DignityAI: Prototyping and Evaluating Community-based Preventive Care Solutions for Ageing Populations in Singapore

Kate Sangwon Lee
National University of Singapore
21 Lower Kent Ridge Rd, Singapore.
katelee@nus.edu.sg

Rafael Calvo
Imperial College London
London SW7 2AZ, UK
r.calvo@imperial.ac.uk

Li Hong Idris Lim
National University of Singapore
21 Lower Kent Ridge Rd, Singapore
lihongidris@nus.edu.sg

Jiang Ying
National University of Singapore
21 Lower Kent Ridge Rd, Singapore
jiangying@nus.edu.sg

Wenru Wang
National University of Singapore
21 Lower Kent Ridge Rd, Singapore
nurww@nus.edu.sg

Raquel Rodríguez-Carvajal
Universidad Autónoma de Madrid
28049 Madrid, Spain
raquel.rodriguez@uam.es

Marco Da Re
Imperial College London
London SW7 2AZ, UK
marco.da-re20@imperial.ac.uk

Pablo Alejandro Fonseca Arroyo
Imperial College London
London SW7 2AZ, UK
p.fonsecaarroyo@imperial.ac.uk

In Singapore, a fast-ageing society, the community care apartment (CCA) is a public-housing model through which the Singapore government facilitates community-based care for older adults. CCA faces challenges including understaffing and the absence of intelligent monitoring, which limits its ambitions to prevent illness and support independent living. We conducted a four-day hybrid hackathon between Singapore and the United Kingdom in which around thirty interdisciplinary students and researchers built four prototypes that pair passive ambient sensing with AI-powered conversational agents, targeting the pre-crisis window across a client–monitor–AI triad. We then evaluated the prototypes through an expert review grounded in Value Sensitive Design (VSD) using five values — autonomy, dignity, well-being, transparency and trust, and care fit — and identified their central value tension. The four prototypes showed complementary, near-mirror value profiles: the monitor-facing digital twin scored highest on care fit but lowest on autonomy, while the resident-facing coaching agent showed the inverse pattern. Three cross-cutting tensions emerged: watching over people versus the feeling of being watched; design restraint that protects dignity at the cost of the residents' own voice; and the inference of social intent from sensor data without explicit consent. We contribute the prototypes, the evaluation protocol, and design recommendations for dignified preventive care.

Ageing in place. Community care. Conversational agents. Ambient sensing. Value Sensitive Design. Hackathon.

1. INTRODUCTION

Singapore is projected to become a "super-aged" society by 2030, with one in four citizens aged 65 or older (Ministry of Health Singapore, 2022). This shift has accelerated policies redirecting healthcare toward community-based preventive models that support ageing in place (Ministry of Health Singapore, 2026). The Community Care Apartment (CCA) operationalises this vision by combining senior-friendly public housing with on-site care services (Thang *et al.*, 2025). However, this model introduces operational friction around human-resource constraints due to the high density of the city (Lee *et al.*, 2022). At our partner CCA A facility, one staff member is responsible for monitoring 189 older residents, relying on manual tracking without intelligent or systematic assistance. This cannot support the goals of preventive care: early indicators of cognitive decline or physical frailty, subtle behavioural shifts such as disrupted sleep, changes in daily routine, or gradual social withdrawal, remain invisible between periodic manual checks (Chen *et al.*, 2024). Consequently, interventions tend to occur only after acute crises occur, such as unattended falls, severe medical decompensation, or prolonged periods of acute social isolation.

To address this gap, a four-day hybrid hackathon with the National University of Singapore (NUS) and Imperial College London in the United Kingdom (UK) was established as a localised solution incubator. Interdisciplinary teams of students and researchers from engineering, human-centred design, and nursing co-created working prototypes that pair continuous passive ambient sensing with AI-powered conversational agents. After the hackathon, an interdisciplinary panel of experts from design engineering, nursing, and health psychology evaluated the prototypes through a review grounded

© Lee et al.

Published by BCS Learning and Development Ltd
Proceedings of BCS HCI 2026, UK

in Value Sensitive Design (VSD) (Friedman *et al.*, 2013), appraising each solution by how well it supports the human values at stake in dignified preventive care. This paper contributes (i) four sensor-integrated conversational-agent prototypes for preventive community care in a CCA context, (ii) a replicable multi-perspective expert evaluation protocol for assessing such prototypes against value-centred criteria, and (iii) findings from that evaluation from seven experts, including three value tensions and design recommendations for the next co-design phase with residents.

2. RELATED WORK

2.1 Community-based and preventive care for ageing populations

Population ageing has prompted a global shift away from acute, institution-centred provision towards preventive, community-based care that enables older adults to remain in familiar surroundings (Hayes *et al.*, 2024). Most older adults prefer to age in place, and international frameworks such as the World Health Organisation's age-friendly cities programme position the community as the primary site of support (World Health Organization, 2007). Singapore has pursued this agenda through a whole-of-nation strategy spanning housing and care services (Ministry of Health Singapore, 2026; Heok, 2021), with the CCA as a central instrument: a public-housing typology that integrates senior-friendly design and tiered care services, enabling residents to live independently (Thang *et al.*, 2025).

Preventive care depends on noticing subtle functional changes before they escalate into a crisis; where staffing is stretched, observation becomes reactive, and intervention follows an incident rather than preventing it (Chen *et al.*, 2024). This tension motivates technological augmentation — but augmentation that supports rather than supplants human care, extending rather than replacing the relational work of carers and the wider community (Sun *et al.*, 2025).

2.2 Human-AI care solutions

Passive ambient sensing offers a continuous, low-burden complement to manual monitoring. Infrared motion, contact, and pressure sensors — and, more recently, non-optical radar — can characterise activities of daily living and surface deviations that act as early markers of cognitive or physical decline. Longitudinal in-home studies have linked reduced outings, slower gait, and disturbed sleep to cognitive impairment (Lussier *et al.*, 2020; Rantz *et al.*, 2013), and a feasibility study with community-dwelling seniors in Singapore showed that such sensing can be deployed unobtrusively and acceptably in real homes (Rawtaer *et al.*, 2020). Critically, radar and ambient sensors capture behavioural signals without cameras, preserving a degree of privacy that camera-based monitoring cannot.

Conversational agents are an increasingly studied channel for older adults' health, supporting well-being check-ins, medication reminders, and companionship, with large language models extending their fluency and contextual range (Espinoza *et al.*, 2025). Yet acceptance is not guaranteed: reviews report concerns about privacy, threats to independence, and emphasise that agents must be controllable, personalisable, and transparent to earn trust (Cho *et al.*, 2025; Sadek, Calvo and Mougenot, 2025). For frail or cognitively impaired users in particular, dignity and autonomy are design requirements rather than afterthoughts. Sensing and conversation are rarely integrated into a single workflow that coordinates the full care triad — residents, family carers, and the professionals who monitor them — without simply adding alerts for an already overloaded monitor to watch. Designing such an integration is as much an organisational and human-centred problem as a technical one, which makes it well suited to hackathons for user-led innovation in digital health that convenes heterogeneous expertise to produce and critique prototypes (Rys, 2026). We adopt this format, paired with a value-centred expert review, to probe how passive sensing and conversational coordination can be combined for preventive community care.

3. METHOD

3.1 Hackathon context and procedures

We report a four-day hybrid hackathon held on 5–8 May 2026, run jointly by a UK design-engineering laboratory and a Singapore human-centred AI design laboratory, with clinical mentorship from a university nursing school. Approximately 30 students and researchers from design engineering, industrial design, architecture, biomedical and electrical engineering, nursing, and the humanities worked in four cross-institutional teams of around five members each. Day 1 was fully facilitated, comprising framing talks, team formation, and a structured AI for empathy workshop (Section 3.2) (Lee and Lim, 2026). Days 2–3 were autonomous build sprints supported by an always-on shared channel for cross-site collaboration. On Day 4, each team presented a working or Wizard-of-Oz prototype, a sensor-integration diagram, and a reflection on monitor workload. Teams designed end-to-end solutions targeting the pre-crisis window across a client–monitor–AI triad: clients (residents and family carers), monitors (care staff and remote clinicians), and an AI layer that mediates between them without increasing monitor load.

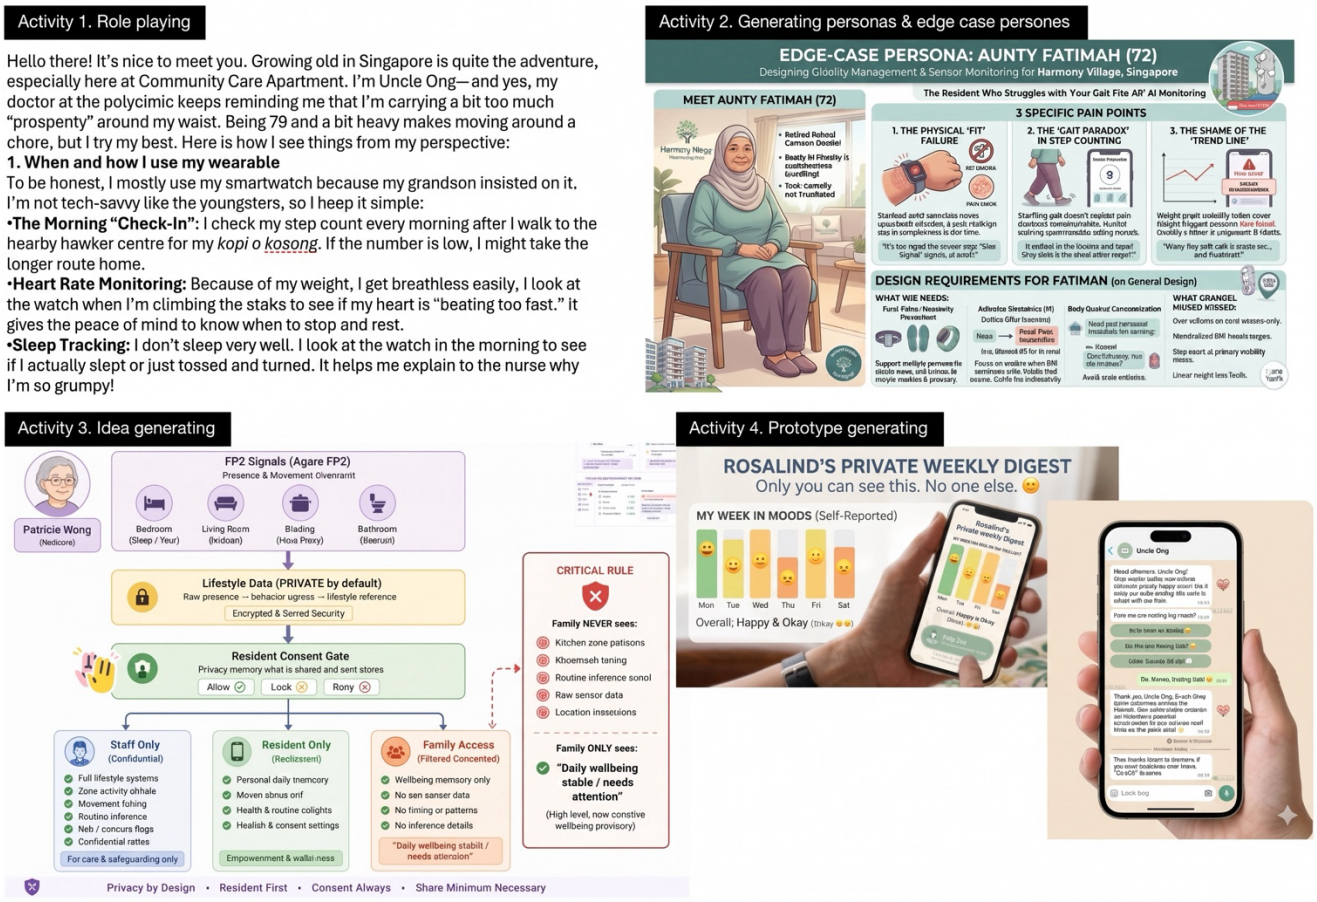

Figure 1: *AI for empathy workshop's four activities.*

Each prototype combined two architectural pillars. The first, passive ambient sensing, used continuous non-optical radar and ambient motion, door, and environmental sensors within simulated apartment units to detect anomalies while preserving privacy.

Secondly, conversational care-coordination interfaces deployed agents built on a shared conversational care platform integrated with a home-automation backend that engage residents in routine well-being checks and relay only triage-relevant signals to the monitoring desk when behavioural shifts are flagged. To structure feedback, prototypes were assessed against six shared delivery criteria: sensor integration, monitor-load reduction, client dignity, empathy fidelity (the traceability of design decisions to a specific user insight), channel appropriateness, and clarity of the Human-AI decision loop. These criteria foreground the project's central concern — relieving the staffing bottleneck without eroding resident dignity or burying carers in noise from monitoring alerts.

3.2 AI for Empathy workshop

Since preventive care hinges on understanding residents as people rather than data points, Day 1 centred on a two-hour workshop designed to build cognitive empathy (Figure 1) — the capacity to reason about another's thoughts and feelings — for the CCA's residents, family carers, and frontline staff. The workshop followed the four phases of the double-diamond design process (Kochanowska and Gagliardi, 2021). In Discover, participants conducted a generative-AI agent role-play to generate stakeholder personas, such as a resident with early dementia, a night-shift care aide responsible for 189 residents, a family carer living abroad, or a remote general practitioner. In Define, teams used AI image and text tools to generate a visual persona and an edge-case "Provotype" that surfaced rarely discussed needs (Lee and Lim, 2026). In Develop, persona-driven ideation with a keep–modify–discard filter generated candidate features, each constrained to reference one sensor data stream and one communication channel. In Deliver, teams produced a rapid UI sketch that made the client–monitor–AI triangle explicit: who sees what, and how overload is avoided. While generative AI was used to accelerate persona ideation, human centricity was safeguarded through clinical mentorship from nursing faculty and grounding prompts in real-world CCA observational data.

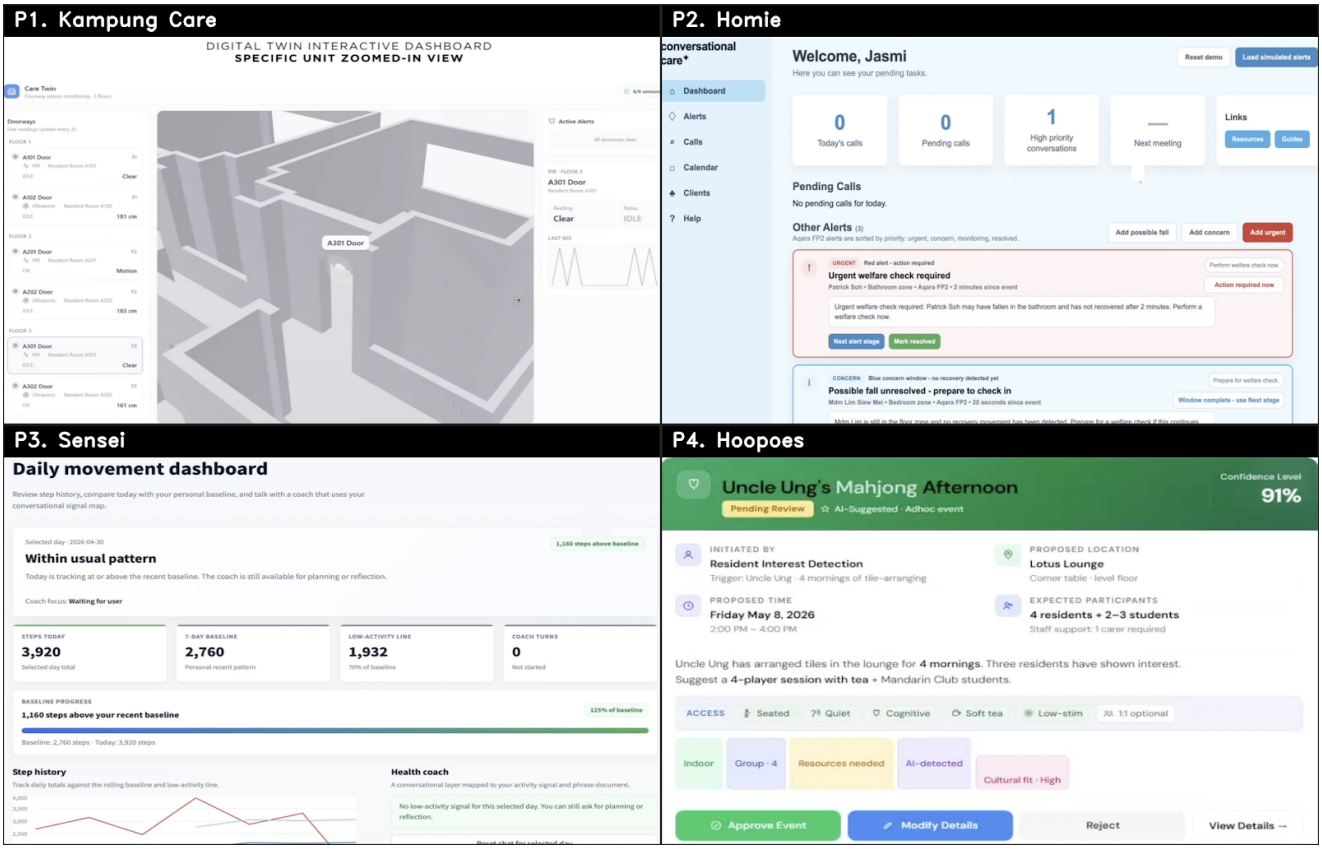

Figure 2: *Four hackathon prototypes*

3.3. Multi-perspective expert evaluation protocol

In June–July 2026, an interdisciplinary panel of seven experts evaluated the four prototypes: design engineering (E2, E3, E4, E5), nursing (E1, E7), and health psychology (E6). Each expert rated systems on a five-point scale (1 = poorly supported, 5 = strongly supported) across five core values (see Appendix): autonomy, dignity, well-being, transparency and trust, and care fit. Five evaluations were completed asynchronously and two through semi-structured interviews (30–40 minutes). Ratings were aggregated into descriptive means and standard deviations, complemented by thematic synthesis of qualitative feedback.

4. FINDINGS

4.1 Four prototypes

Four prototypes were completed and demonstrated on Day 4 (Figure 2), each occupying a different position in the client–monitor–AI triad as follows.

P1: Kampung Care (Monitor-facing triage): ingests ambient radar and environmental data into an interactive 3D "digital twin" dashboard. It spatialises unit alerts to help the sole care worker quickly evaluate emergency urgency.

P2: Homie (Client-facing safety): monitors bathroom duration and movement patterns to identify falls. Alerts escalate through a staged colour-coded timer. It intentionally omits voice output to avoid creating a feeling of constant surveillance.

P3: Sensei (Client-facing coaching): converts wearable activity and sleep data into concise, conversational coaching nudges. It limits conversational turns to encourage physical action over prolonged screen engagement.

P4: Hoopes (Triad coordination): monitors inactivity to match residents based on shared interests, tailoring updates across audiences (e.g., social summaries for family, clinical flags for care staff).

Together, the prototypes cover the full triad: monitor-facing triage (P1), client-facing safety (P2), client-facing coaching (P3), and cross-stakeholder coordination (P4). All four combine passive sensing with a conversational or interpretive layer, and each makes an explicit choice about what to withhold from the monitor in order to limit alert fatigue.

4.2 Quantitative value profiles

Table 1 reports the mean expert rating per value per prototype (n = 7 experts; one expert marked well-being as not applicable for P1). Overall means were closely clustered for the three resident- or triad-

facing prototypes (P2 = 3.87, P4 = 3.86, P3 = 3.84), with the monitor-facing digital twin somewhat lower (P1 = 3.51).

Table 1: *Expert review score summary*

Value	P1 Kampung Care	P2 Homie	P3 Sensei	P4 Hoopoes
Autonomy	3.07 (0.73)	3.57 (0.53)	**4.21** (0.91)	4.00 (0.58)
Dignity	3.29 (0.95)	**4.14** (0.69)	4.14 (0.38)	4.00 (0.58)
Well-being	3.42 (1.11)	3.71 (0.76)	4.14 (0.38)	**4.29** (0.76)
Transparency & Trust	3.71 (0.76)	3.86 (0.38)	3.57 (0.53)	3.57 (0.79)
Care Fit	**4.07** (0.45)	4.07 (0.45)	3.14 (0.90)	3.43 (0.53)
Overall mean	**3.51**	**3.87**	**3.84**	**3.86**

More revealing than the aggregate scores are the near-mirror profiles of values. P1 scored highest of all prototypes on care fit (4.07) but lowest on autonomy (3.07) and dignity (3.29): E6 described it as "*a tool for the care team, not for the resident*", in which residents figure as "*sensor points, not partners*" (E5), with no resident-facing consent or visibility in the current design. P3 showed the inverse profile — the highest autonomy score in the set (4.21) alongside the lowest care-fit score (3.14) — because it serves the resident directly but offers no staff-facing features and its interface complexity was judged to be demanding for older users. P2 achieved the most balanced profile, pairing the joint-highest dignity score (4.14) with the joint-highest care fit (4.07) through deliberate restraint: no camera, no voice, tiered alerts. P4 scored highest on well-being (4.29), reflecting its direct address of relatedness and social connection, but lower on care fit (3.43) and transparency (3.57), where its profile-based matchmaking and cross-centre data flows raised consent concerns. Expert divergence of two or more points concentrated on P1 (well-being ratings ranged from 2 to 5) and P3 (autonomy 2.5–5; care fit 2–4), and largely tracked disciplinary lens: nursing experts foregrounded workflow utility, alert triage, and cost, engineering experts foregrounded data governance and usability, and the psychology expert foregrounded consent, evidence, and the resident's felt experience.

4.3 Cross-cutting value tensions

Three tensions recurred across the panel's lenses and across prototypes.

Watching over versus feeling watched. Every prototype negotiates the same balance between protecting physical safety and the person feeling surveilled (E6). The digital twin (P1) exemplifies it most sharply: "*the more detail the twin shows, the faster staff can react, but the home also becomes more watched*" (E5), and the finer its spatial resolution — down to flat, potentially room-level — the closer it comes to "*following a person's whereabouts at all times*" (E6). Experts converged on the same remedy: the dashboard "*should show what matters, not everything*", with access restricted according to users' roles (E6, E7), and "*residents should know they are on that map, and who can see it*" (E5). Relatedly, experts warned that "*sensor alerts without clinical context become noise*" (E6), as "*the biggest care-fit risk at scale*" (E5); the nursing lens added that alert tiers and colour coding must map onto real staff workflows and specify the follow-up actions expected after an alarm (E1).

Restraint protects dignity — at the cost of the resident's voice. Experts judged that the strongest work on dignity and autonomy came from restraint (E6): P2 Homie gives up the voice channel and stages its alerts; P3 Sensei turns figures into warm, plain messages and keeps exchanges short. But both moves carry a cost. By removing voice, P2 Homie "*protects dignity, but attends only to physical safety*" (E6) — and it also removes the resident's ability to say "I'm fine" when an alert triggers: "*in the most important moment, the system speaks for them*" (E5). By hiding raw numbers, P3 Sensei "*decides what your body is telling you*" (E5), trading transparency for clarity; E5 recommended that "*people should be able to see the numbers if they ask*", while E6 judged the trade-off well balanced by keeping exchanges short and "*offering choices rather than instructions, so the person acts because they want to, not because they are told to*". E6 from health-psychology further noted that the premise underlying P2 Homie's design — that older residents would reject voice and feel watched — is "*an untested assumption, not an established finding*": reviews attribute low acceptance of conversational agents to reliability, threats to independence, and privacy rather than to the voice modality itself, a caution echoed from the nursing lens (E7). The panel therefore recommended treating voice, and a resident-initiated "I'm fine" confirmation, as an open co-design question for the next phase (E6, E7).

Inferring intent without consent. P4 Hoopoes generates social suggestions from sensed inactivity or boredom, matches residents by profile, and shares happy moments with family and across centres. Experts valued its contribution to connection but cautioned that "*the system creates connection, but it also organises it*" (E5), risking turning spontaneous social life into a managed activity, and that boredom sensing, profile-based matchmaking, cross-centre "event intelligence", and photo sharing with families "*all need explicit residents' consent*, *sharing photos with family should be the resident's call, not the system's*" (E5). E6 specified the safeguards: explicit opt-in consent, the ability to decline at any time, and clearly bounded rules about what is shared with whom. The nursing lens raised a complementary reach concern:

the most isolated residents may never engage with an opt-in social system, so "*something more active instead of waiting*" is needed to reach and nudge them (E1). The same consent deficit appeared in P1, where collected data flows to staff without resident control (E3, E5).

A shared recommendation closed the review: across all four prototypes, give residents a lightweight way to see and confirm what the system detects and shares about them, and log when the AI acts versus when a human decides (E5); keep a human "coordinator" in every escalation decision, surface only the alerts that make clinical sense, and make consent and who-can-see-what explicit (E6).

5. DISCUSSION

Our findings are a value see-saw across the care triad. No single prototype supported all five values well: designs that served the monitor scored high on care fit and low on resident autonomy, and designs that served the resident showed the mirror image. This pattern was not imposed by the brief — all teams were asked to design across the triad — but emerged from the anchor each team had chosen for its primary user, as AI-enhanced interventions for older people's care are typically developed and trialled as isolated artefacts with mixed acceptability (Loveys *et al.*, 2022). It suggests that dignified preventive care is unlikely to be delivered by any single artefact and instead requires a portfolio in which resident-facing agency layers (consent gates, confirmation channels, visible data flows) are coupled with monitor-facing triage layers, with the coordination function that P4 Hoopoes gestures toward deciding who sees what. The panel's shared recommendation — a resident-visible record of what is detected and shared, and an explicit log of AI action versus human decision — suggests AI in long-term care as a whole-system approach with stakeholders' diverse needs (Lukkien *et al.*, 2023).

The evaluation also complicates an intentional design strategy: restraint. Withholding a camera reliably earned dignity credit from experts, but each act of withholding also removed a channel through which the resident could participate, confirming an alert, interrogating their own data, declining an inference. Restraint, in other words, can quietly convert the resident from subject to object even while protecting their privacy. This aligns with work on older adults' perceptions of smart-home monitoring, which calls on designers and regulators to support higher levels of user agency over data collection (Percy Campbell *et al.*, 2024).

Methodologically, the VSD expert review identified an untested empirical assumption baked into a design decision — that older adults reject voice interfaces — and reframed it as a falsifiable question for the next co-design phase with residents. This aligns with a weakness of value-sensitive approaches to AI, in which values are often invoked in the abstract but rarely operationalised into concrete, testable design decisions (Sadek, Calvo and Mougenot, 2025).

Several limitations bound these claims. The prototypes were demonstrated, not deployed; the evaluation captures anticipated value support, not experienced value support, and the ultimate arbiters — CCA residents and frontline staff — have not yet rated these systems. The panel was small, drawn in part from the organising teams, and its ratings may exhibit familiarity effects. The four-day format favours demonstrable features over integration and governance work, which may partly explain why transparency and care fit trailed the other values. Also, mean ratings on a five-point scale from seven experts are indicative, not inferential; we treat them as structured qualitative evidence rather than measurement. Furthermore, our study adopted continuous ambient monitoring and conversational AI as baseline paradigms from the outset; future work should critically examine whether alternative, non-surveillance care models better align with older adults' everyday values.

6. FUTURE WORK AND CONCLUSION

Our next step is a resident-facing co-design phase at the CCA A that treats the panel's open questions as hypotheses: whether residents want a voice channel; whether warm summaries or inspectable numbers better support informed choice; and what consent and visibility mechanisms residents find legible. A subset of the prototypes will be integrated on the shared conversational-care platform and evaluated with validated dignity and need-satisfaction instruments alongside workload measures for care staff. Releasing the review sheet and protocol openly would let other design teams adopt the method directly.

This paper presented four conversational ambient-sensing prototypes designed during a hackathon for community preventive care, evaluated using a Value Sensitive Design protocol. The findings demonstrate that while restraint and triage improve workflow feasibility, systematic consent, active resident voice, and human-in-the-loop oversight are critical to prevent depersonalisation. Future work will conduct in-situ participatory co-design at CCA sites to evaluate voice interaction, data granularity, and consent mechanisms directly with older residents and frontline care teams.

ACKNOWLEDGEMENT

We thank all our hackathon participants, and this work is funded by the National University of Singapore-Imperial College London Research Collaboration Fund.

REFERENCES

Chen, Y., Power, M. C., Grodstein, F., Capuano, A. W., Lange-Maia, B. S., Moghtaderi, A., Stapp, E. K., Bhattacharyya, J., Shah, R. C. and Barnes, L. L. (2024) 'Correlates of missed or late versus timely diagnosis of dementia in healthcare settings', Alzheimer's & Dementia, 20(8), pp. 5551-5560.

Cho, M., Re, M. D., Simpson, T., Harrison, M. and Calvo, R. A. 'Conversational AI in Community Care: Preliminary Insights from a Scoping Review'. 38th International BCS Human-Computer Interaction Conference: BCS Learning & Development, 471-480.

Espinoza, F., Cook, D., Da Re, M., Fuentes, M. S. C., Butler, C. R. and Calvo, R. A. (2025) 'Designing AI-Powered Chatbots for Dementia Care in Peru: Stakeholder Engagement and Field Observations', Interacting with Computers, 37(5), pp. 430-443.

Friedman, B., Kahn Jr, P. H., Borning, A. and Huldtgren, A. (2013) 'Value sensitive design and information systems', Early engagement and new technologies: Opening up the laboratory: Springer, pp. 55-95.

Hayes, C., Manning, M., Fitzgerald, C., Condon, B., Griffin, A., O'Connor, M., Glynn, L., Robinson, K. and Galvin, R. (2024) 'Effectiveness of Community-Based Multidisciplinary Integrated Care for Older Adults with General Practitioner Involvement: A Systematic Review and Meta-Analysis', Health & Social Care in the Community, 2024(1), pp. 6437930.

Heok, K. E. (2021) 'Towards ageing well: Planning a future-ready Singapore', Centre for Liveable Cities, Singapore.

Kochanowska, M. and Gagliardi, W. R. (2021) 'The double diamond model: In pursuit of simplicity and flexibility', Perspectives on design II: Research, education and practice: Springer, pp. 19-32.

Lee, J.-J., Ng, T. P., Nasution, I. K., Eng, J. Y., Christensen, R. D. and Fung, J. C. (2022) 'Implementation of a sense of home in high-density multicultural Singapore nursing homes: Challenges and opportunities', International Journal of Environmental Research and Public Health, 19(11), pp. 6557.

Lee, K. S. and Lim, L. H. I. (2026) 'The AI-Human Validation Loop: Fostering Critical Artificial Empathy in Design Education', NUS CTLT AI in Education: Poster Showcase of Teaching Breakthroughs and New Courses Teaching Innovation with GenAI.

Loveys, K., Prina, M., Axford, C., Domènec, Ò. R., Weng, W., Broadbent, E., Pujari, S., Jang, H., Han, Z. A. and Thiyagarajan, J. A. (2022) 'Artificial intelligence for older people receiving long-term care: a systematic review of acceptability and effectiveness studies', The Lancet Healthy Longevity, 3(4), pp. e286-e297.

Lukkien, D. R., Nap, H. H., Buimer, H. P., Peine, A., Boon, W. P., Ket, J. C., Minkman, M. M. and Moors, E. H. (2023) 'Toward responsible artificial intelligence in long-term care: a scoping review on practical approaches', The Gerontologist, 63(1), pp. 155-168.

Lussier, M., Aboujaoudé, A., Couture, M., Moreau, M., Laliberté, C., Giroux, S., Pigot, H., Gaboury, S., Bouchard, K. and Belchior, P. (2020) 'Using ambient assisted living to monitor older adults with Alzheimer disease: single-case study to validate the monitoring report', JMIR medical informatics, 8(11), pp. e20215.

Ministry of Health Singapore 2022. White paper on healthier SG. Ministry of Health Singapore.

Ministry of Health Singapore (2026) THREE MORE AGE WELL NEIGHBOURHOODS TO STRENGTHEN SUPPORT FOR SENIORS IN THE COMMUNITY. Available at: https://www.moh.gov.sg/newsroom/three-more-age-well-neighbourhoods-to-strengthen-support-for-seniors-in-the-community/.

Percy Campbell, J., Buchan, J., Chu, C. H., Bianchi, A., Hoey, J. and Khan, S. S. (2024) 'User perception of smart home surveillance among adults aged 50 years and older: Scoping review', JMIR mHealth and uHealth, 12, pp. e48526.

Rantz, M. J., Skubic, M., Miller, S. J., Galambos, C., Alexander, G., Keller, J. and Popescu, M. (2013) 'Sensor technology to support aging in place', Journal of the American Medical Directors Association, 14(6), pp. 386-391.

Rawtaer, I., Mahendran, R., Kua, E. H., Tan, H. P., Tan, H. X., Lee, T.-S. and Ng, T. P. (2020) 'Early detection of mild cognitive impairment with in-home sensors to monitor behavior patterns in community-dwelling senior citizens in Singapore: Cross-sectional feasibility study', Journal of medical Internet research, 22(5), pp. e16854.

Rys, M. (2026) 'Hackathons as a bridge between academia and industry', Industry and Higher Education, 40(2), pp. 168-179.

Sadek, M., Calvo, R. A. and Mougenot, C. (2025) 'The value-sensitive conversational agent co-design framework', International Journal of Human–Computer Interaction, 41(15), pp. 9533-9564.

Sun, Y., Ankenbauer, S. A., Guo, Z., Chen, Y., Ma, X. and He, L. (2025) 'Rethinking Technological Solutions for Community-Based Older Adult Care: Insights fromOlder Partners' in China',

Proceedings of the ACM on Human-Computer Interaction, 9(2), pp. 1-36.

Thang, L. L., Cao, Y., Heng, C. K. and Kiyota, E. (2025) 'Housing Senior Citizens in Singapore: Ageing in Place in the Community', Architectures of Ageing in Place: Routledge, pp. 79-95.

World Health Organization (2007) Global age-friendly cities: A guide. World Health Organization.

APPENDIX

The five values used in the expert review

Value	What we ask
1. Autonomy	The resident keeps control over the interaction and over what is shared, and with whom.
2. Dignity	The system treats the resident as a person and partner, never as a monitored object or a "problem."
3. Well-being	The design supports emotional and social well-being, not only physical safety.
4. Transparency & Trust	It is clear when the AI acts, when it defers to a human, and why — so residents and carers can trust it.
5. Care Fit	It reduces monitor workload and fits real care work, supporting rather than supplanting human carers.

Scale: 1 = poorly supported · 2 · 3 = partially supported · 4 · 5 = strongly supported. Use N/A if the prototype does not address the value.