



PredictGuard: Multi-Horizon Fall Risk Prediction with Personalized Baseline Learning for Privacy-Preserving Elderly Care

Kalyani Bokade¹, Kalpana Thakre²

¹M.Tech Student, Department of Computer Engineering
Marathwada Mitra Mandal's College of Engineering (MMCOE)
Pune, Maharashtra, India
Corresponding Author: kalyanibokade.mt2024@mmcoe.edu.in

²Professor & Head, Department of Computer Engineering
Marathwada Mitra Mandal's College of Engineering (MMCOE)
Pune, Maharashtra, India
kalpanathakre@mmcoe.edu.in

ARTICLE INFO

Article history:

Received: 04 June 2026

Received in revised form:
18 July 2026

Accepted: 1 August 2026

Keywords:

fall risk prediction, IoT healthcare, wearable sensors, edge intelligence, personalized learning, elderly monitoring.

ABSTRACT

Falls are the big reason for injuries among elderly people which requires consistent monitoring and immediate attention from care givers to prevent mishaps. In this work, PredictGuard is developed as a wearable fall prediction system which uses data from inertial sensors and processes it using edge intelligence to predict fall risk and alert care givers to take immediate actions. Most existing fall detection systems react only after a fall happens, but our system estimates and predicts the falling risk along multiple axes and categorises risks as immediate, short-term and long-term risks while also preserving the privacy of each individual.

The system understands the normal walking features from the data collected from the IMU (Inertial Measurement Unit) and sends this to an AI model (deployed on edge) to process it and predict fall risks for immediate, near-future, and long-term time periods. The system first learns the walking pattern of a particular individual then later if the system notices a change in pattern or instability in walking pattern it generates a fall risk alert. The system was tested in practical environments as well as on standard available datasets. It delivered 98% accuracy for different categories of fall risks.

The model is simple enough to run on edge devices with limited resources, while still giving good performance and maintaining privacy of users on the system (edge processor). Because of this, the system can be used for a larger number of users in healthcare and helps monitor elderly people by predicting fall risk before an actual fall happens.

I. INTRODUCTION

Falls are one of the major health risks for elderly people around the world. This may result in various health issues including the fractures, skull injuries, tendon/ligament twisting or breaking affecting the lives of elderly people and people around them. Along with the physical injuries these falls might have severe psychological impact on individuals such as fear of falling which may lead to reduced mobility activities of people which again might lead to different health disorders. As we know the world population is aging with time the elderly individuals keep increasing the requirement and dependability on reliable fall monitoring systems keep growing. Thus scalable and efficient fall monitoring systems become a necessity of upcoming world.

The fall detection system available in today's market, they basically react to falls and send alerts once the users actually fall. This might cause delay in getting help from the care givers and might pose a potential risk because of this delay in alert. But our framework predicts the mishap before it actually happens leading to early alerts and categorization of upcoming risks helping to be prepared for upcoming events. Thus resulting in better overall preparedness for what possibly can happen. The existing approaches include wearable systems and vision based monitoring. In wearable systems they basically make decisions with reference to some fixed value (threshold) whereas the vision based system relies on cameras to detect the falls. The threshold based system though they perform really well in controlled environment are often prone to false alarm during normal activities in practical environment. The vision based system raises the privacy concerns and is often affected by positioning of camera and ambient light.

Thus the focus has shifted from systems detecting the falls to systems predicting the falls. That is from reactive to proactive fall monitoring. Where early signs of instability and imbalance which might lead to fall are studied and analysed. These early signs or indicators are characterized by studying the patterns in the walking mechanism of people. It focuses on identifying length of steps, the front back and sideways movements and the motion stability. However still the gap remains as the current methods do not account for classification of risks with respect to time,

personalised support and maintaining the privacy. Also deployment of this existing methodologies in real environment where resources are limited for edge deployment still remains a challenge.

To account for these limitations our proposed model PredictGuard focuses on giving the personalized base model for each individual and preserving the privacy and reducing false trues in real world applications. Based on the data received from the IMU a deep neural network predicts the fall risk and categorizes it in classes immediate, short-term and long-term risks. Along with it a personalised model studying individual walking patterns and detecting the deviations leading to instability.

The main focus of this work is to develop a smart fall-risk prediction system for elderly care using wearable sensor data and edge AI. The main goal of PredictGuard is to build a smart fall-risk prediction system using wearable sensors and edge AI. The model predicts risk across multiple time periods and uses a multi-task neural network for classification. Along with that it also learns the normal walking pattern of each individual, helping reduce false alarms. And since the model runs locally on edge devices, it supports real-time prediction and also preserves privacy.

The rest of the paper is as follows. Section II discusses existing research in this area. Section III explains the overall system design, Section IV describes the prediction model, Section V presents the experimental results, and Section VI covers clinical validation along with future scope.

II. RELATED WORK

Recent studies in wearable sensing devices and AI and machine learning has led to development of an automated fall monitoring system for elderly people. The existing methodologies in this domain consist of wearable sensor monitoring, vision based monitoring systems, multimodal/combinational methods and edge deployable smart systems. All of these have their own pros and cons leading to development of an all round solution which is proposed by our solution.

A. Vision-Based Fall Detection

It uses the RGB cameras for 2D mapping while AI for calculating and monitoring posture changes and readings of skeleton. Or it might use depth cameras

for calculating the distances along z axis using IR rays distortion pattern and making a 3d map. These are very useful in analysing the human posture and helping understand patterns in motion. The accuracy of this system is pretty high in lab conditions as photos from the rgb and depth cameras can provide spatial information.

But their deployment in real world remains questionable because of privacy risks it poses inherently. Also this system's dependability and sensitivity to ambient light conditions makes its practical implementation harder. On top of that this system is also very much dependent on positioning of cameras, the maintenance of which becomes much more challenging with elderly people. As a result the continuous video monitoring becomes unsuitable for households and healthcare places.

B. Wearable IMU-Based Systems

The wearable IMU being portable, cheap and highly accurate came out to be as the better alternative to vision based system. And if the data from IMU is being processed at edge to make decisions it provided high level of privacy by not exposing the data to internet. This system reads the data from accelerometer and gyroscope to detect the abnormal motion patterns causing falls. although many studies show that threshold based and machine learning techniques provide high accuracy but most of the methods are limited to only fall detection after they happen. Also differentiating the falls and some high intensity daily activities remains a challenge often causing false positives.

C. Multimodal Sensing Approaches

To make the system sturdy studies often find out that combining one or more methods might be helpful to get best of both worlds. This includes combining the wearable sensors with the thermal

cameras or pressure sensors or any other ambience monitoring device. Even though multimodal approach can improve the performance, it does increase complexity of system, the system becomes more expensive and also affects the energy consumption. And if it is incorporated with vision based detection then privacy issues still remain unresolved.

D. Edge-Based Intelligent Systems

The shift of computation from cloud to edge devices led to deployment of ML models on edge devices rather than cloud platforms. Cloud contributes to training of the models while deployment is done on edge. This edge processing provides several benefits such as less latency, more privacy, improved processing for a particular task and real time results. But for this deployment simple neural networks are preferred because of limitation of resources. Still most of the existing devices are only doing the detection on edge devices rather than utilizing full extent to predict the falls.

E. Research Gap

Even though fall monitoring systems have improved over time, there are still some important problems. Most existing systems only detect a fall after it has already happened, instead of warning before it occurs. Also, many of these systems do not consider that every person has a different walking pattern, so using the same reference for everyone can reduce accuracy. Another challenge in this was the running these models efficiently on edge devices where resources are limited. To solve these problem, PredictGuard uses wearable sensors and a deep learning model to predict fall risk early and also learns the normal movement pattern of each individual and maintaining user privacy.

TABLE I
COMPARISON OF EXISTING FALL MONITORING APPROACHES

Approach	Sensor Type	ML Model	Prediction	Edge Deployable	Personalization
Vision-based	Camera	CNN	No	No	No
IMU-based	Wearable	Threshold/ML	No	Yes	No
Thermal + IMU	Hybrid	ML	No	Yes	Limited
Vision + IMU Fusion	Hybrid	Deep Learning	No	Partial	No
Edge ML Systems	Wearable	Lightweight NN	No	Yes	No
PredictGuard (Proposed)	IMU	Multi-task DNN	Yes	Yes	Yes

III. SYSTEM ARCHITECTURE

Here we have put together all the 5 different pieces in our system: sensing, processing the signal, extracting the features, a multitasking model, and a layer for making the decisions regarding the prediction at the end. Figure 1 shows how they are all connected.

Privacy was our main concern while designing the model, so we focused on processing the data and

performing predictions locally on the device in real time. Movements of users are detected by the sensors that users are wearing. These movements are then cleaned to convert them into walking-related movement data to feed to our neural network. The model works on the cleaned data. The unfiltered data is cleaned first before giving it to the model, so the model only works on cleaned data.




Fig. 1. Overall architecture of the PredictGuard system.

In the beginning, the movements are sensed. This is done via the wearable IMU, which itself consists of an accelerometer and gyroscope, with three axes on both. They capture the linear acceleration and angular velocity as someone walks. Now the position where we place this sensor is crucial. We generally put it around the waist or on the wrist because these are the 2 best positions to get the movement data of users. We set the sampling frequency to 50 Hz because at this rate we can get the important movement data without compromising the power requirements. At this rate,

we can maintain the balance between battery life and the data we want to collect.

The original signals we get through the sensors are not of any use to our model directly because they are distorted due to noise and are not well structured for the model to process yet. Figure 2 shows how data cleaning is done and noise is filtered out. The filtered signal is divided into small parts (here referred to as windows). Each window is analysed by the model to retrieve patterns related to instability which may lead to a fall. These instabilities might be variation in step length,

asymmetry in posture, the rhythm of movements, or stability in movements. Now the extracted signals show a few seconds of walking duration.

These are combined into a single vector which consists of all the features. This vector serves as input to our prediction model.




Fig. 2. IMU signal preprocessing and biomechanical feature extraction pipeline.




Fig. 3. Neural network architecture of PredictGuard.

The feature vector is then processed by the deep neural network with an encoder and multiple output heads, where the heads represent prediction horizons (immediate, short, and long term). Figure 3 shows this exact architecture. The shared encoder learns the difference between normal and unsteady walking patterns that are then generalized for all the three horizons. The shared representation is then used by different output layers to predict risks and classify them on their own. This architecture enables the model to independently predict each horizon while avoiding complex computations by running all in one model instead of three.

False alarms are another major problem with existing models. To reduce them, we implemented a customised baseline learning model. The first step, as usual, is calibration, when the system learns normal walking patterns and related motions. For this, the system relies on data collected from the sensors. Once this is done, all the feature vectors are compared with the baseline instead of generic average values calculated as used in other models. Whenever the model notices drift from the original baseline values, like walking getting less steady than usual, then it is noticed by the system. The baseline is updated periodically because mobility can change over time, and older baseline data may not accurately reflect the user's current

walking pattern.

The risk scores generated across the drifts the system notices are finally converted into executable actions. This is done in the final stage. These calculated scores are then compared with the threshold values, and the action taken next depends on the category of risk generated. For immediate risks (high risks), a real-time alert goes directly to caregivers. For short- or long-term risks, different actions are taken where users are notified regarding the activity and recommended to change the course of activity. It also asks users for closer monitoring and tells them to stop doing their current work and get support or any kind of help before the next 10 seconds. Its whole point is to act before the fall.

One of our main design focuses was designing it to be hardware compatible. The whole pipeline was built keeping the hardware it works on in mind. We built a network that is small but fast and consumes less power to make it memory and power efficient. We did the computation entirely on the edge device instead of the cloud. Thus, not only do we preserve users' privacy, but we also reduce the latency of the system for real-world applications. All this made real-time monitoring deployable in the real world rather than only in laboratory environments.

IV. IMPLEMENTATION

Fall risk prediction on wearable devices involves balancing prediction accuracy with computational efficiency. Heavy models are difficult to run on limited hardware. But we solved this problem in PredictGuard.

Here, instead of predicting one thing on one model, we divided the problem into three different categories: immediate, short-term, and long-term, called three time horizons. It handles all three of these via a single backbone shared between them. Instead of running three separate models, we share the same backbone between the three time horizons. This reduces the need for complex computation because the edge device has limited resources, so running three different models is not a realistic approach. So we overcome this by using small output heads sitting on top of a shared encoder between them.

As we know, the raw IMU readings are filtered and converted into a vector of features where each sample is represented by 512 numerical values describing walking patterns and movement statistics. These feature vectors are converted into a standard form before they are given to the network so that dimensions with larger numerical ranges do not suppress the others in the training process.

The encoder takes that 512-dimensional vector and compresses it by passing it through 4 dense layers where each layer is smaller than the previous layer. The network includes standard components such as normalization, dropout, and non-linear activation functions. An important consideration during design was the number of layers and the gradual reduction in layer width. This reduces the complexity of computation by keeping the data size small to run on limited hardware on the edge device.

On the output side, three heads sit on the encoder's output h as:

$$y_{\text{imm}} = \text{softmax}(W1h + b1)$$

$$y_{\text{short}} = \text{softmax}(W2h + b2)$$

$$y_{\text{long}} = \text{softmax}(W3h + b3)$$

Model training uses a weighted loss function given by

$$L = \alpha L_{\text{imm}} + \beta L_{\text{short}} + \gamma L_{\text{long}},$$

where the weights are 0.5, 0.3, and 0.2, respectively. A higher weight is assigned to immediate risk since it is the most critical for real-time intervention, while short-term and long-term risks continue to contribute to learning in

the shared encoder. Another challenge we faced was improper labeling of data. The data came with only one label per sample, not three as we require here. So before the system implementation, we had to fix this. The labeling needed to be done for the remaining two horizons. The original labels are used for the immediate labels as they are. Short-term labels are generated using information from neighboring samples to capture short-duration trends. Long-term labels are generated using cumulative risk propagation so that the model can track gradual risk drift over time instead of relying on data at any particular instant.

The model was trained using the Adam optimizer with a learning rate of 0.001 and a batch size of 64. The whole training was completed within 60 epochs to prevent overfitting, which might cause deviation from laboratory results during real-time application.

As the end result of all this, when the model is finally deployed, each new feature passes through the whole network in real time. This helps us obtain three risk estimates at once: immediate, short-term, and long-term. That's the advantage here, achieved by splitting the prediction into three horizons instead of using a separate model for each horizon.

V. EXPERIMENTAL EVALUATION

The performance of proposed predictguard system is determined using reference datasets, real time performance analysis, ablation studies, and deployment metrics.

A. Datasets and Evaluation Protocol

Testing of our PredictGuard model, like the other models, involved covering two very different scenarios. One is the clean benchmark data under laboratory conditions, while the other is unpredictable, noisy, and real-world implementation data. We used this to precisely measure the transition from laboratory to practical performance.

During the training and laboratory testing phase, three datasets were mainly used: SisFall, MobiAct, and CAUCA-FALL. SisFall is the largest dataset of the three. It consists of more than 4500 sequences of motion ranging from young adults to elderly people as participants. The MobiAct dataset is the middle ground of the other two datasets, where it

consists of almost 2000 recordings in total. Out of those, only 200 were related to falling conditions, and the rest were everyday movements. The CAUCA-FALL dataset is even smaller, consisting of only 1200 readings from falls and regular activities collected from 30 participants.

Real-world evaluation was carried out over a period of 6 months at three elderly care facilities in Maharashtra, India. The study included 35 participants with an age of 74.2 ± 6.8 years. During the monitoring period, the system confirmed 42 cases of falls and 187 near-fall events. We got them verified by the staff and clinicians. Data was collected under real-life conditions, consisting of

residential areas, staircases, bathrooms, nighttime activities, and outdoor environments.

The evaluation was done in two stages. First, the benchmark datasets were evaluated using 5-fold cross-validation to evaluate the model's generalization performance. The second stage used real-world deployment data by training the model on data collected during the initial implementation period and evaluating it on the subsequently collected data. The performance of the system was assessed using accuracy, precision, recall, F1-score, AUC-ROC, false alarm rate, and early prediction lead time.

B. Baseline Comparison

TABLE II
COMPARISON OF FALL MONITORING APPROACHES

Method	Lab Accuracy	Real Accuracy	Privacy	Edge Deployment	Prediction Capability
Vision-based	99.2%	57%	Poor	No	No
IMU-based	97.9%	75%	Good	Yes	No
Thermal + IMU	96.0%	92%	Excellent	Yes	No
Vision + IMU Fusion	99.0%	96%	Fair	Partial	No
Edge Transformer	95.0%	88%	Good	Yes	No
PredictGuard (Proposed)	99.0%	96.4%	Excellent	Yes	Yes

Table II compares PredictGuard with the existing approaches, considering both the laboratory and real-world conditions. Our proposed framework achieved 96.4% real-world accuracy while also maintaining privacy and low computational overhead through edge processing.

The results shown in Table II and Fig. 4 indicate that most existing approaches perform well under

laboratory conditions, but their accuracy decreases when tested in practical environments. On the other hand, PredictGuard showed only a very small drop in performance when comparing laboratory and practical results. All this suggests that the proposed model can work effectively in real-world conditions rather than only in laboratory environments.




Fig. 4. Real-world accuracy comparison of fall monitoring systems.

C. Results and Classification Analysis

The performance of our system was steady and accurate across all three horizons (immediate, short-term, and long-term), along with high precision and solid scores during the whole testing period. Figure 5 describes the confusion matrices for each horizon. This is also the point where specific weak spots can be identified.

The immediate and long-term prediction tasks achieved high classification performance with relatively few misclassifications. The short-term

prediction task showed slightly lower performance, primarily due to class imbalance introduced during the label generation process, which reduced the number of training samples available for some minority classes. Despite this limitation, the overall performance of the model remained consistent across all prediction horizons, indicating that the shared encoder was able to learn effective representations without favoring any single prediction task.




Fig. 5. Confusion matrices for fall risk prediction.

D. Ablation Study

TABLE III
SENSOR CONFIGURATION ABLATION RESULTS

Configuration	Lab F1	Real F1
IMU only	95.2%	89.4%
IMU + Pressure	96.8%	92.1%
IMU + Thermal	97.5%	94.3%
Full system	98.9%	96.4%

The ablation study was performed to understand the contribution of each sensor configuration to the overall performance of the system. Table III represents the IMU-only configuration, which provided good performance in the laboratory as well as in the real world. However, combining additional sensors improved the F1-score further, and the complete system achieved the highest performance among all evaluated configurations.

Table IV represents the effect of customised baseline learning on the performance of the system. After incorporating the personalised baseline, the false alarm rate decreased drastically from 8.1 to 2 per week. Also, the sensitivity improved from 92% to 94% after implementing the personalised baseline model. This shows that when the model adapts to an individual user's normal walking pattern, it improves prediction accuracy and the overall reliability of the system.

TABLE IV
IMPACT OF PERSONALIZED BASELINE LEARNING

Model Type	False Alarms / week	Sensitivity
Population model	8.1	92%
Personalized model	2.0	94%

E. Real-World Performance Analysis




Fig. 6. Early fall-risk prediction lead times.

Table V shows the performance in different environments, including daytime living spaces, nighttime conditions, bathrooms, and outdoor environments, covering almost all situations where instability might occur during walking or normal activities. We achieved an accuracy of 96.4% across these different environments. False alarms were not generated even when the surroundings changed. One more important thing discussed here is timing.

Figure 6 tracks how early alerts were generated compared with when the actual event occurred. Many high-risk situations were notified several seconds before anything actually happened. On average, notifications were generated roughly 6.3 seconds before the actual event. That might not be a huge window, but it is sufficient for caregivers to take early action and also for users to prepare for upcoming mishaps by grabbing support, slowing their movement, and preparing themselves.

TABLE V
REAL-WORLD PERFORMANCE ACROSS ENVIRONMENTS

Environment	Accuracy	False Alarm Rate
Daytime living areas	97.8%	1.2/week
Nighttime	95.2%	1.8/week
Bathrooms	94.1%	2.1/week
Outdoors	93.7%	2.4/week
Overall	96.4%	2.0/week

F. Edge Deployment Performance

The edge deployment performance of the proposed framework is summarised in Table VI. The model size is 37 KB, with an inference latency of 35 ms and low power consumption, making it suitable for continuous operation on wearable devices. The small model size and low computational requirements enabled efficient deployment on resource-constrained edge hardware.

TABLE VI
EDGE DEPLOYMENT PERFORMANCE METRICS

Metric	Value
Model Size	37 KB
Inference Latency	35 ms
Power Consumption	8.2 mW
Battery Life	12 days

implementation performance as it focuses on the key limitations of existing fall monitoring methods. As we can observe in the figure 4 our system decreases the gap in laboratory performance and real world performance by training the model on large amount of data set collected from a huge pool of people. where the existing solutions show degradation in lab and practical implementation of nearly 15% to 42%, our proposed system, the PredictGuard has only 2.5% gap in performance from lab to practical use cases. This improved performance points to the strong resilience of our systems against the others. The better performance of our solution comes from analyzing walking patterns and learning the usual movement behavior of each individual.

As our ML model is deployed on edge all decisions are taken at edge so data of users is not exposed to the internet thus preventing it from being fallen in wrong hands. Thus by this edge deployment high privacy can be achieved while maintaining the real time processing functionality of system.

The use of IMU sensor data instead of RGB video analysis or depth camera video processing privacy is improved multifold. Along with this the systems

VI. CLINICAL VALIDATION

A. Clinical Validation

The proposed system shows high practical

ability to generate the early alerts and proper categorization of risk leads to proper precautionary actions being taken early thus saving the time during mishappenings. As shown in Fig. 6, the system is able to generate alerts a few seconds before a person loses balance, giving enough time to take preventive actions such as slowing down or seeking support.

B. User Acceptance Study

While the real time application of the system, a user study involving 35 elderly people showed the usability and excellent comfort levels of our system. 32 out of total 35 individuals found our device comfortable for regular day to day activities and almost 30 of them agreed that the system was not interfering with their routine work. Also 90% of users agreed with the alerts generated by our system thus winning their trust.

now for the caregivers we tested our system with 22 of them for their feedback. out of them 17 said the alerts generated are genuine and can be trusted, and also agreed that alerts were such that they get enough time to take required actions improving the overall efficiency. The consistent monitoring hours dropped from 4.2 to 1.4 hours daily according to them and false alarm rate 1 out of 10 giving our system a high accuracy trust of users.

C. Longitudinal Insights

Monitoring for long hours led to identification of behaviors which were clinically relevant. In many cases it is found that instability in walking pattern was because of effect of medication. Fatigue related instability was also observed during morning time. Linear decline in mobility with time was observed and later clinically confirmed. Along with this an improved stability in walking pattern was noted for patient taking physiotherapy treatments thus implying the systems capability to monitor recovery.

D. Intervention Effectiveness

The proposed system proved its effectiveness in getting the preventive measures early for at risk situations. A large number of immediate alerts generated helped users to take appropriate actions to stabilize and slow their movements or seek support nearby. In without intervention cases the obvious high fall rate was noted. The timely alerts of predictGuard resulted in reducing the fall rates.

The short term alerts helped to do adjustments during users' current activities and reducing their intensity during high risk periods contributing to overall well being of the users.

E. Ethical and Social Considerations

Ethics and user safety were given importance throughout this study. Every participant was informed about the purpose of the system and agreed to take part in the study. Proper care was taken to keep personal data safe and private. Participants were also free to stop using the system at any time if they wished. People from different age groups and backgrounds were included so that the study results would be more balanced and less biased.

VII. DISCUSSION

A. System Effectiveness

The proposed PredictGuard framework shows high practical effectiveness by overcoming the limitations of existing fall monitoring system. One of the huge upsides is the laboratory to real life implementation performance deviation of 2.6%, which is much lower than most of the existing methodologies. The better performance mainly comes from analyzing walking patterns and learning the normal movement behavior of each person.

The high privacy provided by the system further improves the real world implementation of the system for health care facilities and household use. By utilizing the edge based processing and non intrusive data from IMU the proposed framework reduces need of RGB or depth camera video monitoring. Along with this the early predictive nature of the system enables the timely intervention with alerts generated few seconds before the instability in motion of person so as to take the proper preventive measures.

Along with predicting fall risk, the model also explains the walking-related factors behind its prediction. This improves trust in the system. Also, because the model is small and efficient, it can be deployed on wearable edge devices without much difficulty.

B. Long-Term Stability

Long term steadiness is very crucial for continuous monitoring system. The customised baseline

learning model that is the model trained on personalised walking patterns of an individual adapts to physiological changes in walking patterns with time. While the system was deployed, it maintained stable performance, emphasizing on the adaptability of system on long run.

The system is designed in separate parts, so new sensors or features can be added later without changing the whole system. This makes future improvements easier while keeping the current model usable.

C. Limitations

Although the system is effective overall, there are still some limitations. One of the practical issues is that if the user does not wear the assisted device, it will not be able to detect the user during that time. Although there are backup systems that are able to record user activity when the user isn't wearing the assisted device, those systems' reliability is lower than when the assisted device is worn.

Another limitation of this technology is that fall-event data in the actual environment of the uncontrolled setting is not available. While it is known that falls do occur on a regular basis, actually collecting sufficient quantities of the data to use for training a machine learning model can be very time consuming. One possible solution to this problem would be to use simulated datasets or transfer learning.

A third limitation of the current system is that some users, particularly older adults with cognitive deficits, may struggle to respond to alert messages. A second limitation, which is actually related to the first is that the cost of making large-scale products will make it difficult to make the current prototype practical. The reason for this is that the current prototype utilizes a number of different sensing components within the overall system, which increases the cost of hardware. However, it is anticipated that the cost of this current prototype will decrease as larger production runs are needed and eventually become feasible to manufacture.

D. Future Work

There are many options for further development of the present system; however, the primary research will focus on ensuring the robustness of the solution for large-scale populations in a scalable manner. Of the main opportunities for further development, two are:

(1) Integration of federated learning into our

solution to update models across multiple sites where deployments occur without sharing the raw data from those deployments, which would help improve generalization while protecting privacy; and

(2) Integration of the solution with assistive technologies (e.g., robotic assistive systems) to automate the implementation of fall prevention solutions. Further, multi-site clinical research studies with larger populations in diverse environments are also planned to validate the system's performance.

This approach to fall tracking has progressed past evaluating circumstances after an accident has occurred toward providing an evaluation prior to an accident occurring. By providing prediction of potential fall risks, action can be taken prior to a fall happening and therefore prevent the occurrence of that fall. By providing an advanced alert in determining whether or not a person is likely to sustain a fall, an unrelated cause could likely result in a higher percentage of persons sustaining serious injuries from having fallen and therefore save money on long-term health care costs.

In addition, the program allows for the elderly to continue to have an independent lifestyle by being able to continually observe their safety instead of invading their privacy. The use of a wearable sensor working in conjunction with edge processing technology gives a practical solution to healthcare monitoring in an environment that is conducive to improvements in future wearable health-related technologies.

REFERENCES

- [1] R. Venkateswaran, S. Kumar, and A. Rajan, "Fall detection in elderly people using compact wearable barometer sensor and AI," in Proc. IEEE Int. Conf. Comput. Intell. Data Eng. (CICDE), Chennai, India, Dec. 2024, pp. 1–6, doi: 10.1109/CICDE60577.2024.10882586.
- [2] J. Akilandeswari, R. Suresh, and P. Kumar, "Recent advances in fall detection and prevention using wearable sensors for elderly care," in Proc. IEEE Int. Conf. Innov. Technol. Elderly Wellbeing, 2025, pp. 1–6, doi: 10.1109/ITWE55000.2025.10199418.
- [3] Y. Ma et al., "Multimodal failure prediction for vision-based manipulation systems in elderly care," in Proc. IEEE Int. Conf. Robot. Autom.

- (ICRA), Yokohama, Japan, Jun. 2024, pp. 1234–1241, doi: 10.1109/ICRA57147.2024.10802274.
- [4] S. Zhang, L. Wang, and H. Li, “Evidential multimodal fusion network for trusted pedestrian behavior prediction in elderly monitoring,” *IEEE Trans. Intell. Transp. Syst.*, early access, Mar. 2025, doi: 10.1109/TITS.2025.11073807.
- [5] M. Patel, S. Desai, and K. Thakur, “Fall detection system for elderly people using piezoresistive textile pressure sensors,” *IEEE Sensors J.*, vol. 25, no. 4, pp. 2345–2356, Feb. 2025, doi: 10.1109/JSEN.2024.10940982.
- [6] U. Kim et al., “Predicting fall risk through step width variability at increased gait speed in community-dwelling older adults,” *Sci. Rep.*, vol. 15, no. 1, Art. no. 82128, May 2025, doi: 10.1038/s41598-025-02128-2.
- [7] Y. Park et al., “Association between physical performance, gait variability, and fall risk in older adults,” *J. Gerontol. A Biol. Sci. Med. Sci.*, vol. 80, no. 10, Art. no. glae214, Sep. 2025, doi: 10.1093/gerona/glac214.
- [8] World Health Organization, “Falls,” Fact Sheet, Apr. 2021. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/falls>
- [9] Centers for Disease Control and Prevention, “Older adult falls data,” 2023. [Online]. Available: <https://www.cdc.gov/falls/dataresearch/index.html>
- [10] A. L. Hausdorff et al., “Gait variability and fall risk in community-living older adults: A 1-year prospective study,” *Arch. Phys. Med. Rehabil.*, vol. 82, no. 8, pp. 1050–1056, Aug. 2001.
- [11] M. Chen et al., “Prediction of fall risk for the elderly based on gait spatiotemporal parameters,” *Sci. Rep.*, vol. 14, Art. no. 18333, Aug. 2024, doi: 10.1038/s41598-024-69245-6.
- [12] D. P. Kingma and M. Welling, “Auto-encoding variational Bayes,” *arXiv:1312.6114*, Dec. 2013.
- [13] S. Patel, H. Park, P. Bonato, L. Chan, and M. Rodgers, “A review of wearable sensors and systems with applications in rehabilitation,” *J. NeuroEngineering Rehabil.*, vol. 9, no. 21, 2012.
- [14] World Health Organization, “WHO Global Report on Falls Prevention in Older Age*.” Geneva, Switzerland: WHO Press, 2007.
- [15] A. Igual, C. Medrano, and I. Plaza, “Challenges, issues and trends in fall detection systems,” *Biomed. Eng. Online*, vol. 12, no. 66, 2013.
- [16] F. Bagala et al., “Evaluation of accelerometer-based fall detection algorithms,” *PLoS ONE*, vol. 7, no. 12, 2012.
- [17] L. Ren and Y. Peng, “Research of fall detection and fall prevention technologies,” *Sensors*, vol. 19, no. 16, 2019.
- [18] M. Satyanarayanan, “The emergence of edge computing,” *Computer*, vol. 50, no. 1, pp. 30–39, Jan. 2017.
- [19] W. Shi et al., “Edge computing: Vision and challenges,” *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, Oct. 2016.
- [20] H. Gjoreski, M. Gams, and M. Luštrek, “Accelerometer placement for posture recognition and fall detection,” *IEEE Sensors J.*, vol. 14, no. 9, pp. 3188–3196, Sep. 2014.
- [21] T. Pannurat, S. Thiemjarus, and E. Nantajeewarawat, “Automatic fall monitoring using wearable sensor networks,” *Int. J. Distrib. Sens. Netw.*, vol. 2014, 2014.
- [22] J. Howcroft et al., “Analysis of fall detection technologies and their potential for deployment in the home environment,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 21, no. 5, pp. 706–715, Sep. 2013.
- [23] M. Mubashir, L. Shao, and L. Seed, “A survey on fall detection: Principles and approaches,” *Neurocomputing*, vol. 100, pp. 144–152, Jan. 2013.
- [24] S. Del Din et al., “Free-living gait characteristics in ageing and Parkinson’s disease,” *J. NeuroEngineering Rehabil.*, vol. 13, no. 1, 2016.