Responsible use of artificial intelligence in the provision of long-term care for older people: a care-centric approach

Caroline Emmer De Albuquerque Green, Tyler Reinmund, Kate Hamblin, Samir K Sinha

Current approaches to the role of artificial intelligence (AI) in the provision of long-term care for older people are largely framed around solving perceived problems in the sector, such as managing workforce shortages by driving greater efficiency through the automation of administrative and care tasks. Although such approaches might highlight some benefits of AI, they tend to overlook broader contextual and ethical implications within the complex structures of care systems. Thus, such narrow approaches can compromise care quality and pose risks for care recipients, caregivers, and care services. In this Personal View, we advocate for an alternative, care-centric approach to AI in long-term care, grounded in co-production and rooted in the view that care is a human need tied to wellbeing, dignity, equality, and human rights. We propose a definition of the responsible use of AI in long-term care with values of care at the forefront. We propose to use this definition as a starting point to drive AI policy and practice, rather than focusing on perceived problems, while also acknowledging and addressing tensions identified during the recent co-creation of responsible AI guidelines for the UK.

Lancet Healthy Longev 2026

Published Online
https://doi.org/10.1016/j.lanhl.2026.100817

Institute for Ethics in AI, Faculty of Philosophy, Schwarzman Centre for the Humanities, University of Oxford, Oxford, UK (C Emmer De Albuquerque Green PhD); **Department of Computer Science, Human Centred Computing, University of Oxford, Oxford, UK** (T Reinmund MSc); **Centre for Care, Centre for International Research on Care, Labour and Equalities (CIRCLE), University of Sheffield, Sheffield, UK** (Prof K Hamblin PhD); **Department of Medicine, University of Toronto, Toronto, ON, Canada** (Prof S K Sinha MD DPhil)

Correspondence to:
Dr Caroline Emmer De Albuquerque Green, Institute for Ethics in AI, Faculty of Philosophy, Schwarzman Centre for the Humanities, University of Oxford, Oxford OX2 6GG, UK
caroline.green@philosophy.ox.ac.uk

Introduction

Long-term care is often defined as supporting people who have disability, chronic illness, or frailty to lead the lives of their choosing as much as possible. Long-term care is multifaceted; it consists of emotional, relational, and practical elements, such as providing support with activities of daily living.[1] The provision of long-term care occurs wherever people who require care and support live or stay, including their own homes, residential care settings, or hospitals, and is offered by paid care providers or unpaid caregivers. It is distinct from, yet complementary to, the provision of health care, including the prevention, diagnosis, and medical treatment of health conditions. Although the provision of long-term care is not limited to older adults, the likelihood that people will require long-term care increases with age.[2]

Various types of artificial intelligence (AI) have been tested and used in the provision of long-term care for older adults for the past 5–10 years in countries including Japan,[3] Germany,[4] and the UK.[5] However, AI enthusiasm across the long-term care sector has increased markedly in recent years with the emergence of new generative AI systems at the forefront of public adoption and the development of many new applications. Generative AI, as found in general-purpose chatbots such as OpenAI's ChatGPT and Google's Gemini or in more specialised care-focused systems, is already being used in care services across the UK to alleviate pressures caused by high administrative workloads.[6] Emerging applications of generative AI in the care sector include the transcription of care assessment conversations,[7] the provision of care-related advice and information,[5] and care-specific voice-activated applications or skills for smart speakers (panel 1).[11]

However, the deployment of AI systems, including generative AI, also raises practical, ethical, and governance considerations that might be linked to the design of the technology itself or to the way people use it in long-term care. For example, generative AI models are known to reproduce biases inherent in the training data, including ageism and ableism.[12] An incident made headlines in the USA, wherein a student had prompted Google's Gemini to generate responses to care-related questions. The output suggested that older people who need care were "not special…not important…and not needed" and should "please die", clearly reproducing shocking bias with potentially harmful impact.[13] Generated outputs can also be factually incorrect or misleading. In Australia, a social worker used a chatbot to generate notes for a child protection case. The report incorrectly reported on key events and the social worker did not verify its factual correctness, which could have led to potentially devastating outcomes for the child. Similar cases are conceivable in the care of older people.[14]

In this Personal View, we introduce a care-centric approach for the use of AI in long-term care. This approach is framed around a definition of the responsible use of AI in long-term care, which offers a foundation on which to build AI policy and implementation and is as follows: "The responsible use of (generative) AI in social care means that the use of AI systems in the care or related to the care of people supports and does not undermine, harm, or breach fundamental values of care, human rights, independence, choice and control, dignity, equality, and wellbeing."[15] The definition was co-produced by different actors across the UK's long-term care sector, also referred to as the social care sector. In this Personal View, we draw on generative AI as an example because of its increasingly wide use in UK long-term care services. However, the approach and definition are independent of specific AI systems.

Limitations of narrow outcome–based AI approaches

Many current UK policy and private sector narratives on AI in long-term care are outcome driven and narrowly framed around particular problems, associated with a care crisis. One commonly identified problem is the heavy administrative workload borne by social care workers, of whom

Panel 1: **High-level definition of generative artificial intelligence**

Generative artificial intelligence (AI) refers to the application of machine learning algorithms to produce new data (eg, text, images, videos, code, and audio) by drawing on patterns learned from training datasets.[8] Applications such as ChatGPT or Claude combine multiple models and elements such as a user interface and prompt engineering into generative AI systems.[9] These systems mark a milestone in AI development in terms of both technical achievement and widespread adoption. Generative AI systems have the potential to be applied across a wide range of tasks, due in part to their ability to operate across multiple data modalities, and, given the ease with which the systems can be accessed, many individuals and organisations have already begun to experiment with generative AI across a diverse assortment of professional activities, based on evidence from a preprint paper by Handa and colleagues.[10] Furthermore, downstream developers can build on pretrained generative AI models or integrate with leading generative AI systems to address specific use cases.

there is a considerable shortage.[15] AI systems are offered as a viable solution to make care provision more efficient and effective[16] by reducing the time spent on administrative tasks. A proposed secondary outcome is that this reduction will enable caregivers to spend more time interacting with the people they care for.[17] Although AI might offer benefits in long-term care, including reduced administrative burden, these problem–outcome-driven approaches to identify the direction of AI in long-term care fall short in contextualising AI deployment in wider care ecosystems and the sociopolitical, legal, and relational contexts of care. The focus of such approaches is therefore narrow. Ethical risks and stakeholder voices frequently take second place to desired outcomes, leading to measures that might not benefit or could even harm those affected. Indeed, how a problem is framed shapes what solutions are proposed.[18]

Furthermore, narrow outcome-driven approaches often do not clearly allocate responsibilities to care providers, regulators, or other actors to ensure that benefits are achieved, risks minimised, or harm remediated. This limitation is evident in the UK, where AI enthusiasm in the provision of long-term care is not matched by equal consideration of governance and regulatory clarity. Broad ethical frameworks for the general use of AI, and even for the use of AI in the provision of health care, are not useful in long-term care contexts, in which there are distinct care-related and regulatory considerations at stake.[19] There is still a paucity of official guidelines on AI-related risks and harms in the provision of long-term care and on how such risks will be identified and responded to under the current or future regulatory frameworks of long-term care service provision or other relevant regulations.

An alternative, responsible framing to the use of AI in the provision of long-term care aims to depart from a set of inconsistently selected and narrowly defined outcomes to a more holistic guiding vision. Responsible AI frameworks recognise the complex societal system of power structures and human decisions on processes such as who develops AI systems, how they are used, and who has access to them.[20] These frameworks seek to foster awareness and reflexivity among multiple actors regarding possible and actual impacts related to the use of AI on individuals and society.[21] Such frameworks allocate responsibilities to key actors and build accountability mechanisms and consider that adverse effects might stem from the technology itself, from how AI is used or misused, and from how AI narratives become part of wider social and political structures.

Owing to the complexity and nuance of many social and ethical challenges surrounding the use of AI in long-term care, scholars have recommended that responsible AI approaches are grounded within specific local and technological contexts.[22] Therefore, in this Personal View, our objective was to collaborate with the UK long-term care community to establish the definition of the responsible use of AI in the UK's long-term care sector as a first step towards identifying a guiding approach to local AI development, policy, and practice.

A cross-sector, co-production model to define the responsible use of AI in long-term care

Cross-sector collaborations can be effective mechanisms for supporting the development of improved solutions to obtain social benefits.[23] Such efforts refer to the strategic collaboration of public, private, and civil society organisations with a shared goal[24] and are built on the recognition that new technologies cannot be successfully embedded across areas such as public services by depending on one entity alone; instead, this requires strategic cooperation between multiple actors.[25] These actors might include companies that develop and deploy technologies, policy makers and government bodies that embed and use these technologies, citizens, and research institutes. Each actor brings their own expertise, which must be combined to create inclusive AI strategies that highlight benefits, address ethics, and tackle roll-out challenges.

Between February, 2024, and March, 2025, we coordinated a cross-sector collaboration to co-produce an approach for the responsible use of AI that is grounded in the priorities and concerns of the long-term care community relating to AI. We convened more than 100 representatives from the long-term care community from across the UK; these representatives are henceforth referred to as collaborators. The collaborators included care and support recipients, unpaid caregivers, professional care workers, care providers, technology providers, policy makers, decision makers, and civil society representatives.

Panel 2: **Social Care Institute for Excellence's guiding principles for co-production**

Equality

No one group or person is more important than another, and everyone has skills and abilities to contribute.

Diversity

Co-production should be as inclusive and diverse as possible, with particular attention given to under-represented groups.

Accessibility

Measures must be taken to ensure that everyone has an equal opportunity to participate fully in an activity in the way that suits them best.

Reciprocity

Participants should receive something in return for their contribution.[28]

Our approach to cross-sector collaboration is aligned with the principles of co-production. Co-production is a collaborative model of research and practice central to the provision of long-term care in the UK using qualitative research methods such as focus group implementation; the approach seeks to rebalance power between decision makers and people with lived experience and to increase the relevance of research findings.[26] The term co-production itself is broad and often encompasses closely related terms such as co-design, co-research, co-creation, and collaboration.[27] Therefore, instead of striving for a universal definition, those wishing to engage in co-production should focus on articulating underpinning values first and then applying them in practice.[27] Our approach uses co-production principles defined by the Social Care Institute for Excellence (SCIE) and Think Local, Act Personal (known as TLAP), two organisations focused on improving the provision of social care services in England (panel 2).

TLAP's approach, inspired by early work on citizen participation in social programmes, conceptualises co-production as a ladder. Ranked lowest on the ladder is coercion, which includes no opportunity to hear the views of stakeholders who typically lack power in relation to services. Ranked higher up the ladder, co-production considers a collaborator's level of education and knowledge, how much they are consulted and engage with the process, and the degree to which they participate in co-design. At the top of the ladder is co-production in which relationships between stakeholders are equal.[29] We also followed the Care Workers' Charity's guide to centring care workers.[30]

Phases of co-production

Our co-production effort consisted of four phases, during which the collaborators engaged in a series of focus group meetings and deliberations drawing on their expertise and experiences. Three hosts from across academia, the private sector, and civil society, involved in long-term care research, coordinated this collaboration, and a steering committee composed of representatives from across the long-term care community was responsible for oversight and deliberation on strategic decisions.

Phase 1: establishing relationships and defining the problem

The first, formative phase of the collaboration included a roundtable discussion on the use of AI in the provision of long-term care. Representatives from across the long-term care community were invited to participate. The goals of this session were to develop a preliminary, shared understanding of the current state of use of AI in long-term care, establish cross-sector relationships, and define the objectives of the co-production project. Based on the contributions made and priorities identified during the roundtable discussion, we published a shared statement on the responsible use of generative AI in social care. This statement highlighted the need for collaboration to address gaps in sector-relevant understanding of and guidance on how to use AI more responsibly in long-term care, given the rapid roll-out of generative AI tools.[31] The statement was endorsed by approximately 30 organisations across the long-term care sector in the UK. Additionally, two aims were defined for the remainder of the co-production process. The first aim was to define the responsible use of AI in long-term care, and the second was to identify gaps and needs for guidance or other outputs to ensure that generative AI, as a specific example of AI, is used responsibly in the provision of long-term care.

Phase 2: working group collaboration

The second phase involved the formation of specific working groups, representing people who received or provided care, paid care professionals, care provider organisations, and technology providers. The groups met three to five times to explore what AI meant for their care, work, and lives and created the outputs and messages that they wanted to share with the public. Two thematic working groups were convened, one to define values and principles of good care and a second to explore an ethical and evaluation framework for the use of AI in long-term care.

Most working groups began their first sessions by exploring what AI in general, and generative AI in particular, meant to them. The participants of these sessions

Panel 3: **Artificial intelligence guidance in long-term care—domains of interest and main concerns**

Improving care and support

Artificial intelligence (AI) systems should be safe and effective to support and improve caregiving, ensuring enhanced outcomes.

Choice and control

People should be able to make informed choices regarding the type of AI systems they use and are subjected to.

Accessibility

AI systems should be accessible to people with different needs, and the costs of AI systems should not create new inequities.

Training

People should be able to learn about the AI systems used in care provision and empowered to use AI responsibly and make choices.

Data privacy

People should have access to information regarding their own personal data and how the data are handled and stored to be able to make informed choices.

Transparency

The use of AI systems in people's care should be transparent and enabled by accessible information.

Human contact and connections

AI systems should not replace human interactions or care provision but support and enhance caring relationships.

Bias and discrimination

Bias and discrimination in AI systems should be actively identified and addressed through proactive measures.

Continuous improvement

User feedback on AI use should be taken seriously, leading to system improvements or risk mitigation when needed.

Co-production

Co-production should be merged with the AI lifecycle and into AI roll-outs in care services and wider policy.

Sustainable technology

People should be informed of how AI in long-term care fits into the broader context and economy of care provision and its social impact.

shared experiences of how AI is already being used at home and places of work, discussed specific AI products (many of which were generative), and explored their hopes and concerns about the use of AI in long-term care. Collaborators agreed that AI can be beneficial in the provision of long-term care, by supporting care recipients, care workers, and care providers. However, the contributors also recognised that AI systems, their use or misuse in the provision of long-term care, and broader AI policy could be harmful.

As these conversations progressed, the working groups reflected on their roles as developers, users, or subjects of AI in long-term care. Each group chose an output to co-create and share publicly. For example, technology providers created a pledge, inviting endorsements from other companies, to commit to developing and marketing products in a fair, transparent, and accountable way. Paid care professionals issued demands for better AI training, stronger accountability, and sustained investment in human caregivers. People who receive care produced a statement highlighting what matters most to them regarding the use of AI in their care. Care providers developed guidance for introducing AI into their services.

Phase 3: deliberation and output refinement

In a third, co-productive and deliberative phase, the outputs of all working groups were subject to an overarching co-production working group. This group scrutinised all outputs from the previous working groups and created a new document, including a definition of the responsible use of AI in long-term care and principles for practice (panel 3).

In the third phase, all collaborators gathered and established a basic shared understanding of long-term care, independent from AI. This understanding assumed that AI has enormous potential to improve the lives of people who need and provide care and support. However, the starting point should not be technological potential but the role of long-term care in the protection and advancement of human rights, independence, choice and control, dignity, and the wellbeing of people who receive care, as enshrined in international and national law and regulation.

Collaborators agreed that AI must benefit or at least not harm people who receive care, as specified in the definition of responsible AI presented in the introduction of this Personal View.

Phase 4: deliberative assembly and finalisation

In the fourth phase, the definition of the responsible use of AI in long-term care and principles for practice were subject to a deliberative assembly, with more than 50 collaborators physically present and deliberating together on the final definitions and practice guidance.[15] This final phase highlighted various points of consensus.

First, collaborators agreed that multiple actors were responsible for ensuring that AI was being used responsibly in long-term care and that accountability systems were necessary for these actors, including technology developers, care providers, and policy makers. Although long-term care is tightly regulated in the UK, collaborators urged for clearer links between long-term care regulation and AI oversight to ensure safety, accountability, and mechanisms for redress.

Second, collaborators shared a mutual understanding that long-term care involves a plurality of aspects that contribute to people's wellbeing, independence, and dignity. AI systems might support some of these aspects, such as administrative tasks, but individuals should still have a choice to receive human-led caregiving, and using AI in the provision of long-term care should be a well-informed and safe choice.

Third, collaborators identified specific domains that require prompt attention (panel 3).[32] Finally, collaborators agreed on the need to look beyond values associated with care to define the responsible use of AI in long-term care, while also examining environmental sustainability and broader social considerations.

We addressed ethical concerns—particularly equity, inclusivity, and transparency—during the collaboration, by recognising the power dynamics within the long-term care system. The work was led by experienced facilitators and civil society groups, ensuring accessible engagement, open communication, and public sharing of outcomes. Ethical clearance for research was granted by the University of Oxford research ethics board (reference number: R93793/RE002).

Concerns and tensions

Throughout the four phases, we gained insights into collaborators' concerns and uncovered key challenges to implementing an approach to the responsible use of AI in long-term care. These challenges were not limited to the technology or how it was being used in the provision of care. A key concern was that government and providers might embrace the use of AI in the provision of long-term care too eagerly, promoting narrow, outcome-driven narratives without sufficient evidence, regulation, or oversight. Collaborators expressed concern that AI adoption could prompt cuts to staffing in residential care homes and other care services as a cost-saving measure, potentially undermining overall care quality. Another concern was that care recipients could face reduced choice and human contact, as they might instead be offered technology-based options that might not suit their individual needs. Collaborators were also concerned about equity of access to AI technologies, with the financial resources for procurement and maintenance of such systems not accessible equally to everyone.

Our co-production exercise furthermore identified tensions between some of the represented groups within the collaboration. Many care recipients and caregivers feel under-represented and unfairly treated in the current UK long-term care policy landscape.[33] The question arose of who should benefit most from AI and whose rights and wellbeing should be the primary concern in policies and regulation. Some groups expressed mistrust in policy makers and care providers, doubting that AI will be used to truly serve people who give or receive care. Furthermore, concerns were raised about how co-production fits into the AI developmental lifecycle, revealing tensions between technology developers and those in the care community. Collaborators noted that co-production principles are rarely followed, limiting meaningful input from people with lived experience. However, openness and dialogue can be challenging for developers under pressure to compete and bring products to market quickly. One approach is to preserve a competitive edge by safeguarding technological innovations through legal mechanisms such as intellectual property rights, although co-production might risk exposure of innovations. Co-production can also be time consuming and expensive.[34]

Discussion

Responsible AI approaches seek to understand, frame, and respond to AI innovation and deployment within existing power structures in society, between organisations and individuals.[35] These approaches recognise that, despite all the incredible developments in and capabilities of AI systems, humans remain responsible for developing, deploying, using, regulating, and creating narratives around AI. Humans also make decisions that can reinforce, create, or indeed challenge power imbalances between groups and individuals.[36] This perspective is particularly important to recognise in regard to the provision of long-term care for older people. Business realities and political and economic interests to manage rising demands by using AI systems and technologies can clash with the needs and human rights of people who receive or provide care. In other words, in the context of development and roll-out of AI in long-term care, the priorities of people who receive care might not match the priorities of policy makers, decision makers, technology developers and providers, or care providers.

By taking a broader look and creating space for reflection, this collaboration helped to clarify what people truly mean by and value in care.

First, the care-centric definition of responsible AI use in long-term care underscores that long-term care has deeper

importance for individuals who use services, extending beyond practical help or emotional support. The need for long-term care is linked to people's physical or mental conditions, helping to maintain independence, autonomy, dignity, and human rights. Aspects of care are rarely one dimensional. Rather, effective care is flexible, empathetic, responsive, and focused on people and their situations in any given moment of time. In this regard, AI might only be a supportive tool, to be put into the service of meeting this need and helping to support people's independence and choices. However, AI must be considered as part of a whole-life context, with a personalised plan that considers at the individual level if, when, and how to integrate AI and other technologies into care at any given time to serve their needs. For example, an AI-powered chatbot might support the needs of a person who receives care for information and advice. However, this advice and information is useless if, for example, the individual needs physical help to get off the floor after having fallen and the chatbot tells them to get someone to support them but there is no human available to help them. Thus, AI would likely only ever be one aspect of care and not replace all of what care provision stands for, similar to how practical support is only one aspect of care and does not represent the whole.

Under the proposed definition of the responsible use of AI in long-term care, care providers, professional care workers, and unpaid caregivers provide essential support, often driven by personal duty and purpose across diverse tasks. Although some of these actions might be supported, enhanced, or even replaced by AI systems, the essential human core of care provision is irreplaceable.[37] The collaboration made clear that the responsible use of AI in long-term care should preserve human interaction and not trade it off for cost or time savings. Instead, investment in human carers and other options to meet people's needs must continue.

Second, this work offered an opportunity to highlight tensions between people and care organisations in long-term care, which exist and unfold independently from AI. Our care-centric approach seeks to empower both people who receive care and care workers. Thereby, it challenges a situation in which these groups are largely subjected to the decisions and interests of more powerful groups such as policy makers, technology developers, and care organisations without any meaningful participation or control over how AI is being developed and placed within the provision of long-term care. Instead, our approach encourages use of what people value in care as a starting point to the development and roll-out of AI systems. This method also clarifies that outcome-driven approaches are not the problem per se. Rather, outcome-driven approaches to AI in long-term care can be problematic when shaped by dominating or influential groups, excluding care recipients, care workers, care providers, and caregivers, lacking oversight of both outcomes and processes behind them. Our care-centric approach offers a guiding vision for setting outcomes and centring care recipients, care workers, and care providers, in defining and monitoring both goals and processes.

However, open questions remain. How can we ensure that a care-centric AI approach is genuinely implemented and not just used as rhetoric or even to legitimise irresponsible practices? How can we pursue the responsible use of AI in long-term care without falling into rigid dogmas that might stifle technological advancements, policy and practical innovation, fresh thinking, and better care solutions? The answers to these questions might lie in practising and deepening co-production and cross-sector collaboration techniques in the provision of long-term care, which involve multiple collaborators who are treated as equals. The collaborators develop bonds of trust, a shared sense, and tools to practise responsibility and accountability mechanisms together.[38,39] Open dialogue based on transparency, evolving evidence, and goodwill to build improved systems will be crucial, as will flexible regulatory mechanisms that allow for innovation but are protective of values of care. Navigating and deconstructing pre-existing conditions of distrust among people, organisations, and institutions and building trust are major challenges, but they can also be viewed as opportunities.[40] Of course, the realities of financial pressures and economic contexts remain, which need to be considered, and for which AI systems and technologies—when developed and deployed correctly—can offer positive solutions.

The UK's cross-sector collaboration is continuing into its second stage, tackling the above-discussed and other open questions together and building actionable good practice solutions.[41]

Conclusion

AI systems and technologies can positively influence the lives of people who receive and provide care. However, such systems cannot transform human needs and values relevant in caring relationships, such as choice and control over one's care, positive human interactions, and personalised solutions that work in individuals' own life contexts. Our care-centric approach for the responsible use of AI in long-term care puts humans first and AI in the service of people. It refocuses on what humans value in care provision, on the human rights and wellbeing of people, and on the impact of AI systems on their lives. Therefore, the approach departs from what drives current narratives around the use of AI in the provision of long-term care, such as narrowly defined economic outcomes, and offers a wider, people-focused starting point. However, many open questions related to putting this definition into practice remain. Next steps include the need to name and address tensions between people, organisations, and contextual realities; allocating responsibilities; and building mechanisms to develop, evaluate, monitor, and respond to harms. A major challenge, and opportunity, will be to protect what we value while embracing AI and what it can offer the

long-term care sector. Importantly, open and democratic dialogues and co-production are necessary, allowing for what might require bold and courageous decisions along the way.

Contributors

CEDAG conceptualised the study. CEDAG, KH, and TR designed the methodology. TR, KH, and SKS validated the data. CEDAG and TR were involved in data curation, analysis, investigation, and project administration. CEDAG and TR wrote the original draft. CEDAG and TR revised the manuscript in response to reviewers' and editors' comments. All authors reviewed and edited the manuscript.

Declaration of interests

CEDAG is Director of Research at the Institute for Ethics in AI, which is funded by the Stephen A Schwarzman Foundation. KH is Director of the Centre for Care, which is funded by the Economic and Social Research Council (ESRC; award ES/W002302/1) with contributions from the National Institute for Health and Care Research (NIHR; Department of Health and Social Care). The views expressed are those of the authors and not necessarily those of the ESRC, UK Research and Innovation, National Health Service, the NIHR, or the Department of Health and Social Care. SKS received salary support to conduct research as a 2024–25 AMS Fellow in Compassion and AI. SKS received salary support and support for living and research expenses from Sinai Health for a research sabbatical at the University of Oxford by Sinai Health. SKS also received non-salary support from the Institute for Ethics in AI to serve as a 2024–25 Visiting Fellow at the Centre for Ethics in AI, University of Oxford and at All Souls College, University of Oxford. TR's doctoral research is co-funded by Newton Europe and the Oxford–Singapore Human–Machine Collaboration Programme, supported by a gift from Amazon Web Services.

Acknowledgments

The study was funded by the Josef und Luise Kraft Stiftung through a grant to the Institute for Ethics in AI. Additional support was provided by the Accelerator Fellowship Programme at the Institute for Ethics in AI. The Care Summit was partly funded by the Josef und Luise Kraft Stiftung and the Accelerator Fellowship Programme at the Institute for Ethics in AI. The funders of the study had no role in study design, data collection, data analysis, data interpretation, or writing of the report. We thank the individuals who led the working groups in this collaboration: Kate Jopling from Think Local, Act Personal (TLAP); Donald Macaskill from Scottish Care; Dr Jane Townson OBE from the Homecare Association; Karolina Gerlich from the Carer Workers Charity; and John Boyle. We also thank the project co-hosts Daniel Casson and Katie Thorn from the Digital Care Hub. We furthermore acknowledge all collaborators, including individuals with lived experience, care workers, technology providers, care providers, and policy makers who shared their perspectives and supported this work. We also acknowledge the organisations that co-led this work, in particular TLAP, the Care Workers' Charity, and the Digital Care Hub. During the preparation of this work the authors used ChatGPT EDU to improve the readability of some sentences. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

1 WHO. Long-term care. Dec 1, 2022. https://www.who.int/europe/news-room/questions-and-answers/item/long-term-care (accessed June 23, 2025).
2 Rahman MM, Rosenberg M, Flores G, et al. A systematic review and meta-analysis of unmet needs for healthcare and long-term care among older people. *Health Econ Rev* 2022; **12:** 60.
3 Takemura S. The challenge to develop and implement artificial intelligence (AI) technologies in health and medical care in Japan. *J Nat Inst Public Health* 2023; **72:** 2–13.
4 Breuer S, Müller R. Digitalization, AI, and robotics for good care and work? German policy imaginaries of healthcare technologies. *Sci Public Policy* 2024; **51:** 951–62.
5 Hamblin K. Sustainable social care: the potential of mainstream "smart" technologies. *Sustainability* 2022; **14:** 2754.
6 Mistry P. More than just hype: how emerging AI use is assisting health and social care. May 16, 2025. https://www.kingsfund.org.uk/insight-and-analysis/long-reads/emerging-ai-use-health-and-social-care (accessed May 26, 2025).
7 Booth R. Social workers in England begin using AI systems to assist their work. Sept 28, 2024. https://www.theguardian.com/society/2024/sep/28/social-workers-england-ai-system-magic-notes (accessed Nov 23, 2025).
8 Feuerriegel S, Hartmann J, Janiesch C, Zschech P. Generative AI. *Bus Inf Syst Eng* 2024; **66:** 111–26.
9 Sarkar A. Will code remain a relevant user interface for end-user programming with generative AI models? Onward! 2023: proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software, Oct 25–27, 2023.
10 Handa K, Tamkin A, McCain M, et al. Which economic tasks are performed with AI? Evidence from millions of Claude conversations. *arXiv* 2025; published online Feb 11. https://doi.org/10.48550/arXiv.2503.04761 (preprint).
11 Wright J. The Alexafication of adult social care: virtual assistants and the changing role of local government in England. *Int J Environ Res Public Health* 2021; **18:** 812.
12 Rickman S. Evaluating gender bias in large language models in long-term care. *BMC Med Inform Decis Mak* 2025; **25:** 274.
13 Clark A, Melissa M. Google AI chatbot responds with a threatening message: "Human ... Please die." Nov 20, 2024. https://www.cbsnews.com/news/google-ai-chatbot-threatening-message-human-please-die/ (accessed July 4, 2025).
14 Office of the Victorian Information Commissioner. Investigation into the use of ChatGPT by a child protection worker. Sept 17, 2024. https://ovic.vic.gov.au/wp-content/uploads/2024/11/DFFH-ChatGPT-investigation-report-20240924-Re-upload.pdf (accessed July 4, 2025).
15 Green C, Thorn K, Boyle J, Jopling K, Casson D. The responsible use of generative AI in social care: a valued led approach. 2025. https://www.oxford-aiethics.ox.ac.uk/sites/default/files/2025-04/AI-in-Social-Care-White-Paper-April-2025-Institute-for-Ethics-in-AI.pdf (accessed Jan 23, 2025).
16 Age UK. State of health and care of older people in England 2024. Sept 24, 2024. https://www.ageuk.org.uk/discover/2024/september/state-of-health-and-care-of-older-people-in-england-2024/ (accessed July 4, 2025).
17 Whitfield G, Wright J, Hamblin K. AI in care: a solution to the 'care crisis' in England? In: Paul R, Carmel E, Cobbe J, eds. Handbook on public policy and artificial intelligence. Edward Elgar Publishing, 2024: 383–96.
18 Whitfield G, Hamblin K. 'Thanks to technology': discourse, care and technology in England. *Ageing Soc* 2025; **45:** 2062–88.
19 Solaiman B. Legal and ethical considerations of artificial intelligence for residents in post-acute and long-term care. *J Am Med Dir Assoc* 2024; **25:** 105105.
20 Vallor S. Defining human-centered AI. In: Régis C, Denis J-L, Axente ML, Kishimoto A, eds. Human-centered AI: A multidisciplinary perspective for policy-makers, auditors, and users. Chapman and Hall/CRC, 2024: 13–20.
21 Dignum V. Responsible artificial intelligence: recommendations and lessons learned. In: Eke DO, Wakunuma K, Akintoye S, eds. Responsible AI in Africa. Challenges and opportunities. Springer International Publishing, 2023: 195–214.
22 Lukkien DRM, Nap HH, Buimer HP, et al. Toward responsible artificial intelligence in long-term care: a scoping review on practical approaches. *Gerontologist* 2023; **63:** 155–68.
23 Mikhaylov SJ, Esteve M, Campion A. Artificial intelligence for the public sector: opportunities and challenges of cross-sector collaboration. *Philos Trans A Math Phys Eng Sci* 2018; **376:** 20170357.
24 Blake M, Lambert C, Burkitt R, et al. NHSX adult social care technology and digital skills review. November, 2021. https://www.ipsos.com/sites/default/files/ct/publication/documents/2021-12/NHSX_Technology_and_Digital_Skills_Review_Main_Report_November_2021.pdf (accessed July 4, 2025).

25 Reinmund T, Kunze L, Jirotka MD. Transitioning towards a proactive practice: a longitudinal field study on the implementation of a ML system in adult social care. CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, May 11–16, 2024.

26 Redman S, Greenhalgh T, Adedokun L, Staniszewska S, Denegri S. Co-production of knowledge: the future. *BMJ* 2021; **372:** n434.

27 Masterson D, Areskoug Josefsson K, Robert G, Nylander E, Kjellström S. Mapping definitions of co-production and co-design in health and social care: a systematic scoping review providing lessons for the future. *Health Expect* 2022; **25:** 902–13.

28 Social Care Institute for Excellence. Co-production: what it is and how to do it. July, 2022. https://www.scie.org.uk/co-production/what-how/ (accessed July 4, 2025).

29 Arnstein SR. A ladder of citizen participation. *J Am Inst Plann* 1969; **35:** 216–24.

30 The Care Workers' Charity. Centring care workers: a guide. https://www.thecareworkerscharity.org.uk/centring-care-workers-a-guide/ (accessed Nov 24, 2025).

31 Institute for Ethics in AI. Oxford statement on the responsible use of generative AI in adult social care. https://www.oxford-aiethics.ox.ac.uk/oxford-statement-responsible-use-generative-ai-adult-social-care (accessed July 4, 2025).

32 Digital Care Hub. Responsible use of generative AI in social care: guidance. https://www.digitalcarehub.co.uk/ai-and-robotics/oxford-project-the-responsible-use-of-generative-ai-in-social-care/responsible-use-of-generative-ai-in-social-care-guidance/ (accessed July 4, 2025).

33 Kenway E. Who cares: the hidden crisis of caregiving, and how we solve it. Wildfire, 2023.

34 Oliver K, Kothari A, Mays N. The dark side of coproduction: do the costs outweigh the benefits for health research? *Health Res Policy Syst* 2019; **17:** 33.

35 Dignum V. Responsible artificial intelligence. How to develop and use AI in a responsible way. Springer, 2019.

36 Vallor S. Carebots and caregivers: sustaining the ethical ideal of care in the twenty-first century. *Philos Technol* 2011; **24:** 251–68.

37 Pugh AJ. The last human job. The work of connecting in a disconnected world. Princeton University Press, 2024.

38 Gentelet K, Mizrahi SK. A human-centered approach to AI governance. Operationalizing human rights through citizen participation. In: Regis C, Denis J-L, Axente ML, Kishimoto A, eds. Human-centered AI. A multidisciplinary perspective for policy-makers, auditors, and users. Chapman and Hall/CRC, 2024: 215–30.

39 Tollon F, Vallor S. The responsible AI ecosystem: a BRAID landscape study. June 18, 2025. https://zenodo.org/records/15195686 (accessed Jan 23, 2025).

40 Afroogh S, Akbari A, Malone E, Kargar M, Alambeigi H. Trust in AI: progress, challenges, and future directions. *Humanit Soc Sci Commun* 2024; **11:** 1568.

41 Digital Care Hub. Ensuring the responsible use of generative AI in social care: a collaborative call to action. March 25, 2025. https://www.digitalcarehub.co.uk/ai-and-robotics/oxford-project-the-responsible-use-of-generative-ai-in-social-care/ensuring-the-responsible-use-of-generative-ai-in-social-care-a-collaborative-call-to-action/ (accessed July 5, 2025).