Investigating the Applicability of Gait-based Health Assessment in a Domestic Environment

Chris Lochhead[1,2][0000−0001−9046−6784], Longfei Chen[1][0000−0002−3935−8021], Robert B. Fisher[1][0000−0001−6860−9371], and Rhona Lochhead

[1] University of Edinburgh, Edinburgh, EH8 9YL, UK
[2] christopher.lochhead@ed.ac.uk

Abstract. At-home health monitoring technologies have extensive use-cases, including gait assessment, disease diagnosis, gait anomaly detection and disease severity quantification. Despite the volume of research into this domain, novel machine learning models and data collection pipelines are often benchmarked on curated datasets, collected in laboratory conditions and annotated by specialists. This separation between the experimental and practical domain is apparent in the lack of inter-disciplinary work in the physical design of these systems. In this work, a series of interviews with 10 healthcare professionals is conducted to co-design an at-home gait monitoring system, with the initial prototype informed by 3 focus groups of older adults. A series of recording experiments are then conducted with the prototype in the homes of 10 older adults to investigate its effectiveness and feasibility when subjected to the conditions of genuine at-home data recording. With novel ST-GCN based models achieving up to 0.93 f1 on the task of gait recognition using only data collected at home, it can be asserted that the developed gait monitor prototype is capable of reliably collecting data of sufficient quality to build accurate profiles of gait. This is achieved despite the design concessions made in the interest of stakeholder acceptability and the challenges presented by at-home gait data.

Keywords: Computer Vision · Gait Assessment · Health Monitoring.

1 Introduction

Gait-based health analysis is far from a new field of study. Before the advent of machine learning, medical science has long known of the descriptive power of gait to indicate a variety of health conditions [Khera and Kumar, 2020]. Research in machine learning-based gait assessment has expanded, with multiple monitoring paradigms from vision-based [Yan et al., 2018, Sabo et al., 2022, Yin et al., 2023], wearables [Fernandes et al., 2018, Jung et al., 2023], and other sensors [Turner and Hayes, 2019, Clever et al., 2022] emerging as a means of developing gait assessment tools to diagnose disease and assess general health.

Despite the advances in this field, with multiple recent studies demonstrating accuracy approaching clinical levels for disease diagnosis [Guo et al., 2022], disease severity assessment [Fernandes et al., 2018] and fall detection [Wang et al., 2020], much of this research is under-utilized. A common shortcoming in much of the literature is the lack of multidisciplinarity [Lochhead and Fisher, 2025]. Most technical papers are designed without any input from healthcare professionals, with their participation often limited to the annotation of datasets [Lochhead et al., 2025, Sabo et al., 2022]. As a result, many studies of end-users of older adult care systems find that issues such as cost, stigma and impracticality are often common features in these systems, with little to no consideration given for these design barriers.

Presented in this work is a prototype at-home gait monitoring system, developed through both public and patient interaction (PPI) workshops with older adults as well as qualitative interviews with healthcare professionals from a variety of backgrounds associated with older adult care. The gait monitor prototype employs a similar data collection and pre-processing method to [Lochhead and Fisher, 2025] to produce more reliable gait data despite the challenges of recording in different homes with a single camera at different viewpoints. The utility of the collected data is then evaluated via a supervised person-recognition task across participants using the state-of-the-art machine learning model from Lochhead and Fisher [2025]. The research question for this work can be defined as "Can a prototype gait monitor system be developed that satisfies the technical requirements to achieve clinical levels of gait recognition accuracy while satisfying the requirements identified by different groups of stakeholders that would be likely to use this system in practice?". The contributions of this work can be summarized as follows:

- A prototype system for ambient at-home vision-based gait monitoring, co-designed through a multidisciplinary approach with healthcare professionals.
- An analysis of the preferences of different healthcare professions (n = 10) when considering the development of at-home health monitors.
- A co-designed output document developed to provide explainability to the gait evaluation and hence greater usability and trustworthiness.
- An experiment and analysis of the prototype in the homes of older adults (n = 10) to assess its effectiveness and utility in practice using the supervised gait recognition task as the evaluation domain.

The literature on gait-based health monitoring is reviewed in Section 2. The method of co-design employed for the design and implementation of the gait monitoring system is then introduced (see Section 3.1), along with the structure of the live experiment of the system (see Section 3.2). Then, the results of the live experiment are discussed along with an analysis of the priorities of the various healthcare professions during the design phase (see section 4). This work concludes with a discussion on the successes and shortcomings of this research and potential future work (see Section 5).

2 Related Work

2.1 Analysis of Existing Live Studies

There is a growing volume of qualitative analysis of the opinions of older adults and other end-users regarding various at-home monitoring technologies. The most comprehensive reviews in this field are typically on the "Smart home"; defined as a technological framework established in a synthetic or real home, in which multiple overlapping monitoring and assistive technologies can be used simultaneously [Liu et al., 2016, Majumder et al., 2017, Ghorayeb et al., 2021, Wang et al., 2020]. [Majumder et al., 2017] investigates overlapping health monitoring methods involving wearable sensors, focusing on the technical metrics of these devices, such as the precision of the measurement. Although they do not consult directly with prospective end users, they do identify issues such as privacy and cost as potential factors that impact uptake and effectiveness.

The smart home review in [Ghorayeb et al., 2021] provides information on the human aspect of the applicability of these technologies. Not only do they identify commonly identified human barriers such as privacy and stigma, they also identify potential solutions. For example, providing manual override buttons for end users to "pause" the system to provide a greater feeling of control and privacy. [Morita et al., 2023] identify that a possible root cause of the lack of end-user analysis is that the majority of technical research into monitoring technologies lack a multidisciplinary approach. Although there exist some reviews in the literature regarding the opinions of end-users, they did not find reviews of articles on the opinions and preferences of healthcare professionals for at-home medical monitoring devices. The findings in [Lochhead et al., 2025] concur with this assessment and are able to qualitatively measure a significant drop in the quality of outcomes of single-discipline versus multi-discipline research where the goal is developing health monitoring systems for use in the real world. Using their metric, they find single-discipline papers score on average 3.06/10 versus 5.43/10 for multidisciplinary work.

A vital component of this type of research is domain-specific experimentation carried out in controlled settings. Although many ML-based approaches employ pre-recorded videos annotated by experts [Sabo et al., 2022, Yan et al., 2018, Guo et al., 2022] or use datasets crafted in a laboratory setting using simulated gait anomalies [Yin et al., 2023, Ortells et al., 2018], there is a movement in the research toward so-called "live studies". These are studies in which the data collected is of participants who belong to the target demographic of the system and tested outside of the laboratory, typically in the participants' own home. For example, researchers in [Pais et al., 2020] conduct a year-long study using ambient and wearable sensors for their 'Domo' system for monitoring various daily events of older adults as a means of allowing nursing staff to more comprehensively monitor their health.

2.2 Machine Learning and Visualization for Gait Analysis

Much of the recent literature on computerized gait analysis rely on ML solutions using different data collection and representation paradigms, with the dominant form of gait analysis using wearable sensors [Wang et al., 2021, Khera and Kumar, 2020]. [Khera and Kumar, 2020] however, identifies a minority (12%) of research in gait analysis for healthcare being carried out using computer vision-based approaches.

One promising method for ML-based gait analysis employs the use of Spatio-temporal Graph Convolutional Networks (ST-GCN) [Yu et al., 2017]. [Yin et al., 2023] make novel contributions to both the data representation and the ST-GCN itself. Their Spatio-temporal joint-adjacency GCN employs 3 streams: one for joint co-ordinates, velocities and a third novel stream using the calculated vectors between connected joints as bone angle representations. They also append the ST-GCN with an attention module to achieve 93.17% accuracy in classifying gait abnormalities on a dataset of healthy participants acting out various common gait conditions (n = 22). Research in [Lochhead and Fisher, 2025, 2024], simplify this general pipeline, reducing model sizes and input streams to achieve state-of-the-art results on gait assessment (94.38% and 97.00% respectively) on systems that are deployed to collect data using constraints not dissimilar to those likely to be found in an at-home setting (single camera viewpoints, cluttered backgrounds, no requirements for clothing, environmental obstacles).

An important component needed for gait analysis frameworks to be effective is explainability. [Gupta et al., 2023] identifies several articles involving Parkinson's diagnosis and severity estimation using interactive online tools that process a variety of input modalities from brain imaging data to voice recordings. The majority of this research provides little focus on visualizations for communicating specific findings, instead focusing on the interactive element of these tools, such as chat-bots and web applications. In [Oliveira et al., 2018], researchers utilize T-SNE to visualize the separation between different classes of data for diagnosing Parkinson's Disease. While this method is potentially more intuitive, it still relies on prior knowledge of T-SNE to be understandable, and lacks input from healthcare professionals regarding the practicality of these visualizations. Researchers in [Jung et al., 2023] develop a novel segmented pie chart style visualization that explains specific abnormalities in gait. While their approach is demonstrably more oriented toward being understandable to lay people, they don't indicate in the development of this methodology any input or analysis from healthcare professionals to justify either their design choices or the data categories they use in their outputs.

3 Methodology

3.1 Healthcare Professional Interviews

To achieve the first three contributions outlined in Section 1, the concerns regarding usability and acceptability must be addressed. Users for this device fall into two broad categories; healthcare professionals and patients. As identified in the previous section, there exists a strong trend in the literature for assessing the opinions and considerations of the latter group, with far less, if any, of the former. To address this shortcoming, a series of interviews were conducted to identify and address the key issues to consider from the point of view of healthcare professionals for this at-home monitoring system. Before undertaking these interviews, a series of 3 Patient and Public Involvement (PPI) workshops were carried out to inform the development of an initial gait monitoring prototype, as well as an initial output document (see figure 1). This prototype served as a demonstration during the interviews for healthcare professionals to base their feedback on. These focus groups were semi-structured group discussions in which gait monitoring technology was discussed, particularly their current technical requirements. From this setting, participants provided opinions and ideas regarding the creation of an acceptable gait monitor prototype. The most prominent themes throughout these workshops were concerns around privacy, stigma and a sense of burden. Each workshop had between 6-12 participants and occurred between June 2022 to August 2023. Participants were recruited from a variety of sources, namely age-related disease charities in the UK as well as those recruited from University networks and local community groups in Edinburgh. All participants were either an older adult (defined as $>$65 years of age) or a caregiver of an older adult with an age-related disease (spouses and children, for example). These initial workshop sessions formed the basis for the initial prototype described in this section.

A total of 10 healthcare professionals were interviewed between August 2023 and February 2024. To maximize the descriptive power of the interview results, people from a variety of disciplines related to older adult care were recruited. All participants were currently working healthcare professionals, working with older adults in a role where assessment of gait was part of their duties. The breakdown of professions was: one Geriatrician, one care home nursing manager, two Nurses, two care workers, one mental health Nurse, one Occupational Therapist, one Physiotherapist and one junior Doctor.

The interviews were carried out 1-to-1 in the place of work of the healthcare professional. The interviews lasted 30-45 minutes and consisted of two phases; a presentation followed by a semi-structured conversation. The presentation provided a layperson breakdown of the gait monitor prototype and what the purpose of the participant's involvement is. They were given a high-level overview of the gait monitoring software's capabilities, concluding with a concept illustration of the output gait analysis file that the gait monitor prototype could produce. See figure 1 for a comparison of the version of the output used during the interviews with the final version made after collating the feedback from the interviews.

The co-designed output document (figure 1) contains six sections of information that could be produced by the data collected by the platform (see Section 4.2) along with two sections for the purposes of interpretability. From top to bottom, left to right these are:

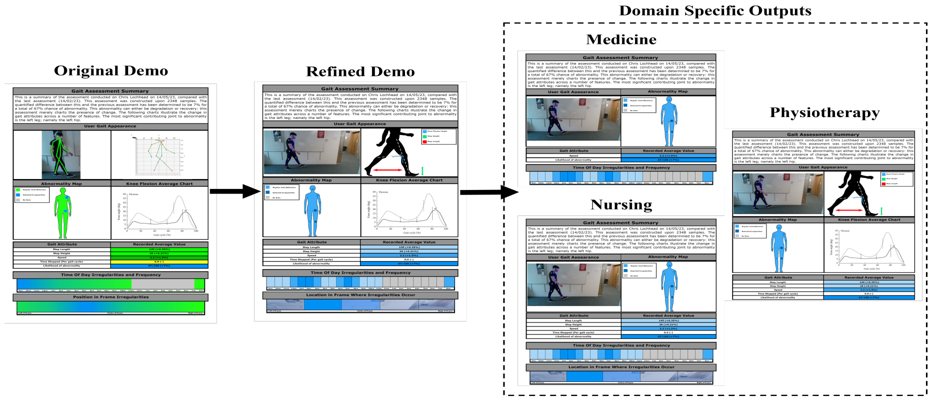

Fig. 1: Comparison of the original demo output synthesized with input from PPI workshops of older adults (left) and the final refined system output following integrating the suggestions made by healthcare professionals (centre). The documents on the right illustrate the profession-specific documents.

1. **Patient description:** contains information about the data, namely the person recorded, the dates of the recordings, the number of instances recorded and a percentage metric summary of the system's belief that the participant's gait has changed from a prior period of monitoring.
2. **Patient image:** an example frame of the patient to confirm the patient is correct.
3. **A legend for gait metrics:** namely knee flexion, step height and width for interpretability of sections 5 and 6.
4. **Abnormality heatmap:** darker spots indicating areas of greatest importance to the change in gait. The technology to produce this type of output given limited graph-based gait input is demonstrated in [Lochhead and Fisher, 2024].
5. **Knee flexion chart:** plotting the angle of the knee across one gait cycle produces a chart which can be directly compared to the baseline.
6. **A table of numerical attributes such as speed and step height:** 'likelihood of abnormality' corresponds to the confidence score produced by [Lochhead and Fisher, 2024] indicating the confidence in the system that the most recent gait assessment is different from the prior session.
7. **Timeline of irregularities:** a heatmap that indicates the time of day that abnormal gait events occur.
8. **Location of irregularities:** A heatmap overlay of a background example. Taken by grouping the mean severity scores from [Lochhead and Fisher, 2024] by their position in the field of view, with the intention of illustrating if abnormal gait events are being influenced by environmental obstacles.

After the presentation, the output file was left open for the participant to refer to and a semi-structured conversation was conducted following the prompts outlined in Table 1 of the supplementary materials. The prompts were designed to address two core themes: usability and acceptability. Users were asked about their own opinions on the system and the output file it produced, specifically the usability of the system and any issues surrounding convenience. They were also asked which sections were and weren't useful; opinions varied depending on profession (see areas of divergence and consensus on this topic in section 4.2). Lastly, the conversation was steered towards barriers to uptake both for them and any theoretical issues they could foresee their patients encountering.

Following the interviews, the output file was updated from the original to a refined version (see Figure 1 center) with the option for bespoke smaller versions of the output (right) which were designed for different healthcare professions. See Section 4.2 for an overview of these compressed output documents.

3.2 Live Study

To address the final contribution outlined in Section 1, a limited live study with 10 older adults across 7 homes was conducted. Recordings were carried out between March 2024 and March of 2025, with the participants all being over 60 and not suffering from any disease where impacted gait is a symptom.

The monitor was placed according to the preferences of the participant. The prototype hardware consists of one Jetson Nano computing device and one Intellisense RGBD camera. The participants were asked to pick somewhere in their home to place the system, with the only requirement being that it was within 1m of an electrical socket. Minor edge-cases were accounted for when deploying the system, such as avoiding large pictures of people/faces within the view of the system. Hanging jackets or other clothing were also avoided from the field of view of the device where possible to reduce the number of incorrect person detection instances.

The recording system is designed to be lightweight and robust to edge cases to satisfy the requirements for diverse recording environments. It runs automatically upon activation of the Jetson Nano, meaning participants had the option to activate/deactivate the system at will without requiring researcher input. Due to the low computational power and storage capabilities of the Jetson Nano (less than 32GB of memory and no on-board GPU), paired with having only a single viewpoint from a single camera, a number of efficiency modifications were made to the system (see Table 2 in the supplementary materials).

After quality assurance and preprocessing, all data are stored as gait graphs with original images discarded to address privacy concerns. Once placed, the system was left to run for between 3 and 7 days. All but 3 of the 10 participants kept the device for the entire 7 days. After the agreed-upon time period, the prototype was retrieved and the collected data transferred to a PC to process the recordings into gait graph data via pose-estimation and outlier removal. The refined data was then applied to a supervised gait-based person recognition task to assess the data quality. The pre-processing stage removed many unsuitable sequences, such as removing any sequences with more than one person in it or sequences where the participant did not walk more than 50% of the way through the frame. This latter requirement was designed to account for camera placement that was in the view of a seating area, to prevent the inclusion of participants sitting down into the final dataset.

The recording system was based on the setup used in Lochhead and Fisher [2025], in which the Jetson Yolov5 model was used for person detection Omidi et al. [2021]. This model was implemented using the yolov5n weights, the smallest available pre-trained configuration. The system is set to detect humans and disregard any detections below 85% confidence. This configuration was found to be sufficient to operate on the prototype gait monitor and account for the majority of false detections. Positive detections of humans were then recorded for 5 seconds and the resulting video was stored as an individual gait instance for later processing into gait graphs.

Further data decimation, pose estimation and feature extraction was not carried out at the edge; rather the collected data was transferred to a PC after recording to maintain the frame rate of the gait monitor.

An automated video clip deletion system removed two prominent edge-cases; if there is more than one person in the video clip and if the person fails to traverse at least 50% of the field of view of the monitor. Despite these edge-case detector algorithms removing around 12.5% of the data, a manual pruning phase was still necessary, in which a researcher manually verifies the collected data to be of sufficient quality. A significant volume of examples were removed (around 85% of the remainder). The most frequent reasons were people walking too close to the camera preventing an accurate pose-estimation, other forms of occlusion, environmental false positives (such as hanging jackets) and people who were not the participant (usually visitors) appearing in frame.

Each model was trained using a 70:20:10 train, validation and test ratio for 120 epochs with mean, f1 and accuracy calculated across 5 stratified folds. This matches the training setup used in [Lochhead and Fisher, 2025] and [Lochhead and Fisher, 2024]. Given that person recognition performance is similar across all datasets, it can be hypothesized that the at-home data is likely to be of sufficient quality that gait abnormalities would be detectable when recording and processing at-home data that exhibited genuine abnormalities.

4 Results

This section first addresses the analysis of the opinions of healthcare professionals on the usability and acceptability of home health monitoring technology. The goal of this analysis is to determine whether the developed prototype is both acceptable from an older adult's point of view whilst still capable of collecting enough quality data to perform accurate gait analysis. An assessment of the collected gait data is then performed using a supervised gait recognition task to establish the quality of the collected data using the prototype.

Table 1: Overview of the most prominent sources of praise and concern identified during the interviews with the healthcare professionals.

Praise	**Concern**
System is likely to be acceptable to a wide range of people within the age demographic.	From a design POV, the initial color scheme is not friendly to the color-blind, and the green-blue range suggests green-good, whereas the model is only supposed to indicate change.
System is clearly unobtrusive and easy to manage, not requiring training or maintenance. This is especially clear upon briefing by a researcher about the nature of the data collected.	There were mutually exclusive metrics of interest to one but not necessarily all groups. For example, knee-flexion was of interest to the Physiotherapist but not the Care workers, leading to unnecessary report content.
The output document is largely intuitive and contains at least some important and easily interpretable data for each healthcare discipline.	The length of the output may make real-life use difficult, as GPs, Nurses, and Care workers all highlighted the limited amount of time they'd have to analyze the results.
The system can be positively conveyed as a means for older adults to look after themselves rather than a device to monitor them for healthcare teams.	More attention should be afforded to developing a system that accommodates those who are less confident using novel technology or more reluctant to use it.

4.1 Analysis of Healthcare Professional Interviews.

Table 4.1 lists the key points of praise and concern for the existing system from the point of view of healthcare professionals. Participants were universally supportive of the robustness of the system using only one camera and not requiring user input or maintenance. In four interviews, comparisons were drawn between this system and monitoring systems that participants already had experience with. A common concern was how often these systems are deemed ineffective due to the burden on the patient to remember to charge or carry the system. Other areas of support included the design of the output document being intuitive, the metrics the collected data can produce and the potential for this system to improve the precision of care.

Multiple common concerns are also identified, for example the color scheme of the output document with respect to color-blind users. Other concerns included the need for additional training to use the system and concerns with uptake due to societal concern about AI and privacy. Different comments were made following a common theme; namely that certain portions of the output file were less useful than others to specific disciplines. One participant made the following recommendation:

> *The knee-flexion part and the numerical metrics aren't terribly important to me, and I would worry that having the time to even glance at this entire document wouldn't be feasible.*

This reflects a trend identified in other interviews that the output document could be optimized and that certain sections were only useful to certain professions. For example, the knee-flexion portion was only viewed as useful by participants with a background in physiotherapy and medical doctors were most prominently in favor of the table of numerical output metrics. One criticism made by a participant with a background in NHS governance and management concerned ethical barriers to uptake in application:

> *NHS technology governance and data sharing regulations could cause problems if you wanted to integrate something like this into current methods of care.*

This concern was borne out of the manner of data collection, being an at-home monitor capable of recording raw images of people's faces and the interiors of their homes. More specifically, there was concern surrounding the first section of the output document revealing a person's face and living space which end-users may wish to opt out of.

The changes in the output document from the one shown to healthcare professionals during the interviews to its final iteration are illustrated in Figure 1. A domain-specific modular design is adopted, producing smaller bespoke outputs produced for different healthcare professions. The color scheme has been changed to be monochromatic to address legibility concerns for color-blindness and communicate that changes in gait are not necessarily negative. For example, changing gait in someone recovering from a broken hip would likely be a positive indicator whereas it would be negative in someone suffering from a dementia-related illness. How to interpret any changes in gait is the domain of the healthcare professional, so it was deemed appropriate design practice to avoid any aspects of design which may impact a users' assessment of the data. The content that appears in each domain-specific output was compiled from consensus from each profession. Knee-flexion charts for example appear only for Physiotherapy where they were explicitly commented on positively. The prototype system itself was modified based on these healthcare professional interviews. Firstly, the minimum original recording time was reduced from 5 days to 3 days from the initial proposal. This was due to the consensus that 3 days is considerably easier for the patient and will produce more than the estimated 10 gait instances required to build a gait profile to support supervised person recognition. When installing the system, the decision was also made to take additional equipment, namely a keyboard and monitor to demonstrate to the end-user in the live study what was being recorded and showing them examples of the data that would ultimately be retained. A universal thread of contention in the interviews was effectively communicating the data recording methodology to make people more comfortable with the system and feel more in control.

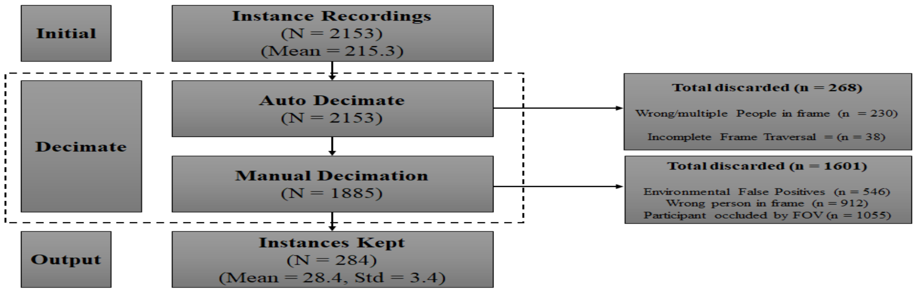

Fig. 2: Flow diagram of the decimation process to obtain high quality gait examples from the domestic environments. The stated mean refers to the mean per participant.

4.2 Utility of the At-home Monitor

The gait monitor system was designed to satisfy a series of requirements derived from the healthcare professional interviews, a summary of which is available in Table 3 of the supplementary materials.

The data collection phase was carried out with the goal of ensuring the collection of quality data with as little intrusion as possible into the lives and comfort of the older adult participants.

Table 2: Mean f1-scores across models on the 5-fold gait assessment (Gait) and person recognition (Person) tasks. Top performance in each column highlighted in bold. Performance on the At-home dataset is comparable to those on the synthetic gait anomaly datasets.

Model	WeightGait		Pathology		Shoe		At-home
	(Gait)	(Person)	(Gait)	(Person)	(Gait)	(Person)	(Person)
K-means	0.34	0.21	0.07	0.18	0.1	0.14	0.09
Logistic Regression	0.53	0.75	0.44	0.89	0.11	0.96	0.63
ST-GCN	0.91	**0.97**	0.89	0.90	0.72	**0.97**	**0.94**
ST-JAGCN	0.83	0.96	0.89	**0.94**	0.72	0.96	0.93
ST-TAGCN	**0.92**	0.96	**0.91**	0.90	**0.74**	0.96	0.93

Across the homes of 10 participants, 2153 example walking instances were detected and deemed of sufficient quality by the monitoring system over a recording period of between 3-10 days (see Figure 2).

After the unusable video clip deletion process was completed, the image data were converted into gait graphs of the format detailed in [Lochhead and Fisher, 2025]. These gait graphs were then normalized, scaled and outliers removed.

Despite the large volume of low-quality or false positive data, a mean 32 instances of regular gait were detected per person (minimum=13, maximum=69).

The right-hand column of Table 2 gives the results on the at-home data on the person recognition task. Person recognition in this context is defined as classifying which person each gait graph instance belongs to. This was carried out using a variety of both simple (K-means, Logistic Regression) and more complex models used in gait recognition and assessment (namely ST-GCN based neural networks). Each model was tested using the format of gait graph they were originally designed for (e.g. ST-GCN uses a joint co-ordinate input stream, ST-TAGCN uses a joint velocity stream). With an average f1 score of 0.93 using a basic ST-GCN, there is a strong indication that different gaits in this dataset fall within the range of discernability (0.90 - 0.97) compared to similar synthetic gait datasets. This demonstrates that at-home data is sufficiently rich to distinguish distinct gait patterns, suggesting the potential for such technology to be applicable to downstream gait assessment tasks. There was concern that the model may prioritize learning the differences in walking pattern in the different environments rather than gait to distinguish between people (e.g. walking direction), however ablations on the 3 sets of 2 participants who live in the same home demonstrated between 0.91-0.93 f1 on the binary person recognition task, supporting the assertion that the model is distinguishing classes by gait patterns rather than environmental patterns.

5 Discussion

This research introduced a prototype system for low-cost, effective and robust at-home gait monitoring, supported by PPI with focus groups of older adults, a co-design phase with a variety of healthcare professionals (n = 10) and a pilot study for data collection conducted on healthy older adults inside their own homes (n = 10). The results indicate that such a system can be acceptable to a broad range of stakeholders, allowing them to integrate such a system into their methods of care. It is also observed that the data collected within the constraints of a single camera, a minimum of 3 days of recording time and limited computational power storage is capable of producing high quality gait graph data. This data is verified as being of sufficient quality using the supervised person recognition task across a total of four datasets, where a variety of models produce comparable results both on controlled gait datasets and the at-home dataset produced in this research. As more data per person (on average) was collected than was used in [Lochhead and Fisher, 2025] for a specific gait profile (32 vs 10 instances per person) as well as achieving 0.94 mean f1 on the supervised gait recognition task, the likelihood of this system being able to detect and quantify changes in a person's gait over time is high. The logical next step is to produce a more longitudinal study featuring multiple 3 day recording sessions per person over a period of months, using participants either recovering from injury or suffering from age-related degenerative diseases like Parkinson's to determine the effectiveness of this system at supporting continuous gait assessment.

There are some shortcomings in this research, the first of which is the limited number of participants for both the co-design phase and live study. Conducting multiple visits to each participant to build multiple gait profiles over time would also further solidify the claim that the data is high enough quality to consistently recognize a person's gait. This could be compared with ground truths of participants recorded in a neutral location to ascertain if the specificity of the person's home plays a part in the high accuracy of person-recognition as opposed to simply their gait. This however is unlikely, given that all gait graphs are 3D and as such should be largely viewpoint invariant.

Acknowledgments. This research was funded by the Legal & General Group (research grant to establish the independent Advanced Care Research Centre at University of Edinburgh). The funder had no role in the conduct of the study, interpretation or the decision to submit for publication. The views expressed are those of the authors and not necessarily those of Legal & General. To the authors knowledge there are no conflicts of interest during the undertaking of this research.

Disclosure of Interests. The authors of this work declare that they have no competing interests.

Bibliography

Henry M Clever, Patrick L Grady, Greg Turk, and Charles C Kemp. Bodypressure-inferring body pose and contact pressure from a depth image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):137–153, 2022.

Carlos Fernandes, Luís Fonseca, Flora Ferreira, Miguel Gago, Lus Costa, Nuno Sousa, Carlos Ferreira, Joao Gama, Wolfram Erlhagen, and Estela Bicho. Artificial neural networks classification of patients with parkinsonism based on gait. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 2024–2030. IEEE, 2018.

Abir Ghorayeb, Rob Comber, and Rachael Gooberman-Hill. Older adults' perspectives of smart home technology: Are we developing the technology that older people want? *International journal of human-computer studies*, 147: 102571, 2021.

Rui Guo, Jie Sun, Chencheng Zhang, and Xiaohua Qian. A self-supervised metric learning framework for the arising-from-chair assessment of parkinsonians with graph convolutional networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9):6461–6471, 2022.

Rohan Gupta, Smita Kumari, Anusha Senapati, Rashmi K Ambasta, and Pravir Kumar. New era of artificial intelligence and machine learning-based detection, diagnosis, and therapeutics in parkinson's disease. *Ageing research reviews*, page 102013, 2023.

Sylvain Jung, Nicolas de l'Escalopier, Laurent Oudre, Charles Truong, Eric Dorveaux, Louis Gorintin, and Damien Ricard. A machine learning pipeline for gait analysis in a semi free-living environment. *Sensors*, 23(8):4000, 2023.

Preeti Khera and Neelesh Kumar. Role of machine learning in gait analysis: a review. *Journal of Medical Engineering & Technology*, 44(8):441–467, 2020.

Lili Liu, Eleni Stroulia, Ioanis Nikolaidis, Antonio Miguel-Cruz, and Adriana Rios Rincon. Smart homes and home health monitoring technologies for older adults: A systematic review. *International journal of medical informatics*, 91:44–59, 2016.

Chris Lochhead and Robert B. Fisher. Towards explainable semi-supervised graph embeddings for gait assessment using per-cluster dimensional weighting. *arXiv preprint arXiv:1812.00169*, 2024.

Chris Lochhead and Robert B. Fisher. Lightweight human gait anomaly assessment using single-stream uniform temporal attention. *Computers in biology and Medicine*, 190, 2025. https://doi.org/https://doi.org/10.1016/j.compbiomed.2025.110076.

Chris Lochhead, Robert B Fisher, et al. On the necessity of multidisciplinarity in the development of at-home health monitoring platforms for older adults: Systematic review. *JMIR Human Factors*, 12(1):e59458, 2025.

Sumit Majumder, Tapas Mondal, and M Jamal Deen. Wearable sensors for remote health monitoring. *Sensors*, 17(1):130, 2017.

Plinio P Morita, Kirti Sundar Sahu, and Arlene Oetomo. Health monitoring using smart home technologies: Scoping review. *JMIR mHealth and uHealth*, 11:e37347, 2023.

Fábio Henrique M Oliveira, Alessandro RP Machado, Adriano O Andrade, et al. On the use of t-distributed stochastic neighbor embedding for data visualization and classification of individuals with parkinson's disease. *Computational and mathematical methods in medicine*, 2018, 2018.

Abbas Omidi, Amirhossein Heydarian, Aida Mohammadshahi, Behnam Asghari Beirami, and Farzan Haddadi. An embedded deep learning-based package for traffic law enforcement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 262–271, 2021.

Javier Ortells, María Trinidad Herrero-Ezquerro, and Ramón A Mollineda. Vision-based gait impairment analysis for aided diagnosis. *Medical & biological engineering & computing*, 56:1553–1564, 2018.

Bruno Pais, Philipp Buluschek, Guillaume DuPasquier, Tobias Nef, Narayan Schütz, Hugo Saner, Daniel Gatica-Perez, and Valérie Santschi. Evaluation of 1-year in-home monitoring technology by home-dwelling older adults, family caregivers, and nurses. *Frontiers in public health*, 8:518957, 2020.

Andrea Sabo, Sina Mehdizadeh, Andrea Iaboni, and Babak Taati. Estimating parkinsonism severity in natural gait videos of older adults with dementia. *IEEE journal of biomedical and health informatics*, pages 2288–2298, 2022.

Alexander Turner and Stephen Hayes. The classification of minor gait alterations using wearable sensors and deep learning. *IEEE Transactions on Biomedical Engineering*, 66(11):3136–3145, 2019.

Ju Wang, Nicolai Spicher, Joana M Warnecke, Mostafa Haghi, Jonas Schwartze, and Thomas M Deserno. Unobtrusive health monitoring in private spaces: The smart home. *Sensors*, 21(3):864, 2021.

Xueyi Wang, Joshua Ellul, and George Azzopardi. Elderly fall detection systems: A literature survey. *Frontiers in Robotics and AI*, 7:71, 2020.

Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

Ziming Yin, Yi Jiang, Jianli Zheng, and Hongliu Yu. Stja-gcn: A multi-branch spatial–temporal joint attention graph convolutional network for abnormal gait recognition. *Applied Sciences*, 13(7):4205, 2023.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.