CareOps: A Multi Agent Control Room for Independent Living with Care

Demonstration Track

Lokesh Singh
University of Southampton
Southampton, United Kingdom
lb5e23@soton.ac.uk

Jayati Deshmukh
University of Southampton
Southampton, United Kingdom
J.Deshmukh@soton.ac.uk

Athina Georgara
University of Southampton
Southampton, United Kingdom
A.Georgara@soton.ac.uk

Tan Viet Tuyen Nguyen
University of Southampton
Southampton, United Kingdom
tuyen.nguyen@soton.ac.uk

Sarvapali D. Ramchurn
University of Southampton
Southampton, United Kingdom
sdr1@soton.ac.uk

ABSTRACT

This paper introduces CareOps, a multi-agent control-room dashboard for Independent Living with Care (ILWC) that coordinates heterogeneous sensors and a Buddy robot across multiple simulated homes. Agents fuse radar, audio, bed, passive infrared (PIR), and gait data into single prioritised incidents with human-readable explanations, and support one-click dispatch of either a robot (via socket) or a human carer (via smartphone notification). The demo presents a single decision-support workflow that orchestrates sensors, software agents, and robots through a unified interface. It provides a human-in-the-loop platform for exploring how professional care staff understand multi-sensor evidence and choose between robot and human responses.

KEYWORDS

Multi-agent Coordination; Assistive robotics; Human-robot interaction; Independent living

ACM Reference Format:
Lokesh Singh, Jayati Deshmukh, Athina Georgara, Tan Viet Tuyen Nguyen, and Sarvapali D. Ramchurn. 2026. CareOps: A Multi Agent Control Room for Independent Living with Care: Demonstration Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/PSHL4649

1 INTRODUCTION

Independent living with care (ILWC) increasingly relies on digital technologies to extend overstretched staff and support ageing-in-place at scale [13, 14]. Care homes now deploy acoustic monitoring [5, 15], radar-based fall detection [4, 10], sleep and continence analytics [17], and ambient or wearable mobility tracking [12, 16]. These systems promise continuous, unobtrusive observation but usually act as isolated alert channels. Each portal raises notifications using its own logic, leaving staff to fuse signals, prioritise incidents, and choose responders. This fragmentation leads to alarm fatigue, cognitive overload, and inconsistent escalation, especially during busy periods such as morning routines or night supervision [3, 11]. From a multi-agent systems (MAS) perspective, ILWC is therefore a distributed decision-making system in which residents, carers, family members, robots, and software services act as agents with partial information and constraints [6, 19], while AI models for health decline prediction and autonomous response must operate under uncertainty and organisational policy [7, 18]. Rather than treating each device or robot as an isolated detector, ILWC can be framed as a coordinated MAS in which software agents fuse evidence, manage information flows, and help allocate scarce human and robotic resources.

Our CareOps dashboard treats ILWC as a human-in-the-loop MAS coordination problem [8]. A single control-room interface lets operators oversee multiple homes, residents, sensors, and robots, while software agents fuse multi-sensor evidence, rank incidents, and coordinate mobile assist robots. Humans inspect explanations, choose responders, and close incidents, turning *alerts everywhere* into coordinated, explainable incident management. An Alert Fusion and Risk Agent combines virtual radar, audio, bed, PIR, and gait into a single incident object with a confidence score, missing-information flags, and a short explanation of why the event matters now [2, 4], and a Decision Agent maintains a cross-home priority queue over severity, recency, and resident risk [1, 9]. False alarms and triage behaviour are made explicit. Each incident is closed with an outcome label and notes. Learning and Audit Agent aggregates live metrics (reliability, false-alarm rate, mean resolution time) as a basis for learning alert-triage and workload-aware notification policies [3, 11]. Because care staff remain legally and ethically accountable for outcomes [7, 18], operators always see sensor level evidence and context, choose whether to send a human or the Buddy robot, can override suggestions with a recorded reason, and complete an audit-ready closure form. In this demo, these logs support descriptive analysis, and future work will use them to study decision patterns and governance style protocols between software agents, robots, nurses, and managers [8].

Demo Video: https://youtu.be/yiVN5xfKq_U

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/PSHL4649

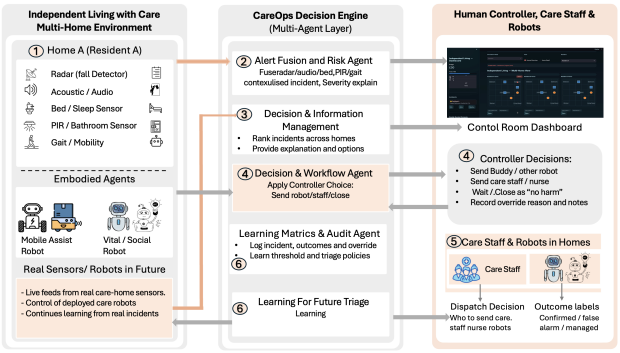

Figure 1: Decision workflow in CareOps: sensors raise incidents, the dashboard supports dispatch or closure, and outcomes feed the Learning & Audit agent.

2 SYSTEM ARCHITECTURE

We instantiate the ILWC Control Room as a compact multi-agent decision engine (Figure 1). Dashboard modules act as agents over a shared incident stream, Alert Fusion & Risk, Decision Priority, Dispatch, and Learning & Audit—forming a closed human-in-the-loop decision cycle from sensor signals to responses and logged outcomes.

2.1 Alert Agent for Fusion and Explanation

The Alert Fusion & Risk Agent analyses radar, bed-exit, PIR/appliance, audio, and gait signals. For each incident i and sensor s, the simulator provides a confidence score $c_{i,s} \in [0, 1]$ and a short "why" explanation; these are fused into an incident-level confidence

$$C_i = \sum_s \alpha_s \, c_{i,s},$$

where α_s encodes the reliability of each modality. The agent renders a structured explanation card (e.g., "possible fall in bathroom: radar impact & lying posture; audio thud & silence; no bed exit spike; high baseline fall risk"), and the fusion layer is modular so alternative evidence combination schemes (e.g., *fall* / *no-fall* / *uncertain*) can be swapped in without changing the operator view.

2.2 Decision Agent for Prioritisation

The Decision Agent maintains a live priority queue of active incidents and drives the incident list and risk indicators in the dashboard. For each incident i, it conceptually computes

$$P_i = w_{\text{sev}} S_i \; + \; w_{\text{age}} A_i \; + \; w_{\text{risk}} R_i \; + \; w_{\text{conf}} C_i,$$

where S_i encodes severity (GREEN/AMBER/RED), A_i is age in seconds, R_i is the resident's risk profile, and C_i is the fused confidence. The weights w_* are initialised from discussions with clinical partners and can be varied between runs to represent different triage policies. In the current demo, the agent also surfaces simple intervention options (e.g., dispatch Buddy robot or send a human carer) with explicit labels and timing, so operators see a ranked, contextualised set of decisions rather than isolated alerts.

2.3 Dispatch and Learning Agents

Once an incident is selected, the Dispatch Agent chooses a responder and logs the decision. Given incident i and responders $j \in \mathcal{R}$ (humans or Buddy), it conceptually selects

$$d(i) = \arg\max_{j \in \mathcal{R}} U_{ij},$$

where U_{ij} encodes grid distance, current workload, and required skill. In the prototype, this utility is rule-based: *Send Buddy Robot* sets Buddy's target room on the 5×5 grid and sends a socket command to the Android client, while *Send Human Carer* triggers a Pushcut notification to the carer's phone. As responders move towards the incident, the dashboard updates the last action, budget, and a simple trust index. On arrival, the operator closes the incident with a label (confirmed / false alarm / managed with no harm). A Learning & Audit Agent aggregates all events and actions into a structured activity log (downloadable as CSV), providing traces from which future MAS algorithms can recalibrate sensor and priority weights and learn workload aware triage policies.

3 DEMONSTRATION SCENARIO

The demo runs a fixed five-event script over multiple simulated ILWC homes (Residents A, B, C). Every 60 seconds, the system injects one random incident from the set of urgent events (possible fall, bleeding, chest pain) and two lower-severity events (unusual inactivity, unusual movement). Each incident is placed in a resident's room on a 5×5 grid (Bedroom, Bathroom, Kitchen, Living Room, Common Space) and appears both in the multi-home overview (red marker, urgency badge) and in the sidebar incident queue with its age in seconds. The video follows a single control-room operator who notices each new alert, clicks through to the relevant home, and inspects the event panel: a fused evidence card summarising radar, audio, bed, PIR, and gait signals with a confidence score, a short "why" sentence, and per sensor breakdown. From this screen, the operator chooses whether to *Send Buddy Robot*, *Send Human Carer*, or close the event. *Send Buddy Robot* sets the robot's path on the grid and issues a socket command (e.g., `SAY:urgent`) to the Android Buddy client; *Send Human Carer* triggers a Pushcut notification with the incident and explanation to the carer's phone. As responders move stepwise towards the incident room, the dashboard updates last actions, budget, and a simple trust index, and the operator records the outcome (confirmed, false alarm, managed with no harm) before closure. All events and actions are written to the Activity Log and can be downloaded as CSV.

4 CONCLUSION AND FUTURE WORK

We presented CareOps, a multi-home ILWC control-room dashboard that treats sensors, software agents, and robots as a multi-agent system with decision logging. In future, we will run a simulation based study with professional care staff from St Monica Trust[1] to compare baseline and AI-assisted use in terms of prioritisation, workload, and trust, informing learned triage policies, personalised risk models, and integration with real sensor feeds and care workflows.

[1]Study protocol "Testing a Care Management Incident Response Dashboard with Professional Care Staff", ERGO/FEPS/110534.

ACKNOWLEDGMENTS

This work is supported by Responsible Ai UK (EP/Y009800/1) (https://rai.ac.uk/).

REFERENCES

[1] Julia Amann, Alessandro Blasimme, Effy Vayena, Dietmar Frey, Vince I Madai, and Precise4Q Consortium. 2020. Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC medical informatics and decision making* 20, 1 (2020), 310.

[2] Ali Azadi and Francisco José García-Peñalvo. 2025. Optimizing Clinical Decision Support System Functionality by Leveraging Specific Human-Computer Interaction Elements: Insights From a Systematic Review. *JMIR Human Factors* 12, 1 (2025), e69333.

[3] Maria Cvach. 2012. Monitor alarm fatigue: an integrative review. *Biomedical instrumentation & technology* 46, 4 (2012), 268–277.

[4] Giovanni Diraco, Alessandro Leone, Pietro Siciliano, Marco Grassi, and Piero Malcovati. 2012. A multi-sensor system for fall detection in ambient assisted living contexts. In *International Conference on Sensor Networks*, Vol. 2. SCITEPRESS, 213–219.

[5] Dorset Council. 2025. Acoustic Monitoring in Care Homes. https://www.local.gov.uk/case-studies/dorset-council-acoustic-monitoring-care-homes. Local Government Association Case Study. Accessed: 2025-12-04.

[6] Adam Eck, Leen-Kiat Soh, and Prashant Doshi. 2023. Decision making in open agent systems. *AI Magazine* 44, 4 (2023), 508–523.

[7] Marzyeh Ghassemi, Tristan Naumann, Peter Schulam, Andrew L Beam, Irene Y Chen, and Rajesh Ranganath. 2020. A review of challenges and opportunities in machine learning for health. *AMIA Summits on Translational Science Proceedings* 2020 (2020), 191.

[8] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. *Advances in neural information processing systems* 29 (2016).

[9] Andreas Holzinger, Chris Biemann, Constantinos S Pattichis, and Douglas B Kell. 2017. What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923* (2017).

[10] Shuting Hu, Siyang Cao, Nima Toosizadeh, Jennifer Barton, Melvin G Hector, and Mindy J Fain. 2024. Radar-based fall detection: A survey [survey]. *IEEE robotics & automation magazine* 31, 3 (2024), 170–185.

[11] Elizabeth Anna Mathilde Michels, Stephen Gilbert, Iulia Koval, and Magdalena Katharina Wekenborg. 2025. Alarm fatigue in healthcare: a scoping review of definitions, influencing factors, and mitigation strategies. *BMC nursing* 24, 1 (2025), 664.

[12] Reuth Mirsky and Peter Stone. 2021. The seeing-eye robot grand challenge: rethinking automated care. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*.

[13] OECD. 2023. *Health at a Glance 2023: OECD Indicators*. OECD Publishing, Paris.

[14] World Health Organization. 2024. *WHO clinical consortium on healthy ageing 2023: meeting report, Geneva, Switzerland, 5-7 December 2023*. World Health Organization.

[15] Jae Mun Sim, Yonnim Lee, and Ohbyung Kwon. 2015. Acoustic sensor based recognition of human activity in everyday life for smart home services. *International Journal of Distributed Sensor Networks* 11, 9 (2015), 679123.

[16] Deepika Singh, Johannes Kropf, Sten Hanke, and Andreas Holzinger. 2017. Ambient assisted living technologies from the perspectives of older people and professionals. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 255–266.

[17] Hannelore Strauven, Chunzhuo Wang, Hans Hallez, Vero Vanden Abeele, Bart Vanrumste, et al. 2024. Unobtrusive nighttime movement monitoring to support nursing home continence care: algorithm development and validation study. *JMIR nursing* 7, 1 (2024), e58094.

[18] Jenna Wiens, Suchi Saria, Mark Sendak, Marzyeh Ghassemi, Vincent X Liu, Finale Doshi-Velez, Kenneth Jung, Katherine Heller, David Kale, Mohammed Saeed, et al. 2019. Do no harm: a roadmap for responsible machine learning for health care. *Nature medicine* 25, 9 (2019), 1337–1340.

[19] Michael Wooldridge. 2009. *An introduction to multiagent systems*. John wiley & sons.