FORTif.ai: A Multimodal Platform for Safer Independent Living

Akil Giri
University of Waterloo
akgiri@uwaterloo.ca

Chris Jackson
University of Waterloo
c8jackson@uwaterloo.ca

Edson Kenzo Takei
York University
etakei@my.yorku.ca

Jessi Huang
University of Waterloo
j658huang@uwaterloo.ca

Lino Kee
University of Waterloo
lino.kee@uwaterloo.ca

Marco Kee
Queen's University
22cvt2@queensu.ca

Michelle Steen
University of Waterloo
mksteen@uwaterloo.ca

Mohammed Elshrief
University of Waterloo
mkelshrief@uwaterloo.ca

Rohan Tuli
University of Waterloo
rtuli@uwaterloo.ca

Sarvesh Sekar
University of Waterloo
s4sekar@uwaterloo.ca

Sidney Liu
York University
sidney1@my.yorku.ca

Spencer Spiegelman
University of Waterloo
smspiegelman@uwaterloo.ca

Sri Meghana Yarlagadda
University of Waterloo
sm2yarlagadda@uwaterloo.ca

Plinio P. Morita
University of Waterloo
plinio.morita@uwaterloo.ca

***Abstract*—Global aging trends present a growing dichotomy: while older adults prefer to age in place, age-related declines in mobility and cognition can make home environments more hazardous and increase the risk of falls. This paper introduces FORTif.ai, an integrated platform that combines (1) LLM-based hazard detection to identify household risks from images and video and (2) an empathetic, LLM-powered chatbot for personalized support. The chatbot uses a retrieval-augmented generation (RAG) pipeline supported by a Weaviate vector database and a four-phase workflow that incorporates hybrid search with thresholded retrieval. We present hazard detection evaluation results from preliminary investigation, across five videos, along with additional image-based testing that emphasizes key behavioral patterns and prompt-driven enhancements in detecting floor-level clutter and common tripping hazards. Collectively, these findings demonstrate a proactive approach to home safety that integrates real-time environmental sensing with conversational support tailored to older adults.**

I. Introduction

Today, approximately 7.6 million Canadians are aged 65 and older, accounting for roughly one-fifth (18.9%) of the total population. This number is expected to grow by 2030, reaching between 21.4% and 23.4%, further establishing Canada as an aging society [1]. Despite this growing population, older adults continue to face significant challenges in maintaining independence, often with limited support. A study conducted by Statistics Canada in 2019 and 2020 found that nearly one in five Canadian seniors (19%) aged 65 and older reported experiencing loneliness, with higher rates among senior women compared to senior men (23% vs. 15%) [2]. Older adults living alone are also at increased risk of fall-related injuries, which are the leading cause of injury-related hospitalizations, with 20–30% of seniors experiencing one or more falls each year [2]. Sensory decline and chronic health conditions are major contributors to falls, as poor vision or hearing can make it more difficult to identify and navigate household obstacles [3]. FORTif.ai is an integrated platform designed to support safer independent living for older adults by combining real-time hazard detection with personalized conversational support. The system represents a comprehensive approach to senior care technology, integrating three core components: an intelligent chatbot, a persistent memory management feature, and a LLM-based hazard detection system. The chatbot component leverages large language models (LLMs) to deliver natural, empathetic interactions, assisting seniors with daily tasks, providing reminders, accessing personal information, and maintaining social engagement. The memory management feature employs a vector database to store and retrieve personalized memories, enabling contextually aware conversations that are particularly beneficial for individuals experiencing cognitive impairments. The hazard detection system utilizes vision-language models to continuously monitor the living environment, identifying common household risks,including floor-level clutter, rugs, and other tripping hazards,and providing actionable real-time alerts and recommendations. FORTif.ai is designed with a multimodal interface, supporting both voice and text inputs, to ensure accessibility for users with varying levels of technological proficiency. By integrating these components into a single platform, FORTif.ai not only aims to reduce the risk of home accidents but also seeks to improve medication adherence, enhance social connectivity, and empower seniors to maintain independence while ensuring safety and well-being. In this paper, we introduce readers to the system and report results from hazard detection testing on five videos, supplemented with additional image-based evaluations. These evaluations include prompt iterations that enhanced detection performance for floor-level hazards. The work addresses the growing need for proactive assistance in senior care, responding to the prevalence and severe consequences of falls in the home, as well as the limitations of existing technologies.

A. Related Works

Prior work on supporting older adults' safety and independence spans environmental sensing, interactive assistance, and integrated smart-home systems. Computer vision approaches use RGB and depth cameras to detect hazards and monitor gait or posture, often reporting strong performance in controlled settings and demonstrating the feasibility of identifying both static obstacles and hazards introduced through daily activity [4], [5]; however, these studies frequently emphasize technical accuracy while offering limited evaluation of how older adults perceive, understand, or respond to automated warnings. In parallel, conversational AI and social chatbots have been explored to support mental health, medication adherence, and social engagement, with evidence that voice-based and natural language interfaces can improve accessibility and adoption when designed with empathy and transparency [6]; nevertheless, these systems are commonly developed without direct integration with environmental sensing, which limits context-aware guidance for daily safety and routines. More recent, smart-home research argues for unified pipelines that couple sensing, analytics, and user interaction to support aging in place [7], while also highlighting challenges around privacy, consent, perceived surveillance, and trust that can undermine real-world adoption even when models perform well [8], [9]. Literature has also shown that although smart home devices hold strong potential for supporting older adults' health, safety, and independent living, existing user studies remain limited in number and often focus primarily on system performance rather than long-term impacts on quality of life, health, or well-being, with overall technology readiness levels still relatively low [10]. At the same time, controlled smart-home experiments demonstrate that non-intrusive ambient sensing can achieve high activity recognition accuracy when compared to wearable approaches highlighting the technical feasibility of zero-effort Ambient Assisted Living systems to support independence [11]. Together, this literature motivates systems that combine proactive user-centered approaches while explicitly addressing barriers to adoption that remain underexplored in many technical evaluations.

B. Problem Definition

Despite advances in home monitoring, a critical gap remains for systems that integrate real-time hazard detection with communication that is trustworthy, autonomy-respecting, and usable for older adults in everyday life. Many existing approaches focus on post-incident detection or emphasize technical performance in controlled settings, providing limited evidence on whether older adults understand, trust, or act upon automated warnings in practice. Consequently, even systems with strong preventative fall detection accuracy may fail to enhance safety or support independence if alerts are misunderstood, perceived as intrusive, or not adopted over time. We address this gap by evaluating an integrated platform that pairs environmental sensing with conversational support, focusing on both technical performance and the human factors that influence adoption and sustained use. Formally, given a stream of visual observations of a home environment, X, the system aims to produce a set of detected hazards, $\hat{H}$, where each hazard $\hat{h} \in \hat{H}$ is tagged with a hazard type and severity level, and is communicated to the user through a conversational interface. The central problem is therefore not only detecting hazards, but presenting hazard information in a way that is interpretable, actionable, and supportive of autonomy for older adults. Therefore, FORTif.ai aims to address the following research questions:

1) When an AI chatbot flags a nearby hazard, what factors shape whether older adults trust or act on the warning?
2) What human conditions (e.g., gait, mobility, vision) and environmental conditions (e.g., lighting, layout, surfaces) do participants identify as high risk, and do these align with what the hazard detection model detects?
3) How might sentiment analysis and conversational interfaces influence comfort and willingness to keep an AI chatbot-based hazard detection and warning system running over time?

II. Methodology

FORTif.ai is designed as an integrated platform combining conversational support, personalized memory management, and LLM-based hazard detection to enhance independent living for older adults. These three components work synergistically: the chatbot provides natural, empathetic interactions to assist with daily tasks and social engagement; the memory system enables persistent, contextually aware storage and retrieval of patient-specific information; and the hazard detection module continuously monitors the living environment to identify and alert users to potential safety risks. In the following sections, we detail the design, implementation, and operational principles of each subsystem, highlighting how they collectively support proactive, personalized, and safe independent living.

A. Chatbot System

The FORTif.ai chatbot is implemented using a Retrieval-Augmented Generation (RAG) architecture, combining the generative capabilities of Google's Gemini 2.5 Flash model with a semantic memory retrieval system to deliver contextually relevant and empathetic responses. The frontend communicates with a Node.js proxy server, which routes requests to a Python FastAPI backend. The FastAPI server handles both text and voice inputs and forwards queries to the RAG API server responsible for conversational intelligence. User queries are classified into four primary action types: calendar management, medication reminders, communication requests, and personal information queries,using the Gemini LLM. Queries concerning personal information invoke the RAG pipeline, which performs a three-stage process: retrieval, augmentation, and generation. Queries are first converted into high-dimensional embeddings, then matched against patient-specific memories stored in a Weaviate vector database using a hybrid semantic and keyword search. To maintain conversational continuity, follow-up queries are rewritten into standalone forms using an LLM-based context summarization pipeline. Retrieved memories are augmented with metadata (topic, emotion, source) and passed to the generation stage, where responses are produced using prompt templates optimized for warmth, validation, and cognitive accessibility. A medical advice detection layer prevents the chatbot from providing clinical recommendations. Queries related to diagnosis, treatment, or medical guidance are automatically redirected with standardized responses advising consultation with healthcare professionals. FORTif.ai supports both voice and text input modalities. Voice input is captured via the browser and processed through speech recognition in the backend, while text inputs follow the same processing pipeline, ensuring consistent response generation across modalities (Figure 1).

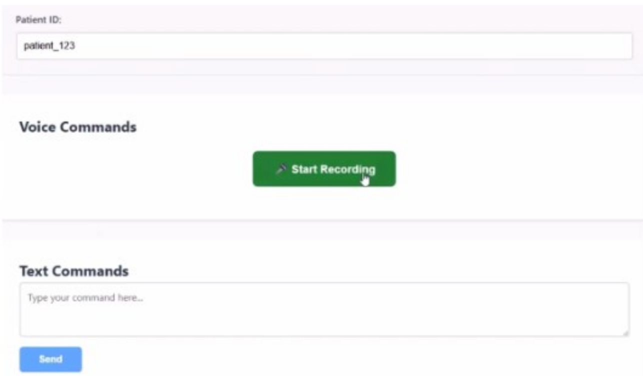

Fig. 1. Screenshot of the FORTif.ai chatbot interface.

B. Memory Management

The memory system enables persistent, patient-specific knowledge storage and retrieval, supporting contextually aware conversation. It is implemented using the Weaviate vector database, enabling hybrid search combining keyword and semantic similarity. Memories are structured with metadata including patient_id, text, topic, emotion, source, sensitivity, and chunking information. Text is embedded using Google's Generative AI embeddings to create high-dimensional vectors capturing semantic meaning. Long texts are intelligently chunked to maintain coherence and allow retrieval within token constraints. Memory retrieval employs a hybrid search strategy with adjustable weighting between semantic similarity and keyword matching. Retrieved results can be filtered by patient ID, sensitivity, emotion, topic, or source to ensure relevance and privacy compliance (Figure 2). Context window management ensures that only the most pertinent memories are included in the chatbot context, preserving critical information while preventing token overflow. A web-based dashboard allows older adults to view, search, filter, and manage memories, including soft or hard deletion, metadata editing, and visualization of memory distribution by topic, emotion, and source. This interface facilitates oversight, quality control, and identification of knowledge gaps.

C. Hazard Detection System

The hazard detection module provides real-time monitoring of the home environment using a vision-language model (Gemini 2.5 Flash Lite) to identify potential safety risks from images and video streams. The system receives images or video

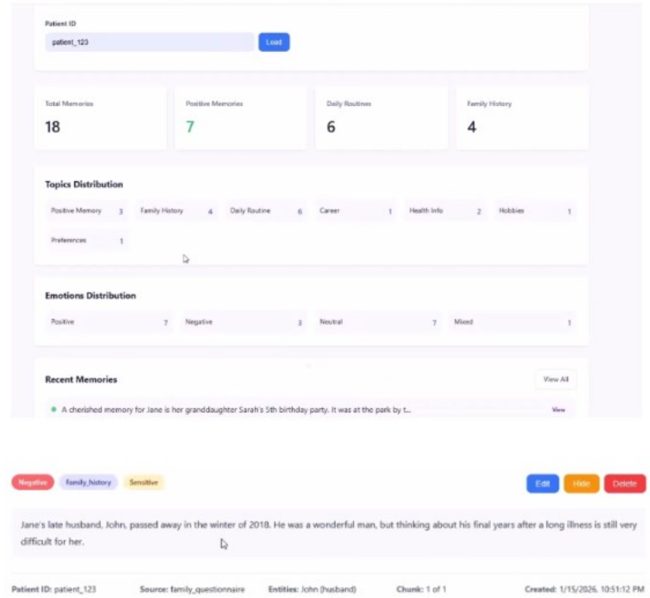

Fig. 2. Screenshot of the memory management dashboard displaying mock-user memories, metadata, and filtering options.

frames and applies structured prompts instructing the model to detect hazards such as wet floors, clutter, sharp edges, or obstructions, independent of visible human presence. Output is returned in structured JSON, including hazard type, location, severity, confidence, and SMS-ready alert text (Figure 3). Video frames are extracted and analyzed at configurable intervals, supporting near real-time safety monitoring. The system implements a signature-based deduplication mechanism to prevent redundant alerts. Hazard notifications are dispatched via SMS using Twilio, constrained to concise, actionable messages. Deduplication ensures that repeated hazards within a cooldown period do not trigger excessive alerts. Hazard detection is integrated with the chatbot interface via an API endpoint, enabling users to trigger analysis and receive structured results within the conversational interface. Severity classification considers both hazard type and spatial context, with risk levels ranging from Safe to Critical based on hazard count and optional area coverage metrics. This allows prioritization of high-risk hazards and ensures actionable, context-aware safety guidance.

III. Results and Discussion

Hazard detection testing for FORTif.ai initially focused on conventional computer vision approaches, evaluating YOLOv8 and SAM3 models across representative hazard classes including rugs, clutter, wires, uneven thresholds, furniture, and staircases. YOLOv8, tested in two iterations with slightly different training datasets, demonstrated low hardware requirements ($\sim$2 GB RAM) and fast inference, making it suitable for deployment on standard laptops. However, its ability to detect small or irregularly shaped hazards, such as scattered wires or cluttered objects, was inconsistent, suggesting that

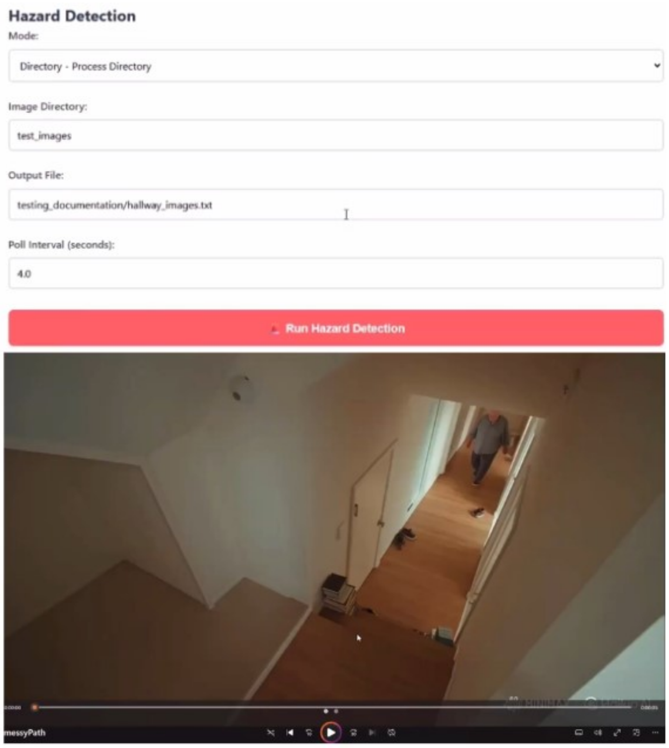

Fig. 3. Screenshot of the hazard detection interface illustrating detected hazards in a AI-Generated video, including type, location, and severity.

additional training data would be necessary to achieve reliable detection. In contrast, SAM3 produced more consistent and accurate hazard identification across all tested classes. Using prompt-based instructions such as “scattered household items,” “thin black power cable,” or “floor carpet,” SAM3 successfully recognized diverse hazards and enabled flexible prompt tuning to handle edge cases. Despite its higher accuracy, SAM3’s computational demands, including the need for a CUDA-enabled GPU and approximately 25 GB of disk space, limit its practical deployment in typical home environments. These limitations motivated the exploration of Gemini 2.5 Flash Lite, a cloud-based large language model capable of multimodal hazard detection with minimal hardware constraints. Testing with Gemini was conducted using both video- and image-based evaluations. Across all tests, the system generated structured outputs indicating whether a person was present, whether hazards were detected, and a detailed hazard list including type, location, severity, confidence, and concise SMS-style notification text, alongside short natural-language summaries. In a clear hallway video, Gemini appropriately reported no hazards while detecting a person with high confidence, demonstrating correct behavior on a negative control. In videos containing hazards, the system successfully detected floor-level clutter, scattered objects, and other tripping risks, producing concise summaries suitable for user-facing alerts. Recurring behaviors were noted: multiple instances of the same hazard type were sometimes not individually distinguished, reducing granularity; pointed furniture edges were occasionally flagged in low-risk contexts, suggesting geometric over-sensitivity; cluttered scenes were correctly identified as hazardous but with less specificity regarding object types; and wet-floor hazards were detected with high frame-level confidence, although overall severity scores occasionally underestimated perceived risk. Prompt design was found to strongly influence detection consistency. Baseline prompts produced variable identification of floor-level hazards, with rugs and small objects such as toys frequently missed and furniture edges sometimes flagged at medium severity despite low risk. Repeated runs on the same scene demonstrated that rugs were undetected in several trials, highlighting inconsistency. Iterative prompt refinement focusing on walkway-relevant hazards improved performance: GoPro image tests with the revised prompt reliably detected floor clutter, bunched or rolled rugs, obstructive items such as vacuum cleaners or toys, and even small animals in the walking path. These observations indicate that prompt specification can substantially enhance alignment between model outputs and intended hazard targets without modifying the underlying model, increasing both reliability and relevance for real-world monitoring. Overall, the progression from YOLOv8 and SAM3 to Gemini highlights the advantages of LLM-based hazard detection in combining dynamic, context-aware interpretation with reduced hardware constraints. Gemini’s multi-modal reasoning and flexible prompt-based approach allow for more nuanced, accurate, and actionable hazard reporting compared to conventional vision-only models. These findings informed the selection of Gemini as the primary model for ongoing hazard detection development, including testing with GoPro video streams and expanded image datasets. Evaluation of the chatbot system will occur in the next phase of the study, focusing on user interaction quality, contextual memory retrieval, and empathetic support.

IV. Conclusion

This paper presented FORTif.ai, an integrated platform designed to support safer independent living by combining real-time hazard detection with an empathetic conversational assistant. The platform integrates a multi-modal hazard detection module capable of analyzing images and video streams to identify household risks and generate structured, actionable warnings, alongside a retrieval-augmented generation chatbot backed by a Weaviate vector database.

Evaluation of the hazard detection component using AI-generated videos and images demonstrated appropriate behavior on negative examples and successful identification of hazards in cluttered and wet-floor scenarios. Iterative prompt refinement improved alignment with floor-level and walkway-relevant hazards, highlighting the importance of prompt design for reliable and context-aware outputs. These results suggest that LLM-based hazard detection offers superior accuracy and flexibility compared to conventional computer vision models.

However, limitations remain: the chatbot system has not yet been formally evaluated, and hazard detection has only been tested in synthetic or curated scenarios rather than real-world settings.

V. Future Work

To address the limitations identified in this study, the next phase of research will focus on controlled, real-world evaluation of the hazard detection module. This evaluation will be conducted in instrumented rooms simulating typical household layouts, including living rooms, bedrooms, and kitchens, with predefined hazards such as wires, rugs, furniture, clutter, spills, and sharp objects. A safety expert panel will define hazard types, locations, and severity, as well as criteria for "person in hazard path" events. Human participants will navigate these layouts along scripted paths while hazard detection outputs and participant observations are recorded and compared using a structured risk metric that accounts for environmental, static, dynamic, and positional hazards. This design will enable assessment of model reliability, consistency across perspectives, and alignment with human perception of risk, providing quantitative and qualitative evidence to refine both the hazard detection module and its scoring methodology.

Parallel to hazard testing, a three-phase methodology has been proposed to assess the chatbot's quality and usability. In Phase 1, external community input from seniors will be gathered to understand preferences, openness to chatbot use, and potential barriers, while internal testing will evaluate edge cases and vulnerability to malicious or nonsensical inputs. Phase 2 involves one-hour in-house sessions with participants interacting with the chatbot, followed by questionnaires to assess response quality, usability, and comfort. Internal evaluations will also be conducted to ensure appropriateness, accuracy, and avoidance of misleading or harmful outputs. Phase 3 implements a "Wizard of Oz" approach, comparing participant interactions with the real chatbot versus a human operator to benchmark performance and identify areas for improvement. Optional refinements include adapting the chatbot to accommodate barriers such as voice tremors, slow speech, or hearing impairments. Together, these phases will provide a comprehensive understanding of the integrated system's effectiveness and inform iterative improvements to support independent living safely and reliably.

In addition, ongoing work includes testing live hazard detection using cameras to capture real-time footage and evaluating different camera angles and positions to optimize coverage. Furthermore, a bathroom-specific modification of the hazard detection tool is being developed to function as a recommendation system in environments where live video footage is not permitted. Collectively, these studies aim to validate the platform in realistic household contexts and expand its applicability across a broader range of living environments.

Parallel to hazard testing, a three-phase methodology has been proposed to assess chatbot quality and usability. Phase 1 gathers external community input from seniors to understand preferences, openness to chatbot use, and potential barriers, while internal testing evaluates edge cases and vulnerability to malicious or nonsensical inputs. Phase 2 involves 1-hour in-house sessions with participants interacting with the chatbot, followed by questionnaires to assess response quality, usability, and comfort. Internal evaluations of response quality will also be conducted to ensure appropriateness, accuracy, and avoidance of misleading or harmful outputs. Phase 3 implements a "Wizard of Oz" approach, comparing participant interactions with the real chatbot versus a human operator to benchmark performance and identify areas for improvement. Optional refinements include adapting the chatbot to account for barriers such as voice tremors, slow speech, or hearing impairments. Together, these studies will provide a comprehensive understanding of the integrated system's effectiveness and inform iterative improvements to support independent living safely and reliably. We also aim to make our codebase publicly available in the near future for use by other developers and researchers.

VI. Acknowledgements

This project was made possible by the support of WAT.ai, the Ubilab at the University of Waterloo and the guidance of Professor Plinio Morita.

Appendices

The following supplementary materials are available online:

- **Appendix A:** Computer Vision Testing for Hazard Detection
- **Appendix B:** Gemini Testing for Hazard Detection
- **Appendix C:** Bathroom Safety Assessment Testing
- **Appendix D:** Example Outputs Hazard Detection Tool - AI Generated Videos

References

[1] Statistics Canada. (2024, Sep. 23) The older people are all right. Statistics Canada. Accessed: 2026-02-28. [Online]. Available: https://www.statcan.gc.ca/o1/en/plus/7059-older-people-are-all-right

[2] Public Health Agency of Canada, *The Safe Living Guide: A Guide to Home Safety for Seniors*, Public Health Agency of Canada, 2015, pDF available on Canada.ca; updated and accessed 2026. [Online]. Available: https://www.canada.ca/en/public-health/services/publications/healthy-living/safe-living-guide-a-guide-home-safety-seniors.html

[3] ——, "Seniors' falls in canada: Second report," Ottawa, ON, Canada, 2014, catalogue No. HP25-1/2014E-PDF. [Online]. Available: https://www.canada.ca/content/dam/phac-aspc/migration/phac-aspc/seniors-aines/publications/public/injury-blessure/seniors_falls-chutes_aines/assets/pdf/seniors_falls-chutes_aines-eng.pdf

[4] Q. R. Khawam, M. Al-Hashimi, and S. K. Abdulateef, "Real-time fall detection for the elderly in home settings using a vision-based yolov8-pose model," *Journal of Al-Qadisiyah Computer Science and Mathematics*, vol. 17, no. 3, pp. Comp153–Comp162, Sep. 2025.

[5] O. Flor-Unda, R. Arcos-Reina, C. Estrella-Caicedo, C. Toapanta, F. Villao, H. Palacios-Cabrera, S. Nunez-Nagy, and B. Alarcos, "A comparative overview of technological advances in fall detection systems for elderly people," *Sensors*, vol. 25, no. 24, p. 7423, Dec. 2025. [Online]. Available: https://doi.org/10.3390/s25247423

[6] S. Yu and T. Chen, "Understanding older adults' acceptance of chatbots in healthcare delivery: An extended utaut model," *Frontiers in Public Health*, vol. 12, November 2024. [Online]. Available: https://doi.org/10.3389/fpubh.2024.1435329

[7] C. VandeWeerd, A. Yalcin, G. Aden-Buie, Y. Wang, M. Roberts, N. Mahser, and D. Fabiano, "Homesense: Design of an ambient home health and wellness monitoring platform for older adults," *Health and Technology*, vol. 10, no. 5, pp. 1291–1309, 2020. [Online]. Available: https://doi.org/10.1007/s12553-019-00404-6

[8] T. Hanjahanja-Phiri, I. Modebelu, G. Morgan, S. Abhari, G. Bin Noon, D. Chumachenko, and P. P. Morita, "Perspectives of older adults on the ethics of active assisted living technologies: A scoping review and conceptual framework," *Digital Health*, vol. 11, Jul. 2025. [Online]. Available: https://doi.org/10.1177/20552076251336004

[9] Ö. Celik, M. Kulla, and J. Stypinska, "Trust formation in healthcare ai: An exploration of older adults' perspectives," in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, vol. 8, no. 1. Washington, DC, USA: The AAAI Press, Oct. 2025, pp. 498–512. [Online]. Available: https://doi.org/10.1609/aies.v8i1.36566

[10] M. Ghafurian, K. Wang, I. Dhode, M. Kapoor, P. P. Morita, and K. Dautenhahn, "Smart home devices for supporting older adults: A systematic review," *IEEE Access*, vol. 11, pp. 47 137–47 158, 2023. [Online]. Available: https://doi.org/10.1109/ACCESS.2023.3266647

[11] K. Wang, A. Saragadam, J. Kaur, A. Dogra, S. Cao, M. Ghafurian, Z. A. Butt, S. Abhari, D. Chumachenko, and P. P. Morita, "A contactless method for recognition of daily living activities for older adults based on ambient assisted living technology," *Internet of Things*, vol. 30, p. 101502, 2025. [Online]. Available: https://doi.org/10.1016/j.iot.2025.101502