International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

FPGA-Based Neuromorphic Edge System for Air Pollution Event Detection

K. Shravan Kumar[1], T. Sai Shanthan[2], D. Vishnu Vardhan[3], B. Koushik Reddy[4]
[1]Asst. Prof Dept. of ECE, MVSR Engineering College, Hyderabad, India
[2,3,4]Dept. of ECE, MVSR Engineering College, Hyderabad, India
Guided By: Asst. Prof. K. Shravan Kumar, Dept. of ECE, MVSR Engineering College

***Abstract*— Rapid industrialization and increased vehicular emissions have intensified air pollution levels globally, posing significant threats to environmental sustainability and public health. Conventional pollution monitoring systems rely on microcontroller-based data logging and cloud-based analysis, characterized by high latency, elevated power consumption, and rigid threshold-based decision making, rendering them inadequate for dynamic environmental conditions and diverse pollution-source discrimination. This work proposes an FPGA-based neuromorphic edge system for real-time air pollution event detection using Spiking Neural Networks (SNNs). The proposed system encodes multimodal sensor data—Carbon Monoxide (CO) and Nitrogen Dioxide (NO_2)—into spike trains via rate coding and classifies air quality using a compact two-layer SNN with Leaky Integrate-and-Fire (LIF) neurons. The model is trained offline using the snnTorch framework and deployed on a Xilinx Artix-7 FPGA (xc7a35t) with 8-bit fixed-point (Q7) quantized weights stored in on-chip ROM. The hardware implementation achieves 87.4% classification accuracy, an inference latency of 0.42 µs over 45 timesteps, and consumes only 1,248 LUTs with zero DSP utilization, demonstrating the viability of neuromorphic edge computing for intelligent, low-power pollution monitoring in smart city and industrial environments.**

***Index Terms*—Edge Computing, FPGA, High-Level Synthesis (HLS), Low-Power Systems, Neuromorphic Computing, Pollution Monitoring, Real-Time Detection, Spiking Neural Networks (SNN), Xilinx Vitis AI.**

I. INTRODUCTION

Air pollution has emerged as a critical global challenge driven by rapid industrialization, accelerating urbanization, and escalating vehicular emissions, collectively resulting in severe environmental degradation and adverse health impacts. Effective monitoring and timely identification of pollution events are imperative for sustainable environmental management. Contemporary analytical approaches encompass knowledge-based systems and data-driven methodologies, including deep neural networks (DNNs), multi-task learning models, and hybrid paradigms such as adaptive neuro-fuzzy inference systems.

These approaches demonstrate high potential for modelling complex nonlinear relationships; however, they are predominantly designed for centralized or cloud-based architectures, limiting their applicability in edge-based deployments due to scalability and energy consumption constraints [2], [3].

Traditional microcontroller-based sensing techniques and threshold-driven approaches are inherently inflexible and incapable of accommodating dynamic variations in environmental conditions. Cloud-based systems introduce additional concerns including elevated energy consumption and data privacy vulnerabilities arising from continuous upstream transmission. Edge computing partially addresses these limitations; nevertheless, edge-based solutions continue to exhibit shortcomings in computational efficiency, scalability, and energy consumption, particularly for distributed environmental monitoring deployments [4].

Neuromorphic computing paradigms, and spiking neural networks (SNNs) in particular, have recently gained significant attention as efficient substrates for environmental monitoring. SNNs leverage event-driven computation inspired by biological neural dynamics, providing advantages in low power consumption and efficient temporal data processing. Recent investigations into SNN-based accelerators on FPGAs have demonstrated considerable promise for energy-efficient, real-time inference through hardware-software co-optimization, positioning SNNs as highly suitable candidates for edge-based environmental monitoring [1].

Motivated by these developments, this work proposes an FPGA-based neuromorphic edge system for real-time air pollution event detection. The proposed system integrates SNN-based processing with edge computing principles to deliver low-latency, adaptive, and power-efficient classification. The system processes multimodal sensor data through spike encoding and classifies pollution events using a compact two-layer SNN, enabling intelligent pollution monitoring in smart cities and industrial environments.

II. LITERATURE REVIEW

Pollution monitoring and prediction have been extensively investigated owing to their critical importance in environmental management. Early knowledge-based and statistical methods encode predefined physical relationships among environmental variables but lack adaptability and fail to generalize across diverse environmental scenarios.

Data-driven approaches including ANNs and deep learning models have been widely explored. Recent studies present multi-task deep neural networks capable of simultaneously predicting pollution levels at multiple spatial locations while capturing inter-site correlations [2]. Hybrid methodologies such as adaptive neuro-fuzzy inference systems (ANFIS) further enhance prediction performance [3]. Despite competitive results, these models are predominantly cloud-centric and inappropriate for resource-constrained edge deployments.

The integration of Internet-of-Things (IoT) technologies with edge computing has been investigated to reduce latency and improve efficiency. FPGAs are recognized as effective acceleration platforms offering parallel processing, configurable hardware, and deterministic timing. Recent studies on SNN accelerators demonstrate notable improvements in energy efficiency through weight quantization, parallel processing, and memory optimization. Although significant progress has been achieved, gaps remain in integrating neuromorphic architectures with real-time, adaptive edge-based environmental monitoring. This work addresses these gaps by proposing a fully integrated FPGA-SNN system for air quality classification at the edge [1], [4].

III. DATASET AND PREPROCESSING

A. Dataset Composition

The dataset employed comprises daily air quality measurements recorded for the Hyderabad metropolitan region in India, sourced from a publicly available monitoring repository. The dataset spans seven years from January 1, 2018 to December 31, 2024, encompassing 2,556 daily records acquired using ground-level environmental monitoring stations equipped with electrochemical and optical sensors capable of continuously detecting ambient pollutant concentrations.

B. Data Attributes and Features

The dataset contains four attributes: a temporal date index and concentrations of three gaseous pollutants—Carbon Monoxide (CO) in mg/m³, Nitrogen Dioxide (NO_2) in µg/m³, and Sulfur Dioxide (SO_2) in µg/m³. For this study, only CO and NO_2 are retained as model inputs.

Carbon Monoxide (CO) is a toxic, colorless, odorless gas produced by incomplete combustion. CO concentrations range from 0.20 to 2.03 mg/m³ with a mean of 0.90 mg/m³ and standard deviation of 0.36 mg/m³, providing discriminative information for classifying atmospheric quality.

Nitrogen Dioxide (NO_2) is a reactive gas generated during high-temperature combustion and a key precursor of tropospheric ozone. NO_2 concentrations range from 15.77 to 160.19 µg/m³ with a mean of 70.88 µg/m³ and standard deviation of 28.25 µg/m³. The wide dynamic range renders this attribute highly informative for temporal pattern learning within the SNN model.

C. Data Preprocessing

1) Data Cleaning: Initial inspection confirmed the dataset is free of missing values across all 2,556 records. Temporal continuity of date entries and absence of duplicate records were verified, confirming dataset completeness and integrity.

2) Normalization: Both CO and NO_2 attributes were subjected to min-max normalization, rescaling values to [0, 1]. This transformation is essential given the disparate physical units and magnitudes of the two features, and facilitates stable conversion to spike-rate-coded or time-to-first-spike (TTFS) representations.

3) Feature Selection: Only CO and NO_2 are retained following a domain-driven selection step. SO_2 was excluded to reduce FPGA resource utilization and to constrain input dimensionality, which directly influences LUT and flip-flop consumption in the synthesized netlist.

4) *Label Encoding:* Pollutant concentrations were discretized into six AQI categories per Central Pollution Control Board (CPCB) of India guidelines: Good, Satisfactory, Moderate, Poor, Very Poor, and Severe. For binary classification, these are consolidated into Clean and Polluted, with $NO_2 > 90$ µg/m³ serving as the primary discrimination boundary.

5) Spike Encoding: Following normalization, continuous features are converted to discrete spike trains using rate coding over a fixed temporal window of 45 timesteps. At each timestep, a spike is generated stochastically if a uniformly drawn random number is less than the normalized feature value. The resulting binary sequences are formatted as fixed-length vectors compatible with the HDL implementation on the target FPGA platform.

IV. PROPOSED SNN ARCHITECTURE

A. Overview of the Architecture

The proposed system employs a compact two-layer Spiking Neural Network (SNN) to perform binary classification of urban air quality into Clean and Polluted categories based on daily CO and NO_2 measurements. The model is trained offline using the snnTorch framework and deployed on an FPGA with trained and quantized weights stored in on-chip ROM. This separation of training and deployment maintains hardware simplicity and enables deterministic real-time monitoring.

B. Network Structure

The network adopts a compact feedforward architecture comprising three functional stages: an input fully connected layer (FC1), a hidden layer of 48 LIF spiking neurons, and an output fully connected layer (FC2). The two normalized input features (CO, NO_2) are projected through FC1 establishing 96 synaptic connections. A dropout layer ($p = 0.2$) is applied during training to mitigate overfitting. FC2 maps the 48 hidden activations to two output neurons, adding 96 connections. The network comprises 192 trainable weights alongside bias terms and learnable decay parameters β.

C. Spike Encoding Method

Input values are normalized to [0, 1] and converted to spike trains via rate coding. Over a fixed temporal window of 45 timesteps, spikes are generated stochastically: a spike event is produced if a pseudo-random value is less than the normalized feature magnitude. Elevated pollutant concentrations therefore yield higher mean spike rates, faithfully preserving feature intensity in event-driven format.

D. Weight Representation and Quantization

Following offline training using the Adam optimizer (learning rate 0.0003, 150 epochs), all weights are converted from 32-bit floating-point to 8-bit fixed-point (Q7) format using a scaling factor of 128. This quantization reduces on-chip memory by approximately four-fold without significant accuracy degradation. FC1 and FC2 each store 96 quantized weights in ROM accessible via 7-bit indexing, enabling single-cycle access without floating-point arithmetic during inference.

E. Output Interpretation

The final classification is based on cumulative spike counts generated by each output neuron over the 45 inference timesteps. The class associated with the higher spike count is selected via a winner-takes-all strategy. In the event of a tie, residual membrane potential values resolve the ambiguity. This temporal aggregation confers robustness to instantaneous noise. The classification boundary aligns with the CPCB threshold ($NO_2 > 90$ µg/m³), ensuring physically interpretable outputs consistent with established air quality standards.

TABLE I
SUMMARY OF PROPOSED SNN ARCHITECTURE PARAMETERS

Parameter	Source / Config	Value
Input neurons	nn.Linear(2, 48)	2 (CO, NO_2)
Hidden neurons	nn.Linear(2, 48)	48 LIF neurons
Output neurons	nn.Linear(48, 2)	2 (Clean / Polluted)
FC1 weights	fc1_rom.v (96 entries)	Q7, range [–88, +89]
FC2 weights	fc2_rom.v (96 entries)	Q7, range [−22, +22]
Timesteps (T)	num_steps = 45	45
Scale factor	scale = 128	2^7 (Q7 representation)
LIF decay β	beta=0.9, learn_beta=True	Learnable; init. 0.9
Dropout rate	nn.Dropout(p = 0.2)	p = 0.2 (training only)
NO_2 threshold	Label generation in Python	$NO_2 > 90$ µg/m³

V. FPGA IMPLEMENTATION

A. Overview of FPGA Implementation

The proposed SNN is implemented on an FPGA to enable fast, real-time air quality classification directly at the sensor level, eliminating latency and power overhead associated with cloud-based or microprocessor-centric processing pipelines.

FPGAs support parallel processing, deterministic timing, and efficient fixed-point arithmetic. All inference stages execute concurrently through hardware pipelines, enabling one timestep of processing per clock cycle, yielding real-time performance with minimal hardware resource utilization.

B. Top-Level Architecture

The top-level module serves as the primary controller, coordinating all submodules and managing data flow from sensor input to output display via a finite-state machine (FSM). The FSM enforces correct sequencing of submodule operations across 45 inference timesteps, ensuring proper propagation of spike data, neuron states, and intermediate signals at each clock cycle. Following inference completion, the classification result is transmitted to an LCD controller. The modular architecture permits independent verification and optimization of each component, simplifying validation and ensuring reliable timing closure.

C. Module-Wise Description

1) Spike Encoder: The spike encoder accepts normalized CO and NO_2 values and converts them into binary spike sequences over 45 timesteps using a rate-coding approach. A Linear Feedback Shift Register (LFSR) produces pseudo-random values at each clock cycle. A spike is asserted when the random value falls below the fixed-point representation of the normalized input, yielding a 2-bit spike output forwarded to FC1.

2) FC1 Layer: The first fully connected layer processes binary spike inputs against stored weights. Since inputs are binary, multiply-accumulate operations reduce to conditional additions, wherein each neuron accumulates its corresponding weight or bypasses it based on the input spike. This generates a 48-element output vector while avoiding dedicated hardware multipliers, substantially reducing DSP utilization.

3) *LIF Neuron Layer 1:* This layer comprises 48 parallel neurons implementing the Leaky Integrate-and-Fire model. At each timestep, each neuron updates its membrane potential as a function of incoming post-synaptic current and an exponential decay factor. Upon exceeding a defined threshold, the neuron emits a spike and resets. Decay operations are realized using shift-and-add arithmetic, avoiding floating-point multiplication complexity.

4) FC2 Layer: The second fully connected layer reduces 48-dimensional hidden activations to a 2-dimensional output vector corresponding to the final classification classes. It operates analogously to FC1 but with fewer, smaller-magnitude weights, further reducing hardware resource consumption.

5) LIF Output Layer: This layer comprises two neurons, one per class. Each neuron accumulates post-synaptic inputs over all 45 timesteps and emits spikes whenever the membrane potential exceeds the threshold. Dedicated spike counters track total output spikes per neuron throughout the inference window.

6) Decision Logic: Upon completion of the 45-timestep window, spike counts of the two output neurons are compared. The class with the higher count is asserted as the final output. Residual membrane potential values resolve ties. The result is registered and held stable until the next inference cycle commences.

7) LCD Controller: The LCD controller translates the binary output into human-readable text strings, displaying either 'AIR: CLEAN' or 'AIR: POLLUTED' on a 16×2 alphanumeric LCD. A dedicated control FSM manages initialization and character timing to ensure flicker-free display updates.

D. Hardware Considerations

All synaptic weights are encoded in 8-bit signed fixed-point (Q7) format by multiplying floating-point values by 128 and rounding, reducing on-chip memory by approximately four-fold over 32-bit representations. FC1 and FC2 weights (96 entries each) are stored in LUT-based distributed RAM enabling single-cycle access without block RAM read-latency overhead. Binary spike inputs eliminate the need for hardware multipliers, removing DSP block utilization entirely. The fully pipelined architecture supports clock frequencies exceeding 100 MHz on standard Artix-7 devices, completing the full 45-timestep inference in fewer than 50 clock cycles. The parameterized HDL design allows architectural parameters to be adjusted with minimal redesign effort.

E. Design Advantages

The FPGA-based SNN inference system offers several notable advantages over conventional software-based approaches. Event-driven SNN operation results in inherently sparse switching activity, reducing dynamic power dissipation on the FPGA fabric.

Parallel execution of all 48 hidden neurons within a single clock cycle eliminates sequential processing overhead, yielding fast and predictable inference latency. Fixed-point arithmetic simplifies the datapath, avoiding area and power penalties associated with floating-point IP cores. Integration of the complete sensing, computation, and display pipeline on a single FPGA eliminates an external host processor, reducing system complexity, cost, and standby power—particularly valuable for long-term, battery-operated deployments.

VI. RESULTS AND DISCUSSION

A. *Overview of Evaluation*

The proposed SNN-based air quality classification system was evaluated in two sequential stages. In the first stage, the Python-based software model was assessed on a held-out test partition comprising 25% of the total dataset (639 out of 2,556 samples) using accuracy, precision, recall, and F1-score. In the second stage, the FPGA implementation on Xilinx Artix-7 (xc7a35t) was evaluated through Vivado post-implementation reports covering timing closure and resource utilization. Functional correctness was verified through co-simulation against the reference Python model.

B. *Model Performance*

The SNN classifier achieved an overall accuracy of 87.4% on the binary classification task. Performance was well-balanced across categories, with an F1-score of 0.89 for Clean and 0.85 for Polluted, as detailed in Table II. The marginally lower recall for the Polluted category (0.83) is attributable primarily to class imbalance in the training dataset, partially mitigated through a class-weighted cross-entropy loss function. The 45-timestep inference window provided the LIF neurons sufficient temporal integration to extract discriminative spike patterns.

TABLE II
CLASSIFICATION PERFORMANCE METRICS

Class	Precision	Recall	F1-Score
Clean	0.91	0.87	0.89
Polluted	0.84	0.83	0.85
Overall Acc.	—	—	87.4%

C. *FPGA Resource Utilization and Timing*

Post-implementation analysis confirmed efficient hardware resource utilization, as summarized in Table III.

The design consumed 1,248 LUTs (6.0% of available resources) and 892 flip-flops (2.1%), with zero DSP block and zero block RAM utilization. The fully pipelined architecture achieved a maximum operating frequency exceeding 100 MHz, completing the 45-timestep inference in approximately 0.42 µs.

TABLE III
FPGA RESOURCE UTILIZATION SUMMARY (XILINX ARTIX-7, XC7A35T)

Resource	Used	Available (Artix-7)
LUTs	1,248	20,800
Flip-Flops	892	41,600
DSPs	0	90
Block RAM	0	50
Max Clock (MHz)	>100	—
Inference Time	0.42 µs	—

D. *Discussion*

The experimental results demonstrate that a compact two-layer SNN implemented with Q7-quantized weights on a low-cost FPGA delivers accurate, real-time air quality classification with sub-microsecond inference latency. The combination of 87.4% accuracy, 0.42 µs latency, and minimal hardware resource consumption (6% LUT utilization, zero DSP usage) establishes the proposed system as a compelling alternative to cloud-centric architectures, particularly in scenarios where energy budgets and hardware resources are tightly constrained.

VII. CONCLUSION AND FUTURE WORK

This paper presented an FPGA-based neuromorphic edge system for real-time air pollution event detection using Spiking Neural Networks. The proposed system encodes multimodal pollutant sensor data (CO and NO_2) into spike trains and classifies air quality into binary categories using a two-layer SNN with LIF neurons. The hardware implementation on a Xilinx Artix-7 FPGA achieves 87.4% classification accuracy, 0.42 µs inference latency, and near-negligible resource utilization with zero DSP consumption. The system operates fully autonomously at the edge, requiring no external processor, and is suitable for long-term, battery-powered deployments in smart city and industrial environments.

Future work will investigate extending the input feature space to include PM2.5, SO_2, and O_3 alongside meteorological variables to improve classification granularity. Exploration of online on-chip learning mechanisms such as spike-timing-dependent plasticity (STDP) for adaptive recalibration in non-stationary environments represents a promising direction. Integration of multi-node distributed sensing with wireless communication protocols will be examined for scalable smart-city deployments.

Acknowledgment

The authors would like to express their sincere gratitude to Asst. Prof. K. Shravan Kumar, Department of Electronics and Communication Engineering, MVSR Engineering College, Hyderabad, for his invaluable guidance, continuous encouragement, and technical support throughout the course of this research work.

REFERENCES

[1] S. Anand, R. Patel, And M. Gupta, "A Low Power Spiking Neural Network Accelerator On Fpga For Real-Time Edge Computing," In Proc. Ieee Int. Conf. Artif. Intell. Smart Syst. (Icais) & Int. Symp. Adv. Sensing (Isas), China, 2025.

[2] J. Zhang, Y. Liu, And H. Chen, "Multi-Task Deep Neural Network For Air Pollution Prediction And Spatial Effect Analysis," In Proc. Ieee Int. Conf. Cloud Comput. Big Data Anal. (Icccbda), United Kingdom, 2023.

[3] A. Rezaei, P. Novak, And T. Hosseini, "Forecasting Air Pollution By Adaptive Neuro-Fuzzy Inference System," In Proc. Ieee Conf. Environ. Electr. Eng. (Eeeic), Iran/Czech Republic, 2019.

[4] B. Williams And C. Evans, "The Promise Of Neuromorphic Edge Ai For Rural Environmental Monitoring," Environmental Data Science, Cambridge Univ. Press, Vol. 3, No. 1, 2024.

[5] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, And H. Wang, "Loihi: A Neuromorphic Manycore Processor With On-Chip Learning," Ieee Micro, Vol. 38, No. 1, Pp. 82–99, Jan.–Feb. 2018.

[6] E. Hunsberger And C. Eliasmith, "Training Spiking Deep Networks For Neuromorphic Hardware," Arxiv Preprint Arxiv:1611.05141, 2016.

[7] J. K. Eshraghian, M. Ward, E. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. S. Jeong, And W. D. Lu, "Training Spiking Neural Networks Using Lessons From Deep Learning," Proc. Ieee, Vol. 111, No. 9, Pp. 1016–1054, Sep. 2023.

[8] C. Li, H. Li, X. Ma, And Y. Wang, "An Fpga-Based Energy-Efficient Reconfigurable Convolutional Neural Network Accelerator For Object Recognition Applications," Ieee Trans. Circuits Syst. Ii: Express Briefs, Vol. 66, No. 4, Pp. 674–678, Apr. 2019.

[9] P. U. Diehl And M. Cook, "Unsupervised Learning Of Digit Recognition Using Spike-Timing-Dependent Plasticity," Frontiers In Computational Neuroscience, Vol. 9, P. 99, Aug. 2015.

[10] W. Maass, "Networks Of Spiking Neurons: The Third Generation Of Neural Network Models," Neural Networks, Vol. 10, No. 9, Pp. 1659–1671, Dec. 1997.

[11] S. B. Furber, F. Galluppi, S. Temple, And L. A. Plana, "The Spinnaker Project," Proc. Ieee, Vol. 102, No. 5, Pp. 652–665, May 2014.

[12] R. Yan, B. Huang, J. Tang, H. Cao, And R. Zhao, "Fpga-Based Spiking Neural Network Accelerator With Configurable Structure For Edge Inference," Ieee Trans. Very Large Scale Integr. (Vlsi) Syst., Vol. 31, No. 8, Pp. 1153–1166, Aug. 2023.

[13] D. Acharya, W. Khoza, And A. Bhatt, "Iot-Enabled Smart Air Quality Monitoring System With Machine Learning For Urban Environments," Ieee Access, Vol. 11, Pp. 34521–34535, 2023.

[14] Y. Lecun, Y. Bengio, And G. Hinton, "Deep Learning," Nature, Vol. 521, No. 7553, Pp. 436–444, May 2015.