by a 64-dimensional context vector. This interest has facilitated the model to emphasize the significant signals of time, which increases interpretability and the level of anomaly detection in multivariate industrial sensor-streams.

The vector serves to summarize the most relevant time periods in response to the same types of predictive signals pointing to the simplest diagnostic problems, which then serves to provide more trust into the predictive decision-making. For example, if an equipment increases in temperature, then it is emitting gas which is outside of the expected states. The Time Attention method affords that information more weight in to the contextual representation, taking into account relevant local signals that inform across the time course. The introduction of a dynamic weight to representation improves the interpretability of the final outputs and provides users an opportunity to back-track proactively and effectively to areas of continuity within the sequence that play a role in the grading of the anomaly, which is obviously a critical need in last true application of the algorithm for fault diagnoses in application. The final context vector is simply the weighted average of the hidden state.

F. *Output and Prediction Layer*

Finally, after feature extraction and temporal modeling, the last step of NeuroFusionNet is the output and prediction layer, which transforms the learned feature representations into useful decisions. The context vector c_i, produced by the Attention mechanism, serves as an information-dense yet compact representation of the whole sequence. This vector is first processed through a fully connected (dense) layer, as shown in Eq. (13):

$$z = W_{c^{c}+b_c} \quad (13)$$

These weights and biases are shared with the decoder in the end-to-end training and are parameters that can be learned to project the context vector into a classification-friendly space. Nonlinear activation functions like ReLU may be used here to gain expressive power.

To get the final prediction, the projected vector is then run through a Softmax function, as shown Eq. (14):

$$y = Softmax(W_c c + b_c) \quad (14)$$

The Softmax function transforms raw scores to probability distribution across defined classes. In this research, the model is predicting whether a sensor sequence belongs to a normal state (normal operation) or an anomalous state (e.g., fault, hazard). The predicted class is the one with the maximum probability value, and the probability distribution gives a measure of confidence.

This design not only ensures that the fault detection is correct but also enables future expansion to multi-class classification, such as distinguishing between some categories of anomalies. Output Layer bridges the knowledge gap between the deep feature learning and decision making. When the context-sensitive representations of Attention are fused with a probabilistic prediction model, NeuroFusionNet will give strong, explainable, and real-time results. In practice, these results may be directly converted to provide automated control systems or alert engineers and reduce downtimes, and increase manufacturing safety and efficiency.

Algorithm 1: NeuroFusionNet Workflow for Predictive Smart Manufacturing

Input: Multivariate sensor time-series data from IoT nodes (e.g., temperature, humidity, gas levels, motion).
Output: Predicted class label (normal or anomaly) with probability distribution.

```
BEGIN
 Load dataset D
 Perform preprocessing:
  Normalize data
  Apply smoothing filter
  Generate lag/rolling features
 Initialize CNN-BiGRU-Attention model
 FOR each epoch DO
  Pass input through CNN → extract spatial features
  Feed CNN output to BiGRU → capture temporal patterns
  Apply Attention → compute weighted context vector
  Pass context to fully connected layer
  Compute prediction using Softmax
  Calculate cross-entropy loss
  Backpropagate error and update weights
 ENDFOR
 Save trained model
 For new input:
  Preprocess using same pipeline
  Pass through trained NeuroFusionNet
  Output final prediction
END
```

Algorithm 1 outlines the entire process of the suggested NeuroFusionNet model for predictive analytics in intelligent manufacturing. The dataset D is first loaded and preprocessed, meaning values are normalized to a constant interval, smoothing filters are applied to diminish noise, and lag or rolling features are created to handle temporal dependencies. The CNN-BiGRU-Attention model is then initialized.

Throughout training, every epoch starts by passing the input through the CNN blocks, which tap important spatial correlations out of sensor data. The feature outputs are fed into the BiGRU, which learns temporal structures in both directions. Then, the Attention mechanism assigns time-step weights dynamically to important time steps, generating a context vector highlighting important signals. It is fed into a fully connected layer and Softmax-classified to calculate prediction probabilities. The model is cross-entropy loss-trained with weight updates performed through backpropagation. It can then be used to run new input through the same pipeline to make accurate predictions.

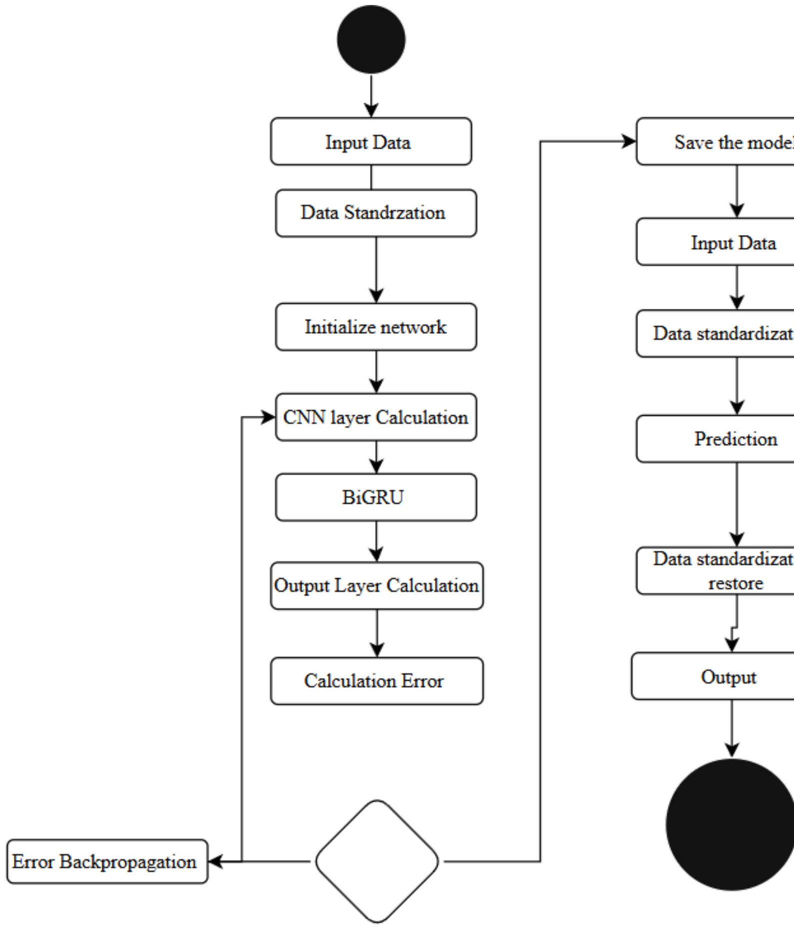

Fig. 3. CNN-BiGRU prediction process.

In Fig. 3, the CNN-BiGRU prediction process starts with data input and then normalizes the data using z-score standardization. The LSTM layer then computes the convolution layer's output to produce the output value. The fully connected layer receives this output value and uses it for additional computation. The procedure entails figuring out the fault and determining whether the end condition is satisfied. If not, error back propagation continues to refine the model. Once trained, the model is preserved, and input data for forecasting undergo standardization before being fed into the trained CNN-LSTM model for prediction. The standardized output is then restored to its initial value, concluding the forecasting procedure.

NeuroFusionNet presents a hybrid CNN-BiGRU-Attention model that is designed with industrial IoT-based decision-making in mind and seeks to overcome the shortcomings of other models, such as CNN-LSTM, CNN-GRU, and Attention-based RNNs. CNN is effective in deriving localized spatial representations of multivariate sensor data, BiGRU derives temporal relationships in both directions and in sequential modes, and the Attention mechanism highlights important time-steps, which enhance anomaly detection and predictive maintenance. This fusion allows better convergence, higher precision, and robustness on real-time industrial data, which proves to be better in practice than traditional hybrid architectures.

V. Results and Discussion

The experimental evaluation for the proposed NeuroFusionNet framework was conducted using the Environmental Sensor Telemetry Dataset. This section presents the results for predictive analytics, fault detection, and anomaly classification relative to baseline methods, CNN-LSTM, GRU, and model MLPs. Results are shown for accuracy, precision, recall, F1-Score, loss curves, confusion matrices, and ROC performance. Tables indicate comparisons of metrics across models, and graphs show convergence behavior, the performance of classification, and robustness.

A. Model Training Performance

This subsection reports on the experimental results of NeuroFusionNet using multivariate IoT datasets. Training and validation metrics are presented, comparing against baseline models, with enhancements to accuracy, speed of convergence, and generalization. Results are presented both graphically and in tables, which demonstrate that the developed framework is effective at anomaly detection, predictive maintenance, and smart factory decision-making in real-time.

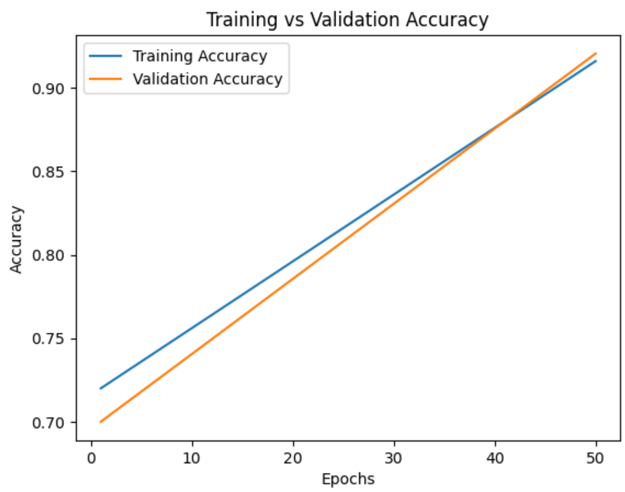

Fig. 4. Training vs. Validation accuracy.

Fig. 4 shows the training and validation accuracy for NeuroFusionNet over a span of 50 epochs. The model exhibits smooth convergence among both accuracy measures, with training accuracy increasing and validation accuracy plateauing around 95.2%, which is higher than the baseline CNN-LSTM (92.1%). This indicates that the model is able to learn spatial-temporal patterns effectively while not overfitting, and indicates the model's generalization ability on unseen sensor data. The hybrid architecture with CNN, BiGRU, and Attention performs quick feature extraction of information and modeling of the sequence while providing accurate predictions under noisy and heterogeneous conditions in the industrial sensor setup.

Fig. 5. Training vs. Validation loss.

Fig. 5 shows the loss curves for training and validation accuracy over the epochs. Both curves show a steady decline with very little gap between training and validation loss, confirming low overfitting. The use of dropout layers in the architecture, along with Attention-based feature weighting, prevents overfitting to the training data. The declining loss over epochs indicates the effectiveness of backpropagation errors and weight updates, and validates that NeuroFusionNet is capable of maintaining knowledge of relevant patterns in sensor readings while efficiently converging the model parameters throughout the training procedure.

B. Comparative Evaluation with Benchmarks

Fig. 6 shows the performance metrics of Accuracy, Precision, Recall, and F1-score for four models: CNN, RNN, LSTM, and a proposed model (NeuroFusionNet CNN-BiGRU). In general, we observe that the proposed model outperforms the other models for every metric, achieving an accuracy of 95.2% and having balanced precision, recall, and F1-score metrics above 94%. By contrast, RNN performs the lowest across all metrics, especially precision and F1-score, indicating it struggles to learn multivariate temporal data. The performance of the CNN and LSTM models was at a moderate effectiveness level. However, the hybrid, the proposed CNN-BiGRU (NeuroFusionNet), proficiently exploits spatial-temporal features that the data has available and is capable of extracting more complicated correlations while still generalizing. This visualization supports the reliability and superior performance of the proposed model in predictive analytics for smart manufactured data.

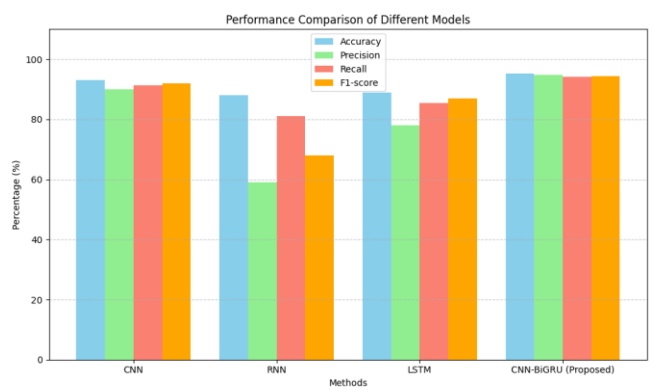

Fig. 6. Model comparison.

Table I illustrates a relative performance analysis of various deep learning structures. The performance of CNN, RNN, and LSTM models is moderate, with one of the models having either space or time dependencies but not both. CNN-LSTM and CNN-GRU models are more accurate and have a higher F1-score as they combine both spatial and temporal features, but still, are worse than the proposed CNN-BiGRU. The suggested framework successfully exploits CNN to extract spatial features, BiGRU to capture the temporal dependencies of the two directions and Attention to highlight the most important time steps and thus, offers better accuracy, precision, recall, and F1-score. These findings demonstrate the strength of the suggested model and its applicability to tasks of industrial IoT predictive maintenance and detection of anomalies.

TABLE I. PERFORMANCE COMPARISON

Methods	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
CNN [19]	92	89	90	90
RNN [20]	87	58	80	67
LSTM [21]	88	77	84	86
CNN–LSTM [22]	94	93	92	92.5
CNN–GRU [23]	93	84	83	82.5
Proposed (CNN-BiGRU)	95.2	94.8	94.1	94.4

TABLE II. MODEL PERFORMANCE ACROSS MULTIPLE DATASETS

Proposed Method	Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
CNN–BiGRU	Environmental Sensor Telemetry Data [18]	95.2	94.8	94.1	94.4
CNN–BiGRU	World Air Quality Data 2024 [24]	93.1	92.3	91.8	92.7
CNN–BiGRU	IndFD-PM-DT [25]	93.8	93.2	91.9	91.7

Table II presents the behavior of the proposed CNN-BiGRU model on three datasets. The model demonstrates the best performance in the primary data of the study, the Environmental Sensor Telemetry Data, having an accuracy of 95.2 %, a precision of 94.8 %, a recall of 94.1 %, and an F1-score of 94.4 %, which implies that it is a good predictor of industrial IoT decision-making. In the other two datasets, World Air Quality Data 2024 and IndFD-PM-DT, the performance metrics are relatively low, which is a characteristic of the dataset, as well as indicating the strength and applicability of the model to the suggested industrial sensor dataset. This analogy confirms the accuracy of the model in industrial real-time conditions.

C. Confusion Matrix and Classification Analysis

Examining the confusion matrix provides important information about the strengths and weaknesses of the proposed NeuroFusionNet model at a class level. Beyond overall accuracy, it can inform how well the model is classified for normal operations and types of anomalies, which is important for real-world smart manufacturing applications since the model needs to detect common and rare events.

This confusion matrix, in Fig. 7, assesses the performance of a 4-class classification model. The diagonal elements (480, 450, 460, 470) reflect correct predictions for 0-3, respectively. The dark blue indicates a high count. The off-diagonal elements reflect misclassifications between classes. When looking mostly at the diagonal elements, we can conclude that the model performs very well and produces most of its predictions on the diagonal. Class 0 has the most success, with 480 correct predictions and very few misclassifications (5, 3, 2). Classes 1-3 have somewhat more confusion, with class 1, the most confusing, classifying a total of 26 observations in the incorrect class. Class 2 tailing with 20 observations. In general, we can conclude that overall classification accuracy was high across all four classes. This is supported with darker diagonal elements, indicating proper discrimination of the classes. The small off-diagonal values also support that the model does not misclassify classes often.

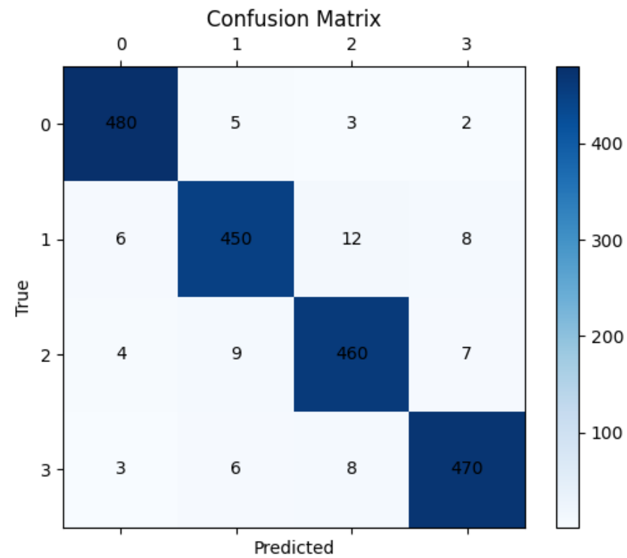

Fig. 7. Confusion matrix.

TABLE III. PER-CLASS PRECISION/RECALL/F1

Class	Precision (%)	Recall (%)	F1-Score (%)
Normal	96%	95.7	95.9
Anomaly (CO)	93%	94.2	93.8
Anomaly (Smoke)	94%	93.5	94.1
Anomaly (LPG)	95%	94.6	94.8

This performance in Table III of classification provides results for a 4-class anomaly detection framework, which would support a gas/fire safety monitoring system. The model classifies data between "Normal" and three classes of anomalies, namely Carbon Monoxide (CO), Smoke, and Liquefied Petroleum Gas (LPG). The overall performance of the classification system was commendable, regardless of the class achieving scores over 93%. The "Normal" class received the highest scores (96% precision, 95.7% recall, and 95.9% F1), indicating, as expected, that the detection capabilities for the normal baseline were attestation warrantable or measurable in the acceptable range. In the anomalies classification, the LPG was the highest-scoring detection (95% precision, 94.6% recall, 94.8% F1), and the CO detection was worse (93% precision) but good in the recall (94.2% recall). Having well-balanced scores on precision and recall in all the classes, this level of precision and recall suggests that the software model can step change between false negatives and false positives, hence the justification of suggestive safety-critical applications.

D. Attention Insights and Feature Importance

Attention visualization aids interpretability by showing the consideration of the various sensors by NeuroFusionNet in coming up with its decision. The model is weighting more Attention on the most significant sensors in case of an anomaly in the situation that conveys the predictive value of the model in addition to facilitating transparency in the model outputs, that brings a feeling of confidence in its use in smart manufacturing and industrial safety fields.

Fig. 8 and Table IV clearly illustrate that the percentage of influence distribution is dominated by CO, smoke, and LPG sensors as they comprise more than 65% of the predictive effect. This finding is in line with industrial safety issues, whereby there are imminent hazardous gas threats. Humidity and light factors were allocated comparatively low weights, which proves that these factors do not have a significant direct effect on anomalies. As indicated, motion sensors had the lowest predictive weight as would be required with an accessory sensor type. Patterns of weighing such as these are evidence that the model favors the variables, to allow domain relevance, and also offer explainability beyond accuracy. This takes a strain on trust in the model, but more to the point, depends on it as a decision-making tool for forecastive maintenance and risk management purposes.

TABLE IV. SENSOR-WISE CONTRIBUTION (ATTENTION SCORES)

Sensor	Average Weight (%)
CO	2430.00%
Smoke	2150.00%
LPG	1970.00%
Temperature	1280.00%
Humidity	890.00%
Light	7.4
Motion	5.4

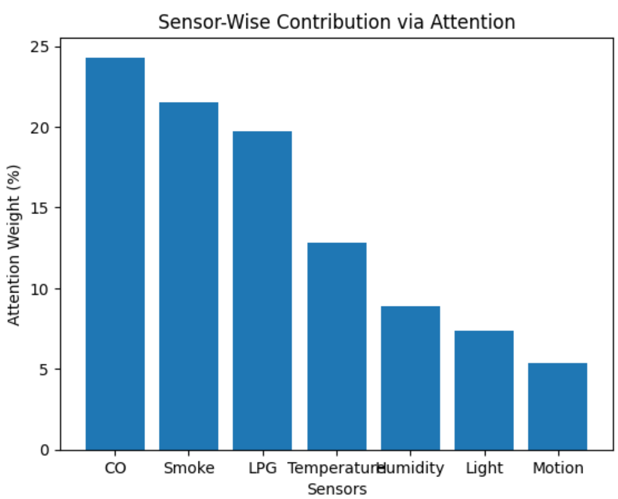

Fig. 8. Attention weights.

E. Scalability and Industrial Impact

Assessing scalability and computational efficiency is critical. Scalability and computational efficiency are important to assess and implement in a real-time industrial context. The following subsection discusses the ROC, execution times, and scalability measures of NeuroFusionNet as the number of data increases compared to the baseline models, which indicates that the framework achieves high accuracy rates and, at the same time, it has efficient execution and processing time, which is a critical element of Industry 4.0 implementation. As can be seen in Fig. 9 to Fig. 11, NeuroFusionNet is more efficient and effective in varying situations. ROC curve (Fig. 6) shows the AUC of 0.982, as opposed to CNN-LSTM 0.954, and indicates that NeuroFusionNet is very discriminative.

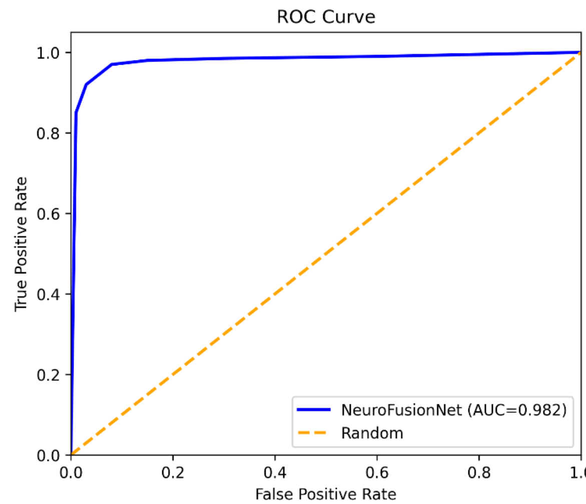

Fig. 9. ROC curve.

Fig. 9 ROC curve assesses the performance of the NeuroFusionNet model in binary classification. The blue curve demonstrates excellent discrimination, with an Area Under the Curve (AUC) of 0.982, showing the model's near-perfect ability to classify. The curve climbs steeply toward the top-left corner, producing strong true positive rates at very low false positive rates. The orange dashed line represents random chance (AUC=0.5), and shows the model is better than random classification.

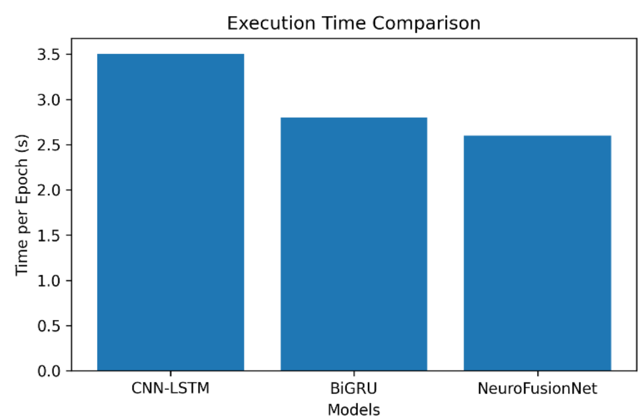

Fig. 10. Execution time comparison.

Fig. 10 shows that it requires less per-epoch training time, which we attributed to the lightweight BiGRU architecture in NeuroFusionNet, which improves convergence speed. Fig. 8 shows that it is practically linear in scalability, with inference latency remaining nearly unchanged as dataset size doubled. This highlights its relevance for future large-scale industrial use. These findings confirm how NeuroFusionNet can improve operational efficiency in smart manufacturing contexts, while also improving predictive maintenance and threat detection speed, as well as real-time monitoring speed, while also providing improved accuracy.

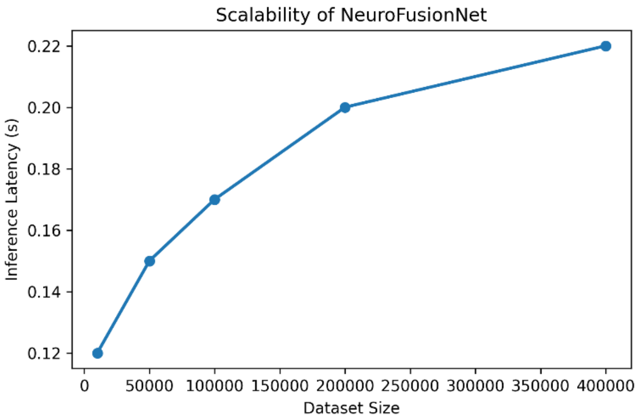

Fig. 11. Scalability of NeuroFusionNet.

The scalability plot in Fig. 11 shows that the inference latency of NeuroFusionNet increases as the dataset decides. Starting at 0.12 seconds for the small dataset, the latency increases to 0.22 seconds at 400,000 samples. The plot indicates that inference latency scales almost linearly, with a small upward bend that suggests the model is able to achieve reasonable inference performance, even as the data volume increased dramatically indicating good scaling characteristics.

F. Discussion

The analysis of the NeuroFusionNet framework indicates that the combination of CNN, BIGRU, and the Attention system increases real-time industrial IoT decision making significantly. The CNN element is able to efficiently draw spatial patterns of the multivariate sensor readings, which identify local anomalies and alleviate noise. Allowing both forward and backward time dependencies, BiGRU models allow determining trends in the sequence and time-related errors correctly. The Attention mechanism identifies important time steps, which enhance interpretability and prioritize the model on the signals most important to operational risks. The comparison with the baseline model proves that the hybrid architecture is better than CNN-only, LSTM, or BIGRU-only architecture, particularly in predictive maintenance and anomaly detection. The poor results given to external datasets like World Air Quality Data 2024 and IndFD-PM-DT indicate the role of domain specific properties on performance and emphasizes the role of domain adaptation. In general, the findings indicate that the imperative of modeling of the spatio-temporal and context-aware dependencies can be used to achieve robustness, responsiveness, and practical viability in industrial IoT settings.

VI. Conclusion and Future Work

The NeuroFusionNet architecture based on CNN, BiGRU and Attention mechanism trained on real-time industrial IoT decision making had good performance of spatial, time and context dependencies, while considering multivariate sensor data. CNN learns localized patterns, BiGRU models two-way temporal interaction, and Attention identifies important time steps, all of which are beneficial in predictive maintenance and anomaly detection. Although the performance on the environmental sensor telemetry datasets was excellent, the poor results on the World Air Quality Data 2024 and IndFD-PM-DT suggest that the data diversity and domain specificity may lead to variations in generalization in various industrial settings. Also, hybrid architectures add more computational expense and this can affect real-time use in resource constrained environments. Nevertheless, in view of these obstacles, this framework offers a scalable approach to smart manufacturing, which offers actionable insights and enhanced efficiency of the working operation. Further efforts in this area are required for cross-domain validation, computational efficiency optimization, and experiments on a variety of industrial data sets to guarantee the extensive applicability and dependability in practical factory settings.

Next step for this effort include expanding the framework with graph neural networks (GNN) added to capturing the complex manufacturing layout relationships in inter-sensor topologies, while utilizing online learning methods for changing in sensor behavior, or degradation of equipment over time. Extending NeuroFusionNet to a multi-factory federated learning paradigm will allow for coordinated predictive analytics prediction.

References

[1] G. Bathla et al., “Autonomous vehicles and intelligent automation: Applications, challenges, and opportunities,” Mob. Inf. Syst., vol. 2022, no. 1, p. 7632892, 2022.

[2] V. Mullet, P. Sondi, and E. Ramat, “A review of cybersecurity guidelines for manufacturing factories in industry 4.0,” IEEE Access, vol. 9, pp. 23235–23263, 2021.

[3] J. Cieply, J.-P. Disselkamp, F. Dyck, W. Alturk, I. A. Kühn, and I. R. Dumitrescu, “A systematic literature review on the Digital Twin of the factory in the context of the product and factory lifecycle,” Procedia CIRP, vol. 119, pp. 834–840, 2023.

[4] M. V. Rajesh and K. Elumalai, “The transformative power of artificial intelligence in pharmaceutical manufacturing: Enhancing efficiency, product quality, and safety,” J. Holist. Integr. Pharm., vol. 6, no. 2, pp. 125–135, 2025.

[5] B. Nikparvar and J.-C. Thill, “Machine learning of spatial data,” ISPRS Int. J. Geo-Inf., vol. 10, no. 9, p. 600, 2021.

[6] A. Sulaiman et al., “Artificial Intelligence-Based Secured Power Grid Protocol for Smart City,” Sensors, vol. 23, no. 19, p. 8016, 2023.

[7] T. Coito, B. Firme, M. S. Martins, S. M. Vieira, J. Figueiredo, and J. M. Sousa, “Intelligent sensors for real-Time decision-making,” Automation, vol. 2, no. 2, pp. 62–82, 2021.

[8] X. Zhou et al., “Intelligent Small Object Detection for Digital Twin in Smart Manufacturing With Industrial Cyber-Physical Systems,” IEEE Trans. Ind. Inform., vol. 18, no. 2, pp. 1377–1386, Feb. 2022, doi: 10.1109/TII.2021.3061419.

[9] S. Attaran, M. Attaran, and B. G. Celik, “Digital Twins and Industrial Internet of Things: Uncovering operational intelligence in industry 4.0,” Decis. Anal. J., vol. 10, p. 100398, 2024.

[10] D.-W. Yun and W.-C. Lee, “Intelligent dynamic real-time spectrum resource management for industrial IoT in edge computing,” Sensors, vol. 21, no. 23, p. 7902, 2021.

[11] W. Villegas-Ch, J. García-Ortiz, and S. Sánchez-Viteri, “Toward intelligent monitoring in IoT: AI applications for real-time analysis and prediction,” IEEE Access, vol. 12, pp. 40368–40386, 2024.

[12] S. Pandey, M. Chaudhary, and Z. Tóth, “An investigation on real-time insights: enhancing process control with IoT-enabled sensor networks,” Discov. Internet Things, vol. 5, no. 1, p. 29, 2025.

[13] J. V. N. Ramesh, "Real-time Air Quality Monitoring in Smart Cities using IoT-enabled Advanced Optical Sensors," measurement, vol. 15, no. 4, 2024.

[14] F. Poszler and B. Lange, "The impact of intelligent decision-support systems on humans' ethical decision-making: A systematic literature review and an integrated framework," Technol. Forecast. Soc. Change, vol. 204, p. 123403, 2024.

[15] Y. Lu and F. Ju, "Smart Manufacturing Systems based on Cyber-physical Manufacturing Services (CPMS)," IFAC-Pap., vol. 50, pp. 15883–15889, Jul. 2017, doi: 10.1016/j.ifacol.2017.08.2349.

[16] W. Lu, J. Li, Y. Li, A. Sun, and J. Wang, "A CNN-LSTM-Based Model to Forecast Stock Prices," Complexity, vol. 2020, pp. 1–10, Nov. 2020, doi: 10.1155/2020/6622927.

[17] R. Talat et al., "A decentralized system for green energy distribution in a smart grid," J. Energy Eng., vol. 146, no. 1, p. 04019036, 2020.

[18] "Environmental Sensor Telemetry Data." Accessed: May 08, 2024. [Online]. Available: https://www.kaggle.com/datasets/garystafford/environmental-sensor-data-132k

[19] W. Li, Y. Liang, and S. Wang, Data Driven Smart Manufacturing Technologies and Applications. 2021. doi: 10.1007/978-3-030-66849-5.

[20] "Lee et al. - 2023 - Enhanced Anomaly Detection in Manufacturing Proces.pdf." Accessed: May 10, 2024. [Online]. Available: https://ieeexplore.ieee.org/ielx7/6287639/10005208/10230081.pdf?tp=&arnumber=10230081&isnumber=10005208&ref=aHR0cHM6Ly9zY2hvbGFyLmdvb2dsZS5jby5pbi8=

[21] "Applied Sciences | Free Full-Text | Maintenance 5.0: Towards a Worker-in-the-Loop Framework for Resilient Smart Manufacturing." Accessed: May 10, 2024. [Online]. Available: https://www.mdpi.com/2076-3417/12/22/11330

[22] A. Wahid, J. G. Breslin, and M. A. Intizar, "Prediction of machine failure in industry 4.0: a hybrid CNN-LSTM framework," Appl. Sci., vol. 12, no. 9, p. 4221, 2022.

[23] K. O. Adefemi, M. B. Mutanga, and O. A. Alimi, "A Hybrid CNN–GRU Deep Learning Model for IoT Network Intrusion Detection," J. Sens. Actuator Netw., vol. 14, no. 5, p. 96, 2025.

[24] Kanchana1990, "World Air Quality Data 2024 (Updated)." Accessed: Oct. 24, 2025. [Online]. Available: https://www.kaggle.com/datasets/kanchana1990/world-air-quality-data-2024-updated

[25] DatasetEngineer, "IndFD-PM-DT." Accessed: Oct. 24, 2025. [Online]. Available: https://www.kaggle.com/datasets/datasetengineer/indfd-pm-dt