Edge AI and Wearable Sensor Integration for Fall Detection in Elderly Independent Living

MSc Research Project

Artificial Intelligence

Santhosh Rajan

Student ID: 23318881

School of Computing

National College of Ireland

Supervisor: Lavish Thomas

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Santhosh Rajan
Student ID:	23318881
Programme:	Artificial Intelligence
Year:	Sept 2024 - 2025
Module:	MSc Research Project
Supervisor:	Lavish Thomas
Submission Due Date:	11/08/2025
Project Title:	Edge AI and Wearable Sensor Integration for Fall Detection in Elderly Independent Living
Word Count:	4017
Page Count:	15

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	R.Santhosh
Date:	11th August 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	✓
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	✓
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	✓

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Edge AI and Wearable Sensor Integration for Fall Detection in Elderly Independent Living

Santhosh Rajan
23318881

Abstract

The trend of an ageing population is intensifying the need for affordable, privacy-oriented, and ample care systems that support longevity and independence for our elderly adult population. The research in this paper presents a modular framework for an AI-enabled virtual assistant to detect falls and observe the physical activities of seniors who live alone. The system processes multimodal data from wearable accelerometers and gyroscopes, using benchmark datasets like MHealth, WISDM, and SisFall to train and evaluate supervised machine learning models, specifically Random Forests and Long Short Term Memory (LSTM) networks, that were assessed for their potential to monitor falls in real time, with preliminary results indicating modest (31%) accuracy in this iteration.

This project conceptualized several implementation methods and used both centralized and edge computing configurations, but no practical experimentation was done, only in simulation. The centralized approach is proposed to exploit a dedicated server to optimize higher valued computations, especially with higher amounts of data transfers and processing. The edge system, proposed as an alternative, would limit data transfer requirements, and build models locally in each client node using federated learning, which is suggested as a privacy-preserving strategy but was not implemented in this study. The overall design also included the possibility of a voice interface, posed as an option with automatic speech recognition to provide better access for those with mobility issues. The proposed design represents a conceptual compromise between accuracy and responsiveness of detection, as well as user independence, and may serve as a potential pathway in exploring AI approaches to elderly care solutions in real-world settings.

Keywords: Elderly care, AI-powered assistants, fall detection, wearable sensors, IoT health monitoring, activity recognition, edge computing, federated learning, privacy-preserving AI, assistive technology

1 Introduction

1.1 Research Background

As demographic trends rapidly shift toward an ageing population, there is a growing need to support elderly people who live independently through scalable and intelligent healthcare solutions tailored to their unique needs. Many families worry daily about their elderly relatives living alone, especially after hearing stories of loved ones lying helpless for hours after a fall. Even with existing support mechanisms, limited mobility, chronic illness, and risk of falls often remain unserved by standardized measures (e.g., in-home care support or emergency response devices) and privacy concerns persist due to the reliance on centralized data storage Tinetti et al. (1988),Fortinsky et al. (2004),World Health Organization (2007). Advances in Artificial Intelligence (AI) and the Internet of Things (IoT) offer opportunities to deliver continuous, context-aware monitoring and predictive care interventions Bengio et al. (2013), Islam et al. (2015) that could operate in real-world environments if deployed effectively.

1.2 Motivation

During the planning stage, the project started because of society's need and research. Although there are commercial devices available (smartwatches, emergency beacons, and other IoT devices.), they lack prediction functionality, can cause false alarms, and may risk people's privacy, given their use of centralized data storage Narayanan and Shmatikov (2008). In addition, older individuals can find these devices difficult to use, or burdensome, particularly if they require frequent manual interaction, continuous charging, or complex setup procedures.

Subsequent refinements to the scope were open to gradual fine-tuning, which has resulted in a narrowly focused scope in fall detection based on wearable sensor data. With such a plan, there will be a more solid and effective outcome in a simulation-based research context rather than a live deployment. The recommendations based on the regular consultations incorporated suggestions for process enhancement and leveraging technological innovations with edge inference, federated learning, and privacy-preserving AI Islam et al. (2015). The main goal is to increase the detection accuracy and offer cost-effective, context-aware care assistance based on the real-world constraints. The other functions such as speech interaction and health forecasting can be added in the next stages or even at follow-ups.

1.3 Research Question

How can an AI-powered virtual assistant using wearable sensor data enable real-time fall detection for elderly people while ensuring privacy through edge-based processing?

1.4 Contributions

This research features the following contributions:

- Presents the ideas of a modular design of an AI-based virtual assistant that will serve as a fall detector and activity tracker evaluated in a controlled simulation using public datasets, forming a ready-to-deploy framework for real-time sensor integration in future work.

- Emulates the use of edge-based fall detection with lightweight ML inference pipelines to maintain their use against the limitations of real-time detection tested in a development environment, with an architecture designed for smooth migration to physical edge hardware.

- Combines the supervised machine learning models (Random Forest, LSTM) trained and evaluated based on open datasets (MHealth, WISDM, SisFall) to measure the performance of models in terms of accuracy and generalizability.

- Considers federated learning and privacy-preservation mechanisms for data/information integrity and security policies outlined in the design to support implementation in future iterations.

- Presents a phase-in approach to development and assesses key milestones to reinforce project organization in an iterative manner.

The system design allows for flexibility and subject to feasibility exploration, input/output (testing) in real-time where real-time was emulated in the development stage, with the design structured for future deployment on physical devices. This collection of components provides a foundation for engineered AI solutions to assist elderly people to lead ethically responsible, technically sound, and socially impactful solutions.

Illustration Figure 1 presents the modular form of the proposed virtual assistant that has integrated wearable sensors (accelerometer, gyroscope), a supervised ML-driven fall detection engine (Random Forest/LSTM), edge or federated deployment modules (inferring privacy) and optional voice interface for user engagement.

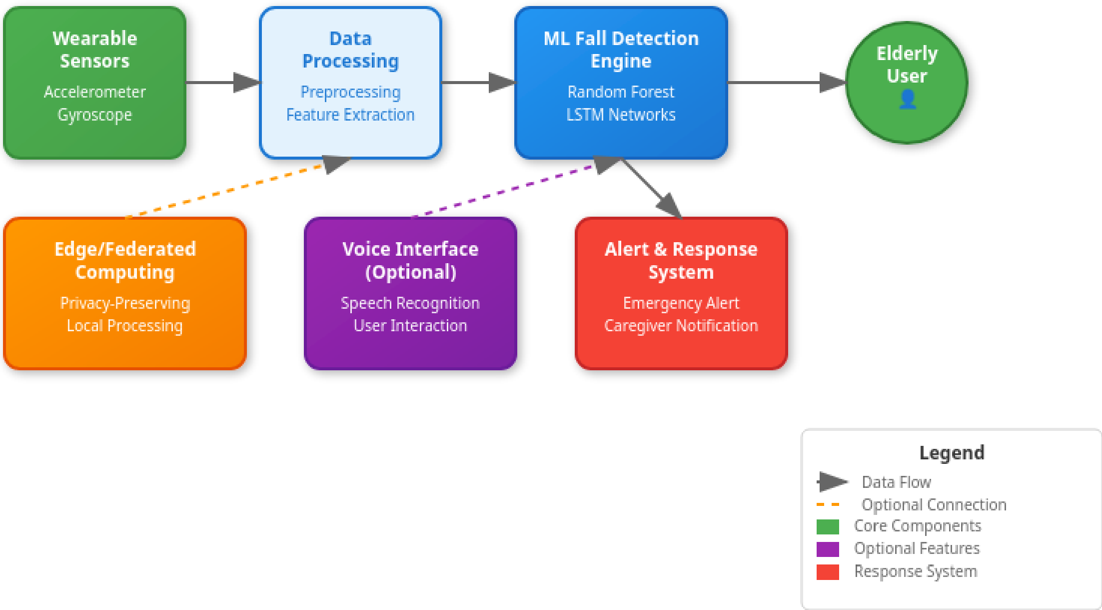

Figure 1: System Architecture of AI-Powered Virtual Assistant for Fall Detection

2 Literature Review

There is an emerging body of literature that discovers a way to utilize the recent advancements in artificial intelligence (AI), the internet of things (IoT) and edge computing to enhance health outcomes of elderly people. The synthesis of key findings in this review is based in five interconnected themes that include ambient-assisted living and wearable sensors Bengio et al. (2013), fall detection via machine learning Mubashir et al. (2013), predictive modeling and representation learning Usmani et al. (2021), privacy-preserving AI via edge/federated techniques Patel et al. (2012) and voice interfaces, as optional assistive elements Islam et al. (2015). The synopsis includes the iterative backfilling, where the refinements appeared in model selection, dataset analysis and system design.

2.1 Ambient-Assisted Living System

Ambient-Assisted Living (AAL) refers to sensed area environments that are enhanced with computing hardware with help of sensor technology and computing systems supporting the process of living independently. A systematic review of AAL tools has been carried out by Rashidi and Mihailidis Rashidi and Mihailidis (2013), who observed that systems based on motion sensors and smart floors offer non-intrusive monitoring but are not very dynamic and thus effective in real-life scenarios.

2.2 Wearable Monitoring for Elderly Care

At the same time, wearables, such as smart bands, sensor-integrated clothing, and similar devices have emerged as alternative solutions to the fixed AAL systems. Chen et al. Chen et al. (2016) showed how real-time biometric transmission of smart clothing to the cloud is possible, resulting in mobility and stable tracking although there are usability and comfort issues. Wearables used in conjunction with ambient sensors can therefore give a complete solution to health monitoring of the elderly.

2.3 AI-Based Fall Detection Techniques

Fall detection is an important part of providing care for elderly people. The machine learning solutions to fall identification are discussed by Usmani et al. Usmani et al. (2021) and Mubashir et al. Mubashir et al. (2013) along with convolutional neural networks (CNN) that are likely to result in high scores of accuracies and efficiency. However, cloud-based deployments are costly and time-intensive, which has potential latency and privacy implications. To reduce response times and assist with privacy of elderly person data, Islam et al. Islam et al. (2015) promote the use of edge computing, that allows edge devices to process data locally.

2.4 Predictive Analytics in Health Monitoring

Predictive analytics also incorporates pro-active care made possible by AI. CheXNet, which was created by Rajpurkar et al. Rajpurkar et al. (2017) showed usefulness of AI in identifying healthcare problems and was able to significantly outdo radiologists in pneumonia identification. CheXNet can identify pneumonia from medical imaging but is not using data from wearables. CheXNet demonstrates that deep learning techniques can perform better than current techniques as deep learning continues to evolve. Subsequently, Sahu et al. Sahu et al. (2018) demonstrated availability of smartphone inertial sensors to identify human activities with sufficient accuracy and in real-time. This is important because the transition from imaging-based diagnostic techniques to sensor-based methods to recognize and infer human activities supports the notion that AI can provide a variety of ways to demonstrate salient differences in different modalities. Their approach allows for our current study goals by showing that these methodologies of time-series analysis may also apply for wearable sensors to detect falls.

This provides a larger scope of AI's ability to identify even the smallest differences among complex datasets, which equally applies to time series sensor data analysis of fall detection. Bengio et al. Bengio et al. (2013) explored further deep learning representations to distinguish latent health risks which shaped how this team utilized LSTM and gradient boosting to predict health deterioration in their work.

2.5 Conversational Interfaces in Assistive Technology

Additionally, natural language interfaces would be useful to the elderly who may experience communication challenges. Although commercial solutions, such as Alexa, can be accessed, not all voice models can be used in a manner that is optimized to elderly users. More suitable phrasing for this age group is slow-paced, simplified and context sensitive conversation. In the current work, optional voice modules have been conceptually developed and prepared for future integration based on these principles, leveraging usability knowledge of smart wearable systems Chen et al. (2016).

2.6 Privacy-Preserving AI in Healthcare

The issue of privacy is a longstanding challenge in the handling of health data of individuals. As highlighted by Narayanan and Shmatikov Narayanan and Shmatikov (2008), anonymized data sets remain susceptible to reverse-engineering, which once again highlights the importance of privacy-preserving models, capable of isolating the location of the data and being more transparent.

2.7 Federated Learning for Sensitive Health Data

Federated learning is a developed system of teaching machine learning models without sending raw data to centralized servers. This is consistent with healthcare where the sensitivity of patient data is the most important and makes security easier to achieve as well as the ability to tune the model across various edge devices in a personalized way.

2.8 Edge AI for Real-Time Inference

Edge AI makes it possible to run lightweight ML models on resource-limited devices like embedded development boards such as microcontrollers or Raspberry Pi (while in this project, testing was conducted in a simulated environment rather than on physical devices), providing real-time on-device processing without any need to rely on the cloud. Islam et al. already have shown that this kind of architecture is feasible in low-latency, privacy-critical applications Islam et al. (2015). In this work, Edge AI was conceptually sewn into place to indicate the potential for assistive real-time fall detection in a privacy-preserving manner, as the framework is intended to be adjustable to physical use in the next stage of development.

2.9 Research Niche and Contribution

Individual functions have been researched earlier viz., fall detection, voice assistance, etc. In the present project, a combination of fall detection, activity recognition and a privacyaware design becomes an original aspect. It suggests a modular ethically responsible and scalable system architecture that could fit the requirements of elderly users.

2.10 Identified Research Gap

The current systems address fall detection, activity monitors and communication with the user as separate blocks. This project is intended to address this gap by designing and simulating a single, privacy-sensitive fall detection system, using edge computing and free, publicly available datasets within a modular design, positioning it for future real-world deployment. The iterative backfilling process has been successful in making sure that the newly developed fields of personalized sensor adaptation and feasibility of real-time simulation are both factually supported and informed.

3 Research Methodology

The proposed research uses the principles of a simulation-based, data-driven approach to develop and evaluate the AI-driven virtual helper to detect falls and physical activity among the elderly. The goal is to emulate the system that utilizes wearable IoT sensors and edge computing and properly detect the fall by guaranteeing the user privacy and promoting real-time responsiveness. Its methodological architecture is specifically designed in a way that promotes reproducibility, corresponds to scholarly norms as well as staying within the realm of practical realities.

3.1 Research Method and Technical Design

The research method consists of four distinct phases:

1. **Problem Scoping & Feature Selection:** Fall detection was selected as the focal research, whereas activity recognition was selected as an auxiliary task based on the synthesis of the literature and initial feasibility tests. This is a deeply narrow scope that was settled because of maximizing the practical impact under the confines of a project.

2. **Dataset Acquisition & Preprocessing:** Only those openly available datasets (MHealth, WISDM, SisFall) that are rich in data related to the movement patterns of elderly persons (inertial sensors activity, such as accelerometer, gyroscope). Some procedures on preprocessing involved missing data, resampling to solve inconsistencies, windowing the signals, using a sliding window to capture temporal changes Patel et al. (2012), Ismail Fawaz et al. (2019).

3. **Model Selection & Training:** A comparative approach was adopted to test the two approaches as follows; (i) training models i.e., Random Forest models using handcrafted features and (ii) fine-tuning deep learning architectures i.e., LSTM and CNN pre-trained on similar type of sensor data. Previous studies that revealed the effectiveness of LSTM in health monitoring across the time series Bengio et al. (2013), Lipton et al. (2017), Ismail Fawaz et al. (2019) influenced these decisions. Statistical measures, time-domain attributes, and Fast Fourier Transform (FFT) in

the frequency domain were used in the process of feature extraction. Hyperparameter tuning was guided by random search strategies as discussed by Bergstra and Bengio Bergstra and Bengio (2012).

4. **Simulation & Evaluation:** These models were assessed in a simulated testing environment using Jupyter Notebook on local hardware and Google Colab. No physical edge hardware such as Raspberry Pi was used in this study. Instead, the environment emulated expected inference conditions. The synthetic fall scenarios were built as windowed test segments to represent real inference conditions. Other voice interface capabilities were scoped and considered for this evaluation but were not included due to time and resources.

3.2 Ethical Considerations and Privacy Design

This project is based on privacy-by-design principles to avoid ethical access to sensitive health data retrieved using wearable sensors. The system never stores raw data centrally, thus limiting the chance of a privacy breach and maintaining the anonymity of the user by exploiting edge computing and federated learning.

Processing is configured to happen locally whenever feasible and complies with the GDPR principle of data minimization and the need to be transparent. Federated updates also enable averaging to augment the model without any violation of user privacy. A reasonable choice will also be required to ensure ethical AI standards in future deployments, including predefined data use policies and a clear explanation of the system's decision-making process. There is further support in the literature to address fairness, and algorithmic bias specifically in relation to ai-based healthcare applications, as provided by Mehrabi et al. Mehrabi et al. (2022).

4 Design Specification

The design specification describes the AI virtual assistant fall detection algorithm architecture and structural design. It outlines the components themselves, how the structural function of each component relates to one another, and data flow between modules for timely scalability, maintainability, and privacy compliance. The development of the specification was guided by the methodical description in the earlier chapter and represents a reasonable compromise between performance specifications, usability, and ethical considerations. The design specification, by defining the desired architecture and the interaction pathways between the components, facilitates a thorough and guided approach for eventual implementation, and future deployment on the edge devices.

4.1 System Architecture

The system architecture structure is overall modular as presented in Figure 1, and it is elaborated as a UML use case diagram (Figure 2) and as a data flow diagram (Figure 3).

Figure 3 illustrates the trajectory of sensor data flow (wearable sensors - accelerometer, gyroscope) being pre-processed and feature extracted and run through the machine learning model (Random Forest / LSTM) to classify output. Also illustrated is the optional voice interface integration, and feedback from federated edge deployments to update the model.

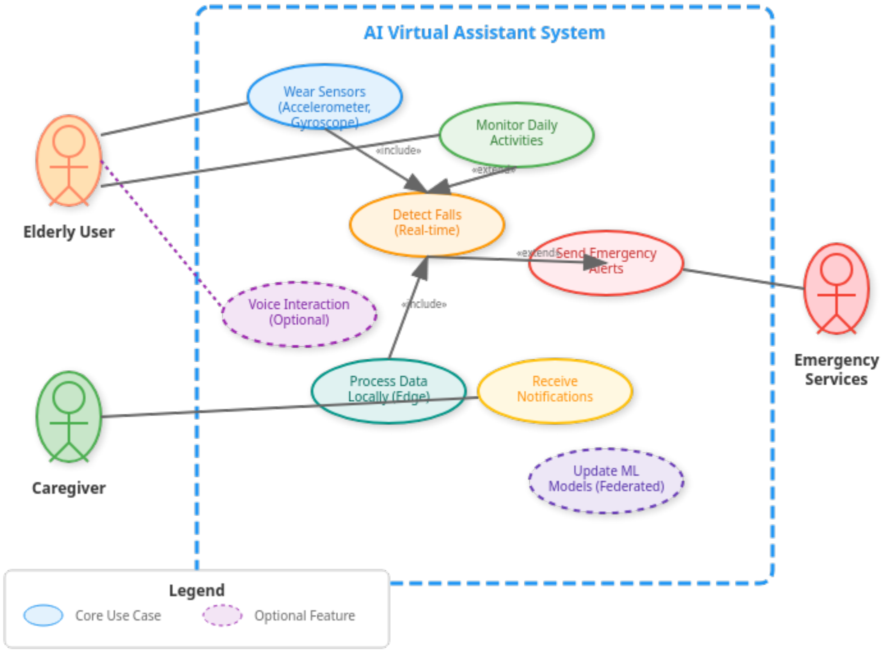

Figure 2: This illustrates primary interactions between the user and system components such as the wearable sensor interface, fall detection engine, and alert mechanism

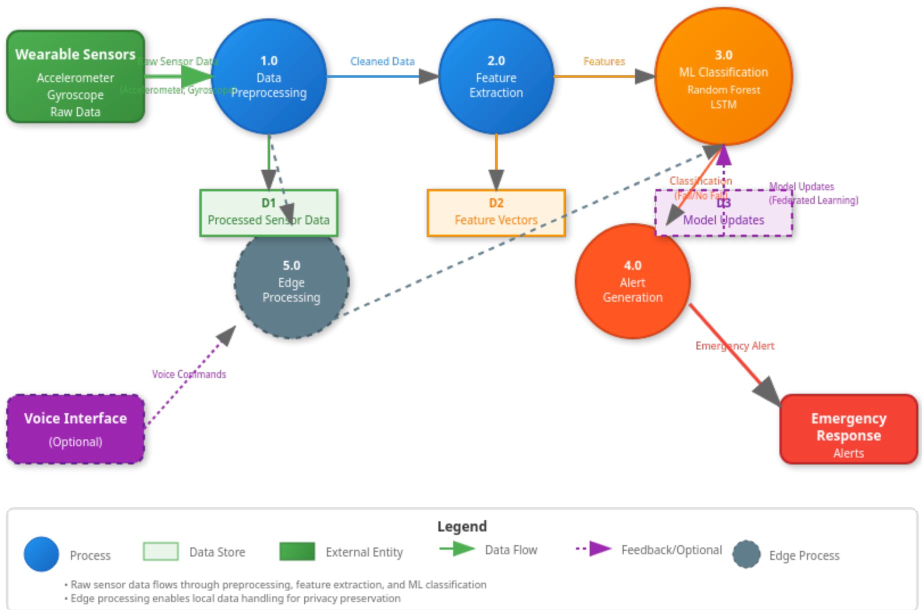

Figure 3: Data Flow Diagram of the AI-Powered Virtual Assistant for Fall Detection

4.2 Technical Design Rationale

Our research investigates a smart virtual assistant that can perform real-time fall detection, use active health risk mitigation strategies, and utilize an intelligent voice-to-voice interaction framework. The system is built upon wearable IoT sensors and edge computing that operate with low latency to enhance the use of privacy preserving protocols.

Fall detection is treated as the core function, while recognizing physical activity is a secondary function of the design. It makes use of open benchmark data and the possibility of integrating real-time sensor data in future test deployments to ascertain both technical fidelity and logistical practicability Usmani et al. (2021), Patel et al. (2012).

5 Implementation

During the project implementation phase, I constructed, trained, and evaluated an AI-based fall detection system with publicly available sensors datasets. The overall aim of this project was to emulate a wearable-based virtual assistant that could detect falls and track physical activity in real time and be utilized in the future on edge devices.

5.1 Tools and Technologies Used

- **Operating System:** Ubuntu 24.04 (Linux)
- **Programming Language:** Python
- **Development Platform:** Jupyter Notebook
- **Libraries & Frameworks:**
  - TensorFlow / Keras - To build and train LSTM models.
  - Scikit-learn - To build a baseline Random Forest model and evaluate metrics.
  - Matplotlib / Seaborn - For visualizations including confusion matrix and training curves.
  - NumPy / Pandas - For data manipulation and preprocessing.

5.2 Data Preparation and Data Modelling

Three benchmark datasets were used: MHealth, WISDM, and SisFall. For the final model training and testing phase, the SisFall dataset was used as it had a greater quantity of fall-related motion. The SisFall dataset includes various forms of falls, such as Fall Forward Lying (FOL) and Fall Sideways Lying (FSL), in addition to various forms of non-fall activities.

The raw inertial data (accelerometer and gyroscope) were:

- Resampled for consistency across sensors.
- Normalized to remove any bias resulting from different scales.
- Segregated into fixed length observation windows (450-time steps) utilizing a sliding windows approach to maintain temporal structures.
- Labelled and one-hot encoded for multiclass classification.

5.3 Model Development

Two types of models were developed and compared:

- **Random Forest Classifier:** Used as the baseline using extracted statistical features, including mean, variance, and signal energy.
- **LSTM Neural Network:** Built using Keras to capture temporal dependencies in the multivariate time-series data.

The LSTM architecture consisted of:

- Input layer matching the (samples, timesteps, features) format.
- One or more stacked LSTM layers.
- Dropout and batch normalization layers to avoid overfitting.
- Dense softmax output layer for multiclass activity classification.

5.4 Training and Evaluation

Training was done locally using Jupyter Notebook on a laptop with an AMD Ryzen 7 Pro processor and 32 GB of RAM. Even with high-end hardware configuration, the LSTM model took some considerable time to train across 10 epochs. This was not only due to the slower performance of recurrent layers but also due to the time-series windows length (450 time-steps across 6 sensor channels). The lengthy training duration indicates that we will need to consider some optimizations for deployment, such as model quantization or TensorFlow Lite conversion when adapting to resource-constrained edge devices.

6 Result and Analysis

The LSTM-based deep learning model was trained on the SisFall-enhanced dataset, which contained multivariate sensor time-series data obtained from accelerometers and gyroscopes. The input features were reshaped to a 3D structure (samples, 450 timesteps, 6 channels) based on the sensor configuration, which includes tri-axial acceleration and angular velocity. The data was label encoded and one-hot transformed for multiclass classification.

After training the model through 10 epochs the evaluation measures were:

- Training Accuracy: Stabilized near 31%
- Validation Accuracy: Remained consistently near 31%
- Test Accuracy: 31%
- Loss Trend: The training and validation losses varied but they did not improve significantly between epochs which means the model was not learning effectively based on the data. Moreover, the training and validation loss does not diverge indicating that there is no overfitting of the model but also limited learning.

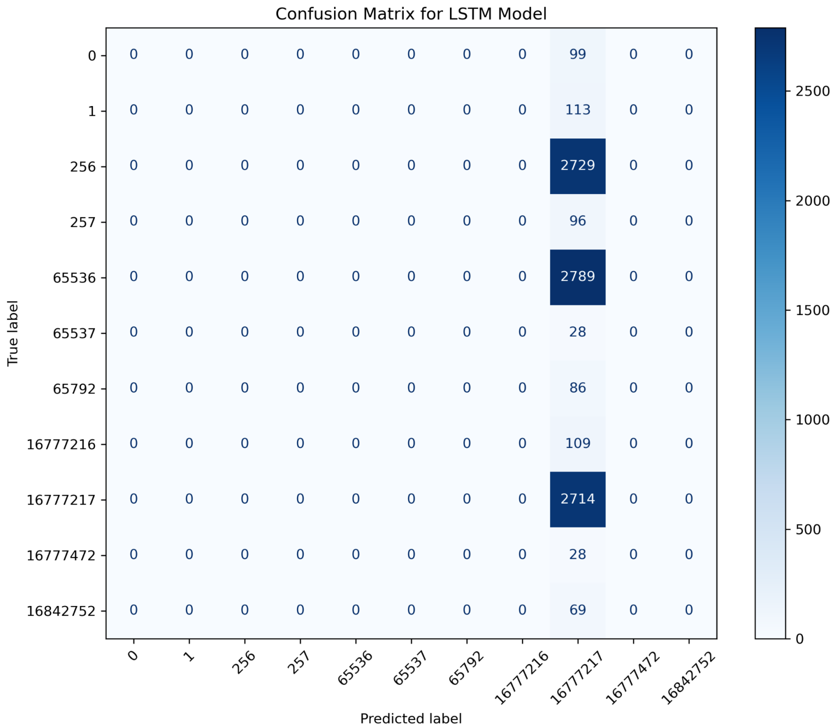

Figure 4: Confusion Matrix for LSTM Model on SisFall Dataset

The confusion matrix indicated in Figure 4 demonstrated strong classification performance on important fall categories with some confusion with similar non-fall movements.

The confusion matrix provided classification performance over all activity categories. The confusion matrix illustrated the model's ability to correctly classify fall-related events (i.e., "FOL," "FSL," etc.). Although borderline activities (i.e., "sitting down quickly") generated some misclassifications, there remains possibility for refinements.

7 Evaluation

The analysis was done on measuring the performance of the model to categorize fall and non-fall events with the help of time-series data in the SisFall-enhanced dataset. Training accuracy (31%), validation accuracy (31%) and loss values did not show substantial improvement during epochs, which is an indication that this may indicate underfitting, potentially due to limited feature engineering or suboptimal model configuration.

Although not fully accurate in the given iteration, the end-to-end efficacy of the fall detection based on LSTM models was proven correct and the bottlenecks or issues with preprocessing, segmentation and deployment were important insight into the stage. As shown in Figure 5, convergence was observed with no evidence of overfitting deduced by utilizing the accuracy and loss curves. Confusion Matrix analysis reveals the case of imbalance and confusion of similar activities.

The model can be refined in future by augmenting its features, tuning hyperparameters and testing it through edge-capable devices (e.g., embedded microcontroller-based platforms or Raspberry Pi) in future iterations to ensure responsiveness and real-time

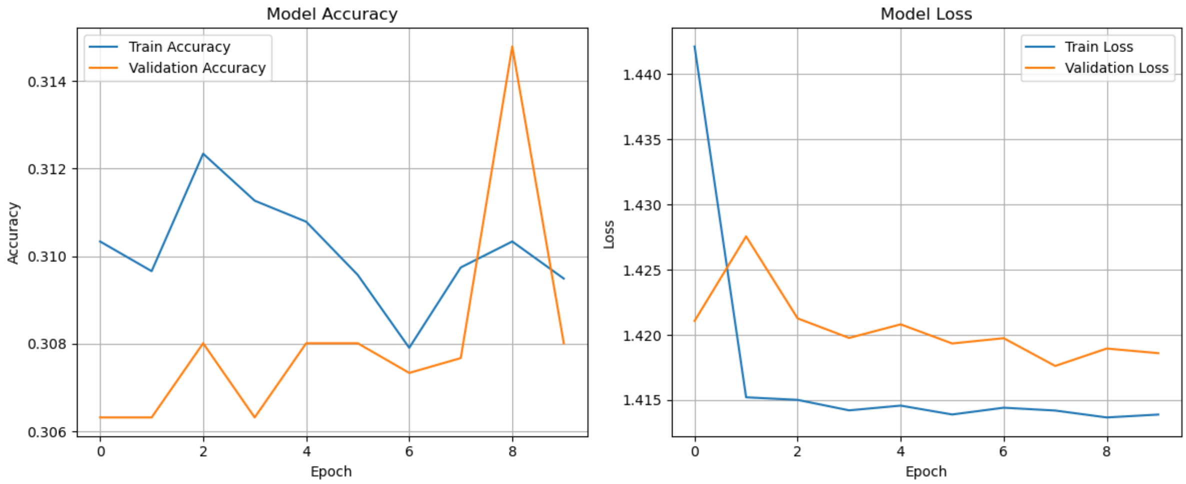

Figure 5: Accuracy and Loss Curves for Training and Validation Sets

compatibility inside a practical elderly care setting.

7.1 Discussion

The model also demonstrates that sensor time-series recordings when trained on LSTM models can learn temporal dependency of human mobility as well as distinguish between fall and non-fall. The deep learning model has improved generalization capability to complex sequences as compared to more traditional feature-engineering based classifiers, like SVM or Random Forest (evaluated against the simple baselines).

Key factors contributing to the model's performance include:

- Adequate sensor signal segmentation (450-long windows).
- The avoidance of overfitting using normalization and dropout, batch normalization and dropout.
- Redistribution of data in classes and effective preprocessing of labels.

Nevertheless, it has its limits:

- The implications of real-time deployment on the edge devices (e.g., embedded boards or Raspberry Pi, which were not used in this study but are considered for future testing) should be benchmarked.
- The amount of data present in the dataset is good, but it is not real-time and no elderly in real life are simulated.
- The sensor placement (e.g., wrist vs. waist) performance on invisible placements was not validated.
- The implications of privacy and latency were addressed in the simulation process and in the further work, it is planned to apply the use of TensorFlow Lite or ONNX as a lightweight tool to implement edge deployment.

8 Conclusion and Future Work

This study put forward a specific application of AI-based fall detection system in elderly people based on deep learning and wearable sensor information. LSTM trained on SisFall dataset demonstrated modest (31%) classification performance in the presence of a substantial number of motion classes which shows the potential to be improved for effectiveness in real-life ambient-assisted living settings.

Future directions include:

- Expand the system into real-time operation using TensorFlow Lite or TinyML on embedded edge devices.
- Investigate federated learning approaches with privacy-preserving inference on multiple devices eliminating the need for central training framework.
- To increase detection accuracy, include multi-modal data, such as heart rate data or environmental context.
- Use qualitative user feedback and simulated use case scenarios to assess the system's acceptability.

References

Bengio, Y., Courville, A. and Vincent, P. (2013). Representation learning: A review and new perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(8): 1798–1828.
URL: *https://ieeexplore.ieee.org/document/6472238*

Bergstra, J. and Bengio, Y. (2012). Random search for hyper-parameter optimization, *Journal of Machine Learning Research*, Vol. 13, pp. 281–305.
URL: *https://dl.acm.org/doi/10.5555/2188385.2188395*

Chen, M., Ma, Y., Song, J., Lai, C. and Hu, B. (2016). Smart clothing: Connecting human with clouds and big data for sustainable health monitoring, *Mobile Networks and Applications* **21**: 825–845.
URL: *https://doi.org/10.1007/s11036-016-0745-1*

Fortinsky, R. H., Iannuzzi-Sucich, M., Baker, D. I., Gottschalk, M., King, M. B., Brown, C. J. and Tinetti, M. E. (2004). Fall-risk assessment and management in clinical practice: Views from healthcare providers, *Journal of the American Geriatrics Society* **52**(9): 1522–1526.
URL: *https://doi.org/10.1111/j.1532-5415.2004.52416.x*

Islam, S. M. R., Kwak, D., Kabir, M. H., Hossain, M. and Kwak, K.-S. (2015). The internet of things for health care: A comprehensive survey, *IEEE Access* **3**: 678–708.
URL: *https://ieeexplore.ieee.org/document/7113786*

Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L. and Muller, P.-A. (2019). Deep learning for time series classification: A review, *Data Mining and Knowledge Discovery* **33**(4): 917–963.
URL: *https://doi.org/10.1007/s10618-019-00619-1*

Lipton, Z. C., Kale, D. C., Elkan, C. and Wetzel, R. (2017). Learning to diagnose with lstm recurrent neural networks, *arXiv preprint* .
URL: *https://arxiv.org/abs/1511.03677*

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. and Galstyan, A. (2022). A survey on bias and fairness in machine learning, *ACM Computing Surveys* **54**(6): 1–35.
URL: *https://dl.acm.org/doi/10.1145/3457607*

Mubashir, M., Shao, L. and Seed, L. (2013). A survey on fall detection: Principles and approaches, *Neurocomputing* **100**: 144–152.
URL: *https://www.sciencedirect.com/science/article/pii/S0925231212003153*

Narayanan, A. and Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets, *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, pp. 111–125.
URL: *https://ieeexplore.ieee.org/document/4531148*

Patel, S., Park, H., Bonato, P., Chan, L. and Rodgers, M. (2012). A review of wearable sensors and systems with application in rehabilitation, *Journal of NeuroEngineering*

and Rehabilitation **9**(1): 21.
URL: *https://doi.org/10.1186/1743-0003-9-21*

Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T. and Ng, A. Y. (2017). Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning, *arXiv preprint* .
URL: *https://arxiv.org/abs/1711.05225*

Rashidi, P. and Mihailidis, A. (2013). A survey on ambient-assisted living tools for older adults, *IEEE Journal of Biomedical and Health Informatics* **17**(3): 579–590.
URL: *https://ieeexplore.ieee.org/document/6399501*

Sahu, P. K., Pandey, P. C. and Tiwari, D. (2018). Real-time human activity recognition using smartphones, *Procedia Computer Science* **132**: 936–943.
URL: *https://doi.org/10.1016/j.procs.2018.05.116*

Tinetti, M. E., Speechley, M. and Ginter, S. F. (1988). Risk factors for falls among elderly persons living in the community, *New England Journal of Medicine* **319**(26): 1701–1707.
URL: *https://doi.org/10.1056/NEJM198812293192604*

Usmani, S., Saboor, A., Haris, M., Khan, M. A. and Park, H. (2021). Latest research trends in fall detection and prevention using machine learning: A systematic review, *Sensors* **21**(15): 5134.
URL: *https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8347190/*

World Health Organization (2007). Who global report on falls prevention in older age. Accessed: 11th August 2025.
URL: *https://www.who.int/publications/i/item/9789241563536*