School and Conference on Edge Intelligence for Scientific Instrumentation and Reconfigurable Computing | (SMR 4200)

14 Sep 2026 - 24 Sep 2026
Outside, Cyberjaya, Selangor, Malaysia

T01 - ALI Endris Mohammed

Efficient Delay-Sensitive Task Offloading to Fog Computing with Multi-Agent Twin Delayed Deep Deterministic Policy Gradient

T02 - BHATTACHARYA Jhilik

Edge Intelligence: Feature Distillation for Object Detection Refinement

T03 - BUDHATHOKI Ram Kaji

A Deep Learning Approach for Identifying Diseases in Lettuce Leaves via CNN

T04 - CHEEKATI Chakradhar

Edge Intelligence for Real-Time Flood Detection: Accelerating SAR Deep Learning with FPGAs

T05 - CHOWDHURY Md Nakib Hayat

Reliability Challenges in Clinical Machine Learning: Data Leakage and Generalization Under Deployment Constraints

T06 - GHONEIM Youmna Sami Mahmoud Khalil

Slow Control GUI based on MasterScada 4D for a Terasic Cyclone V GX Starter Kit. Implementation of image-based self-tests for Cyclone V using designed GUI

T07 - IRWANDI

A Configurable Zynq-7000 FPGA/SoC Platform for Scientific Instrumentation, STEM Education, and Research with FPGA-Based Graphics Acceleration

T08 - JEPTOO Carolyne

Reconfigurable Intelligent Surfaces for Integrated Sensing and Communication in Dense 6G Networks

T09 - KADIR Muhammad Abdul

Development of a Portable Multimodal Sensory Assessment System for Early Detection of Diabetic Peripheral Neuropathy

T10 - KRITO Jauza Akbar

Selection-Aware Machine Learning for Unresolved Stellar Binaries

T11 - MERCADO Reginald Juan

Structural Anomaly Forecasting with Edge-Intelligence and Early Warning System (SAFE-IEWS)

T12 - MOHD SALEH Hariyanti Binti

Adaptive Convolutional Neural Network Training for Semiconductor Image Classification: Toward FPGA Deployment

T13 - NG Charn Loong

BioEdge: Ultra-Low-Power Edge Instrument for Next-Generation Wearable Capacitive Electrophysiological Biomedical Sensor

T14 - RADHY Ahmad

TinyM-Based Electronic Nose for Food Quality Monitoring Using Embedded Rust on ESP32-S3 Edge Devices

T15 - SHARMA Aditya Kumar

SPARSE-BAM: A SPARSITY-AWARE BROKEN ARRAY MULTIPLIER WITH UNIFIED BOUNDARY CORRECTION FOR ENERGY-EFFICIENT NEURAL NETWORK INFERENCE

T16 - SIDDIQUI Amna Sardar

Edge for the Intelligent Grid: A review of Edge Intelligence in Grid Technologies-Current Applications and Emerging Frontiers.

T17 - SOORIARACHCHI Nethulya Aishwarinee

Lightweight LSTM Inference via Knowledge Distillation, Pruning, and Quantization for Real-Time ECG Anomaly Detection at the Edge

T18 - SUHAILI Shamsiah Binti

FPGA-Based Low-Power Cryptographic Hash Architecture for Secure Edge Intelligence Applications

T01

Efficient Delay-Sensitive Task Offloading to Fog Computing with Multi-Agent Twin Delayed Deep Deterministic Policy Gradient

Endris M. Ali[1], **Jemal H. Abawajy**[2], **and Frezwd. Lemma**[2]

[1]*Department CSE, College of Electrical Engineering and Computing, ASTU, Ethiopia;*
[2]*Faculty of Science, Engineering and Built Environment, Deakin University, Australia*

Currently, the network has experienced rapid growth in resource-intensive smart user devices and consumer electronics (CE) designed to run various applications. However, these user devices have limited resources, making it difficult to run real-time applications independently and offload their tasks to the nearest computing nodes [1]. These computing nodes become overloaded and congested in the context of events such as football games, concerts, and contests with high computation and traffic demands [2]. Recently, the adoption of fog-to-fog computing architecture has significantly improved task offloading to deal with these problems while achieving task QoS requirements. In this work, an offloaded horizontal F2F architecture is considered, which needs a coordinated decision process in which tasks are delivered from the user device at the hotspot to optimally allocate sub-tasks to adjacent fog nodes to achieve the latency requirements of user applications. Moreover, due to the dynamic nature of fog systems, updating the offloading decision parameters between fog nodes, such as resource status, current workload, and channel condition, becomes a challenge for real-time applications [3]. To deal with the problem, Multi-agent fully cooperative task offloading and resource allocation (MAFCTORA) algorithms are designed in an F2F cooperative fashion. The main problem in this MADRL-based task offloading is agents' continuous learning space to converge for optimal solutions in the dynamic topology of the Fog network. Moreover, due to DRL agents' heterogeneous and autonomous nature, a decentralized (DecMADRL) partial task offloading approach is designed to minimize task computation delay and energy consumption. In this DecMADRL-based approach, the objective is to optimize overall system latency and energy consumption. We simulate and evaluate the performance of our DecMADRL algorithm with MATD3, MASAC, MAIDDPG, and MAPPO. MATD3 achieved the best overall performance among the evaluated models, yielding the highest average reward of **0.36±0.01**, the lowest average latency of **0.08±0.01**, and the lowest energy consumption of **0.76±0.14**.

Fog Computing; DRL; IoT; Multi-Agent; Task Offloading; TD3

References

[1] Diego Hortelano, Ignacio de Miguel, Ramón J.Durán Barroso, Juan Carlos Aguado, Noemí Merayo, Lidia Ruiz, Adrian Asensio, Xavi Masip-Bruin, Patricia Fernández, Rubén M. Lorenzo, and Evaristo J. Abril. A comprehensive survey on reinforcement-learning-based computation offloading techniques in Edge Computing Systems. *J. Netw. Comput. Appl.*, 216(October 2022):103669, 2023.

[2] Tesfay Zemuy Gebrekidan, Sebastian Stein, and Timothy J. Norman. Combinatorial Client-Master Multiagent Deep Reinforcement Learning for Task Offloading in Mobile Edge Computing. *Proc. Int. Jt. Conf. Auton. Agents Multiagent Syst. AAMAS*, 2024-May:2273–2275, feb 2024.

[3] Hao Hao, Changqiao Xu, Wei Zhang, Shujie Yang, and Gabriel-Miro Muntean. Joint task offloading, resource allocation, and trajectory design for multi-uav cooperative edge computing with task priority. *IEEE Transactions on Mobile Computing*, 23(9):8649–8663, 2024.

T02

Edge Intelligence: Feature Distillation for Object Detection Refinement

Jhilik Bhattacharya [1] , Romina Molina [2] , and Maria Liz Crespo [2]

[1] *Thapar Institute of Engineering and Technology, Patiala, India*
[2] *International Centre for Theoretical Physics, Trieste, Italy*

Edge intelligence plays a key role in autonomous driving systems, which require reliable scene perception and interpretation across heterogeneous environmental conditions. Vision-based systems for surveillance, assisted driving, or fully autonomous vehicles must rely on robust object detection algorithms that function satisfactorily in all weather and lighting conditions. Because the ultimate implementation focuses on an FPGA-based platform, achieving sufficient model compression is a primary requirement while preserving the overall performance.

The criticality of this problem is compounded by the fact that real-time object detection systems for low-power devices are typically implemented using smaller variants of original base networks. These smaller models already exhibit lower detection precision compared to their larger counterparts. For example, the performance of YOLO26-S (9.5M) [1] on the COCO dataset drops by 9% compared to its base model, YOLO26-X (55.7M).

In this research, we highlight that object detection models often struggle with degraded images where object boundaries and textures are blurred due to haze or low light [2]. Through several guided experiments, we infer that the scale-dependent fusion techniques [1,3] used in object detection networks suffer due to poor contrast and obscured features. Conversely, permuting scale inputs at various levels helps increase robustness, even with out-of-distribution data; making a scale-permutable approach an ideal choice for such scenarios. These networks are not only more robust across different environmental conditions, but they are also better suited for knowledge distillation.

Since model pruning and compression are vital steps for deploying tasks on FPGA units, utilizing an appropriate distillation framework is crucial to preserving accuracy. We propose and evaluate a feature distillation technique for object detection under poor visibility conditions and analyse its impact when deployed on an FPGA platform. We formulate a joint training strategy within a scale-permuted student network that learns dehazed features from a dual-teacher network without requiring an explicit dehazing pre-processing step. Through attention transfer, the student network learns to replicate both the teacher's outputs and its underlying decision-making process. This encourages the student to grasp the teacher's reasoning, resulting in a structured distillation framework where the student infers what guides the teacher's outputs rather than blindly matching them. Finally, we emphasize the importance of incorporating continual learning during distillation to prevent catastrophic forgetting. We demonstrate that by using these training techniques, we can successfully distill information from a teacher three times the size of the student, achieving significant performance gains across various scenes without experiencing catastrophic forgetting. Finally, these training techniques though verified for vision-based detection tasks can be applied across any range of edge intelligence problems.

[1] Jocher, G.; Qiu, J. Ultralytics YOLO26 (2026).

[2] Kumar, A.; Chadha, A. From Fog to Failure: The Unintended Consequences of Dehazing on Object Detection in Clear Images. In Proceedings of the ICLR 2025 Workshops: I Can't Believe It's Not Better (ICBINB) (2025).

[3] 7. Tan, M.; Pang, R.; Le, Q.V. EfficientDet: Scalable and Efficient Object Detection, [arXiv:cs.CV/1911.09070].(2020).

T03

A Deep Learning Approach for Identifying Diseases in Lettuce Leaves via CNN

Ram Kaji Budhathoki [1] , and Santosh Shaha [2]

[1,2] *Kathmandu University, Dhulikhel, Kavre, Nepal*

This paper introduced an implementation of deep convolutional neural network for foliar disease detection in lettuce plants using image analysis. By combining images of farms and grocery stores, a comprehensive and diverse datasets were prepared. To address the limitations of the dataset and to increase diversity, data augmentation technique, particularly basic image manipulation (BIM), was employed, resulting in an expanded datasets of 5,700 images in the three classes: healthy leaves and two diseased categories namely: fungal and bacterial. The proposed model was trained for 100 epochs on a personal computer, with hyper-parameters optimized through a random search approach to maximize performance. Evaluation of a test set of 570 images demonstrated the effectiveness of the model, achieving an overall classification accuracy of 96.5%, along with a precision of 96.6% and an F1 score of 96.3%, calculated using weighted averages. These results highlight our model's potential for accurate and reliable lettuce disease detection, contributing to advances in precision agriculture.

[1] J. Pandian, K. Kanchanadevi, V. Kumar, E. R. Jasinska, Z. Leonowicz, and J. Jasinski, Electronics **11**, 1266 (2022).
[2] K. A. Mahlein, U. Steiner, E.-C. Oerke, H. W. Dehne, L. Plümer, and T. Rumpf, Comput. Electron. Agric., **74**, 91 (2010).
[3] C. Wetterich, R. Neves, J. Belasque, R. Ehsani, and L. Marcassa, Appl. Opt., **56**, 15 (2017).
[4] A. Fuentes, S. Yoon, S. Kim, and D. Sun Park, Sensors, **17**, 2029 (2017).
[5] C. DeChant, T. Wiesner, S. Chen, E. Stewart, J. Yosinski, M. Gore, R. Nelson, and H. Lipson, Phytopathology, **107**, 1426 (2017).
[6] Y. Lu, S. Yi, N. Zeng, Y. Liu, and Y. Zhang, Neurocomputing, **267**, 378 (2017).
[7] M. Agarwal, S. K. Gupta, and K. K. Biswas, Sustain. Comput. Inform. Syst., **28**, 100407 (2020).

T04

Edge Intelligence for Real-Time Flood Detection: Accelerating SAR Deep Learning with FPGAs

Chakradhar Cheekati [1]

[1] *(Master's Research Student, Faculty of Electronic Engineering, Universiti Malaysia Perlis, Malaysia.)*

Near-real-time flood mapping using Synthetic Aperture Radar (SAR) imagery is critical for rapid disaster response. However, executing complex deep learning models traditionally relies on cloud infrastructure, creating severe latency and data-transmission bottlenecks during communication outages. To overcome this, we propose an edge intelligence framework that pushes deep-learning flood detection directly to the edge. By leveraging Field-Programmable Gate Arrays (FPGAs), we provide a hardware solution that achieves unmatched energy efficiency and parallel computing capabilities in power-constrained environments. FPGAs significantly outperform standard GPUs for on-device processing, enabling ultra-low-latency, autonomous inference on raw satellite and multimedia data. This edge-native approach bypasses cloud dependency, establishing FPGAs as the superior computational alternative for deploying robust, instantaneous flood analysis in the field.

T05

Reliability Challenges in Clinical Machine Learning: Data Leakage and Generalization Under Deployment Constraints

T06

Slow Control GUI based on MasterScada 4D for a Terasic Cyclone V GX Starter Kit. Implementation of image-based self-tests for Cyclone V using designed GUI

Youmna Ghoneim, Sergey Movchan
Affiliation: Joint Institute for Nuclear Research (JINR), Dubna, Russian Federation
E-mail: youmnasami24@gmail.com; Ghoneim@jinr.ru

Abstract
This work presents an FPGA-based slow-control and reliability-diagnostics framework for the MPD TPC readout subsystems built around a Terasic Cyclone V GX Starter Kit. Verilog firmware implements four LED indication modes—synchronous blink, alternating red/green, wave rotation, and pseudo-random LFSR-driven patterns—controlled by on-board switches and keys that select operating mode, blinking frequency, and global reset under a 50 MHz system clock. The design flow in Quartus Prime Lite 20.1.1 and ModelSim Intel FPGA Edition covers pin assignment for LEDs, switches, push-buttons and UART, synthesis and fitting of the scan_top logic, and functional/timing verification using scripted .do testbenches, yielding a validated slow-control core for integration into the MPD TPC readout chain. Beyond functional correctness, the framework introduces an image-based diagnostic path in which the FPGA stores synthetic or sensor-derived pixel data in on-chip memory, read out to a host PC via Python over JTAG or UART and analyzed for stuck-at faults, missing bits, and localized instabilities. A Python-based OPC UA server exports the derived health metrics to MasterSCADA 4D, enabling real-time visualization, automated test sequencing, and long-term reliability monitoring of FPGA-resident functions in a scalable, extensible architecture.

1. Introduction

The project targets the development of a unified slow-control and reliability-diagnostics framework for the MPD TPC readout subsystems. It addresses the need for validated FPGA control logic that can be continuously monitored and quantitatively assessed under realistic operating conditions.

2. Methodology

A Terasic Cyclone V GX Starter Kit hosts Verilog firmware implementing four LED indication modes—synchronous, alternating, wave, and pseudo-random LFSR patterns—controlled by on-board switches and keys and clocked at 50 MHz. The design is implemented in Quartus Prime Lite 20.1.1 with ModelSim Intel FPGA Edition for scripted functional and timing verification of all LED modes and UART communication. An image-based diagnostic path writes synthetic or sensor-derived pixel frames into on-chip memory, which are read out to a Python environment via JTAG or UART, converted to arrays, and processed to extract fault and stability metrics.

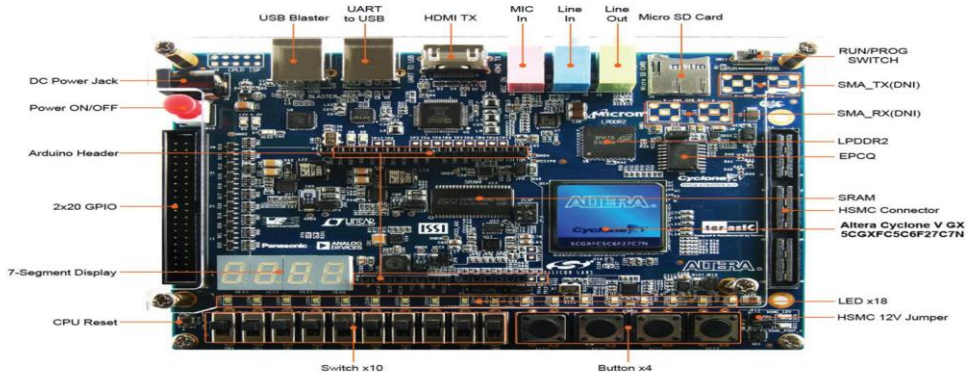

Figure 1 Cyclone V GX Starter Kit (Board Overview & Architecture)

3. Results

The verified scan_top design provides a stable slow-control core suitable for integration into the MPD TPC readout chain. Image analysis detects stuck-at patterns, missing bits, and spatially localized instabilities, yielding per-mode reliability indicators and global health metrics for the FPGA logic.

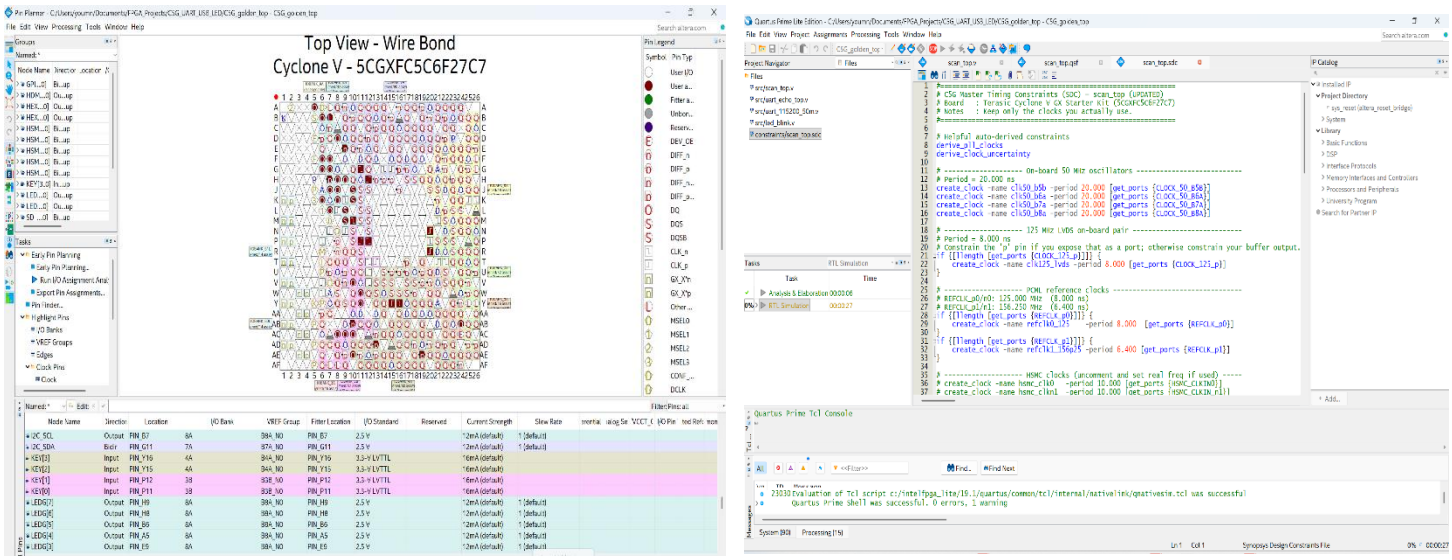

Figure 2 a) Pin Planner view of the Cyclone V device showing LED and key assignments for the FPGA slow-control prototype. b) Timing-constraint setup for on-board clocks in Quartus Prime Lite for the Cyclone V FPGA design.

4. Discussion and Conclusion

A Python-based OPC UA server publishes these metrics as standardized process tags to MasterSCADA 4D, which offers real-time visualization, trend analysis, and initiation of new test cycles. The resulting FPGA–Python–SCADA chain forms a dual-purpose platform that simultaneously delivers operational slow-control and scalable online reliability monitoring for current and future MPD TPC FPGA subsystems.

5. References

- P. R. W. F. Falk et al., “Development of slow control systems for the MPD TPC,” J. Instrum., vol. 15, no. 03, C03010, 2020.
- M. Nicolaidis, “Design for soft-error mitigation,” IEEE Trans. Device Mater. Rel., vol. 5, no. 3, pp. 405–418, 2005.

T07

A Configurable Zynq-7000 FPGA/SoC Platform for Scientific Instrumentation, STEM Education, and Research with FPGA-Based Graphics Acceleration

Irwandi[1], Yuwaldi Away[1], Ihsan Hariadi[2], Fadjar Rahino[3]

[1]STEM Research Center, Universitas Syiah Kuala
[2]School of Electrical Engineering and Informatics, Institut Teknologi Bandung
[3]National Research and Innovation Agency (BRIN)

Extended Abstract

Scientific instrumentation plays a fundamental role in STEM education, engineering training, and experimental research. However, the high cost of commercial platforms such as National Instruments ELVIS and Digilent Analog Discovery often limits their adoption, particularly in developing countries. To address this challenge, we are developing a configurable and scalable Zynq-7000 FPGA/SoC platform for scientific instrumentation, designed to provide affordable access to laboratory-grade measurement, signal generation, data acquisition, and visualization capabilities for education and research.

The proposed platform is envisioned as a multi-tier architecture consisting of three performance levels. The educational tier is based on ESP32 microcontrollers and has been successfully implemented for introductory experiments in physics, electronics, sensor systems, and programming. The research tier, currently under development using STM32 microcontrollers, aims to provide improved sampling performance, enhanced signal processing capabilities, and greater flexibility for laboratory applications. The long-term objective is to develop a professional-grade platform based on the **Zynq-7000 FPGA/SoC architecture**, enabling high-speed data acquisition, real-time signal processing, and hardware acceleration for advanced scientific instrumentation.

Current development utilizes a Zybo Zynq-7000 development board integrated with a 3.5-inch ILI9486 LCD display. To minimize power consumption, thermal load, and software complexity, wireless communication is implemented through an external ESP8266 module that bridges serial communication to UDP over Wi-Fi, avoiding the need to run a full operating system and networking stack on the Zynq platform. Development is performed using Vivado 2025. Initial milestones have been achieved, including successful implementation of XADC-based data acquisition and LCD visualization through the Processing System (PS). In parallel, an FPGA-based graphics accelerator has been developed in the Programmable Logic (PL) to reduce the processing burden associated with graphics generation and visualization.

Ongoing work focuses on further migrating time-critical functions from the Processing System to the Programmable Logic to exploit the parallel processing capabilities of FPGA hardware. The graphics accelerator is intended to handle computationally intensive visualization tasks in the PL, while the PS remains responsible for system control, communication, and other sequential operations. The integration of XADC-based data acquisition with the FPGA-based graphics accelerator is currently under development. This transition involves the development of finite state machine (FSM)-based architectures and careful timing optimization, which present significantly different design challenges compared to conventional embedded software development. Although AI-assisted development tools are increasingly useful, FPGA implementation remains highly dependent on timing-aware hardware design methodologies.

At the current stage, the platform utilizes the internal XADC resources available within the Zynq-7000 device. Future development will integrate external high-speed ADC and DAC components with target sampling rates of approximately 250 MSPS or higher, enabling applications in digital signal processing, communication systems, high-speed electronics, and experimental physics. Because such performance levels require careful attention to signal integrity, future professional

versions are expected to employ dedicated custom PCB designs that tightly integrate the Zynq-7000 SoC, analog front-end circuitry, and high-speed data converters.

By combining an open architecture, configurable firmware and software, FPGA-based hardware acceleration, embedded visualization, and scalable hardware platforms, this work aims to establish an affordable and flexible scientific instrumentation platform for STEM education, research, and innovation. The ongoing development is directed toward a configurable FPGA-enabled instrumentation platform capable of bridging educational laboratory applications and more advanced scientific and engineering applications.

Keywords: FPGA, Zynq-7000, Scientific Instrumentation, STEM Education, Embedded Systems, Data Acquisition, Signal Processing, Reconfigurable Computing.

T08

Reconfigurable Intelligent Surfaces for Integrated Sensing and Communication in Dense 6G Networks

C. Jeptoo [1] , D. Waswa [1,2] , and J. Jena

[1]*Centre for Broadband Communication, Department of Physics, Nelson Mandela University*

Sixth-generation (6G) networks are expected to deliver ultra-high data rates, sub-millisecond latency, and massive device connectivity. Among the defining use cases is integrated sensing and communication (ISAC), where the same infrastructure simultaneously performs environmental monitoring and user data transmission. However, dense urban deployments introduce a fundamental challenge as millimeter-wave and sub-terahertz signals that enable high bandwidth are highly susceptible to blockages from buildings, vehicles, and human movement, causing frequent link failures that degrade both communication reliability and sensing coverage.

Existing solutions address blockages through beam switching or network handovers, but these introduce latencies ranging from tens to hundreds of milliseconds, orders of magnitude above 6G requirements. Furthermore, conventional ISAC systems treat sensing and communication as separate functions that either share time slots or use independent beams, failing to exploit the spatial degrees of freedom available in dense deployments. This creates a critical gap, how can a network maintain continuous communication while simultaneously sensing its environment under dynamic blockage conditions, without compromising either function.

This work proposes a reconfigurable intelligent surface (RIS)-assisted ISAC framework for dense 6G networks. The RIS is spatially partitioned such that selected cells perform wide-beam environmental scanning (humidity, temperature, device presence) while remaining cells execute narrow-beam user tracking and data transmission using the same waveform. An AI-driven controller optimizes this partition dynamically based on channel conditions and traffic demand. The primary objective is to demonstrate that RIS-assisted ISAC can maintain link reliability above 99% under NLoS conditions while achieving sensing accuracy.

The system employs OTFS modulation for Doppler-resilient joint sensing and communication. A hierarchical AI architecture uses federated reinforcement learning at edge servers for global RIS configuration, while lightweight event-triggered inference runs locally on RIS controllers to avoid computational saturation. Simulation evaluates Bit Error Rate (BER), spectral efficiency, missed detection probability, and end-to-end latency under varying blockage densities.

This work will quantify the reliability improvement achievable through RIS-assisted ISAC compared to conventional beam-switching baselines, establish design parameters for RIS element count and phase resolution in dense 6G scenarios, and demonstrate an AI control framework that improves resource utilization without introducing prohibitive computational overhead. These findings aim to inform RIS hardware specifications and ISAC waveform standardization for 6G.

Reference

Aldirmaz-Colak, S., Namdar, M., Basgumus, A., Özyurt, S., Kulac, S., Calik, N., Yazici, M.A., Serbes, A. and Durak-Ata, L., 2025. A comprehensive review on ISAC for 6G: Enabling technologies, security, and AI/ML perspectives. *IEEE Access*, *13*, pp.97152-97193.

Kaushik, A., Singh, R., Li, M., Luo, H., Dayarathna, S., Senanayake, R., An, X., Stirling-Gallacher, R.A., Shin, W. and Di Renzo, M., 2024. Integrated sensing and communications for IoT: Synergies with key 6G technology enablers. *IEEE Internet of Things Magazine*, *7*(5), pp.136-143.

Björnson, E., Wymeersch, H., Matthiesen, B., Popovski, P., Sanguinetti, L. and De Carvalho, E., 2022. Reconfigurable intelligent surfaces: A signal processing perspective with wireless applications. *IEEE Signal Processing Magazine*, *39*(2), pp.135-158.

Sharma, S., Singya, P.K., Deka, K., Adjih, C. and Sharma, M., 2025. Terahertz communication: State-of-the-art and future directions. *IEEE Open Journal of the Communications Society*.

T09

Development of a Portable Multimodal Sensory Assessment System for Early Detection of Diabetic Peripheral Neuropathy

Muhammad A. Kadir [1] , **Arpita Roy** [1] , **Md M. Jomaddar** [1] , **Md M. Hossen** [1] , and **Hridoy C. Barman** [1]

[1]*Department of Biomedical Physics and Technology, University of Dhaka, Dhaka, Bangladesh*

Diabetic peripheral neuropathy (DPN) remains a major complication of diabetes and a leading cause of ulcers and amputations. Conventional diagnostic tools such as the 10g monofilament, biothesiometer, and nerve conduction studies either lack sensitivity to early small-fiber dysfunction or are unsuitable for widespread use in resource-limited settings [1, 2]. This study aimed to design and develop a low-cost, portable multimodal sensing system capable of quantifying sensory foot pressure threshold, tactile reaction time (TRT), and temperature discrimination threshold (TDT), for early DPN detection. A custom sensory probe incorporating a strain-gauge load cell and HX711 amplifier enabled accurate measurement of just-detectable pressure by recording the load applied to the plantar surface until participant response. For thermal assessment, a MOSFET driven Peltier module and thermistor-controlled feedback was developed, allowing rapid and stable bidirectional temperature modulation. TRT was captured through a skin impedance–based sensing interface designed to detect tactile responses with millisecond resolution. A preliminary study involving 25 participants (15 healthy and 10 diabetic) was conducted to evaluate the system's performance. Healthy participants exhibited low sensory thresholds at distal toes (0.4–0.5 g, equivalent load) and higher thresholds at the heel (≈1.0 g), while diabetic participants showed markedly elevated thresholds (>80 g at several sites). TRT was prolonged in diabetics (185–194 ms) relative to healthy participants (166–167 ms), while TDT values were consistently higher in diabetics, reflecting reduced thermosensory acuity. Statistical analysis revealed strong group separation across modalities ($p < 0.05$). The proposed system enables quantitative, site-specific assessment of sensory function and demonstrates potential for accessible, early DPN detection. Its design emphasizes affordability, scalability, and usability in low-resource healthcare environments, with future work focused on clinical validation in larger cohorts.

[1] D. Selvarajah et al. The lancet Diabetes & endocrinology, **7**, 938-948 (2019).
[2] J. Røikjer et al. Journal of diabetes science and technology, **20**, 403-418 (2026).

T10

Selection-Aware Machine Learning for Unresolved Stellar Binaries

Jauza Akbar Krito[1] **and Salsabila Zahirah Pranida**[2]

[1]*Department of Physics, Universitas Gadjah Mada, Indonesia*
[2]*Independent Researcher, Indonesia*

Edge intelligence is increasingly attractive for scientific instruments that must reduce data volume and make useful decisions close to the point of measurement. A central design question is whether a more complex learned model actually improves scientifically meaningful detection, and how the resulting selection effects should be calibrated. We investigate this question through unresolved stellar binaries using Gaia DR3 and 2MASS photometry and parallax [1, 2]. A PARSEC-based forward model [3] generates synthetic F/G/K main-sequence single and binary systems with exact flux addition and a heteroscedastic observational-noise model. We compare a classical colour–magnitude overluminosity criterion, Random Forest, XGBoost, LightGBM, a plain multilayer perceptron, and a physics-informed neural inverse model, then combine the photometric experiment with a Gaia astrometric forward model to map detectability across mass ratio, primary mass, distance, and orbital period.

In a joint experiment with 50,000 synthetic binaries, the photometric and astrometric detection fractions are 7.8% and 34.4%, respectively; 6.1% are photometric-only, 32.7% astrometric-only, 1.7% detected by both, and 59.5% missed by both. The overlap is smaller than expected if the two channels were independent, demonstrating structured complementarity. Tree-based baselines outperform both neural models on the present synthetic benchmark, while the added physical constraints do not yield a measurable classification gain. Transfer to real Gaia non-single-star data is also poorly calibrated: a threshold designed for 1% synthetic false positives produces 25.2% false positives in the real control sample.

These results suggest a practical principle for edge scientific intelligence: model efficiency should be evaluated together with physically interpretable selection functions and domain calibration, not accuracy alone. For resource-constrained scientific inference, a simpler model can be preferable when its detection regime is measurable, complementary sensors are understood, and failure modes are explicit. The framework provides an astronomy case study for hardware–software co-design of compact, selection-aware scientific inference pipelines.

[1] Gaia Collaboration, A&A **674**, A1 (2023).
[2] M. F. Skrutskie et al., AJ **131**, 1163 (2006).
[3] A. Bressan et al., MNRAS **427**, 127 (2012).

T11

Structural Anomaly Forecasting with Edge-Intelligence and Early Warning System (SAFE-IEWS)

For School and Conference on Edge Intelligence
for Scientific Instrumentation and Reconfigurable Computing

14 – 24 September 2026,
University Malaysia Perlis (UniMAP) Kuala Lumpur, Malaysia

Reginald Juan Magpantay Mercado [1]

[1]*GTek Enterprise*

This Industrial IoT - Smart Predictive Maintenance System project, is targeting to design, develop, and build an early warning system (EWS) integrating a wireless sensor network with edge-intelligence utilizing Polar Fire SOC or PSoC Edge. This embedded machine learning (ML) system will detect/forecast the structural integrity or failure of structures such as columns, walls, buildings, platforms, towers, and bridges after a severe ground shaking or earthquake event. The SAFE-IEWS will provide alert data/information to structural engineers and safety authorities (performing preventive maintenance) if strengthening of structures are necessary to ensure public safety. The final task of the SAFE-IEWS is to provide early warning data/informaion to emergency authorities and vulnerable public if there is a need to evacuate to safer grounds away from imminent danger from failing/falling structures.

This SAFE-IEWS solution is a structure anomaly detector/predictor system that will integrate the following sensors:

- A 3-axis (+/- 2g, 20-bit) precision accelerometer to monitor structures for vibrations.
- A 2-axis (+/- 90 degress) precision inclinometer to monitor tilting of structures resulting to unsafe angle or displacement.

The resulting data from these sensors will be collected in real-time, stored, and processed by a reconfigurable instrument based on a Polar Fire SOC or PSoC Edge. Anomaly prediction will be inferred by the ML algorithm running in the Polar Fire SOC or PSoC Edge. Early warning data or information will be sent via Wi-Fi to a data center (remote computer) where structural engineers and safety officers are monitoring.

The SAFE-IEWS will be in a portable package and will operate in stand-alone mode. It will utilize a 12V, 2A solar-battery power system designed to operate 24hrs-7days continuously.

The project is currently at the prototyping phase, the hardware and sensors are being evaluated using either a Polar Fire SOC (FPGA-SOC from Microchip, originally from Actel/Microsemi) or PSoC Edge (Arm-based Reconfigurable SoC from Infineon, originally from Cypress).

T12

Adaptive Convolutional Neural Network Training for Semiconductor Image Classification: Toward FPGA Deployment

Hariyanti Mohd Saleh[1]

[1]*Faculty of Electronic Engineering & Technology, Universiti Malaysia Perlis, 02600 Arau, Perlis, Malaysia*

Reliable convolutional neural network (CNN) training is relevant to semiconductor image classification and to possible deployment on resource-constrained hardware. This work presents the Cyclical Local Adaptive Learning Rate Strategy (CLARA), an enhanced learning-rate scheduler for CNN training. CLARA uses epoch-level training-loss feedback to adjust the lower learning-rate bound while retaining iteration-level cyclic updates. Earlier work examined CNN optimization using a cyclical learning rate for wafer defect image classification. In the reported ResNet-50 experiment on wafer defect maps, CLARA achieved 99.39% validation accuracy. This result applies only to the evaluated dataset and model configuration. Because CLARA is applied only during training, it does not change the CNN architecture or add computational overhead during inference. Future work will investigate whether the trained model can be compressed through quantization or pruning and implemented as a proof of concept on an FPGA board. The planned work will draw on prior experience with Xilinx SDK, standalone and FreeRTOS board support packages, register control, and data transfer between software and programmable logic. FPGA resource utilization, latency, power consumption, and real-time performance have not yet been evaluated.

[1] H. M. Saleh, N. A. Mat Isa, N. H. Saad, and M. I. F. Maruzuki, "Toward Solving the Learning Rate Problem in Convolutional Neural Networks Using an Enhanced Cyclical Learning Rate Strategy," IEEE Access, 106615-106633 (2026), doi: 10.1109/ACCESS.2026.3710474.
[2] C. C. Nelson, H. M. Saleh, and N. H. Saad, "Wafer defect image classification optimization using CNN with cyclical learning rate," in Proc. IEEE Int. Conf. Artif. Intell. Eng. Technol. (ICAIET), 2025, doi: 10.1109/ICAIET67524.2025.11265332.

T13

BioEdge: Ultra-Low-Power Edge Instrument for Next-Generation Wearable Capacitive Electrophysiological Biomedical Sensor

Charn Loong Ng [1]

[1]*Celestica ATS Engineering, Malaysia.*

Abstract:
Multichannel capacitive electrophysiological sensing offers a non-contact, continuous paradigm for digital healthcare [1]; however, conventional wearable architectures are severely constrained by high telemetry power draw, severe motion artifact corruption, and data privacy vulnerabilities inherent to cloud-dependent computing [2 – 5]]. To overcome these critical barriers, **BioEdge** introduces a novel, hardware-software co-designed edge instrument that performs autonomous, end-to-end signal conditioning and physiological event decoding directly on resource-constrained microcontrollers (MCUs) or field-programmable gate arrays (FPGAs). By deploying ultra-lightweight, quantized deep neural networks at the point of acquisition, the system dynamically suppresses non-stationary capacitive motion noise and classifies multi-modal electrophysiological patterns with >95% accuracy and an end-to-end latency of <10 ms. Operating within a strict sub-50 mW total power budget, this on-chip intelligence eliminates the bandwidth overhead and energy penalty of continuous raw data streaming while preserving biometric privacy. EdgeBio establishes a scalable, generalized foundation for next-generation ambulatory monitors, closed-loop therapeutics, and privacy-preserving human-machine interaction.

[1] C. L. Ng *et al*., "A Versatile and Wireless Multichannel Capacitive EMG Measurement System for Digital Healthcare," in *IEEE Internet of Things Journal*, vol. 11, no. 11, pp. 20120-20137, 1 June1, 2024, doi: 10.1109/JIOT.2024.3370960.

[2] S. Roy, A. N. M. W. Azad, S. Baidya, M. K. Alam and F. Khan, "Powering Solutions for Biomedical Sensors and Implants Inside the Human Body: A Comprehensive Review on Energy Harvesting Units, Energy Storage, and Wireless Power Transfer Techniques," in *IEEE Transactions on Power Electronics*, vol. 37, no. 10, pp. 12237-12263, Oct. 2022, doi: 10.1109/TPEL.2022.3164890.

[3] M. Tang and Y. Xin, "Efficient Energy Consumption Optimization for Wireless Sensor Health Monitoring System in Mobile-Edge Computing," in *IEEE Internet of Things Journal*, vol. 11, no. 5, pp. 7948-7955, 1 March1, 2024, doi: 10.1109/JIOT.2023.3317830.

[4] R. Govindarajan, R. V, K. Malathy, M. Gopikrishnan, A. A. R and P. M, "A Novel Design of Edge-AI Powered Real-Time Health Monitoring System Using Energy-Efficient Embedded Sensors," *2026 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES)*, Chennai, India, 2026, pp. 1-9, doi: 10.1109/ICSES66558.2026.11479024.

[5] J. Yoo, "Energy-Efficient AI at the edge for Biomedical Applications," *2023 20th International SoC Design Conference (ISOCC)*, Jeju, Korea, Republic of, 2023, pp. 202-202, doi: 10.1109/ISOCC59558.2023.10396488.

T14

TinyML-Based Electronic Nose for Food Quality Monitoring Using Embedded Rust on ESP32-S3 Edge Devices

Ahmad Radhy [1], Akhmad Maulvin Nazir Zakaria [1], Moreno Rizkihakiki [1] and Muhammad Abigail Novarianto [1]

[1]Department of Instrumentation Engineering, Institut Teknologi Sepuluh Nopember, Surabaya, Indonesia.

Food quality monitoring increasingly requires intelligent instrumentation capable of performing real-time analysis directly at the sensing node. Conventional cloud-centric approaches often suffer from latency, bandwidth limitations, and privacy concerns, making them less suitable for resource-constrained industrial and agricultural environments. Edge Intelligence has emerged as a promising paradigm that enables artificial intelligence (AI) processing close to the data source, improving responsiveness, reliability, and operational efficiency [1].

This work presents the development of a TinyML-based electronic nose (e-nose) platform for food quality monitoring, with a particular focus on aroma profiling of Arabica coffee. The proposed system integrates a custom-built gas sensor array, signal conditioning circuitry, and an ESP32-S3 edge computing platform programmed entirely in Embedded Rust. The use of Rust provides memory safety and robustness for embedded AI deployment, while recent advances in Rust-based TinyML inference engines demonstrate the feasibility of executing neural network models on highly resource-constrained devices [2].

The e-nose acquires multivariate volatile organic compound (VOC) responses and performs on-device inference without reliance on cloud infrastructure. Machine learning models are trained to classify and characterize coffee aroma profiles, enabling the generation of digital aroma fingerprints for traceability and quality assessment. Similar approaches have demonstrated the potential of electronic noses for discriminating coffee origins and establishing extensible aroma databases [3].

Beyond food quality monitoring, this work contributes to the broader vision of intelligent scientific instrumentation by combining sensor systems, TinyML, and edge computing within a low-power embedded platform. Future work will investigate hardware-software co-design and reconfigurable edge architectures to further improve inference efficiency, scalability, and energy performance, in line with emerging trends in reconfigurable edge intelligence systems [4].

Keywords: TinyML, Electronic Nose, Edge Intelligence, Scientific Instrumentation, Embedded Rust, ESP32-S3, Food Quality Monitoring, Coffee Aroma Profiling.

[1] C. Savaglio, P. Mazzei and G. Fortino, "Edge Intelligence for Industrial IoT: Opportunities and Limitations," *Procedia Computer Science,* vol. 232, pp. 397-405, 2024.

[2] M. Carnelos, F. Pasti and N. Bellotto, "MicroFlow: An Efficient Rust-Based Inference Engine for TinyML," *Internet of Things,* vol. 30, 2025.

[3] C.-H. Lee, I.-T. Chen, H.-C. Yang and Y. J. Chen, "An AI-Powered Electronic Nose System with Fingerprint Extraction for Aroma Recognition of Coffee Beans," *Micromachines (MDPI),* vol. 13, 2022.

[4] W. Foudhaili, A. Nechi, C. Thermann, M. A. Johmani, R. Buchty, M. Berekovic and S. Mulhem, "Reconfigurable Edge Hardware for Intelligent IDS: Systematic Approach," in *20th International Symposium: Applied Reconfigurable Computing. Architectures, Tools, and Applications (ARC 2024)*, Aveiro, Portugal, 2024.

T15

SPARSE-BAM: A SPARSITY-AWARE BROKEN ARRAY MULTIPLIER WITH UNIFIED BOUNDARY CORRECTION FOR ENERGY-EFFICIENT NEURAL NETWORK INFERENCE

Aditya Kumar Sharma[1] Geetika Srivastava[1*] Raghvendra Tiwari[1] Sandeep Kumar Pandey[1] Pratham Agrawal[2]

[1]*Department of Physics & Electronics, Dr. RML Avadh University, Ayodhya (U.P.), India*

[2]*Department of Electrical Engineering, IIT Delhi, India*

*Corresponding author: Geetika Srivastava (geetika_gkp@rediffmail.com)

Edge intelligence on reconfigurable hardware platforms demands energy efficient arithmetic units which capable of sustaining real-time throughput within tight resource budgets. Approximate computing intentionally soothing arithmetic exactness in exchange for reduced logic demand and dynamic power is well suited to error-resilient inference workloads such as image processing and On device classification. Multipliers are among the most resource concentrated arithmetic primitives; their optimisation directly impacts LUT resource, DSP block consumption, and on-chip power consumption. Conventional broken array multipliers apply fixed partial-product pruning regardless of operand content, ignoring the structural sparsity introduced by weight quantisation and ReLU activation functions in deployed neural networks.

This work proposes SPARSE-BAM, a sparsity-aware broken array multiplier described in C++ for synthesis with Vitis HLS and implemented on a Xilinx Zynq-7000 SoC (XC7Z020) using Vivado. Three co-designed innovations are introduced. A Sparsity-Conditioned Vertical Cut (SCVC) will adapt the partial-product pruning depth dynamically as a function of operand pop count, synthesised as a shallow LUT tree with no DSP inference, allowing sparser operands to sustain deeper cuts without exceeding a target error budget. A Unified Boundary Correction (UBC) will fuse the output-bit bias strategy of PAAM01 with the sign-bit recycling principle of SIBAM into a single adder stage, jointly minimising Mean Relative Error Distance (MRED) for signed 16-bit operands beyond what either technique achieves independently. Zero-Row Clock Gating (ZRCG) will suppress switching activity on zero partial-product rows via dataflow pragmas, directly harvesting ReLU-induced activation sparsity commonly exceeding 50% in deployed convolutional networks at zero accuracy cost. A closed-form MRED model under non-uniform sparsity-weighted input distributions will be derived, replacing the uniform-input assumption of prior work and enabling application-specific error budgeting for FPGA instrumentation systems.

Based on benchmarking against the existing FPGA approximate multiplier literatures and specific bit pattern sparsity in different applications, the SPARSE-BAM consumes approximately 100 to 140 LUTs and zero DSP48E1 blocks on the Zynq-7000 fabric, compared to 1 to 2 DSP48 blocks required by an HLS-inferred exact multiplier. On-chip dynamic power savings of 40 to 70% under 50% ReLU sparsity, as estimated from Vivado switching-activity injection (SAIF-based analysis). MRED is expected remains below 5 to 10% for typical quantised neural network weight distributions and classification accuracy loss under 5% for an INT8 LeNet-5 evaluated on MNIST.

[1]. J. Han and M. Orshansky, "Approximate Computing: An Emerging Paradigm for Energy-Efficient Design," in *Proc. 18th IEEE European Test Symp. (ETS)*, pp. 1–6, Montpellier, France (2013).

[2]. H. Jiang, J. Han, and F. Lombardi, "Approximate Arithmetic Circuits: A Survey, Characterization, and Recent Applications," *ACM J. Emerg. Technol. Comput. Syst.* 13, Article 60 (2017).

[3]. H. R. Mahdiani, A. Ahmadi, S. M. Fakhraie, and C. Lucas, "Bio-Inspired Imprecise Computational Blocks for Efficient VLSI Implementation of Soft-Computing Applications," *IEEE Trans. Circuits Syst. I* 57, 850 (2010).

[4]. P. Balasubramanian, R. Nayar, and D. L. Maskell, "Approximate Array Multipliers," *Electronics* 10, 630 (2021).

[5]. A. S. Roy and A. S. Dhar, "SIBAM—Sign Inclusive Broken Array Multiplier Design for Error Tolerant Applications," *IEEE Trans. Circuits Syst. II, Exp. Briefs* 67, 2881 (2020).

[6]. F. Farshchi, M. S. Abrishami, and S. M. Fakhraie, "New Approximate Multiplier for Low Power Digital Signal Processing," in *Proc. 17th CSI Int. Symp. Comput. Arch. Digital Syst. (CADS)*, pp. 25–30 (2013).

T16

Edge for the Intelligent Grid: A review of Edge Intelligence in Grid Technologies Current Applications and Emerging Frontiers

Amna Sardar Siddiqui [1]

[1] *Department of Electrical Engineering, Faculty of Engineering, University of Malaya, 50603 Kuala Lumpur, Malaysia*
22080364@siswa.um.edu.my

Abstract

Rapid proliferation of distributed energy sources, deregulated bidirectional power flows over vast areas, and emerging technologies of AI and IoT has revolutionized the conventional grids to self-healing smart grids. Centralized Control architectures such as Cloud Computing are outpaced by the dynamic capabilities required for quick operational decision making, cost effectiveness and mitigating environmental concerns. Fusion of three technological advancements of Artificial Intelligence, Internet of Things and Edge computing has led to advent of the "Edge Intelligence" resulting in a dynamic, evolving architecture capable of autonomous decision making on real time data at the source itself, thus reducing the latency, improving the grid's responsiveness, resilience and reliability, and mitigating the environmental impact of intensive water consumption and carbon emission by the cloud data centres. In terms of economics, the Edge Intelligence for smart grid market size is expected to expand to $ 138.2 Billion by 2034 from $ 15.2 Billion in 2024, with the growth at a Compund Annual Growth Rate (CAGR) of 24.70% during the forecast period from 2025 to 2034 [1]. Altough several reviews have examined edge computing for power systems, the practical deployment across full spectrum of grid technologies is limited. This systematic review bridges the gap through analysing the practical applications, industrial deployment and challenges spanning power system stability and control [2], advanced grid management [3-4], proactive maintenance [5], cyber resilience [6], resource optimization and allocation [7], and fault detection and isolation [8]. The review also aims to provide the academia and industry with potential and emerging frontiers in the field of Grid Edge AI such as federated learning [9] and quantum computing [10] for transforming the conventional grid architectures into resilient, self-optimizing networks capable of autonomously responding to dynamic operating conditions.

Keywords

Grid Edge, Edge Intelligence, Edge Computing, Grid Management and Control.

References

[1] Market.us, Edge AI for Smart Grid Market Size Report (2025).
[2] Y. Xiang,Z. Yi,X. Lu, Z. Yu, D. Shi, C. Xu, X. Li, Z. Wang, Elec. Pow. Sys. Res.199,107448 (2021).
[3] F. Yıldırım, Y. Yalman, K. Ç. Bayındır, E. Terciyanlı, Appl. Sci. 15, 4592 (2025).
[4] J. Huang, S. Zhou, G. Li, Q. Shen, Sci. Rep. 15, 24890 (2025).
[5] Y. Xu, N. M. M. Nascimento, P. H. F. de Sousa, F. G. Nogueira, B. C. Torrico, T. Han, C. Jia, P. P. R. Filho, Appl. Soft Comput. 101, 107053 (2021).
[6] W.-T. Lin, G. Chen, Y. Huang, Appl. Energy 314, 118828 (2022).
[7] Y. Chi et al., IEEE Access 11, 6541–6550 (2023).
[8] S. Netsanet, D. Zheng, Z. Wei, G. Teshager, Front. Energy Res. 10, 826915 (2022).
[9] Y. Li, D. Qin, H. V. Poor et al., Nat. Commun. 15, 9044 (2024).
[10] H. Tran-Dang, D. S. Kim, J. Supercomput. 82, 273 (2026).

T17

Lightweight LSTM Inference via Knowledge Distillation, Pruning, and Quantization for Real-Time ECG Anomaly Detection at the Edge

Nethulya A. Sooriarachchi [1] , Vishalya P. Sooriarachchi [1]

[1]SLTC Research University, Meepe, Padukka, Sri Lanka 1
[2] COPELABS, Lusofona University,Campo Grande, Lisbon, Portugal 1

Deploying Long Short-Term Memory (LSTM) networks on resource-constrained edge hardware is a key challenge in real-time scientific instrumentation [1]. Large LSTM models achieve strong accuracy on sequential sensor data but are impractical for Field Programmable Gate Array (FPGA) deployment due to their parameter counts and floating-point arithmetic [2]. While model compression techniques such as knowledge distillation, pruning, and quantization have been studied individually, their joint application to LSTM architectures targeting FPGA-based scientific instrumentation remains underexplored. This work addresses that gap by proposing an integrated three-stage compression pipeline for recurrent models, producing a compact student LSTM that closely matches its larger teacher on Electrocardiogram (ECG) arrhythmia classification. A two-layer teacher LSTM (128 hidden units, 199,555 parameters) is trained on the Massachusetts Institute of Technology–Beth Israel Hospital (MIT-BIH) Arrhythmia Database [3] for five-class beat classification, achieving 97.83% accuracy (macro F1: 0.611). Knowledge distillation [4, 5] transfers the teacher's soft output distributions to a compact one-layer student (32 units, 4,579 parameters) via temperature-scaled KL divergence, followed by L1 magnitude pruning at 50% sparsity and dynamic int8 quantization. The full pipeline achieves a 97.7% parameter reduction while retaining 96.67% accuracy, a drop of only 1.16% from the teacher, confirming that joint distillation and structured sparsity preserve model quality far more effectively than training a small LSTM from scratch. Future work will target FPGA synthesis via HLS4ML for on-device latency and power evaluation. We also propose extending the framework to support continual edge adaptation, where the deployed student periodically synchronises with an updated teacher LSTM to incorporate new labelled patterns without full retraining while also enabling real-time model tracking in evolving instrumentation environments.

Keywords: Long Short-Term Memory (LSTM), Field Programmable Gate Array (FGPA), Knowledge Distillation, Electrocardiogram (ECG)

[1] H. Li et al., Comput. Math. Methods Med. 2022, 8039281 (2022).
[2] B. Ramhorst, G. A. Constantinides, V. Loncar, arXiv:2308.05170 (2023).
[3] G. B. Moody and R. G. Mark, IEEE Eng. Med. Biol. Mag. 20, 45 (2001).
[4] G. Hinton, O. Vinyals, J. Dean, arXiv:1503.02531 (2015).
[5] A. M. Mansourian et al., Trans. Mach. Learn. Res., arXiv:2503.12067 (2025).

T18

FPGA-Based Low-Power Cryptographic Hash Architecture for Secure Edge Intelligence Applications

Shamsiah Suhaili [1], Norhuzaimin Julai [1], Maimun Huja Husin [1], Asrani Lit [1] and Nordiana Rajaee [1]

[1]1Electrical & Electronic Engineering Department, Faculty of Engineering, UNIMAS, 94300 Kota Samarahan, Sarawak, Malaysia

Edge intelligence systems rely heavily on secure and energy-efficient hardware to protect real-time data processed at the network edge. This work presents a conventional, high-performance FPGA-based cryptographic hash architecture specifically designed to provide hardware-level security for these edge applications [1]. To meet the strict power and resource constraints of edge environments, the proposed architecture utilizes Verilog HDL implementation integrated with hardware optimization techniques, including Gray encoding and architectural unfolding. These methods successfully reduce dynamic switching activity [2] while maximizing throughput efficiency [3]. Experimental evaluation conducted via FPGA development tools demonstrates optimized area utilization, high processing speeds, and low dynamic power consumption. The resulting design provides a robust, low-power hardware security solution essential for safeguarding data within secure edge intelligence ecosystems.

Reference:

[1] Sideris, A., & Dasygenis, M. (2023). Enhancing the Hardware Pipelining Optimization Technique of the SHA-3 via FPGA. Computation, 11(8), 152. https://doi.org/10.3390/computation11080152

[2] Vignesh, S., Tripathi, A. N., Chaudhari, S., S, C., & Padhy, J. B. (2026). *Low-power, high-throughput FPGA implementation of pipelined AES with real-time performance and scalability analysis*. *Arabian Journal for Science and Engineering*. https://doi.org/10.1007/s13369-026-11299-7

[3]Welte, B., & Zambreno, J. (2023). An FPGA Implementation of SipHash. 2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 63-70. https://doi.org/10.1109/ipdpsw59300.2023.00022