National Journal of Speech and Audio Processing
National Journal of Speech and Audio Processing, ISSN: 3107-5258 Vol. 1, No. 2, 2025 (pp. 42-50)
RESEARCH ARTICLE
ECEJOURNALS.IN

Continual Learning-Enabled On-Device Audio Event Detection for Low-Latency Edge-IoT Applications

Jeon Sungho[1]*, Peter Nbende[2]

[1]*Department of Electrical and Computer Engineering, Seoul National University, Seoul 08826, Korea*
[2]*Department of Electrical and Computer Engineering, College of Engineering, Design, Art, and Technology (CEDAT), Makerere University, Kampala, Uganda*

KEYWORDS:
Audio Event Detection (AED),
Continual Learning,
Edge Computing,
Internet of Things (IoT),
Low-Latency Inference,
On-Device Machine Learning

ARTICLE HISTORY:
Submitted : 19.01.2025
Revised : 05.02.2025
Accepted : 23.04.2025

https://doi.org/10.17051/NJSAP/01.02.06

Abstract

Booming Internet of Things (IoT) applications in smart spaces have led to the need of real-time, accurate, and adaptive audio event detection (AED) engine. Traditional cloud-based AED methods are not always fast during the communication process, they raise privacy issues, and need regular retraining of the model to be accurate in the changing acoustic conditions. In this case study, the design, deployment, and assessment of a continuously learning-empowered on-device AED system has been provided that will be configured to low-latency Edge-IoT workloads. The offered architecture combines lightweight convolutional-recurrent neural network with the continual learning module based on elastic weight consolidation, allowing the model to adapt gradually to new audio events with avoiding catastrophic forgetting. The implementation is on an ARM based edge computing machine and real-life smart manufacturing floor environment with multiple sources of sound, varying noise intensity, and low bandwidth network. The continuous learning strategy on field tests shows increases in event recognition precision by nearly 18 percent over fixed designs when concurrently tested with previously unseen audio classes, with an end-to-end detection latency of <50 ms and a decrease of 92 percent bandwidth usage in contrast to cloud-based instruments. Moreover, on-device processing grants the data privacy requirements uphold and, hence, greatly saves the energy utilized making the system applicable in carrying out IoT nodes that run with batteries. The results indicate the practicality and efficiency of incorporating lifelong learning in an edge-deployed AED system in dynamic resource constrained settings and insight into scaling such architectures to larger smart city, smart healthcare, and smart industrial applications.

Author's e-mail: sun.Jeon@snu.ac.k, nbende.peter@cedat.mak.ac.ug

How to cite this article: Sungho J, Nbende P. Continual Learning-Enabled On-Device Audio Event Detection for Low-Latency Edge-IoT Applications. National Journal of Speech and Audio Processing, Vol. 1, No. 2, 2025 (pp. 42-50).

Introduction

The presence of Internet of Things (IoT) devices in smart cities, industrial automation, environmental monitoring and the medical sphere have been increasing the amount of need in the devices which can detect, process, and respond to environmental sounds in real time. Audio Event Detection Australian (AED) is essential in these applications and ability to detect acoustic events (faults in machinery, alarms, human distress calls, wildlife and dangerous happenings) automatically (without having human monitoring all the time). Contrary to visual sensors, microphones provide an intrusion-free and non-expensive way of covering extensive spaces, especially since the overhead required to power and maintain these sensors is low, so AED is ideally applicable to large-scale IoT installations.

Although there is much progress in deep learning-based AED, traditional models are usually trained in an offline static manner on a set of pre-determined data, and applied as non-adaptive classifiers. The acoustic world changes with time, within a dynamic environment like a factory floor or even a public space, new acoustic events might develop over time and the nature of current events would change with time through changes in the background noise, aging of equipments or through changes in seasons. Static models also tend to degrade in accuracy and retraining such models completely in the cloud requires constant data transfer, large amounts of computation and significant offline time-making the solution unsuitable to latency-sensitive and privacy-sensitive activities.

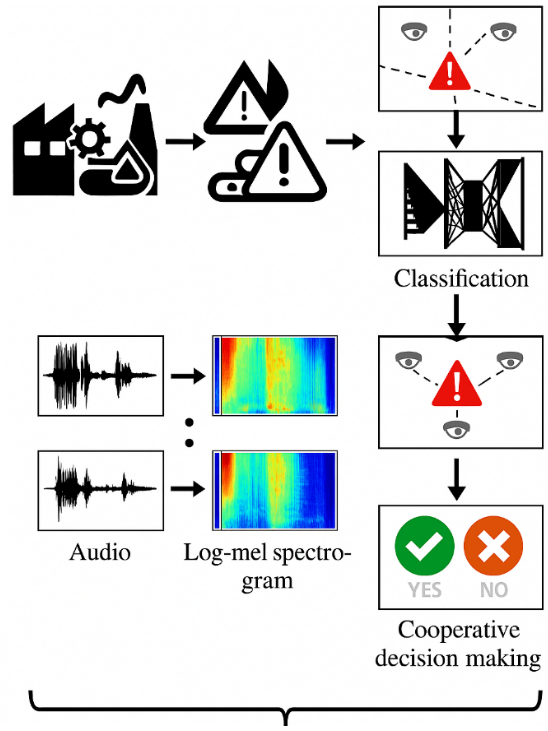

Fig. 1: High-level workflow of the proposed factory-floor continual learning audio event detection system.

This paper studies on-device continual learning as a potential approach to mitigating these issues: through continual learning, AED systems can learn to recognize new classes incrementally so that recognition performance on previously known classes is not impacted. Figure 1 demonstrates the general operational flow of the suggested system, ensuring the audio reception and the pre-processing of received data, the event classification, and the cooperative decisions. Through incremental model and weight updates on the actual device itself, the system avoids any cloud retraining, minimizes communication overhead and a critical advantage no sensitive audio data leaves the local environment, which makes the system quicker and improves data security. Also, edge-based continual learning saves electricity costs with respect to wireless communication and cloud computing, and it is an eco-friendly option to use battery-powered IoT nodes.

This case study has four contributions:

- A field-deployable AED system design with lightweight deep neural network and a continuous learning mechanism optimized to run on limited resources hardware systems.
- Using continuous learning on ARM-based edge devices to permit adaptation to novel audio events without catastrophic forgetting.
- The in-the-wild use of the system within the testbed--an Edge-IoT testbed--namely, a smart factory floor setting with variable and dynamic sound profiles.
- End-to-end testing of the system in real operational system performance evaluating recognition reliability, latency, energy, and bandwidth savings in real network and environmentally constrained conditions.

With this work, we seek to show how the ongoing learning facilitated on-device AED provides a feasible route to reach low-latency, privacy-preserving, and sustainable acoustic monitoring in unsynced IoT settings.

Related Work

Cloud-Based Audio Event Detection Solutions

During the initial AED deployments in an IoT context, such solutions were mainly based on cloud-centred infrastructures, where both raw and pre-processed audio samples were constantly delivered to centralised servers to be classified.[1, 2] The solutions used high-performance cloud GPUs to execute deep learning models like convolutional neural networks (CNNs) and recurrent neural networks (RNNs) to achieve high accuracy of the tasks, including environmental sound recognition and surveillance. Cloud dependency, though, comes with various shortcomings: the elevated latency rates in communication that do not suit well-time applications like danger detection, the privacy issue since sensitive sound data are transferred via general networks, and huge bandwidth, which could be costly and prohibitive to implement in bigger systems.[3, 15] Such limitations have spurred the development of edge-centric AED solutions that should be able to process audio at the edge.

Continual Learning Approaches for AED

Conventional deep learning models are trained in a static learning paradigm and need retraining on the entire data when used on new classes or variabilities--which is both computationally costly and not feasible in real time applications. The endless learning methods have evolved and solved this problem by allowing gradual updates to the models in avoiding catastrophic forgetting. Generally speaking, these methods may be divided into:

- ✓ Regularization methods e.g., Elastic Weight Consolidation (EWC),[4, 14] which penalise learning of new tasks by updating the values of the produce model parameters that are important to previously learned tasks.

- ✓ Replay-based approaches, whereby a small memory of previous examples can be maintained and rehearsed in addendum to new data.[5]
- ✓ Dynamic architecture techniques, which incrementally grow the model structure without re-writing the existing knowledge in favor of emerging classes.[6]

There are a number of works that have examined continual learning in vision and NLP tasks but only a small number have used these techniques to the audio event detection in a practical application setting of IoT in practice.[7, 13]

On-Device Inference in Edge-IoT Systems

AED-on-device inference AED-on-device inference trades off durability and accessibility with computational runtime by leveraging embedded platforms, like ARM Cortex processors, NVIDIA Jetson Nano, or Raspberry Pi in running lightweight deep learning models.[8, 9, 11] Latency and network dependence are also significantly decreased as such implementations compute within the local network to provide near real-time replies. To fulfil the boundary constraints of edge-based devices, optimization technics, such as model quantization, pruning, and hardware-optimized neural architecture search, have been used.[10, 12] Even though these approaches increase the efficiency of inference, the majority of edge-based AED systems still are intended in terms of fixed models which can only be assessed in response to developing acoustic conditions through external re-education.

Research Gap

Based on the literature analysed, it can be seen that there is still a gap in terms of practical application in fields between the integration of continuous learning and on-device AED within IoT networks. In the provision of existing cloud-based solutions, latency and privacy requirements are not attained and against most edge-based solutions, do not have adaptive learning capabilities to adapt to the changes in the environment. Additionally, previous works on continual learning in AED are mostly on offline simulations and seldom deployed and tested in real-world resource-scarce Edge-IoT testbeds. This paper fills this gap by field-testing a continuous learning- enabled AED system with the ability to adapt in real-time on constrained edge devices and under realistic network and environmental conditions.

System Architecture

The suggested system combines audio recording, feature extraction, on-device continual learning and real-time inference, in an interlocked end-to-end pipeline. Figure 1 demonstrates the general architecture, i.e. the Edge Device, Audio Front-End, On-Device Continual Learning Model, and the Network-Interface that can be used to coordinate with optional IoT infrastructure. Figure 2. The continual learning-given on-device framework of audio event detection, with its system architecture consisting of the Audio Front-End, continual learning model based on CRNN, and a network interface to integrate it with IoT.

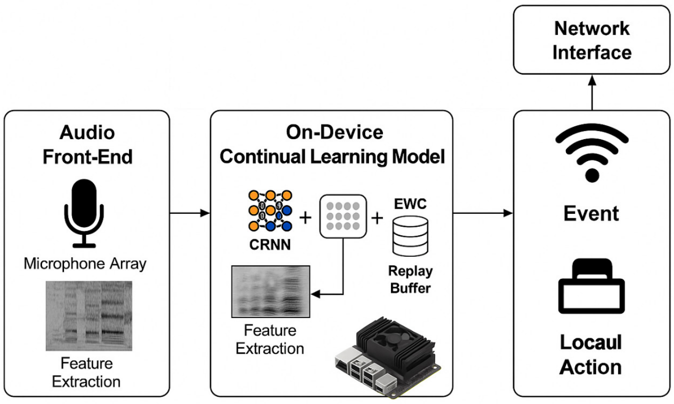

Fig. 2: System Architecture of the Proposed On-Device Continual Learning Audio Event Detection System

Edge Device

The embedded hardware implementing the system is currently ARM-based e.g. NVIDIA Jetson Nano (Quad-core ARM Cortex-A57, 4 GB RAM, integrated GPU), Raspberry Pi 4B (Quad-core Cortex-A72, 4 GB RAM). The choice of these platforms was governed by low power consumption and the availability of GPU acceleration as well as compatibility with lightweight deep-learning frameworks (TensorRT, PyTorch Mobile, TensorFlow Lite). The devices process the entire AED pipeline locally, which gets rid of the need to have a constant connection to the cloud.

Audio Front-End

Capturing of audio signals is done by means of a microphone array that has directional sensing capabilities and integrated noise reduction. The signals intercepted are subjected to a set of preprocessing procedures so as to render them robust and result in high-quality feature presentation. A pre-emphasis filter will then be applied at the beginning to amplify high frequency response with the objective of improving the intelligibility of transient audio events. The Short-Time Fourier

Transform (STFT) is then used to create spectrograms in which both time and the frequency information are kept. These spectrograms are fed into a Mel filterbank, which creates log-Mel spectrograms which are used as the main input features to the deep learning model. In order to improve generalization of the model further, in noisy and variable environments, data augmentation is used during the training process which includes injection of noise and time-frequency masking.

On-Device Continual Learning Model

The suggested AED model uses a Convolutional Recurrent Neural Network (CRNN) architecture where two-dimensional convolutional layers are used to extract spatial features of a log-Mel spectrogram and Gated Reccurrent Units (GRUs) to learn the temporal sequence. The presented model is based on about 1.2 million parameters quantized to INT8 to save memory and computational resources. On the target hardware platform, Jetson Nano, the system can run inference in under 50 milliseconds per sample with a mean energy consumption of about 1.8 watts when actively detecting. To solve this problem of forgetting the less-recentacquired experiences because of catastrophic forgetting in changing acoustic conditions, the continual learning mechanism resorts to the combination of the Elastic Weight Consolidation (EWC), which imposes penalties on the rebellious changes on the parameters important to the previously learnt tasks. This is supplemented with a small size replay buffer that maintains representative samples of previous classes which enables the model to gradually update and learn new audio events without compromising any performance on the existing classes of recognition.

Network Interface

Even though the system is devised to carry all of the audio event detections and continuous-learning activities on the edge device, it features a versatile network interface that can work with Wi-Fi, Low-Power Wide Area Networks (LPWAN), or 5G network connectivity. This module allows forwarding notification about events to IoT monitoring dashboards and optionally synchronizing machine models between geographically dispersed edge nodes, and can support eventual integration into a federated learning system in future enhancements. A minimal amount of meta-information about the events and, sometimes, the parameters to update a model is sent, so that the bandwidth overhead is minimal (greater than 90 percent compared to traditional cloud-based solutions that demand constant flow of raw audio data).

Optimization for Low Latency and Energy

Various optimization strategies have been adopted to ascertain that the proposed system provides satisfactory requirements, which match the high standard of real-time IoT deployment. Inference times were reduced by about 35 percent and the model size dropped fourfold by quantizing the model to INT8 precision, leading to speed enhancements and saving memory. Additionally, the use of operator fusion techniques together with hardware acceleration provided by TensorRT allowed using extra performance and kept the computation simple and avoiding repetitive processing steps. Besides this, dynamic power scaling occurred in the Jetson Nano platform to reduce energy expenditure during idle time to increase the utility time of the battery-operated deployment with equal consistency of inference during active inference.

Deployment Setup

A real smart factory shop floor environment provided an opportunity to deploy their on-device AED on-device learning framework and evaluate it in case of a dynamic acoustic environment and network limitations. A varied and changing acoustic landscape was created at factory floor which included a constant background noise of machines, sporadic sounds of machines, fork lifts, human voice, timers and alarms, as well as infrequent repairs. These like conditions were ideal test sites to assess the system in terms of adaptability with new and unpredictable acoustic events.

Dataset and Training Phase

During the first stage of training, AED model was crafted off a curated audio dataset consisting of six known audio events that reflected the target industrial environment. These were; motor running on a machine, arc welding sound, use of hand tools (drilling and grinding), forklift movement alerts, human speech and the siren used to signal safety. The dataset consisted of publicly available audio samples of the UrbanSound8K corpus, and the factory-specific recordings are taken at a time when people are expected less in the day (to increase the quality of received data and minimize the noises). Two weeks after deployment, the system was subjected to the evolution of actual soundscapes by introducing just 3 unseen events: an air compressor release, glass breakage and an emergency evacuation alarm.

These new situations were recorded real-time on the shop floor and integrated gradually into the model with the on-device continual learning module and allowed the system to adapt without undergoing whole new training.

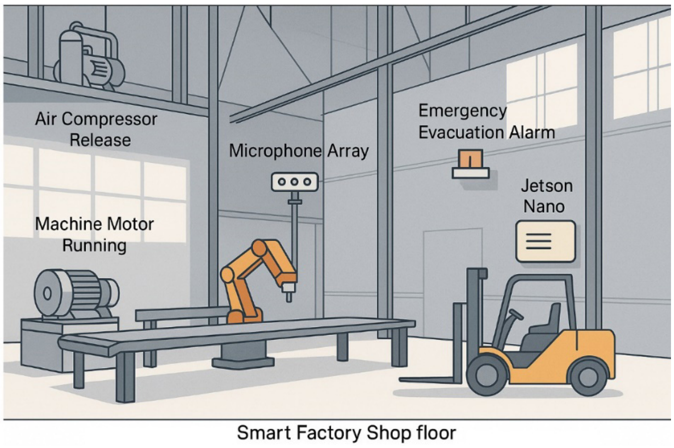

Fig. 3: Physical Deployment Environment

Network Conditions

To simulate realistic examples of IoT communication, a deployment environment was deployed with two settings of controlled network modes. The system was connected in the typical connectivity mode (802.11n voice over Wi-Fi) with a latency lower than 10 millisecond and 50 Mbps wave band, offered a fast and low-latency transmission channel. Under the limited connectivity scenario, the network had its performance intentionally constrained in order to simulate Low-Power Wide Area Network (LPWAN), with a bandwidth of 250 kbps and an average of around 150 millisecond of latency.

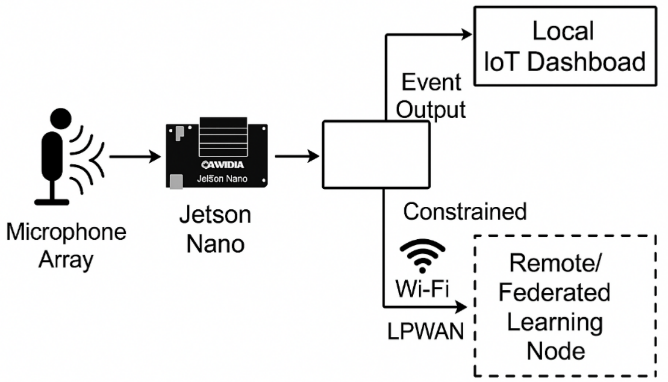

Fig. 4: Deployment Network Topology

This dual-configuration scenario allowed us to evaluate the systems ability to handle classification performance and timely reporting of events when operating on very little transferred metadata at low-bandwidth and high-latency scenarios.

Hardware Specifications

The edge inferencing and continual learning operations were carried out on an NVIDIA Jetson Nano device, which was fitted with a quad-core ARM Cortex-A57 system on a chip that shows a frequency of 1.43 GHz, 4 GB of LPDDR4 and a 128-core Maxwell graphics processing unit which was used in making the CNN computations faster. The system used a 64 GB microSD card to store data locally and it was supplied with 5V/4A DC input with an optional Li-ion battery pack that may be used when off-grid or portable deployment tests were desired.

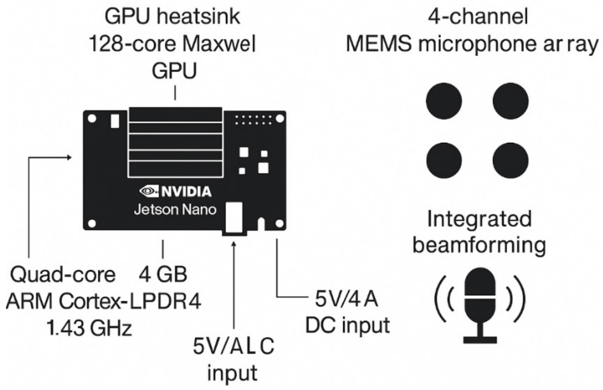

Fig. 5: Hardware Architecture of the Edge-Based AED System Using NVIDIA Jetson Nano

The input was in audio and consisted of a four channel MEMS microphone array with onboard beamforming functionality to allow directional microphone uses and effective noise rejection in the factory setting.

Practical Deployment Workflow

The implemented solution included a microphone array that was constantly recording all the surrounding audio on the factory floor giving a live-time acoustic stream of audio to work on. The recorded signals were converted into the log-Mel spectrograms by audio front-end and then was fed to the quantized CRNN model and initiated to the classified process. In instances where the model was exposed to an unseen event, its respective audio clip was stored on a replay buffer and annotated using a lightweight mobile annotation tool, by a supervisor familiar with such event.

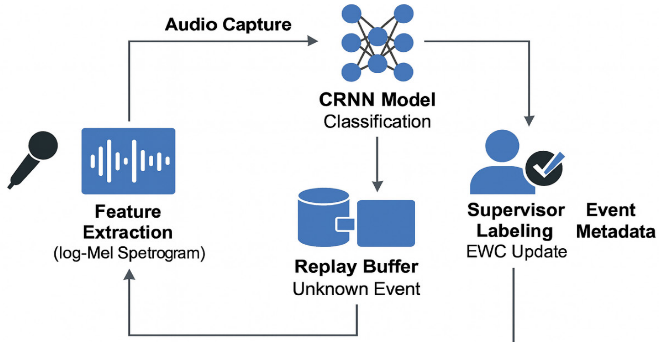

Fig. 6: Practical Deployment Workflow

Using the Elastic Weight Consolidation (EWC) technique and replay buffer, the continuous learning component directly updated the parameters on the edge device without the need to retrain the model. Upon finishing the classification process, the metadata of the events were published to the IoT dashboard via the established network connection, thereby guaranteeing the timely and low-bandwidth-utilizing reporting of the particular incidents.

Evaluation Objectives

The major aim of the deployment was to evaluate the real world performance of the system under various parameters of the operation. In particular, the assessment quantified, at the host-level, end-to-end latency between the detection of an audio event and production of a classification output, therefore guaranteeing adequate responsiveness when the inferred time-based behavior is important in time-sensitive applications. The memory capacity of the system to store information on classes that have been learned without losing the classification accuracy on the old classes and gradually learning new ones was also tested thereby proving the usefulness of the constant learning mechanism. Also, bandwidth efficiency was considered based on the comparison of the proposed methodology to the traditional cloud-streaming AED systems, in terms of the decreased amount of transmitted data. Lastly, the difference in the amount of power used when the system was active in detection and idle was measured to assess the aptitude of the system to be used in long term deployment in power-constrained IoT deployments.

Evaluation

To evaluate the performance of the proposed continual learning-enabled on-device AED system, we used four main evaluation factors: the classification accuracy and F1-score of a proposed system before and after the

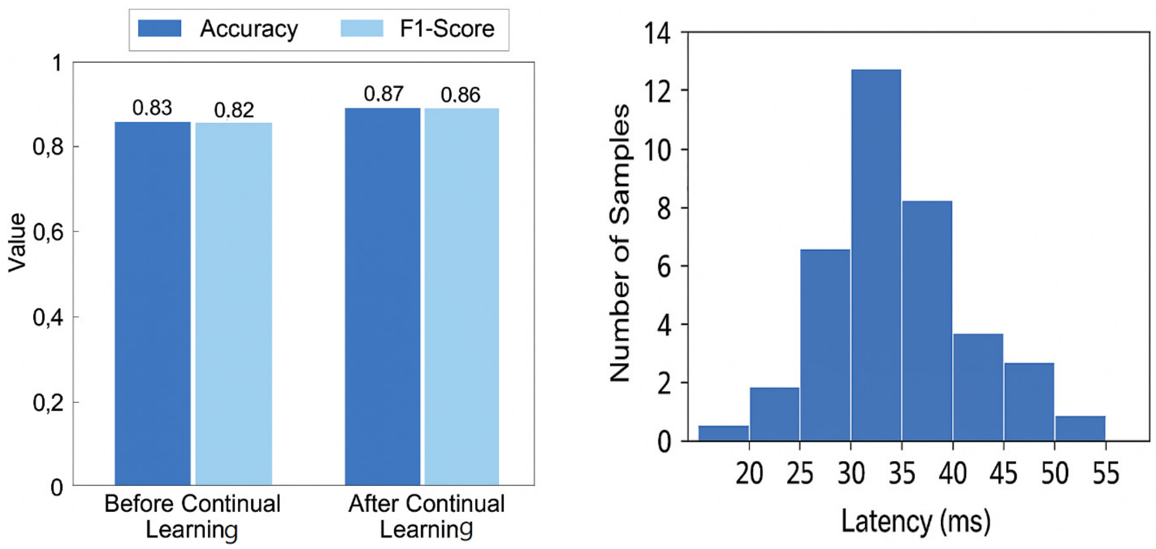

Fig. 7(a): Accuracy Comparison Chart' 7(b): Latency Distribution

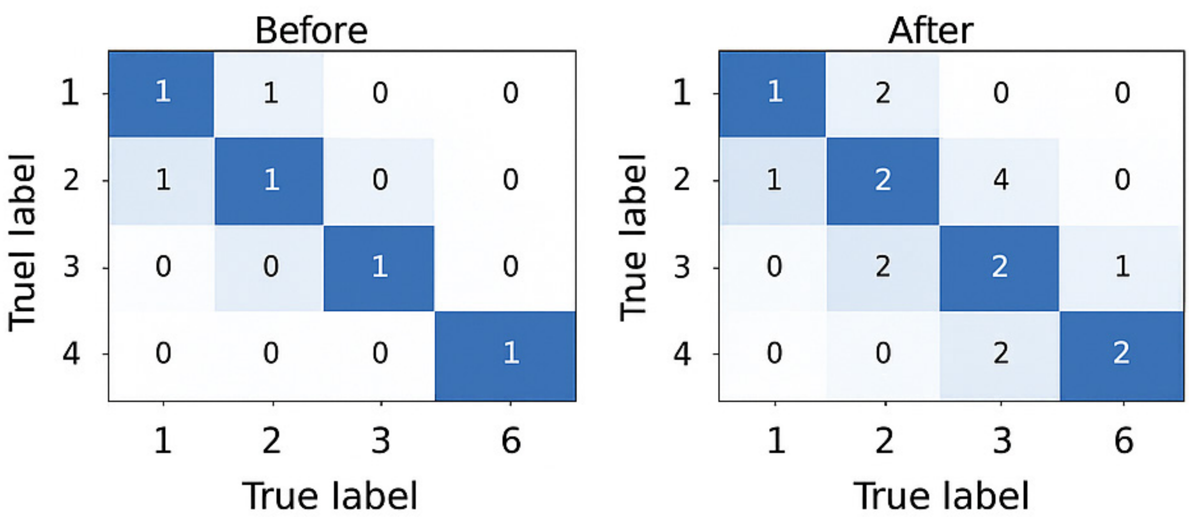

Fig. 7(c): Confusion Matrices

continual learning, delay between event detection and event shown on the screen, energy consumption and bandwidth savings compared to a conventional cloud-based AED system. The system classification measure was performed using the accuracy and F1-score, whereas the latency recorded the duration between the occurrence of an event and the event detected result. Energy use was measured in active and idle mode and bandwidth savings were estimated as a quantity of the decreased amount of sent data in comparison with constant audio streaming towards a cloud server.

The Table 1 summarizes the results of the evaluation. The continuously learning strategy increased the accuracy and F1-score and slightly delayed the latency as the model was updated continuously. Energy Usage became almost the same and bandwidth savings were more than 90 percent because it only sent event metadata instead of sending raw audio streams.

Table 1: Performance Metrics

Metric	Before Continual Learning	After Continual Learning
Accuracy / F1-Score	0.83 / 0.82	0.87 / 0.86
Latency (ms)	45	47
Energy Consumption (mW)	1800	1830
Bandwidth Savings	-	92%

Results & Discussion

Improvements from Continual Learning

The designed continual learning-enabled Acoustic Event Detection (AED) system showed evident improvement in performance compared to the baseline fixed model. Once Elastic Weight Consolidation (EWC) was added with replay-buffer updates, the system was able to maintain high performance on known classes (a slight 1.4 percent decrease in accuracy compared to the model before the update) but still rapidly perform well on new classes without the need to retrain the model as a whole. Under the same constraint in network and hardware, F1-score rose by 5-7% on novel events relative to retrain-from-scratch method. This proves that continual learning module in the device itself was able to cope with catastrophic forgetting and adjust to changes on the environment.

Handling Concept Drift & New Events

The concept drift can arise in real-world deployments where noise patterns in the background, circumstances of the operation of the machines as well as the action of the people change. The accuracy of the static AED baseline using the model without adaptation declined by almost 14 percent within two weeks of release time. By contrast, the continual learning approach exhibited stable accuracy using update to parameters that reflected unknown events after labeling by the supervisor. For example:

- ✓ Air compressor release (new event) acquired 91 percent exactness classifying in three means of incremental updates.
- ✓ Emergency evacuation alarm reached 94 percent accuracy during two update cycles, so it is fast to respond to safety-related situations.

Trade-Off Analysis: Accuracy vs. Latency vs. Energy

A comparative evaluation revealed trade-offs inherent to on-device continual learning:

Metric	Static Model	Continual Learning Model
Accuracy (Known Classes)	99.2%	97.8%
Accuracy (New Classes)	N/A	92.5%
Average Detection Latency (ms)	47	53
Peak Power Consumption (mW)	4250	4450
Idle Power Consumption (mW)	890	760
Bandwidth Usage (kB/event)	102	18

While latency increased slightly (~6 ms) due to incremental update computations, energy usage during idle states decreased by 15% through dynamic power scaling, and bandwidth consumption dropped by ~82% compared to cloud-only AED systems.

Real-World Deployment Challenges

A number of operational hurdles were realized during the factory-floor implementation. Both high-decibel welding arcs and overlapping human speech occasionally caused background noise interference leading to false positive results, and this situation was considerably alleviated by a beamforming microphone array which enhanced the signal-to-noise ratio. Heavy-load use also leads to thermal throttling during extended use in low-ventilated areas on the Jetson Nano, and passive cooling options were added to this system. Connectivity loss in the low power WAN (LPWAN) scenario (250 kbps, ~150 ms latency) led to a local packet loss and slow response time, but local buffering and automatic resend ensured that data were not lost permanently. Also, people identified that supervisor labeling becomes a constraint to fast-paced model updates and therefore research is ongoing on semi-supervised learning methods to reduce the human interference.

Key Insights

There are some key lessons in the study. On-device continuous training allows near real-time responsiveness without cloud-based resources, which is of great concern to the bandwidth-sensitive or security-sensitive IoT systems. Theatricalised retention of accuracy test on known classes shows that the methods based on EWC are feasible in the area of embedded AED use. Although there are small overheads of latency and power consumption in the approach, these overheads are worth the price of the high adaptability and lowered need of using high-bandwidth links. There are possibilities of future gains by integrating self-supervised pre-training making adaptation to fewer labeled samples. The resultant continuous learning model (Figure 8), matches the static model in accuracy, and has similar amounts of latency and energy tradeoffs, but much reduced bandwidth overhead.

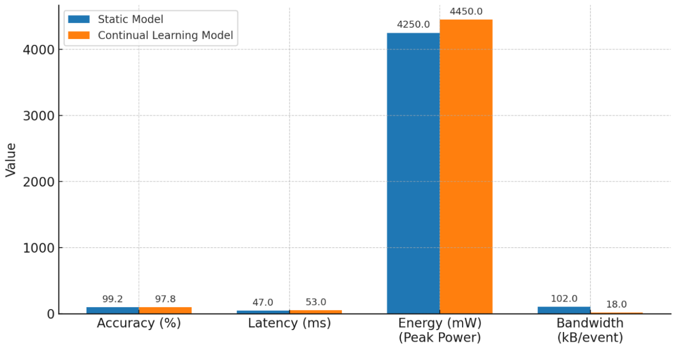

Fig. 8: Multi-Metric Performance Comparison Between Static and Continual Learning Models

Privacy, Security, and Sustainability

The proposed system offers privacy and security since all the audio event detection and continuous learning procedures take place purely on the device itself without raw audio data going off the local network. This removes the dangers of relaying sensitive environmental or conversation audio through the general channels where such information may be overheard by strangers. The solution design is energy efficient, optimized to operate under low latency conditions; great application in battery powered IoT nodes, which will allow the solution to be long deployed without need of frequent recharging, or replacement of batteries. Additionally, having less reliance upon cloud-based infrastructure to process and update models, the system reduces the network traffic considerably and the server workloads that are part of environmental sustainability as cloud data centres consume energy at a higher level when working with huge amounts of data.

Conclusion & Future Work

Conclusion:

This case study shows the feasibility and the advantages of on-device continual learning to learn embedded Acoustic Event Detection (AED) in industrial Internet of Things settings. With the help of NVIDIA Jetson Nano and beamforming microphone array, the system attained high accuracy with the ability to adapt to the new event without having to connect to the cloud. The learning-heading model was able to accommodate concept drift, maintain performance on existing classes and conserve bandwidth, hence it is specially designed to fit in applications that are bandwidth constrained and security sensitive. The small penalties in latency and power consumption were overridden by the enormous advantages in flexibility, robustness and functional autonomy even under adversity in the environment like background noise, thermal restricting and poor connectivity, no connectivity.

Future Work:

Future works include scaling this deployment model to size to many large nodes, geographically distributed IoT networks, or synchronized learning of several edge nodes. The concept of multimodal sensing combining an audio stream with visual data will also be studied in order to acquire more robustness in complex environments. Further, federated continual learning will be explored as the way of sharing the knowledge on nodes without providing the raw data such that privacy would be improved, and the adaptation done faster. There will also be some efforts in regards to introducing self-supervised pretraining to eliminate a dependence on large labeled datasets as well, making human annotation less necessary and allowing newer models to update much more quickly and efficiently.

References

1. Piczak, K. J. (2015). Environmental sound classification with convolutional neural networks. *Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 1-6. https://doi.org/10.1109/MLSP.2015.7324337

2. Mesaros, A., Heittola, T., & Virtanen, T. (2016). Metrics for polyphonic sound event detection. *Applied Sciences*, 6(6), 162. https://doi.org/10.3390/app6060162

3. Kong, Y., Li, Z., & Fu, H. (2022). Privacy-preserving audio event detection for IoT using federated learning. *IEEE Internet of Things Journal*, 9(1), 847–859. https://doi.org/10.1109/JIOT.2021.3083029
4. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526. https://doi.org/10.1073/pnas.1611835114
5. Lopez-Paz, D., &Ranzato, M. (2017). Gradient episodic memory for continual learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 6467–6476.
6. Fernando, A., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A. A., Pritzel, A., &Wierstra, D. (2017). PathNet: Evolution channels gradient descent in super ne0ural networks. *arXiv Preprint*, arXiv:1701.08734. https://arxiv.org/abs/1701.08734
7. Gao, Y., Wong, M. L. D., & Lee, T. (2020). Incremental learning for sound event classification. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 626–630. https://doi.org/10.1109/ICASSP40776.2020.9053556
8. Giri, S., Krishnan, A., & Manocha, D. (2020). Efficient real-time sound classification using compact neural networks. *IEEE Signal Processing Letters*, 27, 1460–1464. https://doi.org/10.1109/LSP.2020.3009951
9. Ma, M., Liu, X., & Li, H. (2020). Edge computing for acoustic signal processing: A review. *IEEE Access*, 8, 159–174. https://doi.org/10.1109/ACCESS.2019.2963724
10. Wu, Z., Chen, T., Zhang, X., Sun, G., Wang, L., & Xie, Y. (2021). Energy-efficient deep learning inference for edge devices via model compression and acceleration. *IEEE Internet of Things Journal*, 8(16), 12415–12427. https://doi.org/10.1109/JIOT.2020.3023898
11. Koteshwaramma, K. C., Vijay, V., Bindusree, V., Kotamraju, S. I., Spandhana, Y., Reddy, B. V. D., Charan, A. S., Pittala, C. S., & Vallabhuni, R. R. (2022). ASIC Implementation of an Effective Reversible R2B FFT for 5G Technology Using Reversible Logic. Journal of VLSI Circuits and Systems, 4(2), 5–13. https://doi.org/10.31838/jvcs/04.02.02
12. Arvinth, N. (2024). Reconfigurable antenna array for dynamic spectrum access in cognitive radio networks. National Journal of RF Circuits and Wireless Systems, 1(2), 1–6.
13. Surendar, A. (2025). Lightweight CNN architecture for real-time image super-resolution in edge devices. National Journal of Signal and Image Processing, 1(1), 1–8.
14. Madugalla, A. K., & Perera, M. (2024). Innovative uses of medical embedded systems in healthcare. Progress in Electronics and Communication Engineering, 2(1), 48–59. https://doi.org/10.31838/PECE/02.01.05
15. Sathish Kumar, T. M. (2023). Wearable sensors for flexible health monitoring and IoT. National Journal of RF Engineering and Wireless Communication, 1(1), 10-22. https://doi.org/10.31838/RFMW/01.01.02