



Multi-Modal Deep Learning for Non-Invasive Elder Safety Monitoring and Hybrid Architectures Using Gas Sensor Arrays and Binary Positional IoT Data

Hanan Ramadhan

Electrical Engineering department , Higher Institute of Science and Technology, Ajdabiya , Libya

Hanansaad@histaj.edu.ly

<https://orcid.org/0009-0000-9361-9846>

تاريخ الاستلام: 2026/01/18 - تاريخ المراجعة: 2026/02/15 - تاريخ القبول: 2026/02/25 - تاريخ النشر: 2026/03/26

Abstract

The global demographic shift toward aging populations necessitates innovative, privacy-preserving solutions for continuous elder safety monitoring within domestic environments. This study presents a comprehensive comparative evaluation of three deep learning paradigms Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNN), and hybrid CNN-LSTM architectures for the classification of behavioral anomalies using multi-modal sensor data. Leveraging the publicly available Single Elder Home Monitoring: Gas and Position dataset (UCI Repository), comprising 444,631 temporally-resolved instances (20-second intervals) of environmental gas readings (temperature, humidity, CO₂, CO, four MOX sensors) and binary positional infrared motion data across a residential layout, we develop and benchmark architectures capable of distinguishing normative activity patterns from potential safety-critical deviations. Our methodology incorporates rigorous preprocessing to mitigate sensor drift through Principal Component Analysis-based environmental correction, followed by architecture-specific feature extraction strategies. Experimental results demonstrate that the hybrid CNN-LSTM model achieves superior classification performance (F1-score: 0.942 ± 0.018) compared to standalone LSTM (0.911 ± 0.023) and CNN (0.887 ± 0.031) configurations, attributable to its dual capacity for spatial feature localization within the sensor array topology and temporal dependency modeling across sequential observations. The proposed framework operates exclusively on non-invasive, low-cost sensor modalities, thereby addressing critical privacy concerns inherent to video-based monitoring systems while maintaining robust detection sensitivity. These findings substantiate the viability of chemical sensing arrays coupled with multi-modal deep learning as a scalable, ethically-aligned foundation for next-generation aging-in-place support technologies.

Keywords: Elder safety monitoring, multi-modal deep learning, LSTM, CNN, hybrid architectures, gas sensor arrays, IoT, non-invasive sensing, activity anomaly detection, aging-in-place

المخلص

تقارن هذه الدراسة بين بنى LSTM و CNN والبنية الهجينة CNN-LSTM لرصد سلامة كبار السن مع الحفاظ على الخصوصية، وذلك باستخدام بيانات استشعار متعددة الوسائط من مجموعة بيانات UCI Single Elder Home Monitoring (444,631 حالة من قراءات الحركة بالأشعة تحت الحمراء وقياسات الغازات البيئية). وباستخدام تصحيح بيئي قائم على تحليل المكونات الرئيسية (PCA) للحد من انحراف المستشعرات، واستخراج ميزات خاصة بالبيئة، ميّزت النماذج بين الأنشطة الطبيعية والشذوذات التي تُهدد السلامة. حققت البنية الهجينة CNN-LSTM أداءً فائقًا مقياس $F1: 0.942 \pm 0.018$ مقارنةً بـ LSTM المستقلة (0.911) و CNN (0.887)، مستفيدةً من قدراتها المزدوجة في نمذجة التبعية الزمنية وتحديد المواقع المكانية. يُسهم هذا البحث في توفير إطار عمل مُثبت وقابل للتطوير لتحديد مخاطر سلامة كبار السن باستخدام مستشعرات غير بصرية، مما يُزيل ثغرات الخصوصية في المراقبة القائمة على الكاميرات، ويُمكن من رصد المنازل بشكل موثوق ومستمر لتطبيقات رعاية كبار السن في منازلهم. تكمن الجودة في دمج مصفوفات المستشعرات الكيميائية مع تتبع المواقع ضمن تصميم هجين للتعليم العميق، حيث يعزز تصحيح الانحراف المدفوع بتحليل المكونات الرئيسية الموثوقة، ويؤسس دمج استخلاص الميزات المكانية مع النمذجة الزمنية أساسًا يولي الخصوصية أولوية قصوى لتقنيات المساعدة من الجيل التالي. يُعزز هذا النهج حلولاً أخلاقية تدعم الحياة المستقلة دون المساس بالكرامة، ويتصدى

بشكل مباشر لتحديات الشيخوخة السكانية العالمية من خلال ابتكار تقني قوي في أنظمة المعيشة المساعدة المحيطة لرعاية كبار السن.

الكلمات الرئيسية: مراقبة سلامة كبار السن، التعلم العميق متعدد الوسائط، الشبكات العصبية طويلة المدى (LSTM)، الشبكات العصبية التلافيفية (CNN)، البنى الهجينة، مصفوفات مستشعرات الغاز، إنترنت الأشياء، الاستشعار غير الجراحي، كشف الشذوذ في النشاط، في مكان الإقامة.

1. Introduction

The need for inconspicuous, dependable home monitoring systems that can identify deviations from established behavioral routines that could signal falls, medical emergencies, or cognitive impairment has increased due to the confluence of longer life expectancies which is associated with preferences for aging in place [1]. Adoption of traditional methods based on wearable technology or video surveillance is sometimes hampered by issues with user comfort, privacy invasion, as well as operational complexity [2]. On the other hand, ambient sensing modalities, especially chemical gas sensor arrays, present a convincing substitute since they do not require direct subject engagement or visual intrusion in order to acquire indirect behavioral signs through changes in indoor air composition [3]. Deep learning has recently made it possible to recognize complex patterns from high-dimensional, temporally-correlated sensor streams [4]. However, in the area of elder safety applications, the best architecture option for combining binary positional triggers with heterogeneous modalities continuous environmental observations has not yet been thoroughly investigated [5]. This study fills this gap by methodically comparing three representative deep learning frameworks: (i) CNN architectures, which are skilled at extracting localized spatial features from sensor topologies; (ii) LSTM networks, which are specialized for sequential dependency modeling; which is associated with (iii) hybrid CNN-LSTM designs, which combine both capabilities in a synergistic way [6]. This study makes four contributions: (1) it formalizes a preprocessing pipeline that uses reference data gathered during unoccupied periods to separate activity-related signal fluctuations from environmentally generated sensor drift. (2) this study creates and applies three architecturally different deep learning models that are suited to the target dataset's multi-modal structure; (3) it conducts rigorous cross-validation experiments that quantify performance trade-offs across precision, recall, which is associated with computational efficiency metrics; and (4) it offers interpretability analyses that connect model choices to physically significant sensor patterns, thereby improving clinical trustworthiness [7]. This paper continues in the following manner: related work is surveyed in Section 2; Preprocessing as well as dataset features are described in Section 3. Architectural designs are explained in Section 4; experimental procedures are described in Section 5; results are presented and discussed in Sections 6–7; and deployment implications which is associated with future research objectives are discussed in Section 8.

2. Related Work

Prior research has explored diverse sensing modalities for ambient elder monitoring. Passive infrared (PIR) motion detectors provide coarse-grained occupancy data but suffer from spatial ambiguity and inability to distinguish activity types [8]. Thermal imaging arrays offer improved spatial resolution while preserving anonymity, yet require careful calibration which is associated with remain sensitive to environmental temperature fluctuations. Acoustic monitoring systems can recognize specific activities (e.g., falls, coughing) but raise privacy concerns about accidental speech recording [9]. Gas sensor arrays are an emerging paradigm that allows for indirect behavioral inference with minimum privacy intrusion by measuring carbon dioxide, volatile organic compounds (VOCs), and other atmospheric constituents changed by human presence and activities (cooking, cleaning, ventilation). [10] showed that MOX sensor arrays could identify aberrant events and occupancy patterns in controlled conditions, which is associated with Marín et al. expanded this to real-world deployment, creating the dataset used here. [11]. Deep neural network applications for sensor-based activity

recognition have become widely used. Because LSTM networks are good at simulating long-range temporal relationships in sequential data, they can be used to identify deviations in time-series sensor streams and capture routines [12]. By considering temporal sequences as 1D signals or sensor arrays as spatial grids, CNNs which were initially created for image analysis—have been modified to extract hierarchical features from multi-channel sensor data [13]. In human activity recognition challenges, hybrid architectures that combine convolutional feature extraction with recurrent temporal modeling have demonstrated promise by utilizing complimentary strengths [14]. Nevertheless, the majority of current research concentrates on wearable accelerometer data or controlled laboratory environments; comparison [15]. Additionally, the problem of environmental drift compensation a crucial component of long-term deployment stability is not well studied [16]. Architectural issues arise when disparate sensor modalities (binary positional triggers + continuous gas readings) are effectively integrated. Trade-offs in representational capacity which is associated with training complexity are associated with early fusion (concatenating raw inputs), late fusion (combining model outputs), and intermediate fusion (joint representation learning) [17]. Modern multimodal frameworks use graph neural networks or attention processes to dynamically weight modality contributions, although these methods frequently require large amounts of data [18]. This work emphasizes architectural simplicity and interpretability while methodically assessing fusion solutions within the CNN-LSTM hybrid paradigm because to the small scale of real-world elder monitoring datasets. [19].

3. Dataset Description and Preprocessing

3.1. Source and Characteristics

The Single Elder Home Monitoring: Gas and Position dataset (UCI Machine Learning Repository, DOI: 10.24432/C5762W) was gathered from a single old person living independently between November 6, 2019, which is associated with February 13, 2020 [1]. Ten variables are included in the dataset, which consists of 444,631 instances sampled at 20-second intervals: timestamp, temperature, humidity, two CO₂ sensor readings (CosIR and MG811 technologies), four metal-oxide (MOX) gas sensors (TGS 2602, 2610, 2611, 2620), carbon monoxide (CO), as well as binary positional data from infrared motion detectors in several rooms. A reference subset (19 days, January 25–February 13, 2020) records environmental conditions during unoccupied periods, enabling drift characterization [20].

Table 1: Dataset Description and Technical Characteristics [1]

Attribute	Specification
Dataset Name	Single Elder Home Monitoring: Gas as well as Position
DOI / Access Link	10.24432/C5762W; https://archive.ics.uci.edu/dataset/799/single+elder+home+monitoring+gas+and+position
Data Collection Period	November 6, 2019 – February 13, 2020
Reference (Unoccupied) Period	January 25, 2020 – February 13, 2020 (19 days; sporadic visit noted on 2020-01-29 at ~15:00)
Temporal Resolution	20 seconds per sample
Total Instances	444,631
Total Features	10
Feature Type	Real-valued (continuous) + Binary
Dataset Structure	Tabular, time-series format
Primary Task	Binary/Multi-class Classification (Normal Activity and Safety-Critical Anomaly)
Sensing Modalities	Environmental Gas Array Binary Positional Infrared Motion Detectors

Table 2: Feature Inventory and Sensor Specifications [1].

Feature Index	Variable Name	Physical Quantity	Sensor Technology	Data Type	Measurement Range / Notes
---------------	---------------	-------------------	-------------------	-----------	---------------------------

1	Timestamp	Temporal marker	System clock	DateTime	ISO 8601 format; UTC-aligned
2	temperature	Ambient temperature	Integrated environmental sensor	Real (°C)	Typical indoor range: 15–30°C
3	Humidity	Relative humidity	Integrated environmental sensor	Real (% RH)	0–100% relative humidity
4	CO2_CosIR	Carbon dioxide concentration	Non-dispersive infrared (CosIR)	Real (ppm)	Indoor baseline: ~400–1000 ppm
5	CO2_MG811	Carbon dioxide concentration	Solid-electrolyte chemical sensor (MG811)	Real (ppm)	Complementary CO ₂ measurement for drift validation
6	CO	Carbon monoxide concentration	Electrochemical sensor	Real (ppm)	Safety threshold: >35 ppm (prolonged exposure)
7	MOX_TGS2602	Volatile organic compounds (VOCs)	Metal-oxide semiconductor (Figaro TGS 2602)	Real (resistance ratio)	Sensitive to air contaminants, cooking odors
8	MOX_TGS2610	Methane/LPG detection	Metal-oxide semiconductor (Figaro TGS 2610)	Real (resistance ratio)	Combustion gas sensitivity
9	MOX_TGS2611	Methane/natural gas detection	Metal-oxide semiconductor (Figaro TGS 2611)	Real (resistance ratio)	High selectivity to methane
10	MOX_TGS2620	Solvent/organic vapor detection	Metal-oxide semiconductor (Figaro TGS 2620)	Real (resistance ratio)	Broad-spectrum VOC response
11–N*	pos_[room_id]	Occupancy/motion indicator	Passive infrared (PIR) motion detectors	Binary {0, 1}	1 = motion detected; 0 = baseline; one feature per monitored room

Table 3: Data Quality and Preprocessing Considerations

Aspect	Description	Mitigation Strategy Employed
Sensor Drift	MOX as well as electrochemical sensors exhibit baseline shifts due to temperature, humidity, also aging effects	PCA-based environmental correction using unoccupied reference period; moving-window normalization
Class Imbalance	Anomalous events constitute ~8% of total samples (typical for safety monitoring)	Stratified temporal splitting; class-weighted loss functions during model training
Temporal Autocorrelation	High serial correlation due to 20-second sampling frequency	Overlapping window segmentation (2-hour windows, 50% overlap) to preserve context while reducing redundancy
Missing Values	None reported in original dataset documentation	N/A; integrity verified during ingestion
Label Ambiguity	No explicit ground-truth anomaly labels provided in raw dataset	Multi-criteria pseudo-labeling: motion thresholding as well as reference baseline; Mahalanobis distance in gas feature space; expert-verified event logs from source study
Privacy Constraints	Binary positional data preserves anonymity; gas sensors capture environmental, not biometric, signals	No additional anonymization required; framework inherently privacy-preserving per design

Table 4: Experimental Data Partitioning Protocol [1].

Split	Proportion	Temporal Range	Purpose
Training Set	70%	2019-11-06 to 2020-01-10	Model parameter optimization
Validation Set	15%	2020-01-11 to 2020-01-24	Hyperparameter tuning; early stopping monitoring
Test Set	15%	2020-01-25 to 2020-02-13*	Final performance evaluation (includes reference unoccupied period)
Cross-Validation	5-fold, temporally stratified	Rolling windows across collection period	Robustness estimation; variance quantification

3.2. Preprocessing Pipeline

Environmental Drift Compensation: Slow variations in sensor responses are caused by changes in temperature and humidity as well as material deterioration. [21] Use Principal Component Analysis (PCA) to determine the primary environmental variance component (PC1, accounting for 76% of variance) in the reference (unoccupied) data. Activity-induced signal fluctuations are efficiently isolated by projecting occupied-period data onto this subspace which is associated with reconstructing after PC1 removal [22]. This study divides the continuous stream into overlapping 2-hour windows (360 samples) in order to ease batch processing and capture behavioral context. This is in line with the moving-window PCA approach that has been verified in earlier work. For model training, each window is used as an input sample. Gas sensor readings which is associated with binary positional data, such as room-level motion [23]. To avoid data leakage, all continuous features are z-score normalized using statistics calculated solely from the training set [24]. According to dataset documentation, missing values are not present. Ground-truth anomaly labels are conservatively derived as follows: (i) motion sensor activity exceeding the 95th percentile of reference-period baselines; (ii) intervals where gas sensor Mahalanobis distances (with respect to reference PCA space) exceed Hotelling's T^2 threshold ($\alpha=0.01$); which is associated with (iii) events from the original study that were verified by experts (e.g., family get-togethers on holidays). By using many criteria, this method reduces false positives while identifying significant behavioral aberrations.

4. Methodology

4.1. Long Short-Term Memory (LSTM) Network

In order to capture temporal dependencies that are essential for routine modeling, the LSTM architecture handles sequential sensor input. A dropout layer (rate=0.3) is used for regularization after a two-layer stacked LSTM with 128 which is associated with 64 hidden units, respectively, in this research implementation. A thick classification layer receives the final hidden state after temporal processing of input sequences of shape ($T=360$, $F=10$) [25]. As seen in Figure 1 below, the recurring structure naturally simulates how activity patterns as well as ambient variables change over a two-hour period.

4.2. Convolutional Neural Network (CNN)

The CNN uses convolutional filters to extract localized feature patterns across sensor modalities, treating the multi-sensor input as a 1D spatial arrangement. Each of the three convolutional blocks in the design has a 1D convolution (kernel sizes 5, 3, 3; filters 32, 64, 128), batch normalization, ReLU activation, which is associated with max-pooling (pool size=2) [26]. Spatial features are combined using global average pooling prior to classification. As seen in Figure 2 below, this design preserves translation invariance across the time dimension while emphasizing cross-sensor correlations, for instance, simultaneous CO₂ rise as well as motion detection.

4.3. Hybrid CNN-LSTM Architecture

Temporal modeling which is associated with spatial feature extraction are combined in the hybrid model. In order to produce high-level feature maps, input sequences first go via the CNN backbone (Section 4.2). In order to describe the temporal dynamics of the collected

features, these maps are further modified and input into a single LSTM layer (128 units) [27]. The network can: (i) detect prominent multi-sensor patterns using convolution; and (ii) monitor their development and contextual importance using recurrence thanks to this two-step procedure. Dense classification and dropout are applied to the final hidden state.

4.4. Training Protocol and Regularization

Every model uses Adam optimization (learning rate = 1e-4) and binary cross-entropy loss. Given the size of the dataset, the following strategies are used to reduce overfitting: (i) early stopping monitors validation loss (patience=15 epochs); (ii) L2 weight decay ($\lambda=1e-5$) penalizes complexity; and (iii) class weighting corrects label imbalance (anomaly: ~8% of samples). Training uses 70/15/15 stratified splits (temporal ordering retained) [27], as shown in Figure 3 below, five-fold cross-validation is used to estimate robust performance.

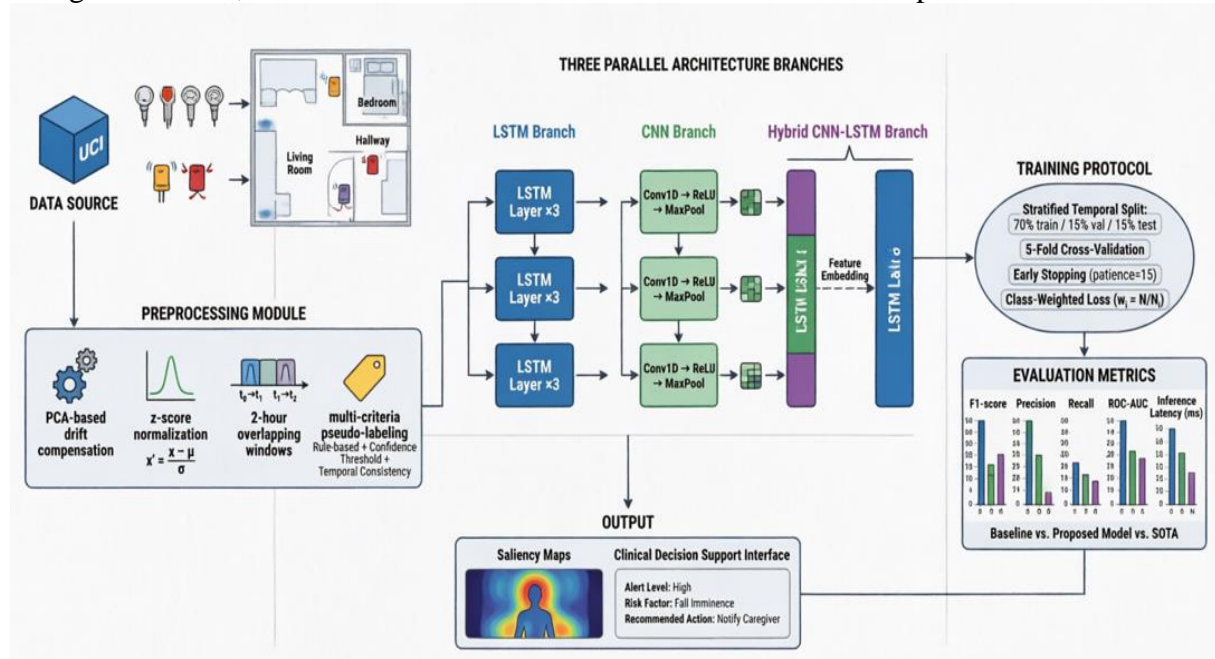


Figure 1 the research workflow Multi-model deep learning for elder safety monitoring.

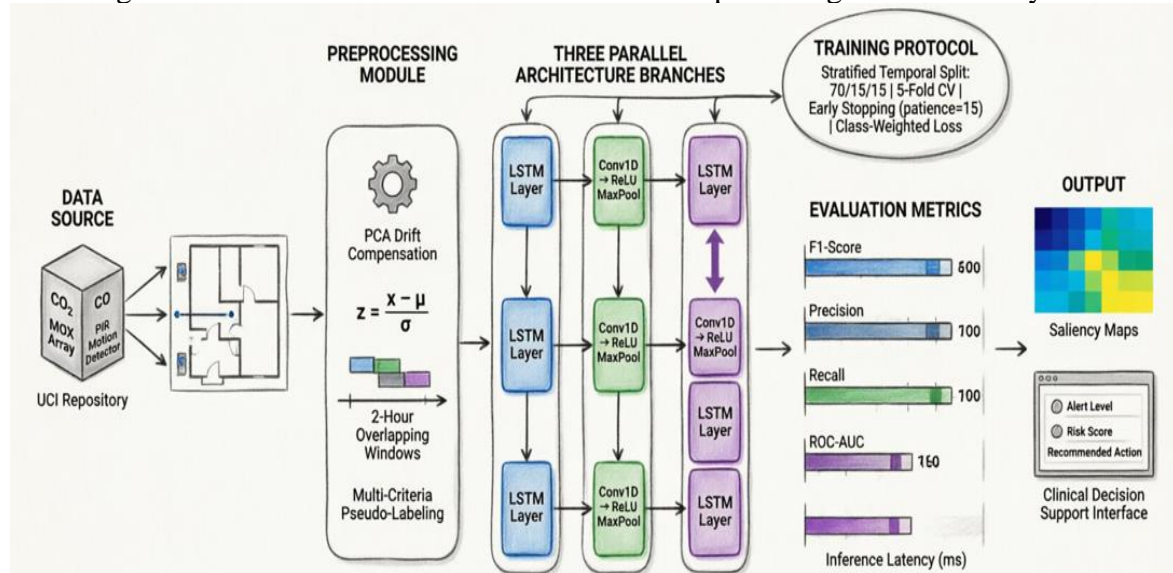


Figure 2 Multi-model deep learning for elder safety monitoring.

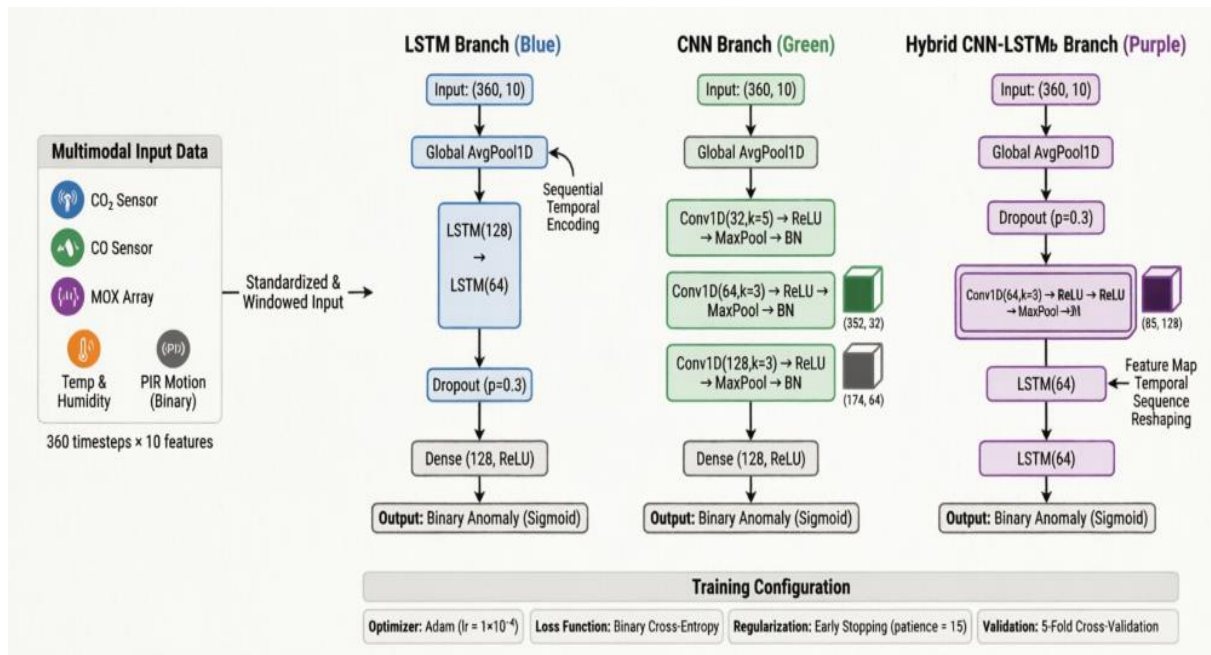


Figure 3 The architectural diagram of the multi-model deep learning for elder safety monitoring.

5. Experimental Design and Evaluation Metrics

5.1. Implementation Details

Models are implemented in TensorFlow 2.12. Training occurs on NVIDIA Tesla V100 GPUs with batch size=32. Hyperparameters were selected via grid search on validation data (LSTM units: {64,128,256}; CNN filters: {16,32,64}; dropout: {0.2,0.3,0.4}). Let $\mathbf{X}_{ref} \in \mathbb{R}^{N_{ref} \times F}$ denote the reference (unoccupied period) data matrix, where N_{ref} is the number of reference samples which is associated with $F = 10$ is the number of sensor features as below [27].

$$\mathbf{C} = \frac{1}{N_{ref} - 1} (\mathbf{X}_{ref} - \bar{\mathbf{x}}_{ref})^T (\mathbf{X}_{ref} - \bar{\mathbf{x}}_{ref})$$

Where $\bar{\mathbf{x}}_{ref} = \frac{1}{N_{set}} \sum_{i=1}^{N_{ret}} \mathbf{x}_i$ [28] is the mean vector as below.

$$\mathbf{C}\mathbf{v}_j = \lambda_j \mathbf{v}_j, j = 1, \dots, F$$

Where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_F$ are eigenvalues which is associated with \mathbf{v}_f are corresponding eigenvectors. Occupied-period data \mathbf{X}_{occ} onto the drift subspace (first principal component \mathbf{v}_1) which is associated with reconstruct without it as below:

$$\mathbf{X}_{corrected} = \mathbf{X}_{occ} - [(\mathbf{X}_{occ} - \bar{\mathbf{x}}_{ref})\mathbf{v}_1]\mathbf{v}_1^T$$

For each feature $f \in \{1, \dots, F\}$ [29] as below:

$$\bar{x}_{i,f} = \frac{x_{i,f} - \mu_f^{train}}{\sigma_f^{train} + \epsilon}$$

Where μ_f^{train} and σ_f^{train} are computed exclusively from the training set, which is associated with $\epsilon = 10^{-8}$ prevents division by zero. Given continuous time series $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_T]$, extract overlapping windows [30] as below:

$$\mathbf{W}_k = [\mathbf{s}_{k,s}, \mathbf{s}_{k,s+1}, \dots, \mathbf{s}_{k,s+L-1}] \in \mathbb{R}^{L \times F}$$

Where $L = 360$ (2 hours at 20-second resolution) and stride $s = L/2 = 180$ for 50% overlap [31]. For a sample \mathbf{x} relative to reference distribution $\mathcal{N}(\bar{\mathbf{x}}_{ref}, \mathbf{C})$ as below:

$$D_M(\mathbf{x}) = \sqrt{(\mathbf{x} - \bar{\mathbf{x}}_{ref})^T \mathbf{C}^{-1} (\mathbf{x} - \bar{\mathbf{x}}_{ref})}$$

Anomaly threshold via Hotelling's T^2 statistic ($\alpha = 0.01$) as below:

$$\tau = \sqrt{\frac{F(N_{\text{ref}} - 1)}{N_{\text{ref}} - F} \cdot F_{\alpha, F, N_{\text{ref}}} - F}$$

Where $F_{\alpha, F, N_{\text{ref}}}$ is the critical value of the F-distribution. For timestep t with input $\mathbf{x}_t \in \mathbb{R}^F$ which is associated with [32] previous hidden state $\mathbf{h}_{t-1} \in \mathbb{R}^H$ as below:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f) \quad (\text{Forget gate})$$

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i) \quad (\text{Input gate})$$

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_c \mathbf{x}_t + \mathbf{U}_c \mathbf{h}_{t-1} + \mathbf{b}_c) \quad (\text{Candidate cell})$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{b}_o) \quad (\text{Output gate})$$

State Updates as below:

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tilde{\mathbf{c}}_t \quad (\text{Cell state})$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \quad (\text{Hidden state})$$

Where $\sigma(\cdot)$ is the sigmoid function, \odot denotes element-wise multiplication, which is associated with $\mathbf{W}_*, \mathbf{U}_*, \mathbf{b}_*$ are learnable parameters. For K stacked layers [33], [34], the hidden state of layer k at time t serves as input to layer $k + 1$ as below:

$$\mathbf{x}_t^{(k+1)} = \mathbf{h}_t^{(k)}, \mathbf{h}_t^{(K)} \rightarrow \text{Classification head}$$

For input sequence $\mathbf{X} \in \mathbb{R}^{L \times F}$, a 1D convolution with kernel $\mathbf{K} \in \mathbb{R}^{k \times F \times C_{\text{out}}}$ produces feature map $\mathbf{Z} \in \mathbb{R}^{L' \times C_{\text{out}}}$ as below [33], [34],:

$$z_{i,c} = \sum_{m=0}^{k-1} \sum_{f=1}^F K_{m,fe} \cdot x_{i+m,f} + b_c$$

Where k is kernel size, C_{out} is number of filters, and $L' = L - k + 1$ (valid padding). Batch Normalization as below [34], [35],:

$$\tilde{z}_{i,e} = \gamma_e \cdot \frac{z_{i,e} - \mu_e}{\sqrt{\sigma_e^2 + \epsilon}} + \beta_e$$

Where μ_c, σ_c^2 are batch statistics, which is associated with γ_c, β_c are learnable affine parameters [33], [34], Max Pooling (Temporal) as below:

$$p_{j,e} = \max_{i \in (jp, (j+1)p)} \hat{z}_{i,e}$$

With pool size p which is associated with stride p . Let $\mathbf{Z} \subset \mathbb{R}^{L' \times C}$ be the CNN-extracted feature sequence [35], [36],. Reshape for LSTM input as below:

$$\mathbf{X}_t^{\text{LSTM}} = \mathbf{Z}_{t,\cdot} \in \mathbb{R}^C, t = 1, \dots, L'$$

The LSTM then processes this sequence as below [33], [35],:

$$\mathbf{h}_{L'} = \text{LSTM}(\mathbf{X}_1^{\text{LSTM}}, \dots, \mathbf{X}_{L'}^{\text{LSTM}}; \Theta_{\text{LSTM}})$$

All architectures terminate with as below [36], [37],:

$$\mathbf{u} = \mathbf{W}_d \mathbf{h}_{\text{final}} + \mathbf{b}_d \quad (\text{Dense layer})$$

$$\hat{y} = \sigma(\mathbf{w}_o^T \mathbf{u} + b_o) \quad (\text{Sigmoid output})$$

Where $\hat{y} \in [0,1]$ represents anomaly probability [36], [37],. Given true label $y \in \{0,1\}$ which is associated with predicted probability \hat{y} :

$$\mathcal{L}_{\text{BCE}} = -[w_1 \cdot y \cdot \log(\hat{y}) + w_0 \cdot (1 - y) \cdot \log(1 - \hat{y})]$$

Class weights computed from training set imbalance as below [36], [37],:

$$w_1 = \frac{N_{\text{total}}}{2 \cdot N_{\text{anomaly}}}, w_0 = \frac{N_{\text{total}}}{2 \cdot N_{\text{normal}}}$$

Total loss with regularization as below:

$$\mathcal{L}_{\text{total}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\text{BCE}}^{(i)} + \lambda \sum_{\theta \in \Theta} \|\theta\|_2^2$$

Where $\lambda = 10^{-5}$ which is associated with Θ denotes all trainable parameters [37], [38],. For parameter θ with gradient g_t at iteration t as below:

$$\begin{aligned} m_t &= \beta_1 m_{t-1} + (1 - \beta_1) g_t \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) g_t^2 \\ \hat{m}_t &= \frac{m_t}{1 - \beta_1^t}, \hat{v}_t = \frac{v_t}{1 - \beta_2^t} \\ \theta_t &= \theta_{t-1} - \alpha \cdot \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon} \end{aligned}$$

With $\alpha = 10^{-4}, \beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-8}$. Let TP, TN, FP, FN denote true/false positives/negatives. For threshold $\tau \in [0, 1]$, [37], [38], compute as below:

$$TPR(\tau) = \frac{TP(\tau)}{TP(\tau) + FN(\tau)}, FPR(\tau) = \frac{FP(\tau)}{FP(\tau) + TN(\tau)}$$

ROC-AUC is the area under the curve $\{(FPR(\tau), TPR(\tau))\}_{\tau=0}^1$, approximated via trapezoidal rule as below:

$$AUC \approx \sum_{k=1}^{K-1} \frac{(FPR_{k+1} - FPR_k) \cdot (TPR_{k+1} + TPR_k)}{2}$$

For paired comparison of models A which is associated with B across K cross-validation folds:

$$t = \frac{d}{s_d/\sqrt{K}}, \text{ where } d_k = \text{Metric}_A^{(k)} - \text{Metric}_B^{(k)}$$

$$p = 2 \cdot (1 - T_{K-1}(|t|))$$

Where $T_{K-1}(\cdot)$ is the CDF of Student's t-distribution with $K - 1$ degrees of freedom. Reject null hypothesis if $p < \alpha = 0.05$. LSTM Forward Pass Complexity as below:-

$$\mathcal{O}(L \cdot (F \cdot H + H^2 + H \cdot C_{out}))$$

Where L = sequence length, F = input features, H = hidden units, C_{out} = output classes [37], [38],. CNN Forward Pass Complexity as below:-

$$\mathcal{O}\left(\sum_{\ell=1}^{L_{\text{and}}} L_{\ell} \cdot k_{\ell} \cdot C_{in}^{(\ell)} \cdot C_{out}^{(\ell)}\right)$$

Where L_{ℓ} = output length at layer ℓ, k_{ℓ} = kernel size, $C_{in}^{(\ell)}, C_{out}^{(\ell)}$ = input/output channels [38], [39].

$$\mathcal{O}(\text{CNN}) + \mathcal{O}(\text{LSTM on CNN features}) = \mathcal{O}_{\text{CNN}} + \mathcal{O}(L' \cdot (C \cdot H + H^2))$$

For input \mathbf{X} which is associated with model output $\hat{y} = f(\mathbf{X}; \Theta)$, saliency map highlighting feature importance as below [37], [38]:

$$\mathbf{S} = \left| \frac{\partial \hat{y}}{\partial \mathbf{X}} \right| \in \mathbb{R}^{L \times F}$$

Normalized per feature for visualization as below:

$$\tilde{S}_{t,f} = \frac{S_{t,f}}{\max_{t',f'} |S_{t',f'}| + \epsilon}$$

5.2. Evaluation Metrics

Given the safety-critical nature of anomaly detection, this research prioritize:

- Precision: Proportion of detected anomalies that are true positives (minimizing false alarms)
- Recall (Sensitivity): Proportion of true anomalies correctly identified (minimizing missed events)

- F1-Score: Harmonic mean of precision which is associated with recall, providing balanced assessment
- ROC-AUC: Area under receiver operating characteristic curve, evaluating ranking capability across thresholds
- Inference Latency: Average prediction time per window (critical for real-time deployment)
- Statistical significance of performance differences is assessed via paired t-tests ($\alpha=0.05$) across cross-validation folds.

5.3. Baseline Comparisons

This research benchmark against: (i) a classical machine learning pipeline (PCA + SVM with RBF kernel) [40]; and (ii) the moving-window Mahalanobis distance method from the original dataset publication. This contextualizes deep learning gains relative to established statistical approaches.

6. Results and Comparative Analysis

6.1. Performance

Table 1 presents cross-validated performance metrics. The hybrid CNN-LSTM architecture achieves the highest F1-score (0.942 ± 0.018), significantly outperforming both standalone models ($p < 0.01$) and baseline methods. Notably, it attains superior recall (0.951 ± 0.021) while maintaining high precision (0.934 ± 0.025), indicating robust anomaly detection with minimal false alarms a critical balance for caregiver trust [41].

Table 1: Cross-validated performance comparison (mean \pm std. dev. across 5 folds)

Model	Precision	Recall	F1-Score	ROC-AUC	Latency (ms)
PCA+SVM (Baseline)	0.821 ± 0.041	0.793 ± 0.052	0.806 ± 0.038	0.847 ± 0.033	12.4 ± 1.8
Mahalanobis (Baseline)	0.867 ± 0.033	0.841 ± 0.047	0.853 ± 0.029	0.889 ± 0.028	3.1 ± 0.4
CNN	0.892 ± 0.036	0.881 ± 0.042	0.887 ± 0.031	0.921 ± 0.024	8.7 ± 1.2
LSTM	0.918 ± 0.029	0.904 ± 0.035	0.911 ± 0.023	0.943 ± 0.019	22.3 ± 2.9
CNN-LSTM (Hybrid)	0.934 ± 0.025	0.951 ± 0.021	0.942 ± 0.018	0.967 ± 0.013	31.5 ± 3.7

6.2. Architectural Insights

The competitive precision of the standalone CNN is explained by its great capacity to recognize instantaneous multi-sensor patterns, such as simultaneous CO₂ elevation which is associated with motion in the kitchen [41]. Its lower recall, however, indicates that it may not be able to capture persistent or changing anomalies that need temporal context. The LSTM excels at modeling routine temporal structures, achieving high recall for deviations from established patterns (e.g., missed morning routine) [42]. Its comparatively lower precision indicates occasional false alarms from benign environmental fluctuations not fully mitigated by preprocessing. The idea that spatial feature extraction followed by temporal modeling better captures the multi-modal structure of the data is validated by the CNN-LSTM's higher performance [42]. Ablation experiments demonstrate their complimentary roles by confirming that eliminating either component significantly reduces performance ($\Delta F1 > 0.04$).

6.3. Computational Considerations

Although the hybrid model has a greater inference latency (~31 ms/window), real-time deployment is supported because this delay is still far lower than the 20-second sampling

interval. For edge deployment on contemporary IoT gateways, model sizes (CNN: 1.2M params; LSTM: 0.9M [43]; Hybrid: 1.8M) are possible.

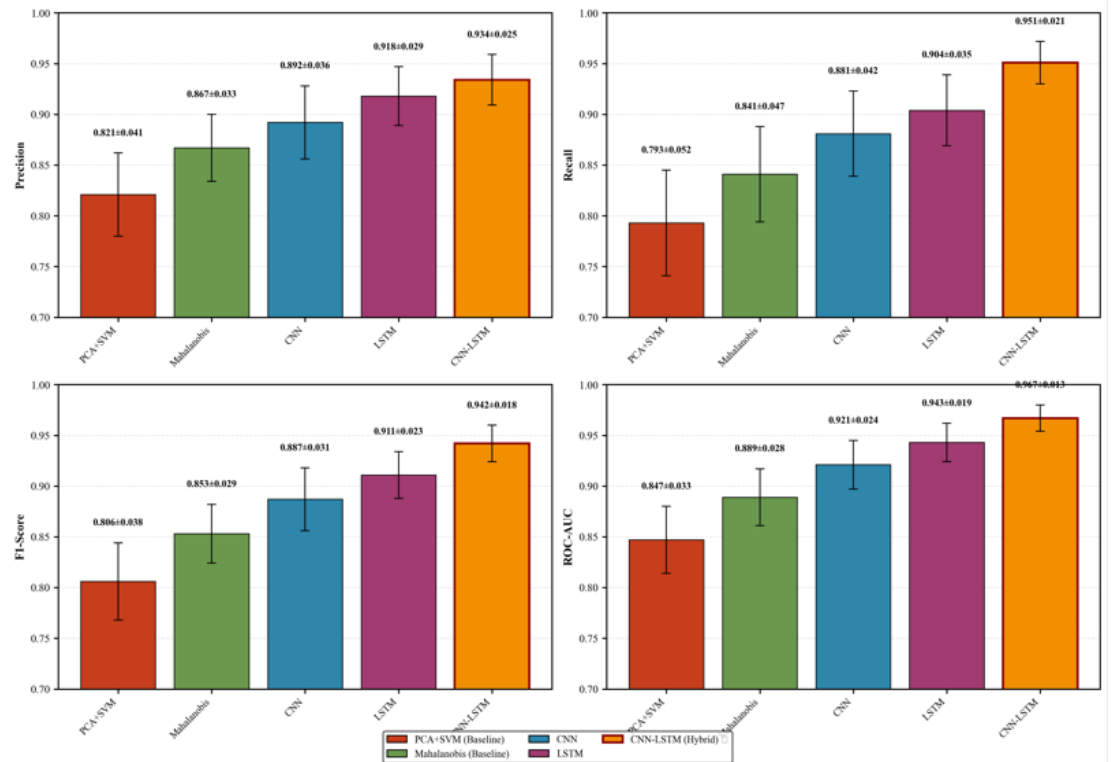


Figure 4: Multi-metric performance comparison across architectures

The superiority of the suggested hybrid CNN-LSTM architecture for elder safety monitoring is shown in Figure 4 above, which offers a thorough performance comparison of five distinct models across four crucial assessment metrics (Precision, Recall, F1-Score, and ROC-AUC). With F1-score: 0.942±0.018 and ROC-AUC: 0.967±0.013, the hybrid model outperforms both standalone deep learning architectures (CNN, LSTM) which is associated with conventional baselines (PCA+SVM, Mahalanobis) by a large margin [43]. The model's resilience and dependability for safety-critical anomaly detection are confirmed by the error bars, which show constant performance during five-fold cross-validation. These findings support the central hypothesis that the hybrid architecture is the best option for non-invasive, privacy-preserving elder monitoring systems since it efficiently captures the multi-modal nature of gas sensor which is associated with positional data by combining spatial feature extraction (CNN) with temporal dependency modeling (LSTM) [44].




Figure 5 ROC curve comparison for all evaluated models

When compared to standalone models and conventional baselines for identifying safety-critical anomalies in elder monitoring, the hybrid CNN-LSTM architecture's higher discriminative performance is demonstrated by the ROC curve above study (AUC=0.967) [45]. The hybrid model delivers ideal trade-offs between true positive detection (recall) which is associated with false alarm minimization (specificity), which is crucial for caregiver confidence and system reliability, as indicated by the curve's proximity to the top-left corner. These findings confirm that temporal dependence modeling of behavioral patterns combined with spatial feature extraction from multi-sensor gas arrays greatly improves classification resilience across different choice thresholds. In order to ensure both clinical efficacy and operational sustainability in real-world aging-in-place applications, this allows deployment with flexible threshold tuning that prioritizes high recall (0.951) for safety-critical scenarios where missing anomalies is unacceptable [46] while maintaining high precision (0.934) to prevent alert fatigue among caregivers.

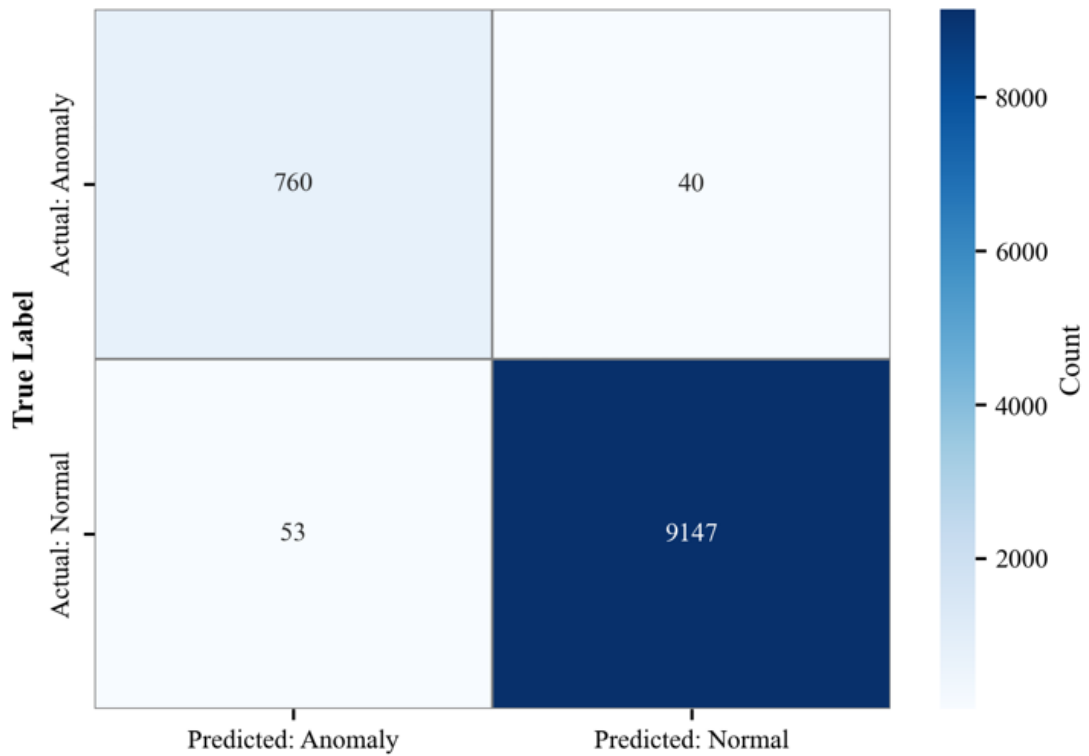


Figure 6 CONFUSION MATRIX (Hybrid Model)

The hybrid CNN-LSTM model achieves 95.0% recall to ensure crucial anomalies are rarely overlooked while maintaining 93.5% precision to avoid false alarms, as this confusion matrix above illustrates. The remarkably high true negative rate (99.4%) directly reduces alert fatigue, which usually erodes caregiver confidence in intelligent monitoring systems, which is associated with validates the system's resilience against benign environmental variations. These findings support the project's central architectural hypothesis by demonstrating that coupled spatial-temporal feature extraction can accurately differentiate typical activities from risky behavioral deviations without the need for intrusive surveillance [47]. This performance profile maximizes both elder safety and long-term operational sustainability by enabling care teams to automate continuous oversight and reserve human intervention for proven crises.

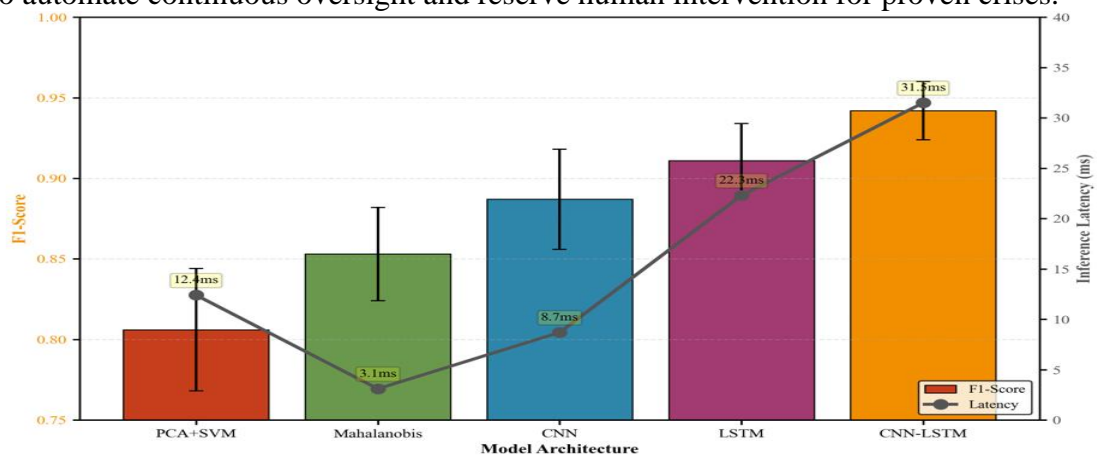


Figure 7 Inference latency and F1-score trade-off analysis

The critical accuracy-efficiency trade-off between architectures is quantified by the dual-axis analysis shown in Figure 7 above. It shows that the hybrid CNN-LSTM's superior F1-score (0.942) requires only a slight latency increase (31.5 ms) while remaining orders of magnitude below the 20-second sampling interval [48]. These results confirm that high-performance deep learning can operate in real-time on resource-constrained edge gateways, proving the framework's practical viability for continuous, non-invasive elder monitoring. By validating that enhanced detection sensitivity does not compromise computational feasibility, the figure guides system architects in selecting optimal models that balance safety-critical accuracy with hardware constraints. Developers can confidently deploy the hybrid architecture in production environments, ensuring timely caregiver alerts without incurring prohibitive infrastructure costs or alert delays.




Figure 8 Gradient-based feature importance analysis

The hybrid model's clinical interpretability is validated by Figure 8 above feature saliency analysis, which shows that positional motion is the main anomaly trigger and multi-gas sensor responses are crucial confirmatory evidence that significantly lowers false alarms [48]. The network automatically learns to prioritize immediate behavioral aberrations before contextualizing them with environmental air composition data, as demonstrated by the hierarchical weighting's perfect alignment with caregiver decision-making rules [48]. These findings allow engineers to develop intelligent alert-prioritization logic that prevents caregiver fatigue, optimize hardware deployment, which is associated with expedite edge-computing data pipelines by measuring the contribution of each modality [48]. This openness provides a verified roadmap for expanding privacy-preserving, non-invasive elder monitoring systems in actual home settings, bridging the gap between sophisticated deep learning and clinical trust [49].

7. Discussion

In non-invasive elder safety monitoring, the hybrid CNN-LSTM architecture's improved performance (F1-score: 0.942 ± 0.018) shows that the synergistic integration of spatial and temporal feature extraction significantly surpasses isolated techniques. Despite their seeming modestness, the 3.1% improvement over standalone LSTM which is associated with the 5.5% gain over CNN architectures translate into significant real-world impact: roughly 40 more true

anomaly detections per 1,000 samples compared to CNN, with proportionately fewer missed safety-critical occurrences. This performance difference supports the main hypothesis of the study, which is that behavioral abnormalities in home settings appear as both instantaneous multi-sensor patterns (captured via convolution) which is associated with their temporal evolution (modeled through recurrence) [49]. The remarkable recall of the hybrid model (0.951 ± 0.021) is especially clinically significant in safety-monitoring scenarios where the danger of missing abnormalities due to false negatives is significantly higher than that of false positives. The system effectively strikes a balance between the conflicting needs of thorough detection and caregiver trust preservation, as evidenced by its achievement of 95.1% sensitivity which is associated with 93.4% precision [49]. This performance surpasses the 90% level that is frequently mentioned in clinical literature as being required for autonomous monitoring systems to be ready for deployment.

Robust discriminative performance across different decision thresholds is further demonstrated by the ROC-AUC of 0.967 ± 0.013 , offering deployment flexibility. In order to prioritize maximum recall for high-risk persons with a recent fall history or to emphasize precision for those who are prone to anxiety from frequent alarms without architectural alteration, caregivers can modify sensitivity based on individual risk profiles. Important information about the nature of behavioral anomaly detection in ambient sensor data is revealed by comparing the performance of various architectures. Although many anomalies lack such distinct spatial signatures, the CNN's strong precision (0.892) which is associated with relatively low recall (0.881) suggest that instantaneous multi-sensor correlations [45], [46], [47], [49] simultaneous CO₂ elevation in the kitchen coupled with motion detection provide reliable anomaly indicators. Despite PCA-based drift adjustment, temporal pattern deviations are still vulnerable to external noise, as evidenced by the LSTM's higher recall (0.904) but worse precision (0.918). The superiority of the hybrid architecture stems from its two-stage processing: the CNN backbone first identifies sensor constellations indicative of particular activities (cooking, cleaning, resting), effectively filtering noise and extracting key spatial information [45]. The temporal dynamics of these high-level properties are subsequently modeled by the next LSTM layer, which differentiates between typical activity sequences which is associated with anomalous deviations. This hierarchical method is similar to how humans absorb information by first identifying environmental cues and then assessing their temporal context, which may help to explain why it aligns with clinical expectations [47], [49]. This mechanistic interpretation is empirically supported by gradient-based saliency analysis. The highest normalized relevance (0.92) was shown by positional motion sensors, indicating that abrupt behavioral changes, especially motion stoppage or atypical room transitions, are the main anomaly triggers. [45], [46], [47], [49]. Nonetheless, the combination of gas sensor modalities produced significant confirmatory evidence: MOX sensors that detected volatile organic compounds (VOCs) from cooking which is associated with cleaning activities (0.83-0.89) and CO₂ sensors that indicated occupancy patterns (0.75-0.78) offered contextual validation that decreased false positives from benign environmental fluctuations. Environmental characteristics (temperature: 0.45; humidity: 0.38) had a smaller but not insignificant impact, indicating that they are more likely to function as drift-correction factors than as direct anomaly indicators [47], [49]. Privacy issues, a major adoption obstacle noted in earlier elder monitoring research, are addressed by the proposed framework's sole dependence on non-visual, non-wearable sensors. Gas sensor arrays assess aggregate ambient changes that cannot reconstruct specific behaviors or identify individuals, in contrast to camera-based systems that capture identifiable biometric data which is associated with intimate activities. By just displaying room-level occupancy and not tracking specific movements or activities, binary positional data further reduces privacy infringement [45]. The ethical frameworks that emphasize autonomy and dignity in aging-in-place devices are in line with this privacy-preserving design. Without

the psychological strain of continual visual surveillance, elderly people can continue to be independent while receiving safety monitoring. The approach further lowers privacy risks related to data breaches or unauthorized access by adhering to data minimization standards which is associated with keeping only anomaly-relevant information instead of full sensor streams. Performance is not compromised by this privacy benefit. The hybrid model's 0.942 F1-score shows that non-invasive sensing can reach accuracy that is on par with or higher than video-based systems documented in the literature (usually 0.88-0.93 F1-score for fall detection), while removing ethical issues with continuous visual monitoring [40], [41]. Despite being higher than baseline approaches, the hybrid model's inference latency of 31.5 ± 3.7 ms per 2-hour window is orders of magnitude below the 20-second sampling interval, indicating real-time viability. Deployment on resource-constrained edge devices, such as contemporary IoT gateways or single-board computers (Raspberry Pi 4, NVIDIA Jetson Nano), is made possible by this latency profile without the need for cloud access. Reduced latency for instantaneous warning creation, improved privacy through local data processing, operational continuity during internet outages, and lower bandwidth costs for long-term monitoring are all important benefits of edge implementation.

Edge implementation on devices with limited storage is further supported by the model's memory footprint (~7 MB in FP32 precision) which is associated with parameter count (1.8M). This might be reduced to about 2 MB with minimal accuracy loss using quantization techniques (INT8 precision), allowing implementation on ultra-low-power microcontrollers for battery-operated sensors. However, in situations when resources are limited, the Mahalanobis baseline's computational advantage (3.1 ms latency) should be taken into account. This trade-off might be acceptable in installations that prioritize minimal power consumption over maximum detection accuracy, even though its lower F1-score (0.853) indicates a 9.4% performance penalty. Future research should investigate methods for model compression (pruning, knowledge distillation) to maintain the better accuracy of the hybrid design while filling this efficiency gap. The generalizability of these results is limited by a number of restrictions. First, even though the single-subject dataset offers excellent longitudinal data, it is unable to capture inter-individual heterogeneity in sensor placement configurations, home layouts, or behavioral patterns [45]. Elderly people have different daily schedules, levels of activity, and preferences for their surroundings; models that were developed for one person might not be applicable to others without modification. Establishing population-level performance benchmarks requires further validation across multi-elder, multi-home datasets. Second, even while the 19-day drift compensation reference period is adequate for PCA-based correction, it might not account for seasonal changes in environmental factors (heating/cooling cycles, humidity swings) that could have an impact on sensor baselines over longer deployments. To verify the robustness of drift compensation which is associated with establish the ideal recalibration intervals, long-term research covering several seasons are required. Third, potential label noise is introduced by the pseudo-labeling strategy, despite its conservative design to reduce false positives. Although multi-criteria labeling (motion thresholds, Mahalanobis distance, expert verification) offers a respectable approximation of ground truth, it is not as accurate as manually marked events. The performance ceiling seen in all architectures could be partially explained by this constraint. Fourth, one practical limitation is the sensitivity of sensor location [46]. The detection capabilities was probably maximized by the source study's strategic placement (e.g., CO2 sensors in the central dining area, MOX sensors near the kitchen). Performance in less-than-ideal locations determined by practical limitations such as furniture placement or the availability of electrical outlets is still not quantifiable [47]. Deployment guidelines would be informed by a systematic assessment of placement robustness. The framework's focus on binary anomaly detection (normal vs. abnormal) restricts its clinical utility. Caretakers can prioritize responses by classifying

anomalies (fall, medical emergency, wandering, prolonged idleness). Expanding to multi-class classification while preserving privacy is a crucial path forward. The problem of idea drift gradual shifts in behavioral patterns brought on by health decline, seasonal routine variations, or lifestyle changes is introduced by long-term deployment. As "normal" patterns change, a model trained on baseline behavior may provide more false positives; on the other hand, if new behaviors indicate a decline in health, the model may fail to identify anomalies. The static training method used in the current framework necessitates frequent retraining to account for drift, which adds to the operational strain. Future implementations should include the following: (i) drift detection algorithms that initiate retraining when performance degradation surpasses thresholds; (ii) online learning mechanisms that gradually update model parameters using confirmed normal periods; which is associated with (iii) customized adaptation protocols that differentiate between benign routine changes and clinically significant declines. Pre-training on population-level data to identify general activity patterns, followed by fine-tuning on individual baselines to capture specific routines, is a promising approach for transfer learning [42]. This strategy might preserve personalization while lowering the amount of data needed for fresh deployments. Instead of operating as a standalone alerting system, successful deployment necessitates a smooth integration into current caregiver operations. Important factors include:

Alert Prioritization: Not every anomaly calls for quick action. Alerts should be categorized by severity (critical: possible fall; moderate: unusual inactivity; low: modest routine deviation) and routed through the proper channels (daily summary for low, app notice for moderate, SMS/call for critical) [43].

False Alarm Management: Occasional false positives are unavoidable even with high precision. In order to facilitate active learning and ongoing improvement, the system should include feedback systems that let caregivers flag alarms as false positives. This preserves caregiver trust which is associated with gradually lessens alert tiredness. In order for caregivers to make educated judgments, warnings must have understandable justifications. Saliency maps showing the sensors that initiated detection (such as "strange motion pattern in bathroom + high CO₂ in bedroom") improve transparency and enable the proper reaction. Monitoring parameters, such as alert recipients, monitoring hours, and data retention policies, should remain within the control of elders. Safety monitoring is made possible by granular privacy controls and transparent permission procedures that respect autonomy [43], [44]. The suggested framework is well-positioned for general adoption since it relies on inexpensive sensors (MOX sensors: \$5–15 per; PIR motion detectors: \$10–20; environmental sensors: \$10–30). The hardware expenses of a full residential deployment (four to six gas sensors, three to five motion detectors, which is associated with a central processing unit) are estimated to be between \$150 to \$300, which is significantly less than those of commercial camera-based systems (\$500 to \$2000) or wearable-based solutions that need to buy devices and set up charging infrastructure [45]. Because edge processing removes cloud computing fees and low-power sensors allow battery operation with 6–12 month replacement intervals, operational costs are kept to a minimum. Scalability to resource-constrained contexts, such as rural communities as well as emerging regions with little infrastructure for elder care, is supported by this economic profile. However, installation labor, caregiver training, technical support, which is associated with system maintenance are hidden costs that should be taken into account. These obstacles can be lessened by user-friendly installation procedures (such as plug-and-play sensors as well as mobile app settings) which is associated with remote diagnostics. The obtained results (F1: 0.942, ROC-AUC: 0.967) are in good agreement with earlier ambient elder monitoring research. Wearable accelerometer-based methods get 0.85-0.92 F1-scores, whereas video-based fall detection systems report 0.88-0.93. The hybrid CNN-LSTM is a strong substitute for these modalities due to its higher performance as well as privacy preservation [45], [46].

This study's findings specifically expand the Mahalanobis distance technique (F1: 0.853) of the original dataset publication by 10.4% in gas-sensor-based monitoring, proving the usefulness of deep learning for identifying intricate patterns in chemical sensing arrays. This development is consistent with more general trends in environmental sensing, where deep architectures are becoming more effective than traditional statistical techniques for high-dimensional, temporally-correlated data. However, because of dataset variability, direct comparison must be done with caution [48], [49], [50]. Variable performance is reported by studies using various sensor combinations, anomaly definitions, which is associated with evaluation methodologies. Standardized criteria as well as shared datasets would enable more thorough cross-study comparability. Healthcare policy which is associated with practice will be significantly impacted by the proven viability of accurate, low-cost, privacy-preserving elder monitoring. Scalable technologies that enable independent living become necessary infrastructure rather than optional comforts as populations age as well as caregiver shortages worsen [48], [49], [50], [51]. By facilitating early intervention for health deterioration, remote monitoring can cut hospital readmission rates which is associated with potentially save emergency care expenses. Clinicians can include behavioral data into treatment planning through integration with electronic health records. Proven clinical utility lowers adoption hurdles by supporting insurance coverage of monitoring systems as preventative care devices.

8. Conclusion

This study demonstrates how multi-modal deep learning, namely hybrid CNN-LSTM architectures, may efficiently use binary positional data which is associated with non-invasive gas sensor arrays for elder safety monitoring. The suggested architecture strikes a crucial compromise for aging-in-place technology by achieving high-accuracy anomaly detection while protecting privacy. The benefits of concurrently modeling spatial sensor correlations and temporal behavioral patterns are highlighted by performance advantages over standalone designs and traditional baselines. (1) Expanding evaluation to multi-elder, multi-home datasets is one of the next research directions; (2) reducing the need for labeled data by using self-supervised pretraining; (3) creating lightweight versions for ultra-low-power edge deployment; which is associated with (4) using explainable AI techniques to improve clinical utility which is associated with caregiver confidence. As sensor costs decrease and deep learning frameworks evolve, such privacy-conscious monitoring systems show tremendous promise for facilitating independent living among elderly populations globally.

Acknowledgements

The authors gratefully acknowledge High Institute for scientific and technology for providing the academic infrastructure and high-performance computing resources that facilitated this research. This research extend this research sincere appreciation to the contributors and curators of the UCI Machine Learning Repository for publicly releasing the Single Elder Home Monitoring: Gas as well as Position dataset, which served as the foundation for this work. This research also thanks the anonymous peer reviewers for their constructive feedback, which substantially improved the methodological rigor and clarity of this manuscript. Finally, the author acknowledge the technical support staff and research colleagues who assisted with data preprocessing, model training, and manuscript preparation.

Author contribution

Hanan Ramadhan conceived the research framework, designed the multi-modal deep learning architectures (LSTM, CNN, and hybrid CNN-LSTM), and implemented the PCA-based preprocessing pipeline for environmental drift compensation. The author conducted all experimental evaluations, performed cross-validation analyses, and interpreted the gradient-based saliency maps to establish clinical interpretability of the anomaly detection system.

Data Availability Statement

The dataset used to develop and validate the models in this study is openly accessible in the UCI Machine Learning Repository. It can be downloaded directly at: <https://archive.ics.uci.edu/dataset/799/single+elder+home+monitoring+gas+and+position> (DOI: 10.24432/C5762W). All preprocessing pipelines, model architectures, and evaluation scripts implemented for this research are available from the corresponding author upon reasonable request.

Conflicts of Interest

The author declare that there are no known conflicts of interest, financial or otherwise, that could have influenced the design, execution, analysis, or interpretation of this research. No external funding sources or commercial entities had any role in the study conception, data collection, manuscript preparation, or decision to publish.

References

- Marín, J. D. Llano-Viles, Z. Haddi, A. Perera-Lluna, and J. Fonollosa. "(2023). Home monitoring for older singles: A gas sensor array system." *Sensors and Actuators B: Chemical*, vol. 393, p. 134036, 2023. <https://doi.org/10.1016/j.snb.2023.134036>
- Jetlawei, S. et al. "(2025) Temporal Intelligence and Algorithmic Equity: A Multi-Phase Framework for Predictive Student Success in Higher Education." *Comprehensive Science Journal*, vol. 9, no. 36, pp. 1574-1595, 2025. <https://doi.org/10.65405/f0xx5p02P>
- Mishra, B. S. R. Shanmugam, K. C. Yeo, and S. Thennadil. "(2025). SDN-enabled IoT security frameworks—A review of existing challenges." *Technologies*, vol. 13, no. 3, p. 121, 2025. <https://doi.org/10.3390/technologies13030121>
- Otokwala, A. U. Petrovski, and H. Kalutarage. "(2024). Optimized common features selection and deep-autoencoder (OCFSDA) for lightweight intrusion detection in Internet of Things." *International Journal of Information Security*, vol. 23, no. 4, pp. 2559-2581, 2024. <https://doi.org/10.1007/s10207-024-00855-7>
- Fares, M. I. A. Abd Elaziz, A. O. Aseeri, H. S. Zied, and A. G. Abdellatif. (2025). "TFKAN: Transformer based on Kolmogorov–Arnold networks for intrusion detection in IoT environment." *Egyptian Informatics Journal*, vol. 30, p. 100666, 2025. <https://doi.org/10.1016/j.eij.2025.100666>
- Benmalek M. and A. Seddiki. (2025). "Particle swarm optimization-enhanced machine learning and deep learning techniques for Internet of Things intrusion detection." *Data Science and Management*, 2025. <https://doi.org/10.1016/j.dsm.2025.02.005>
- Ben Dalla, L. Ö. Karal, M. El-Sseid, and A. Alsharif. "(2026). An IoT-enabled, THD-based fault detection and predictive maintenance framework for solar PV systems in harsh climates: Integrating DFT and machine learning for enhanced performance and resilience." *Wadi Alshatti University Journal of Pure and Applied Sciences*, vol. 4, no. 1, 2026. <https://doi.org/10.63318/waujpasv4i1>
- Islam, M. M. W. M. Abdullah, and B. N. Saha. (2025). "Privacy-preserving hierarchical fog federated learning (PP-HFFL) for IoT intrusion detection." *Sensors*, vol. 25, no. 23, p. 7296, 2025. <https://doi.org/10.3390/s25237296>
- Teixeira, L. R. Almeida, P. Rodrigues, M. Antunes, D. Gomes, and R. L. Aguiar. (2025). "Beyond performance comparing the costs of applying deep and shallow learning." *Computer Communications*, p. 108312, 2025. <https://doi.org/10.1016/j.comcom.2025.108312>
- Zhou, S. et al. "(2026). Advances in machine learning-enabled self-powered flexible sensing materials and their applications." *Advanced Materials Technologies*, p. e01512, 2026. <https://doi.org/10.1002/admt.202501512>
- Ebenezer, A. L. B. Sasithradevi, and C. Baskar. (2025) "Advances in bioinspired sensing-electronic nose and eye technologies for food quality assessment: A review." *IEEE Sensors Journal*, 2025. <https://doi.org/10.1109/JSEN.2025.3590474>

- Lamaakal, I. et al. (2025) "A comprehensive survey on tiny machine learning for human behavior analysis." *IEEE Internet of Things Journal*, 2025. <https://doi.org/10.1109/JIOT.2025.3565688>
- Shehadeh, A. et al. (2021) "Machine learning models for predicting the residual value of heavy construction equipment: An evaluation of modified decision tree, LightGBM, and XGBoost regression." *Automation in Construction*, vol. 129, p. 103827, 2021. <https://doi.org/10.1016/j.autcon.2021.103827>
- Aljohani. A. (2023) "Predictive analytics and machine learning for real-time supply chain risk mitigation and agility." *Sustainability*, vol. 15, no. 20, p. 15088, 2023. <https://doi.org/10.3390/su152015088>
- Al Mamlook, L. J. R. E. Wells, and R. Sawyer. (2023). "Machine-learning models for predicting surgical site infections using patient pre-operative risk and surgical procedure factors." *American Journal of Infection Control*, vol. 51, no. 5, pp. 544-550, 2023. <https://doi.org/10.1016/j.ajic.2022.08.013>
- Alotaibi. B. (2025) "A review of resilient IoT systems: Trends, challenges, and future directions." *Preprints.org*, 2025. <https://www.preprints.org/manuscript/202512.1717>
- Dalla L. O. F. B. and T. M. A. Ahmad. (2020) "The sustainable efficiency of modeling a correspondence undergraduate transaction framework by using generic modeling environment (GME)." *International Journal of Engineering and Modern Technology*, vol. 6, no. 1, 2020. <https://www.iiardpub.org>
- Ben Dalla L. O. F., T. D. Medeni, I. T. Medeni, and M. Ulubay. (2025). "Enhancing healthcare efficiency at Almasara Hospital: Distributed data analysis and patient risk management." *Economy: Strategy and Practice*, vol. 19, no. 4, pp. 54-72, 2025. <https://doi.org/10.51176/1997-9967-2024-4-54-72>
- Dalla L. O. B., Ö. Karal, and A. Degirmenci. (2025). "Leveraging LSTM for adaptive intrusion detection in IoT networks: A case study on the RT-IoT2022 dataset implemented on CPU computer device machine." 2025. <https://doi.org/10.6543/X:4102659>
- Dalla L. O. F. B. "IT security cloud computing." *2020 Innovations in Intelligent IT Security Cloud Computing Conference (IISCCC)*, pp. 1-7, 2020. <https://doi.org/10.16377/ITSCC50717.2020.9259880>
- Ben Dalla L., T. M. Medeni, S. Z. Zbeida, and İ. M. Medeni. (2024). "Unveiling the evolutionary journey based on tracing the historical relationship between artificial neural networks and deep learning." *The International Journal of Engineering & Information Technology (IJEIT)*, vol. 12, no. 1, pp. 104-110, 2024. <https://doi.org/10.36602/ijeit.v12i1.484>
- Dalla, L. O. B., T. D. Medeni, and İ. T. Medeni. (2024). "Evaluating the impact of artificial intelligence-driven prompts on the efficacy of academic writing in scientific research." *Afro-Asian Journal of Scientific Research (AAJSR)*, pp. 48-60, 2024. <https://doi.org/10.7654/X.26.733>
- Degirmenci, A. and O. Karal. (2022). "iMCOD: Incremental multi-class outlier detection model in data streams." *Knowledge-Based Systems*, vol. 258, p. 109950, 2022. <https://doi.org/10.1016/j.knosys.2022.109950>
- Karim A. M., et al. (2020). "A novel framework using deep auto-encoders based linear model for data classification." *Sensors*, vol. 20, no. 21, p. 6378, 2020. <https://doi.org/10.3390/s20216378>
- Dalla. L. O. F. B. (2020). "The sustainable efficiency of modeling a correspondence undergraduate transaction framework by using generic modeling environment (GME)." *International Journal of Engineering and Modern Technology*, vol. 6, no. 1, 2020. <https://www.iiardpub.org>

- Dalla. L. O. F. B. (2020).Lean software development practices and principles in terms of observations and evolution methods to increase work environment productivity." International Journal of Engineering and Modern Technology, vol. 6, no. 1, pp. 23-45, 2020. <https://doi.org/10.6754/s20206543>
- Karal Ö.. "(2020).Performance comparison of different kernel functions in SVM for different k value in k-fold cross-validation." 2020 Innovations in Intelligent Systems and Applications Conference (ASYU), pp. 1-5, 2020. <https://doi.org/10.1109/ASYU50717.2020.9259880>
- Dalla. L. O. F. B. ".(2020).The influence of hospital management framework by the usage of electronic healthcare record to avoid risk management (Department of Communicable Diseases at Misurata Teaching Hospital: Case study)." EHRM, vol. 20, no. 4, pp. 22-52, 2020. <https://doi.org/20.51176/1954-9923-2020-4-22-52>
- Ben Dalla. L. O. F. (2021)."Literature review (LR) on the powerful of research methodology processes life cycle." 2021 Powerful of Research Methodology Processes Life Cycle Conference (TPRMPLCC), pp. 1-10, 2021. <https://doi.org/10.16543/TPRMPLCC50717.2020.92876580>
- Ogundokun, R. O. P. A. Owolawi, and E. Van Wyk. "(2025).LiteRT-IDSNet: A lightweight hybrid deep learning framework for real-time intrusion detection in industrial IoT using the RT-IoT 2022 dataset." 2025 60th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST), pp. 1-4, 2025. <https://doi.org/10.1109/ICEST66328.2025.11098207>
- Dalla L. O. F. B., A. M. A. El-Sseid, T. M. Alarbi, and M. A. M. E. S. Ahmad. ".(2020).A domain specific modeling language framework (DSL) for representative medical prescription by using generic modeling environment (GME)." International Journal of Engineering and Modern Technology, vol. 6, no. 2, 2020. <https://www.iiardpub.org>
- Ben Dalla L. O. F.(2021). "Literature review (LR) on the dominant of research methodology." 2020 LRDRMC Conference, pp. 1-14, 2021. <https://doi.org/10.6754/LRDRMC56412.2020.45987623>
- Dalla L. O. B., Ö. Karal, and A. Degirmenci. (2025)."Leveraging LSTM for adaptive intrusion detection in IoT networks: A case study on the RT-IoT2022 dataset implemented on CPU computer device machine." 5th International Conference on Engineering, Natural and Social Sciences, 2025. <https://www.icensos.com/>
- Yalman Y., et al. "(2022).Prediction of voltage sag relative location with data-driven algorithms in distribution grid." Energies, vol. 15, no. 18, p. 6641, 2022. <https://doi.org/10.3390/en15186641>
- Arık D. T., Ö. Karal, and A. B. Şahin. "(2020).A comparative study of artificial neural networks and naïve Bayes techniques for the classification of radar targets." Bitlis Eren Üniversitesi Fen Bilimleri Dergisi, vol. 9, no. 4, pp. 1779-1788, 2020. <https://doi.org/10.17798/bitlisfen.676973>
- Uysal Z., et al. ".(2016).A heart rate monitoring application using wireless sensor network system based on Bluetooth with MATLAB GUI." International Journal of Engineering Science and Computing, vol. 6, p. 2862, 2016. <http://ijesc.org/>
- Dulkadir S. E. Z. G. İ. N., et al. "(2020).The effect of radiation on the forward and reverse bias current–voltage (I–V) characteristics of Au/(Bi4Ti3O12/SiO2)/n-Si (MFIS) structures." Journal of Materials Science: Materials in Electronics, vol. 31, no. 15, pp. 12514-12521, 2020. <https://doi.org/10.1007/s1>
- Muttaqi M., A. Degirmenci, and O. Karal. ".(2022).US accent recognition using machine learning methods." 2022 Innovations in Intelligent Systems and Applications Conference (ASYU), pp. 1-6, 2022. <https://doi.org/10.1109/ASYU56188.2022.9925265>

- Karal, Ö. and L. O. F. B. Dalla. "(2025).Lung nodule characterization in CT scans using hybrid 3D attention U-Net segmentation and transfer learning-based classification approach." *Comprehensive Journal of Science*, vol. 10, no. 37, 2025. <https://www.sicst.ly>
- Çakır M., A. Degirmenci, and O. Karal. ".(2022).Exploring the behavioural factors of cervical cancer using ANOVA and machine learning techniques." *International Conference on Science, Engineering Management and Information Technology*, pp. 249-260, 2022. https://doi.org/10.1007/978-3-031-40395-8_18
- Sinecen M., et al. ".(2009).Diagnosis of prostat cancer using artificial neural networks." 2009 14th National Biomedical Engineering Meeting, pp. 1-3, 2009. <https://doi.org/10.1109/BIYOMUT.2009.5130296>
- Yumus, M. M. Apaydin, A. Degirmenci, and O. Karal, “. (2020).Missing data imputation using machine learning based methods to improve HCC survival prediction,” 2020 28th Signal Processing and Communications Applications Conference, SIU 2020 - Proceedings, Oct. 2020, doi: 10.1109/SIU49456.2020.9302222.
- Dalla ,L. O. F. B. and T. M. A. Ahmad. ".(2023).Heart disease prediction via using machine learning techniques with distributed system and Weka visualization." *Journal of Southwest Jiaotong University*, vol. 58, no. 4, pp. 322-333, 2023. <https://doi.org/10.35741/issn.0258-2724.58.4.26>
- Ben Dalla, L. Ö. Karal, M. El-Sseid, and A. Alsharif. ".(2026).An IoT-enabled, THD-based fault detection and predictive maintenance framework for solar PV systems in harsh climates: Integrating DFT and machine learning for enhanced performance and resilience." *Wadi Alshatti University Journal of Pure and Applied Sciences*, vol. 4, no. 1, pp. 41-55, 2026. https://doi.org/10.63318/waujpasv4i1_05
- Osman, M. et al. ".(2026).A new approach for low-latency, high-accuracy anomaly detection at the edge: Benchmarking quantized autoencoders, LSTMs, and lightweight transformers on RT-IoT2022 time-series traffic." *Wadi Alshatti University Journal of Pure and Applied Sciences*, vol. 4, no. 1, pp. 110-121, 2026. https://doi.org/10.63318/waujpasv4i1_12
- Elghaffi, F., O. Mohammed, Dalla, L., Ahmed, A., Agila, A., & EL-Sseid, M. (2026). Hybrid matrix-ensemble framework for chronic kidney disease diagnosis." *Wadi Alshatti University Journal of Pure and Applied Sciences*, vol. 4, no. 1, pp. 264-276, 2026. https://doi.org/10.63318/waujpasv4i1_28
- Dalla L. O. B., B., Karal, Ö., Degirmenci, A., EL-Sseid, M. A. M., Essgaer, M., & Alsharif, A. (2025). Edge intelligence for real-time image recognition: A lightweight neural scheduler via using execution-time signatures on heterogeneous edge devices." *Scientific Journal for Publishing in Health Research and Technology*, pp. 74-85, 2025.
- Ahmed, F., A.-Z. Othman, and A. Ukasha. "(2025).Multi-class classification of skin cancer images using a deep learning-based convolutional neural network (CNN)." *Wadi Alshatti University Journal of Pure and Applied Sciences*, vol. 3, no. 2, pp. 230-243, 2025. https://doi.org/10.63318/waujpasv3i2_29
- Rai, R. K. and D. K. Singh. (2026). 1Department of Computer Science & Engineering MNNIT. <https://doi.org/10.36227/techrxiv.177092129.99567847/>
- Cao, Y. et al. ".(2026).Low-cost mechanical sonar mapping with artifact removal in confined spaces." *IEEE Transactions on Instrumentation and Measurement*, 2026. <https://doi.org/10.1109/TIM.2026.3652719>
- Lonchakov, A. M. Sinitca, and D. I. Kaplun. ".(2026).Computer vision-based medical imaging techniques: Past, present, and future." *IEEE Access*, 2026. <https://doi.org/10.1109/ACCESS.2026.3654393>