ORIGINAL RESEARCH **Open access**

Federated Continual Learning Framework for Adaptive Maintenance of Fall Risk Prediction Models across Aging Populations Using Wearable Accelerometer Data from Smartwatches

Igor Petrov[1*], Alexei Smirnov[1]

Abstract

Fall risk in aging populations is a modifiable health concern, with mobility patterns changing over time due to factors like frailty, comorbidities, and medication. Smartwatch accelerometers provide a privacy-sensitive way to monitor gait and movement outside clinical settings. However, federated learning, which supports privacy by keeping sensor data local, faces challenges in aging populations due to concept drift from gradual mobility decline, which can invalidate static models. This article proposes a federated continual learning framework to adaptively maintain fall risk prediction models using smartwatch data. The system includes local models that combine feature extraction with temporal sequence modeling, continual learning to prevent forgetting, and a federated server for privacy-preserving coordination. It aims to support personalized fall risk monitoring, reduce concept drift, and enable scalable deployment in senior care settings, with clinical validation necessary for real-world assessment.

Keywords Federated learning, Continual learning, Fall risk prediction, Smartwatch accelerometer, Ageing populations, Concept drift

*Correspondence:
Igor Petrov
igor.petrov@gmail.com
[1] Department of Artificial Intelligence in Healthcare, Saint Petersburg State University, Saint Petersburg, Russia

Introduction

Falls among older adults remain a persistent clinical and public health concern because they can lead to injury, functional decline, loss of confidence, institutionalisation, and increased care dependency. Static fall risk assessments are useful snapshots, yet they may fail to capture gradual changes in gait stability, postural control, and activity regularity that occur between clinical encounters. Wearable sensor studies have shown that movement-derived features can support prospective fall-risk modelling and elderly faller classification, suggesting that fall risk should be understood as a dynamic trajectory rather than a fixed category [1,2]. Systematic work on wearable sensors further indicates that continuous sensing can complement clinical assessment by capturing everyday movement patterns that are difficult to observe in clinic-based testing [3].

Smartwatch accelerometers and inertial sensors provide a practical foundation for remote fall risk monitoring because they can capture walking, sit-to-stand transitions, postural changes, and daily activity variation in free-living environments. Deep learning and machine learning approaches using inertial data have been proposed to represent spatio-temporal gait patterns and fall-risk signatures, while more recent reviews highlight the growing role of wearable sensors in constructing fall-risk prediction models for community-dwelling older adults [4,5]. At the same time, federated learning offers a privacy-preserving paradigm in which models can be trained across distributed devices without transferring raw accelerometer streams to a central server [6,7]. This combination is particularly relevant for senior care because movement data may reveal sensitive information about health status, frailty, and daily routines.

Age-related mobility decline creates a modelling challenge because the meaning of a movement pattern may change as an individual transitions from robust mobility to prefrailty or frailty. A gait signature that once represented ordinary daily variability may later indicate instability, fatigue, or increased fall susceptibility, generating concept drift in wearable sensor data. Continual learning has been proposed as a mechanism for maintaining adaptive clinical models while limiting catastrophic forgetting, but clinical artificial intelligence systems still require careful monitoring and updating after deployment [8,9]. In federated settings, this challenge is intensified by non-independent and non-identically distributed wearable data, user-specific behaviour patterns, and device-level heterogeneity [10].

This article proposes a conceptual federated continual learning framework for adaptive maintenance of smartwatch-based fall risk prediction models across ageing populations. The framework builds on privacy-preserving federated health learning, personalised wearable modelling, and continual learning principles to support local adaptation while retaining historically relevant fall-risk knowledge [6,11,12]. It is presented as an AI Systems/Frameworks article rather than an experimental study, and therefore it does not report dataset sizes, training epochs, or performance estimates. The remainder of the manuscript outlines the background, system architecture, local smartwatch model, federated continual learning mechanism, mobility-shift handling strategy, privacy and personalisation design, clinical integration pathway, evaluation strategy, and limitations.

Background

Wearable accelerometers and fall risk

Wearable accelerometers and inertial measurement units can capture gait rhythm, postural transition dynamics, step irregularity, turning behaviour, and daily activity patterns that are clinically relevant to fall risk. Studies using wearable sensors have examined prospective fall-risk prediction, elderly faller classification, and feature selection from movement signals, indicating that accelerometer-derived measures can represent subtle mobility deficits [1,2]. Reviews of wearable fall risk assessment further show that inertial sensors are useful for continuous monitoring because they capture behaviour in naturalistic contexts rather than isolated clinical tasks [3]. More recent work using accelerometry and clinical assessment has reinforced the conceptual value of combining movement signals with individual health information for fall-risk prediction [13,14].

Federated learning for health monitoring

Federated learning enables distributed model training by allowing local devices or institutions to compute model updates without transferring raw health data to a central server. In wearable healthcare, frameworks such as FedHealth illustrate how transfer learning and federated learning can be combined to support privacy-preserving learning from distributed sensor data [6]. Broader reviews of federated learning in medical contexts and mobile health describe its relevance for privacy-sensitive domains, including applications where data are fragmented across devices, institutions, or individuals [12,15]. Cross-institutional medical federated learning also demonstrates how collaborative learning can scale across sites while reducing dependence on centralised data pooling [7,16].

Continual learning and catastrophic forgetting

Continual learning addresses the problem of updating models over time while retaining previously acquired knowledge. In clinical artificial intelligence, continual updating is particularly important because disease patterns, patient behaviour, device use, and care pathways can change after deployment [8,9]. Systematic reviews of continual learning in medicine and medical imaging describe common strategies such as rehearsal, regularisation, and architectural adaptation, each of which attempts to reduce catastrophic forgetting under sequential learning conditions [17,18]. These strategies provide a conceptual basis for fall-risk models that must adapt to new movement patterns without discarding earlier representations of stability, instability, and recovery.

Federated continual learning

Federated continual learning combines the privacy-preserving coordination of federated learning with the sequential adaptation capacity of continual learning. This combination is important for wearable fall-risk monitoring because each smartwatch may observe a unique stream of evolving movement data, while the broader system must still learn from population-level patterns across heterogeneous users. Studies on non-IID data and continual processes in federated learning show that distributional heterogeneity and sequential tasks remain major challenges for federated systems [10]. Federated and online human activity recognition approaches further demonstrate the relevance of continual, personalised, and semi-supervised learning in distributed wearable environments [19,20].

Table 1 distinguishes standard federated learning from the proposed federated continual learning framework by showing how temporal adaptation and forgetting prevention change the logic of smartwatch-based fall risk prediction.

Table 1. Conceptual Differentiation between Standard Federated Learning and Federated Continual Learning for Smartwatch-Based Fall Risk Prediction

Analytical dimension	Standard federated learning in wearable health	Federated continual learning in the proposed framework	Added conceptual value for ageing populations
Primary modelling assumption	User data are distributed but relatively stable over time	User data are distributed and temporally evolving	Treats fall risk as a changing mobility trajectory rather than a fixed classification problem
Main privacy mechanism	Raw smartwatch data remain local while model updates are shared	Raw data remain local, and historical movement knowledge is retained through privacy-preserving local mechanisms	Protects both current sensor streams and longitudinal mobility histories
Adaptation target	Population-level model improvement across devices	Joint population-level learning and individual ageing-trajectory adaptation	Supports both shared fall-risk signatures and personalised mobility decline detection
Main technical risk	Non-IID data may weaken aggregation stability	Non-IID data plus sequential drift may cause catastrophic forgetting	Addresses instability caused by heterogeneous and ageing-related movement changes
Local device role	Computes updates for a global model	Maintains personalised model, detects drift, and performs continual adaptation	Gives the smartwatch an active role in lifelong model maintenance
Historical knowledge handling	Earlier patterns may be overwritten during retraining	Earlier movement-risk representations are protected through rehearsal or regularisation	Preserves baseline mobility knowledge needed to interpret later decline
Clinical relevance	Supports privacy-preserving collaborative model development	Supports adaptive, longitudinal fall-risk monitoring	Better aligns with frailty progression, medication

			changes, and evolving daily activity routines
Evaluation priority	Accuracy, aggregation performance, and privacy	Accuracy, calibration, forgetting, drift response, communication cost, and alert burden	Expands evaluation from model performance to clinical sustainability

Personalisation versus generalisation

Fall risk prediction requires a balance between population-level generalisation and individual-level personalisation. Population-level models can learn common fall-risk signatures across older adults, but personalised models are needed because baseline gait, activity routines, comorbidities, and device-wearing patterns vary substantially between individuals. Federated personalisation approaches for in-home health monitoring and wearable activity recognition suggest that local adaptation can preserve user-specific information while still benefiting from collaborative learning [11,21,22]. In elderly monitoring, this trade-off is central because excessive generalisation may miss individual deterioration, whereas excessive personalisation may reduce robustness across changing care contexts [23,24].

Framework Overview

High-level architecture

The proposed framework consists of a local smartwatch model, an on-device continual learning module, and a federated aggregation server. Each smartwatch processes accelerometer-derived movement windows locally and updates a fall-risk model using recent data together with mechanisms designed to retain historical knowledge. The federated server receives privacy-preserving model updates rather than raw accelerometer streams, allowing population-level coordination without centralising personal movement traces [6,7]. This architecture extends existing federated wearable health concepts toward adaptive maintenance of fall-risk models in ageing populations [11, 25].

Figure 1 presents the proposed federated continual learning architecture for adaptive smartwatch-based fall risk prediction across ageing populations.

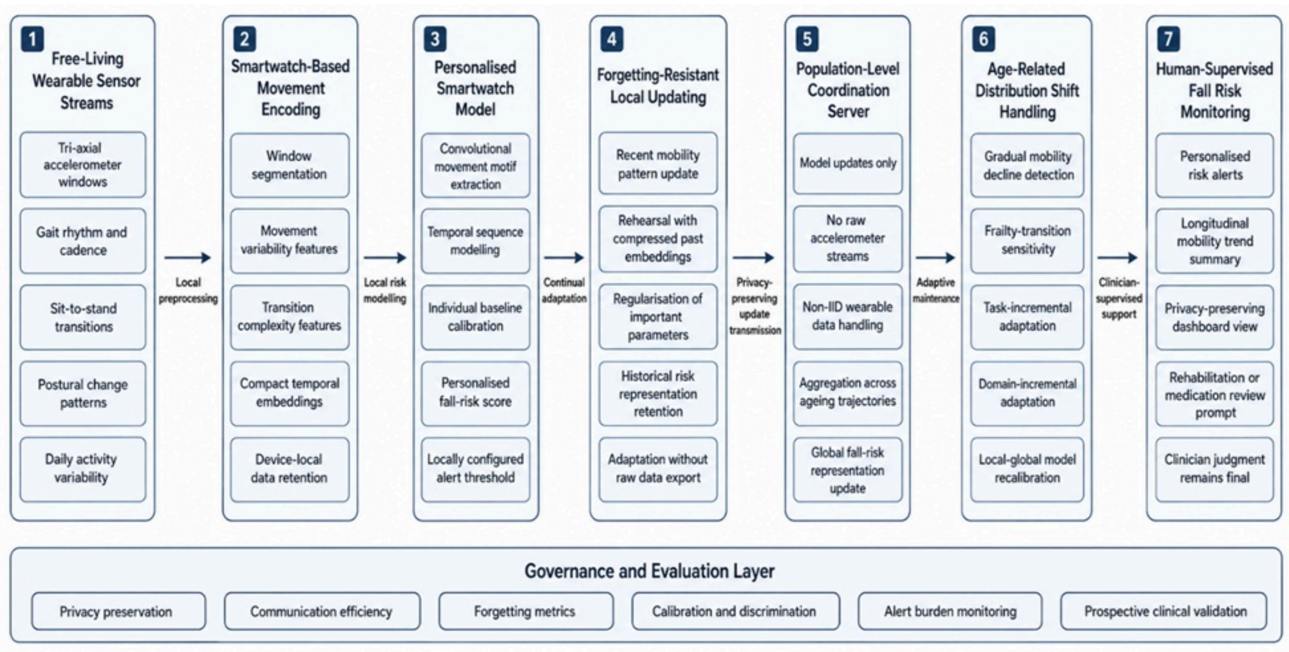

Figure 1. Federated continual learning architecture for adaptive smartwatch-based fall risk prediction

Core assumptions

The framework assumes that smartwatches can collect continuous or periodic accelerometer data during ordinary daily activity, perform lightweight preprocessing, and store a limited local representation of historically informative movement segments. It also assumes that communication with a federated server occurs periodically rather than continuously, allowing local learning to continue under intermittent connectivity. These assumptions align with mobile health federated learning, where devices may contribute updates while preserving raw data locally and operating under resource constraints [15,26]. They also reflect wearable activity-recognition settings in which non-IID user behaviour and local adaptation are expected rather than exceptional [27,28].

Design principles

The framework is guided by five design principles: prevention of algorithmic forgetting, personalisation to individual ageing trajectories, communication efficiency, privacy preservation, and scalability across heterogeneous elderly cohorts. Forgetting prevention is necessary because the model should retain earlier fall-risk signatures even as mobility changes, while personalisation is necessary because each user's baseline movement pattern may differ. Communication efficiency and privacy preservation are essential because smartwatch systems must operate within battery, bandwidth, and data-governance constraints [26,29]. Scalability is supported by federated aggregation, which enables distributed learning across many devices without requiring centralised sensor repositories [7, 12].

Local Smartwatch Model

Input representation

The local smartwatch model would receive short windows of tri-axial accelerometer data and transform them into representations of gait rhythm, step regularity, postural transitions, and activity intensity. Raw signal windows could be complemented by derived features such as cadence-related patterns, transition complexity, and movement variability, provided that these features remain on the device. Prior fall-risk studies using wearable sensors support the conceptual relevance of movement-derived inputs for distinguishing fall-prone and non-fall-prone profiles [1,2]. Deep learning work on inertial sensors further suggests that temporal gait structure and domain-informed spatio-temporal features can support fall-risk assessment [4].

Neural architecture

A compact local architecture could combine convolutional layers for extracting local movement motifs with recurrent or sequence-oriented layers for representing temporal dependencies across activity windows. The architecture should be designed for on-device inference and incremental updating, with computational requirements appropriate for smartwatch battery and processor limitations. Federated human activity recognition studies have explored feature extraction, dynamic layer sharing, and prototype-guided personalisation, all of which are conceptually relevant to compact wearable models [21,27,29]. Such a model would not need to expose raw movement streams, because local feature learning and update generation could occur directly on the smartwatch [6].

Output fall risk score

The local model would produce an individualised fall-risk estimate that may be represented as a continuous risk score or as clinically interpretable risk categories. The threshold for alert generation should be personalised because older adults differ in baseline gait, habitual activity, frailty trajectory, and tolerance for false alerts. Wearable fall-risk prediction studies and reviews suggest that sensor-derived risk estimates are most meaningful when interpreted alongside clinical context rather than treated as universal labels [3,5,14]. In a federated continual learning framework, the smartwatch would therefore maintain local calibration while still contributing privacy-preserving updates to the wider model [11,22].

Federated Continual Learning Mechanism

Local continual update

In the proposed framework, each smartwatch performs local continual updates as new movement patterns are observed over time. A rehearsal-based module could retain a small privacy-preserving memory of representative past movement embeddings or compressed exemplars, allowing new data to be interleaved with historical patterns during adaptation. Continual learning reviews in medicine identify rehearsal as a practical strategy for reducing catastrophic forgetting when models encounter sequentially changing clinical data [17,30]. For fall risk prediction, this mechanism would be expected to help the model adapt to emerging frailty-related movement changes without losing earlier knowledge of the user's stable baseline.

Regularisation-based option

As an alternative or complement to rehearsal, the framework could use regularisation-based continual learning to protect model parameters that are important for previously learned fall-risk patterns. Elastic weight consolidation and related approaches conceptually penalise excessive changes to influential parameters, thereby reducing the risk that adaptation to new movement phases overwrites older knowledge. Continual learning surveys in medical imaging and clinical analysis describe regularisation as one of the central families of forgetting-prevention methods [18,31]. In smartwatch-based fall risk modelling, regularisation would be especially useful when memory buffers are constrained or when storing even compressed historical movement examples is undesirable.

Federated aggregation for FCL

After local continual updates, the smartwatch would transmit model updates, importance summaries, or privacy-preserving adaptation signals to the federated server rather than raw accelerometer windows or exact rehearsal buffers. The server would aggregate these updates to maintain a population-level model that reflects distributed ageing trajectories while allowing local devices to preserve personalised knowledge. Prior work on federated health learning, mobile health applications, and privacy-preserving clinical prediction supports the feasibility of collaborative learning without direct data centralisation [12,15,32]. Conceptually, the aggregation process should be designed to tolerate non-IID wearable streams and sequential mobility shifts, which remain key challenges in federated continual learning [10,33].

Handling Age-Related Mobility Shifts

Concept drift in fall risk factors

Age-related mobility shifts can alter the relationship between accelerometer patterns and fall risk because gait speed may decline, postural transitions may become slower, and step-to-step variability may increase as frailty progresses. These changes imply that a fall risk model trained on earlier movement behaviour may become less reliable when the user's mobility state changes. Wearable sensor research has shown that accelerometry can capture gait stability, symmetry, and transition-related features that are relevant to fall-risk assessment, making it suitable for detecting drift in movement-derived risk patterns [4,13]. A federated continual learning system would therefore treat fall risk prediction as an adaptive process rather than a one-time model deployment.

Task-incremental versus domain-incremental scenario

The framework distinguishes between task-incremental and domain-incremental adaptation in ageing-related mobility monitoring. A task-incremental scenario would arise when a clinically meaningful phase, such as a transition from prefrail to frail mobility, is explicitly recognised and treated as a new learning phase. A domain-incremental scenario would arise when movement distributions shift gradually without a new label or explicit clinical phase boundary. Continual learning approaches for clinical systems and electronic health records illustrate why models must adapt across changing feature spaces and temporal care contexts while retaining earlier knowledge [9,34].

Personalised drift detection

Personalised drift detection would operate locally on the smartwatch by monitoring changes in movement variability, transition complexity, activity regularity, and model uncertainty over time. When the system detects a sustained shift, it could trigger local adaptation while preserving earlier risk representations through rehearsal or regularisation. Federated human activity recognition studies highlight the importance of user-specific adaptation because wearable sensor distributions differ substantially across people and contexts [19,20,28]. In fall risk monitoring, this local trigger would reduce dependence on central supervision while allowing adaptation to the individual's evolving mobility trajectory.

Privacy and Personalisation

Privacy guarantees

The proposed framework is designed so that raw accelerometer segments, exact temporal routines, and locally stored rehearsal examples remain on the smartwatch. The federated server would receive only model updates, compressed adaptation signals, or differentially private summaries of parameter importance, depending on the chosen implementation. Privacy-preserving edge federated learning and federated digital health frameworks support this design principle because they emphasise collaborative model improvement without direct centralisation of sensitive health data [7,26]. For older adults, this is especially important because movement traces may indirectly reveal frailty, daily living patterns, and care dependency.

Personalised risk thresholds

Personalised risk thresholds would combine local model output with user-specific clinical context, such as baseline mobility, prior fall history, medication burden, or clinician-defined tolerance for alert sensitivity. These threshold parameters should remain locally controlled, because they may encode sensitive medical and behavioural information. Personalised federated learning for in-home health monitoring and wearable sensor activity recognition shows how user-specific adaptation can coexist with broader collaborative learning [11,22]. In this framework, the global model would provide shared fall-risk representations, while the smartwatch would preserve individual calibration for real-time alerting.

Clinical Integration – Real-Time Adaptive Risk Alert

Local alert generation

The local smartwatch model would generate fall-risk alerts when the personalised risk estimate exceeds a clinically configured threshold. To support interpretability, the alert could describe the type of movement change that contributed to increased risk, such as reduced postural transition stability, altered gait rhythm, or increased irregularity in daily activity. Wearable fall-risk studies suggest that movement-derived features are clinically meaningful when linked to gait, stability, and functional behaviour rather than presented as opaque signal outputs [1,14,35]. Such alerts should be treated as decision-support prompts rather than deterministic predictions of imminent falls.

Lifelong learning dashboard

A clinician-facing dashboard could summarise longitudinal fall-risk trends, mobility-phase changes, and population-level drift indicators derived from federated statistics. The dashboard should avoid exposing raw movement data and instead present privacy-preserving summaries that support care planning, rehabilitation referral, or medication review. Federated learning in mobile health and medical settings has been proposed as a way to coordinate learning across distributed data sources while preserving privacy and governance boundaries [12,15]. In senior care facilities, the dashboard would be expected to support proactive monitoring while leaving clinical judgement and patient communication under human oversight.

Evaluation Strategy

Table 2 provides an evaluation matrix that connects technical performance, clinical usefulness, privacy protection, and deployment feasibility in the proposed federated continual learning framework.

Table 2. Evaluation Matrix for Federated Continual Fall Risk Prediction across Technical, Clinical, and Governance Domains

Evaluation domain	Core evaluation question	Candidate assessment criteria	Why this strengthens the framework
Forgetting resistance	Does adaptation to new mobility phases degrade earlier fall-risk knowledge?	Backward transfer, retention of baseline-risk representations, performance before and after drift	Tests whether continual learning actually protects historical mobility information
Drift responsiveness	Does the model detect and respond to ageing-related mobility shifts?	Change in gait variability, transition instability, uncertainty increase, domain-shift sensitivity	Links technical adaptation to clinically meaningful mobility decline
Personalisation quality	Does the model reflect individual baseline movement patterns?	Local calibration, personalised threshold stability, user-specific sensitivity and specificity	Prevents population-level averaging from obscuring individual deterioration
Federated robustness	Can aggregation remain stable under heterogeneous smartwatch data?	Non-IID tolerance, client variability, update divergence, aggregation stability	Addresses the central challenge of distributed elderly populations
Privacy protection	Are raw movement traces and sensitive routines protected?	Raw-data locality, compressed update sharing, privacy-preserving summaries, buffer governance	Ensures that longitudinal movement data are not unnecessarily centralised
Device feasibility	Can the framework operate within smartwatch constraints?	Battery burden, memory footprint, update frequency, communication cost, on-device computation	Connects the conceptual framework to realistic deployment conditions
Clinical utility	Do alerts support meaningful care decisions without replacing clinicians?	Calibration, actionable alert explanations, false-alert burden, referral or review usefulness	Keeps the system aligned with decision support rather than autonomous diagnosis
Prospective validation	Does the framework remain useful in real-world senior care settings?	Longitudinal pilot studies, frailty subgroup analysis, care-setting comparison, clinician acceptance	Establishes the pathway from conceptual architecture to clinical evaluation

Forgetting metrics

The evaluation of the proposed framework should examine whether adaptation to new mobility phases causes degradation in previously learned fall-risk patterns. Conceptual forgetting metrics could include backward transfer, retention of earlier risk representations, and comparison between federated continual learning and standard federated learning after distributional shifts. Continual learning reviews in medicine emphasise that preventing catastrophic forgetting is central when clinical models are updated sequentially over time [17,30,31]. In this framework, forgetting evaluation should focus on whether the model remains sensitive to both historical baseline risk and newly emerging mobility decline.

Predictive metrics for fall risk

Predictive evaluation should assess whether the model supports clinically useful fall-risk stratification across age brackets, frailty states, and heterogeneous activity routines. Metrics such as discrimination, calibration, sensitivity at clinically acceptable specificity, and decision-support utility could be considered, but they should be interpreted in relation to care context rather than as isolated technical outputs. Reviews of wearable fall-risk assessment and wearable prediction models emphasise the need to connect sensor-based modelling with clinically meaningful validation [3,5]. Because this article is conceptual, such metrics are proposed as future evaluation criteria rather than reported results.

Convergence and communication efficiency

A practical evaluation should also assess whether the federated continual learning process can operate within smartwatch constraints, including communication cost, update frequency, memory use, and on-device computation. Federated wearable activity recognition has explored strategies such as dynamic layer sharing, prototype-guided personalisation, and similarity-aware clustering, all of which are relevant to reducing unnecessary communication while preserving adaptation [21,28,29]. Privacy-preserving mobile health learning further highlights the need to balance model updating with edge-device limitations [26]. The framework should therefore be evaluated not only for predictive behaviour but also for sustainability in real-world device ecosystems.

Limitations

Technical limitations

The proposed framework depends on smartwatch resources that may be limited by battery life, processor capacity, memory availability, and intermittent connectivity. A rehearsal buffer that is too small may fail to preserve sufficient historical mobility information, while a buffer that is too large may be impractical or raise privacy concerns. Non-IID wearable data may also reduce the stability of federated aggregation, especially when users differ in device placement, activity routines, frailty trajectories, and adherence to wearing the smartwatch [10]. These constraints mean that implementation should prioritise lightweight continual learning, robust aggregation, and careful privacy-preserving representation of historical movement patterns.

Clinical limitations

Fall risk prediction remains inherently uncertain because falls depend on environmental hazards, acute illness, medication changes, cognition, footwear, home layout, and chance events that may not be fully captured by wrist accelerometry. Alerts may therefore generate false positives or miss clinically relevant risk changes, and older adults may experience alert fatigue or discomfort with continuous monitoring. Clinical artificial intelligence quality improvement literature stresses the need for ongoing monitoring, governance, and validation after deployment [9]. Prospective clinical evaluation would be necessary before the framework could be treated as a reliable component of senior care decision support.

Conclusion

The proposed federated continual learning framework offers a conceptual pathway for adaptive maintenance of smartwatch-based fall risk prediction models across ageing populations. It treats fall risk as a dynamic process shaped by changing mobility, frailty progression, and evolving daily activity patterns. By combining local smartwatch inference, continual updating, and federated aggregation, the framework aims to support lifelong risk monitoring without centralising raw accelerometer data.

A key strength of the framework is its emphasis on continuous adaptation to age-related mobility shifts. Instead of assuming that a model trained at one point in time remains valid indefinitely, the system is designed to update as the user’s movement profile changes. Its continual learning component is intended to reduce algorithmic forgetting, while its federated structure supports privacy-preserving learning across distributed older adult populations.

Several technical and clinical challenges remain. Smartwatch memory, battery, connectivity, and computational limits may constrain how much continual learning can occur locally. Clinical deployment would also require transparent alert design, careful threshold calibration, user acceptance, clinician oversight, and prospective validation in real-world settings.

Future work should evaluate this framework through pilot studies using public fall detection resources and prospective deployments in senior care facilities. Such studies should examine whether federated continual learning can maintain adaptive, personalised, and privacy-preserving fall risk models over time. The long-term goal is not to replace clinical judgement, but to provide a continuously learning decision-support layer for earlier recognition of mobility decline and preventable fall risk.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 19 Jan 2026 Revised: 22 Jan 2026 Accepted: 03 Mar 2026
Published online: 20 July 2026

References

Howcroft J, Kofman J, Lemaire ED. Prospective fall-risk prediction models for older adults based on wearable sensors. IEEE Trans Neural Syst Rehabil Eng. 2017;25(10):1812-20.
https://doi.org/10.1109/TNSRE.2016.2639794.

Howcroft J, Kofman J, Lemaire ED. Feature selection for elderly faller classification based on wearable sensors. J Neuroeng Rehabil. 2017;14(1):47.
https://doi.org/10.1186/s12984-017-0255-9.

Bet P, Castro PC, Ponti MA. Fall detection and fall risk assessment in older person using wearable sensors: a systematic review. Int J Med Inform. 2019;130:103946.
https://doi.org/10.1016/j.ijmedinf.2019.103946.

Tunca C, Salur G, Ersoy C. Deep learning for fall risk assessment with inertial sensors: utilizing domain knowledge in spatio-temporal gait parameters. IEEE J Biomed Health Inform. 2020;24(7):1994-2005.
https://doi.org/10.1109/JBHI.2019.2948879.

Wang B, Liu Y, Lu A, Wang C. Application of wearable sensors in constructing a fall risk prediction model for community-dwelling older adults: a scoping review. Arch Gerontol Geriatr. 2025;129:105689.
https://doi.org/10.1016/j.archger.2024.105689.

Chen Y, Qin X, Wang J, Yu C, Gao W. FedHealth: a federated transfer learning framework for wearable healthcare. IEEE Intell Syst. 2020;35(4):83-93.
https://doi.org/10.1109/MIS.2020.2988604.

Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3(1):119.
https://doi.org/10.1038/s41746-020-00323-1.

Lee CS, Lee AY. Clinical applications of continual learning machine learning. Lancet Digit Health. 2020;2(6):e279-e281.
https://doi.org/10.1016/S2589-7500(20)30102-3.

Feng J, Phillips RV, Malenica I, Bishara A, Hubbard AE, Celi LA, et al. Clinical artificial intelligence quality improvement: towards continual monitoring and updating of AI algorithms in healthcare. NPJ Digit Med. 2022;5(1):66.
https://doi.org/10.1038/s41746-022-00584-6.

Criado MF, Casado FE, Iglesias R, Regueiro CV, Barro S. Non-IID data and continual learning processes in federated learning: a long road ahead. Inf Fusion. 2022;88:263-80.
https://doi.org/10.1016/j.inffus.2022.07.003.

Wu Q, Chen X, Zhou Z, Zhang J. FedHome: cloud-edge based personalized federated learning for in-home health monitoring. IEEE Trans Mob Comput. 2022;21(8):2818-32.
https://doi.org/10.1109/TMC.2020.3047205.

Pfitzner B, Steckhan N, Arnrich B. Federated learning in a medical context: a systematic literature review. ACM Trans Internet Technol. 2021;21(2):1-31.
https://doi.org/10.1145/3418292.

Lien WC, Ching CT, Lai ZW, Wang HM, Lin JS, Huang YC, et al. Intelligent fall-risk assessment based on gait stability and symmetry among older adults using tri-axial accelerometry. Front Bioeng Biotechnol. 2022;10:887269.
https://doi.org/10.3389/fbioe.2022.887269.

González-Castro A, Benítez-Andrades JA, González-González R, Prada-García C, Leirós-Rodríguez R. Predicting fall risk in older adults: a machine learning comparison of accelerometric and non-accelerometric factors. Digit Health. 2025;11:20552076251331752.
https://doi.org/10.1177/20552076251331752.

Wang T, Du Y, Gong Y, Choo KKR, Guo Y. Applications of federated learning in mobile health: scoping review. J Med Internet Res. 2023;25:e43006.
https://doi.org/10.2196/43006.

Pati S, Baid U, Edwards B, Sheller M, Wang SH, Reina GA, et al. Federated learning enables big data for rare cancer boundary detection. Nat Commun. 2022;13(1):7346.
https://doi.org/10.1038/s41467-022-33407-5.

Bruno P, Quarta A, Calimeri F. Continual learning in medicine: a systematic literature review. Neural Process Lett. 2025;57(1):2.
https://doi.org/10.1007/s11063-024-11634-4.

Wu X, Xu Z, Tong RK. Continual learning in medical image analysis: a survey. Comput Biol Med. 2024;182:109206.
https://doi.org/10.1016/j.compbiomed.2024.109206.

Yu H, Chen Z, Zhang X, Chen X, Zhuang F, Xiong H, et al. FedHAR: semi-supervised online learning for personalized federated human activity recognition. IEEE Trans Mob Comput. 2023;22(6):3318-32.
https://doi.org/10.1109/TMC.2021.3131778.

Presotto R, Civitarese G, Bettini C. Semi-supervised and personalized federated activity recognition based on active learning and label propagation. Pers Ubiquit Comput. 2022;26(5):1281-98.
https://doi.org/10.1007/s00779-021-01574-3.

Cheng D, Zhang L, Bu C, Wang X, Wu H, Song A. ProtoHAR: prototype guided personalized federated learning for human activity recognition. IEEE J Biomed Health Inform. 2023;27(8):3900-11.
https://doi.org/10.1109/JBHI.2023.3268172.

Chai Y, Liu H, Zhu H, Pan Y, Zhou A, Liu H, et al. A profile similarity-based personalized federated learning method for wearable sensor-based human activity recognition. Inf Manage. 2024;61(7):103922.
https://doi.org/10.1016/j.im.2024.103922.

Yu Z, Liu J, Yang M, Cheng Y, Hu J, Li X. An elderly fall detection method based on federated learning and extreme learning machine (Fed-ELM). IEEE Access. 2022;10:130816-24.
https://doi.org/10.1109/ACCESS.2022.3227088.

Ghosh S, Ghosh SK. FEEL: federated learning framework for elderly healthcare using edge-IoMT. IEEE Trans Comput Soc Syst. 2023;10(4):1800-9.
https://doi.org/10.1109/TCSS.2022.3211504.

Qi P, Chiaro D, Piccialli F. FL-FD: federated learning-based fall detection with multimodal data fusion. Inf Fusion. 2023;99:101890.
https://doi.org/10.1016/j.inffus.2023.101890.

Aminifar A, Shokri M, Aminifar A. Privacy-preserving edge federated learning for intelligent mobile-health systems. Future Gener Comput Syst. 2024;161:625-37.
https://doi.org/10.1016/j.future.2024.07.019.

Xiao Z, Xu X, Xing H, Song F, Wang X, Zhao B. A federated learning system with enhanced feature extraction for human activity recognition. Knowl Based Syst. 2021;229:107338.
https://doi.org/10.1016/j.knosys.2021.107338.

Ouyang X, Xie Z, Zhou J, Huang J, Xing G. ClusterFL: a similarity-aware federated learning system for human activity recognition. In: Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys). 2021. p. 54-66.
https://doi.org/10.1145/3458864.3466628.

Tu L, Ouyang X, Zhou J, He Y, Xing G. FedDL: federated learning via dynamic layer sharing for human activity recognition. In: Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems (SenSys). 2021. p. 15-28.
https://doi.org/10.1145/3485730.3485946.

Kumari P, Chauhan J, Bozorgpour A, Huang B, Azad R, Merhof D. Continual learning in medical image analysis: a comprehensive review of recent advancements and future prospects. Med Image Anal. 2025;103730.
https://doi.org/10.1016/j.media.2025.103730.

Qazi MA, Hashmi AU, Sanjeev S, Almakky I, Saeed N, Gonzalez C, et al. Continual learning in medical imaging: a survey and practical analysis. ACM Comput Surv. 2026;58(8):1-25.

Verma T, Jin L, Zhou J, Huang J, Tan M, Choong BC, et al. Privacy-preserving continual learning methods for medical image classification: a comparative analysis. Front Med (Lausanne). 2023;10:1227515.
https://doi.org/10.3389/fmed.2023.1227515.

Jagdeesh K, Kanimozhi N, Sardar TH, Naveenkumar N, Mahalakshmi B, Chandrasekar A, et al. Federated learning with continual update for privacy-preserving clinical event prediction across distributed hospitals using MCN-GNN. Sci Rep. 2026;13(1).

Zhao X, Liu Z, Ji B, Xi P, Peng S. TransEHR: alignment-free electronic health records continual learning across feature spaces. Expert Syst Appl. 2025;129020.
https://doi.org/10.1016/j.eswa.2025.129020.

Haescher M, Chodan W, Höpfner F, Bieber G, Aehnelt M, Srinivasan K, et al. Automated fall risk assessment of elderly using wearable devices. J Rehabil Assist Technol Eng. 2020;7:2055668320946209.
https://doi.org/10.1177/2055668320946209.