AJJ Research

Frontier Systems Technical Paper

Company-Led Technical Preprint / Technical Disclosure Version

Version 1.0 | 2026

RR-Ethics™: A Human-Centred Ethical Readiness Index for Responsible Deployment of Embodied AI Humanoid Eldercare Robotics

A Systems-Oriented HERI Framework for Legal Accountability, Safety Boundaries, Public Health Protection and Institutional Deployment Governance

Xin Zhao[1†], Weijie Tan[1†], Tian Shen[1†], Wentao Zhao[1], Qiang Huang[1], Hon Hsiang Ong[1], Wenqun Guo[1], Tingting Shen[2], Chunqiu Yan[2], and Jian Zhang[1,*]

[1] AJJ Healthcare Management Pte. Ltd., Singapore.

[2] Hangzhou Huaxi Intelligent Technology Co. Ltd., Hangzhou, China.

[†] These authors contributed equally to this work and share first authorship.

* Correspondence: zhangjian@ajjmedtech.com.sg

Document Type: Company-led technical preprint / technical disclosure for DOI registration

Status: Public research version; not independently peer-reviewed; not a journal article or final published paper

DOI: 10.5281/zenodo.21002067, https://zenodo.org/records/21002067 |
Website: https://www.ajjmedtech.com.sg/publications-3 | DOI metadata type: technical report

Research Programme: Embodied AI Humanoid Eldercare Robotics

Important Notice

This public version presents the Humanoid Eldercare Robotics Ethical Readiness Index (HERI™) Version 1.0 for transparency, academic discussion and future development purposes. HERI Version 1.0 represents an initial framework contribution in an emerging field and should be interpreted together with its stated research boundaries and limitations. Future development may include broader expert review, multi-site evaluation, inter-rater reliability assessment and additional validation activities.

Publication and Evidence Boundary

HERI Version 1.0 is presented as a framework-definition contribution developed within a bounded pioneer-developer context involving AJJ Healthcare Management and Hangzhou Huaxi Intelligent Technology. The framework is informed by literature review, system-level analysis and preliminary deployment-related observations. It should not be interpreted as an independently validated clinical instrument, regulatory standard or universal deployment certification system. Evidence discussed in this version should be interpreted as framework-support evidence intended to inform future refinement and evaluation activities.

Prior Public Version and Future Journal Submission Boundary

This document is released as a public framework-definition version intended to support transparency, discussion and institutional awareness. It is not presented as a peer-reviewed journal publication and should not be interpreted as the final scholarly version of HERI. Future journal submissions may contain expanded methodology, additional evidence, independent validation results, refined scoring approaches and updated framework components.

Abstract

The rapid emergence of embodied AI humanoid robots in eldercare settings is creating new opportunities for supporting ageing populations, workforce shortages and long-term care delivery. However, the deployment of such systems also raises significant ethical, legal, governance and institutional readiness challenges involving vulnerable older adults, caregivers, healthcare organisations and public-interest obligations. Although existing literature provides important contributions in AI ethics, care robotics, dignity protection, accountability and human–robot interaction, most frameworks remain principle-oriented and do not provide a practical mechanism for evaluating whether a robot is ethically ready for deployment in a real-world care environment. To address this gap, RR-Ethics™ introduces the Humanoid Eldercare Robotics Ethical Readiness Index (HERI™), a structured and evidence-oriented readiness framework designed to support institutional deployment assessment of embodied AI humanoid eldercare robotics. HERI Version 1.0 consists of ten assessment dimensions covering dignity and personhood protection, consent and autonomy safeguards, privacy and data proportionality, safety and medical boundary control, human oversight and override, legal accountability, auditability, public health protection, emotional dependency management and caregiver fairness. The framework combines evidence-based scoring, weighted readiness assessment, critical-dimension floors and readiness-level interpretation to support transparent and systematic deployment review. HERI is intended as a decision-support framework rather than a regulatory, certification, legal or clinical approval mechanism. As an initial framework contribution developed within a bounded pioneer-developer context, HERI Version 1.0 should be regarded as a foundation for future expert review, independent evaluation and continued framework refinement.

Keywords: Humanoid eldercare robotics; ethical readiness; embodied AI; healthcare governance; responsible AI; long-term care; robotics ethics; institutional deployment

1. Introduction

Population ageing is rapidly reshaping healthcare and long-term care systems worldwide[1]. Many countries are experiencing increasing demand for eldercare services while simultaneously facing persistent workforce shortages, rising care costs and growing pressure on institutional care environments . These challenges have stimulated significant interest in the development and deployment of advanced robotics, artificial intelligence (AI) and embodied AI systems capable of supporting care delivery[2].

Recent advances in embodied AI humanoid robotics have expanded the potential role of robots beyond industrial automation and logistics. Humanoid eldercare robots are increasingly being developed for functions such as companionship, information assistance, monitoring, mobility support, activity facilitation and selected care-related tasks. As these systems become more capable and socially interactive, they are also becoming more deeply integrated into environments involving vulnerable older adults, caregivers, healthcare organisations and public-interest responsibilities.

The deployment of such systems creates opportunities to improve care capacity, operational efficiency and service accessibility. However, it also introduces a complex set of ethical, legal, governance and institutional challenges. Questions concerning dignity, autonomy, privacy, accountability, transparency, safety, emotional dependency and human oversight become increasingly important as robots move from experimental settings into real-world care environments[2-8].

Over the past decade, substantial research has been conducted in fields including AI ethics, healthcare AI governance, social robotics, human–robot interaction and responsible innovation[2-8]. Existing literature has

contributed valuable insights regarding ethical principles, risk identification and normative guidance for care-related technologies. Nevertheless, most existing approaches remain principle-oriented and do not provide institutions with a practical mechanism for evaluating whether a specific robotic system is ethically prepared for deployment.

In practice, healthcare providers, long-term care operators and technology developers often face a more operational question:

Is a humanoid eldercare robot ethically ready for deployment within a real-world care environment?

Answering this question requires more than the identification of ethical concerns. Institutions must be able to assess evidence, evaluate safeguards, review governance structures and determine whether appropriate protections have been established before deployment occurs.

To address this implementation gap, RR-Ethics™ introduces the Humanoid Eldercare Robotics Ethical Readiness Index (HERI™). HERI is designed as a structured and evidence-oriented framework for assessing ethical readiness in institutional eldercare settings. Rather than focusing solely on ethical principles, HERI seeks to translate ethical expectations into observable deployment requirements that can be reviewed, documented and assessed.

HERI Version 1.0 consists of ten assessment dimensions covering dignity and personhood protection, consent and autonomy safeguards, privacy and data proportionality, safety and medical boundary control, human oversight, accountability, auditability, public health protection, emotional dependency management and caregiver fairness. Together, these dimensions form a readiness architecture intended to support transparent, systematic and human-centred deployment evaluation.

As an initial framework contribution developed within an emerging field, HERI Version 1.0 should be understood as a foundation for future refinement, independent evaluation and broader validation efforts.

2. Existing Ethical Assessment Gap

The growing body of literature on AI ethics, healthcare governance and social robotics has significantly advanced understanding of the ethical challenges associated with robotic care technologies. Existing studies have explored important issues including dignity protection, autonomy, accountability, privacy, emotional dependency, trust, human–robot interaction and responsible deployment within care environments[2-8].

These contributions provide valuable ethical foundations and have helped identify many of the risks that may emerge when robots are introduced into eldercare settings. However, most existing approaches focus primarily on ethical principles, conceptual analysis or risk discussion. Relatively few provide a structured mechanism for determining whether a specific robotic system is ethically prepared for deployment within a real-world institutional environment.

From an operational perspective, care providers, healthcare organisations and deployment partners require more than ethical awareness. They must also be able to evaluate evidence, assess safeguards, review governance mechanisms and determine whether minimum readiness conditions have been achieved before deployment occurs.

As a result, a practical gap remains between ethical principles and deployment decision-making. Existing frameworks often help answer:

What ethical issues should be considered?

What risks may arise?

What principles should guide responsible innovation?

However, institutions frequently require an additional question to be addressed:

Has sufficient ethical readiness been demonstrated to support deployment?

HERI was developed to address this implementation-oriented gap. Rather than replacing existing ethical frameworks, HERI seeks to translate key ethical expectations into a structured readiness architecture supported by evidence requirements, weighted assessment, critical-dimension safeguards and deployment readiness interpretation.

Within the reviewed comparator literature and public-source search scope, HERI is positioned as a first-definition-type, deployment-oriented ethical readiness framework that operationalises established care-robot ethics, healthcare AI governance and AI robot accountability principles into a structured, evidence-linked assessment architecture for embodied AI humanoid eldercare robotics.

Table 1 summarises several influential contributions that informed HERI development and highlights the distinction between principle-oriented ethics frameworks and deployment-oriented readiness assessment.

Table 1. Comparator literature, HERI leverage points and boundary of use.

Boundary of use	HERI leverage point	Authority / scope	Comparator literature
Use as gap evidence; do not treat as readiness-index validation.	Shows care-robot ethics is well established but mainly principle-based.	Systematic review of argument-based aged-care robot ethics.	Vandemeulebroucke et al. (2018) [3]
Do not overstate as a humanoid deployment index.	Supports consent, dependency, emotional replacement and caregiver justice dimensions.	Long-term care social-robot ethics, including dementia, consent and substitution.	Hung et al. (2025) [4]
Use as literature-landscape evidence, not as a scoring architecture.	Supports the claim that ethical issues are rich, broad and fragmented.	Broad review of ethical aspects of social robots in elderly care.	Leineweber et al. (2026) [5]
Avoid leaving dignity as abstract prose; convert it into criteria.	Anchors dignity and personhood as operational readiness conditions.	Dignity-based bioethics framework for robotic assistance.	Felber et al. (2022) [6]
Use for accountability logic; do not present as eldercare-specific.	Supports responsibility chains, legal accountability and evidence trails.	AI robot accountability and organisational responsibility framework.	Toth et al. (2022) [7]

3. HERI V1.0 Methodology and Scoring Architecture

3.1. Framework Purpose

The Humanoid Eldercare Robotics Ethical Readiness Index (HERI™) was developed to provide a structured, evidence-oriented mechanism for assessing ethical readiness prior to the deployment of embodied AI humanoid robots in institutional eldercare environments. Unlike principle-oriented ethics frameworks that primarily identify ethical concerns or normative obligations, HERI focuses on deployment readiness. The framework is designed to support institutional review by translating ethical expectations into observable evidence requirements, structured assessment criteria and readiness-level interpretation. HERI does not assess technological capability, engineering performance or commercial competitiveness. Instead, the framework focuses on whether sufficient ethical, governance, accountability and safety safeguards have been established before deployment occurs.

The framework is intended to support:

- institutional deployment review;
- governance and oversight assessment;
- deployment-risk identification;

- documentation and evidence evaluation;
- readiness benchmarking across care environments.

HERI should be interpreted as a decision-support framework rather than a regulatory approval mechanism, legal opinion, certification programme or clinical authorisation system. To operationalise these objectives, HERI integrates ethical principles, governance safeguards, evidence requirements and deployment-readiness criteria into a unified assessment architecture.

A HERI assessment should be tied to a specific robot configuration, deployment function, institution, software version and assessment date, and should not be generalised across different deployment contexts without further review.

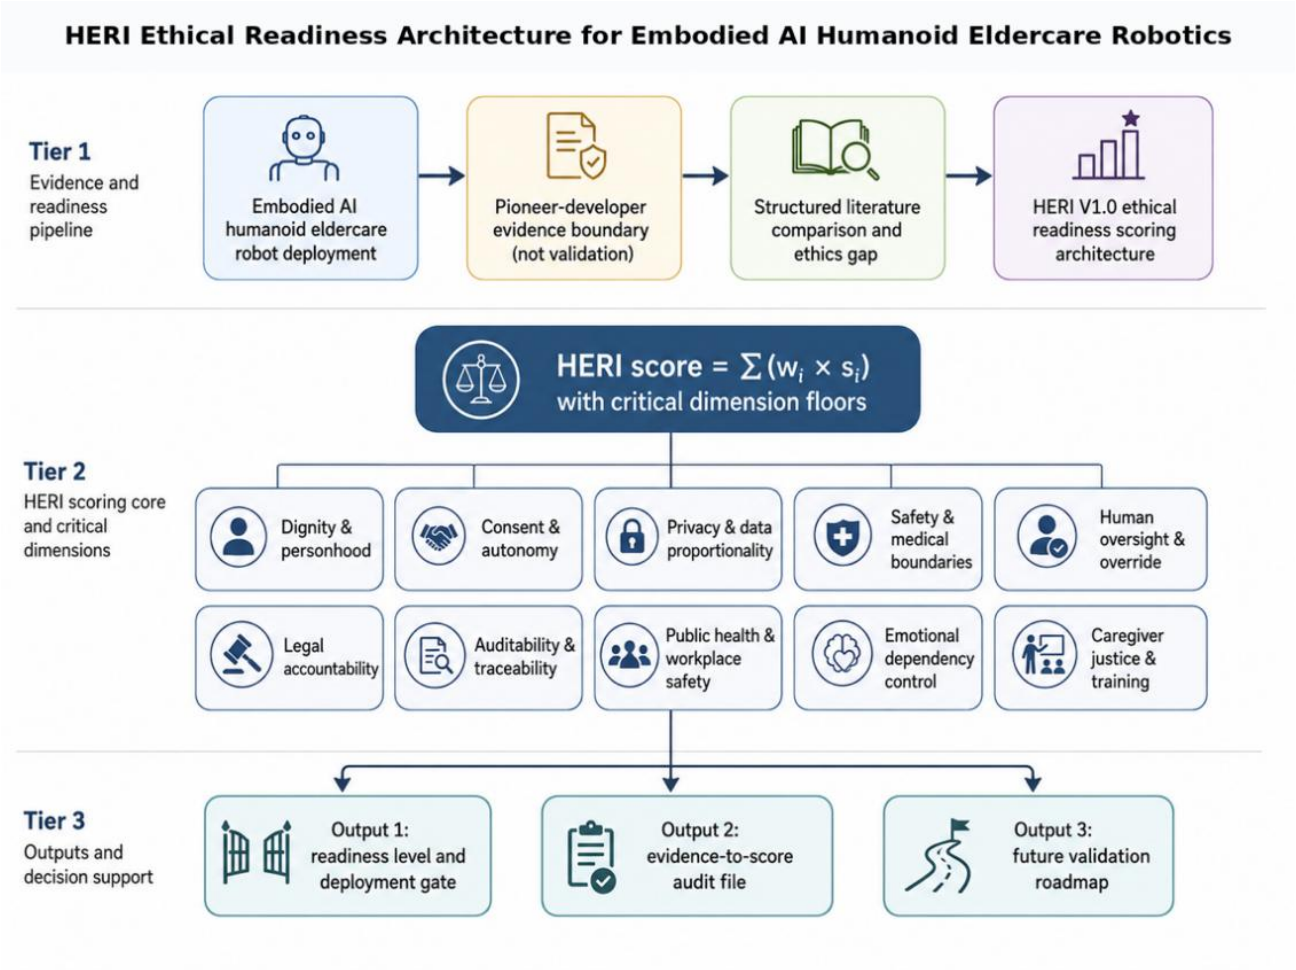

Figure 1. HERI ethical readiness architecture for embodied AI humanoid eldercare robotics.

The architecture links the pioneer-developer evidence boundary, structured literature comparison, ten HERI dimensions, critical floors, deployment gates and future validation roadmap.

3.2. Composite Formula, Scoring Scale and Weighting Principle

Each HERI dimension is assessed using a 0–5 evidence-maturity scoring scale. This scale is designed to distinguish between absence of evidence, initial awareness, partial documentation, defined controls, evidence-supported implementation and mature auditable readiness.

The weighted raw HERI score is calculated as follows:

$$HERI_{raw} = \sum_{i=1}^{10} w_i s_i, \text{ where } \sum_{i=1}^{10} w_i = 1 \text{ and } 0 \le s_i \le 5$$

where s_i represents the **0–5** score of HERI dimension i, and w_i represents the assigned weight of that dimension.

For readiness interpretation, the raw score is normalised to a **0–1** scale:

$$HERI_{norm} = \frac{HERI_{raw}}{5}$$

Using the proposed HERI V1.0 dimension set, the weighted formula may be expanded as follows:

$$HERI_{raw} = w_1 S_{HD} + w_2 S_{IC} + w_3 S_{PDP} + w_4 S_{SMB} + w_5 S_{HO} + w_6 S_{LA} + w_7 S_{AT} + w_8 S_{PHW} + w_9 S_{EDR} + w_{10} S_{CJT}$$

where S_{HD} , S_{IC} , S_{PDP}, S_{SMB} , S_{HO}, S_{LA} , S_{AT}, S_{PHW}, S_{EDR} and S_{CJT} represent the **0 – 5** evidence-maturity scores of the ten HERI dimensions.

where the abbreviations are defined as follows:

- ***HD*** = Dignity and personhood protection;
- ***IC*** = Consent, autonomy and proxy safeguards;
- ***PDP*** = Privacy and data proportionality;
- ***SMB*** = Safety and medical boundary control;
- ***HO*** = Human oversight and override;
- ***LA*** = Legal accountability and responsibility chain;
- ***AT*** = Auditability and traceability;
- ***PHW*** = Public health and workplace safety;
- ***EDR*** = Emotional dependency and substitution control;
- ***CJT*** = Caregiver justice, training and workflow fairness.

The proposed weighting structure is initial and evidence-informed. It reflects the higher deployment significance of safety and medical boundary control while preserving balanced attention to dignity, consent, privacy, oversight, accountability, auditability, public health, emotional dependency and caregiver justice. Future expert review, inter-rater reliability testing, Delphi-style refinement, sensitivity analysis and multi-site validation may further refine these weights. HERI should not be interpreted as an automatic mathematical decision tool. A deployment may achieve a relatively high aggregate score but still remain ethically unready if critical safeguards are insufficient.

3.3. Ten HERI Dimensions, Literature Anchors and Evidence Requirements

HERI Version 1.0 evaluates ethical readiness through ten dimensions representing key ethical, legal, governance, public-health and operational considerations associated with the deployment of embodied AI humanoid robots in eldercare environments. The dimensions were derived from the comparator literature reviewed in Section 2 and translated into deployment-oriented assessment criteria through the HERI framework architecture. Rather than functioning as isolated ethical categories, the dimensions are intended

to operate as an integrated readiness system in which dignity, autonomy, privacy, safety, oversight, accountability and organisational safeguards collectively contribute to deployment readiness. Each dimension is associated with a defined literature anchor, evidence requirements, core scoring question and deployment risk rationale. Together, these dimensions provide the substantive assessment foundation of HERI Version 1.0. Table 2 presents the ten HERI dimensions, associated literature anchors and evidence requirements.

Table 2. HERI dimensions, weights, literature anchors, evidence requirements, scoring focus and risk if absent.

No.	HERI dimension and initial weight	Literature anchor	Evidence requirements and scoring focus	Core scoring question	Risk if absent
D1	Dignity and personhood protection Weight: 0.10	Felber et al. [6]; Vandemeulebroucke et al. [3]	Role communication, avoidance of infantilisation, privacy-sensitive assistance and preservation of meaningful human care.	Does deployment preserve dignity and avoid reducing residents to workflow objects?	Residents may be reduced to workflow objects or exposed to dignity harm.
D2	Consent, autonomy and proxy safeguards Weight: 0.10	Hung et al. [4]; care-robot ethics [3]	Consent scripts, proxy-consent logic, resident preference records, opt-out procedures and communication of robot limitations.	Are residents or authorised representatives able to understand and control participation?	Residents or authorised representatives may lack meaningful control over participation.
D3	Privacy and data proportionality Weight: 0.10	AI ethics and healthcare AI governance [2]	Data map, access controls, minimisation policy, retention rules, anonymisation safeguards and purpose limitation.	Is data collection proportionate to defined care, safety and institutional purposes?	Data collection may exceed care and safety purposes or become insufficiently governed.
D4	Safety and medical boundary control Weight: 0.15	Robot safety and healthcare AI governance [2]	Risk assessment, permitted-task list, restricted-function list, safety testing, escalation triggers and non-clinical boundary statements.	Are safety limits, task boundaries and non-clinical boundaries documented and enforced?	Physical, public health or medical-adjacent risks may be underestimated.
D5	Human oversight and override Weight: 0.10	Responsible AI and HRI literature [2]	Supervision plan, override logs, escalation pathways, staff authority rules and fallback procedures.	Can humans supervise, intervene and reverse robot-supported workflows when required?	Robot-supported workflows may proceed without effective human control.
D6	Legal accountability and responsibility chain Weight: 0.10	Toth et al. [7]	Responsibility matrix, contract/institutional role allocation, incident ownership, maintenance responsibility and dispute pathway.	Who is accountable when robot-supported deployment causes harm, confusion or operational failure?	Accountability may become dispersed or unclear after harm, confusion or incident.
D7	Auditability and traceability Weight: 0.10	Toth et al. [7]; governance literature	Evidence repository, system logs, scoring files, review records, version records and source-record controls.	Can readiness claims be independently reviewed through evidence, logs and source records?	Readiness claims may be unverifiable and unable to support external review.

No.	HERI dimension and initial weight	Literature anchor	Evidence requirements and scoring focus	Core scoring question	Risk if absent
D8	Public health and workplace safety Weight: 0.10	Healthcare AI governance, public-health protection and workplace-safety literature [2]	Hygiene protocols, infection-control compatibility, staff safety procedures, workplace risk assessment and training records.	Does deployment protect residents, staff and institutional public health obligations?	Deployment may create avoidable risks for residents, staff or institutional operations.
D9	Emotional dependency and substitution control Weight: 0.075	Hung et al. [4]; social-robot ethics [5]	Dependency risk review, human-care substitution controls, disclosure of robot limitations and monitoring for over-reliance.	Does the robot avoid inappropriate emotional replacement or weakening of human care?	Companionship functions may inappropriately replace or weaken human care.
D10	Caregiver justice, training and workflow fairness Weight: 0.075	Hung et al. [4]; LTC implementation literature	Training records, workload assessment, staff feedback, role-protection policies and digital-proficiency support.	Does deployment support rather than burden, displace or deskill caregivers?	Automation may shift burden, deskill staff or create unequal adaptation pressure.

3.4. Evidence-Maturity Scoring Levels

To support consistent assessment across institutions and deployment contexts, HERI adopts a structured evidence-maturity scoring approach. The scoring framework is intended to evaluate not only whether a safeguard exists, but also the extent to which it is documented, implemented, evidenced and auditable. The evidence-maturity model recognises that ethical readiness develops progressively. Organisations may demonstrate awareness of a requirement without having fully implemented controls, while more mature deployments are expected to provide documented procedures, supporting records and evidence of operationalisation. Accordingly, HERI uses a five-level evidence-maturity framework ranging from absence of evidence to mature and auditable implementation. Table 3 summarises the evidence-maturity scoring levels used throughout HERI Version 1.0.

Table 3. Evidence-maturity scoring guide for each HERI dimension.

Score	Evidence maturity	Evidence meaning	Dimension-level interpretation
0	Absent	No evidence, no assigned responsibility or issue not assessed.	No dimension-level readiness demonstrated.
1	Initial awareness	Concern recognised but without documented controls.	Dimension recognised, but controls are not established.

Score	Evidence maturity	Evidence meaning	Dimension-level interpretation
2	Partial documentation	Some policies or records exist, but implementation evidence is incomplete or unclear.	Partial dimension-level preparation; implementation evidence remains incomplete.
3	Defined controls	Controls, roles and procedures are documented and can be reviewed.	Documented dimension-level controls; aggregate stage-gate assessment remains required.
4	Evidence-supported implementation	Controls are implemented, logged, trained and supported by source records.	Evidence-supported dimension-level implementation; not an overall deployment decision.
5	Mature and auditable	Evidence is complete, traceable, reviewed and subject to continuous improvement.	Mature dimension-level evidence; aggregate stage-gate assessment and formal approvals remain required.

3.5 Readiness Levels and Critical Floors

The overall HERI score is intended to support readiness interpretation rather than provide a simple pass-or-fail outcome. Ethical readiness exists along a continuum, and institutions may demonstrate varying levels of preparedness depending on governance maturity, evidence quality and deployment safeguards. To facilitate deployment review, HERI groups overall scores into readiness categories that reflect increasing levels of ethical preparedness[8]. These categories are designed to support institutional decision-making, remediation planning and future deployment benchmarking. However, ethical readiness cannot be determined solely by aggregate scores. Certain dimensions represent foundational safeguards whose absence may create unacceptable deployment risks regardless of overall performance in other areas. For this reason, HERI incorporates critical-dimension floors to prevent high aggregate scores from masking significant deficiencies in core areas such as safety, oversight, accountability or privacy protection. Table 4 presents the HERI readiness-level interpretation framework.

Table 4. HERI readiness levels and deployment stage-gate.

$HERI_{norm}$ score	Overall readiness level	Deployment interpretation
$HERI_{norm} \geq 0.85$	Level 4	Threshold-positive ethical readiness for controlled institutional deployment, provided all Level 4 critical-floor requirements are met.
$0.70 \leq HERI_{norm} < 0.85$	Level 3	Acceptable for limited ethically supervised pilot deployment if all Level 3 critical-floor requirements are met and remediation actions are defined.
$0.55 \leq HERI_{norm} < 0.70$	Level 2	Conditional remediation required before deployment.
$HERI_{norm} < 0.55$	Level 1	Not ethically ready for deployment.

Table 5 presents the critical-dimension floor requirements applied within HERI Version 1.0.

Table 5. Critical dimension floors and stage-gate logic.

Critical area	Suggested floor	Rationale	Gate effect
Safety and medical boundary control	Level 3 pilot: ≥3 Level 4 deployment: ≥4	Deployment cannot proceed if task limits, restricted functions, safety testing, escalation triggers and non-clinical boundaries are undefined or not implemented.	Failure overrides aggregate score and prevents positive readiness interpretation.
Human oversight and override	Level 3 pilot: ≥3 Level 4 deployment: ≥4	Deployment cannot proceed without documented and implemented human supervision, intervention authority, escalation pathways and fallback procedures.	Failure requires remediation before pilot or controlled deployment.
Privacy and data proportionality	Level 3 pilot: ≥3 Level 4 deployment: ≥4	Deployment cannot proceed if data collection, access, retention, purpose limitation and privacy safeguards are not governed and evidenced.	Failure prevents threshold-positive interpretation.
Legal accountability and responsibility chain	Level 3 pilot: ≥3 Level 4 deployment: ≥4	Deployment cannot proceed if incident responsibility, maintenance ownership, escalation duties and dispute pathways are unclear or unsupported by records.	Failure blocks responsible institutional use.
Auditability and traceability	Level 3 pilot: ≥3 Level 4 deployment: ≥4	Deployment cannot proceed if readiness claims cannot be reviewed through evidence, logs, source records, scoring files and version records.	Failure prevents auditable readiness assessment.

Note. Level 3 critical floors support limited supervised pilot interpretation only. Level 4 controlled institutional deployment requires all critical dimensions to score at least 4 and no non-critical HERI dimension below 3. These proposed floor requirements remain subject to future expert review, inter-rater reliability testing and sensitivity validation.

4. Illustrative Evidence Requirements and Example Application

HERI is intended to support evidence-based ethical readiness assessment rather than principle-based review alone. Accordingly, each dimension is associated with observable evidence requirements that may be examined during institutional assessment. Examples of supporting evidence may include governance documents, deployment protocols, safety procedures, consent processes, audit records, staff training materials, accountability structures, privacy safeguards and operational review mechanisms. The purpose of evidence collection is not to demonstrate perfection, but to establish whether sufficient safeguards have been implemented and documented to support responsible deployment. To illustrate how evidence may be mapped to HERI dimensions, Version 1.0 includes an example evidence-reference structure. The example is intended for explanatory purposes and should not be interpreted as a mandatory documentation standard. Table 6 provides an illustrative mapping between HERI dimensions and representative evidence categories that may be considered during readiness assessment.

Table 6. Illustrative evidence-to-score mapping for the HERI dimensions.

HERI dimension	Illustrative evidence file	How evidence supports scoring	Auditability note
Dignity and personhood	Resident-role communication, privacy-sensitive assistance notes and care relationship safeguards.	Score is supported if records show the robot's role and limits are communicated and access to meaningful human care remains preserved.	Evidence should show resident-facing communication and protection against infantilising or objectifying interaction.
Consent, autonomy and proxy safeguards	Consent/proxy-consent workflow, opt-out pathway and resident preference file.	Score is supported if participation, refusal and proxy decision pathways are documented in a resident-sensitive way.	Records should be dated, attributable and linked to institutional consent procedures.
Privacy and data proportionality	Data map, access logs, anonymisation controls, retention procedure and purpose-limitation statement.	Score is supported if data collection is necessary, proportionate, controlled and linked to defined care or safety purposes.	Access, retention and anonymisation records should be reviewable.
Safety and medical boundary control	Permitted-task list, restricted-function list, safety testing and incident-response plan.	Score is supported if robot functions do not exceed documented boundaries and escalation is defined.	Evidence should identify safety tests, medical-adjacent boundaries and escalation ownership.
Human oversight and override	Supervisor assignment, override records, escalation logs and fallback procedures.	Score is supported if humans remain able to supervise, interrupt and reverse robot-supported workflows.	Override and escalation events should be logged and reviewable.
Legal accountability and responsibility chain	Responsibility matrix, incident ownership, maintenance responsibility and dispute pathway.	Score is supported if institutional and developer responsibilities are traceable before deployment.	Accountability records should clarify ownership for incidents, maintenance and corrective action.
Auditability and traceability	Evidence repository, scoring worksheet, source-record register, version history and access record.	Score is supported if each readiness claim can be audited through source records.	Readiness claims should be traceable to identifiable records rather than narrative assertion.
Public health and workplace safety	Cleaning procedure, infection-control notes, workplace risk assessment and staff training records.	Score is supported if deployment protects residents, staff and institutional public health obligations.	Records should link public health controls to site workflow and staff responsibilities.
Emotional dependency and substitution control	Disclosure of robot limitations, monitoring for over-reliance and human-care substitution safeguards.	Score is supported if companionship or interaction functions do not replace required human attention.	Evidence should show limitation disclosure and monitoring of dependency or over-reliance risks.
Caregiver justice, training and workflow fairness	Training records, workflow feedback, workload-impact notes and staff role-protection records.	Score is supported if deployment reduces burden without unfairly shifting risk, workload or deskilling pressure to caregivers.	Evidence should include staff feedback, training completion and role clarity.

5. Limitations and Future Development

HERI Version 1.0 should be interpreted as an initial framework contribution developed within an emerging field of embodied AI humanoid eldercare robotics. While the framework seeks to translate established ethical principles into a structured readiness architecture, several limitations should be acknowledged. First, HERI Version 1.0 was developed within a bounded pioneer-developer context and therefore does not constitute an independently validated assessment instrument. The framework is intended to provide an initial operational model rather than a finalised standard. Second, the proposed weighting structure reflects evidence-informed judgement and literature-supported prioritisation. Although the weighting approach was designed to align with deployment-risk considerations, future refinement may further improve consistency and generalisability. Third, HERI Version 1.0 has not yet undergone large-scale inter-rater reliability testing across multiple independent institutions. Additional evaluation is required to assess scoring consistency under different deployment conditions and assessment teams.

Future development of HERI may include:

- broader expert review and feedback;
- Delphi-based refinement of dimensions and weighting structures;
- inter-rater reliability assessment;
- multi-site evaluation across diverse eldercare environments;
- sensitivity analysis of scoring and weighting assumptions;
- comparison with emerging international governance frameworks;
- future HERI version updates informed by empirical deployment experience.

HERI is intended to evolve through ongoing review, practical application and independent evaluation. Future versions may incorporate additional evidence, refined scoring approaches and expanded deployment guidance as the field of embodied AI eldercare robotics continues to mature.

6. Conclusions

The emergence of embodied AI humanoid robotics in eldercare environments presents significant opportunities for addressing workforce shortages, supporting care delivery and improving service accessibility. At the same time, these technologies introduce complex ethical, legal, governance and deployment challenges that extend beyond technical performance alone. Existing literature has established important ethical principles and identified many of the risks associated with care-related robotics. However, institutions often require a more operational mechanism for determining whether sufficient ethical safeguards have been established prior to deployment. To address this gap, RR-Ethics™ introduces the Humanoid Eldercare Robotics Ethical Readiness Index (HERI™), a structured and evidence-oriented framework designed to support ethical readiness assessment for embodied AI humanoid eldercare robotics. HERI Version 1.0 integrates ten assessment dimensions, evidence-maturity scoring, weighted evaluation, critical-dimension safeguards and readiness-level interpretation into a unified deployment-review architecture. The framework is intended to support transparent, systematic and human-centred assessment while remaining distinct from regulatory approval, clinical certification or legal authorisation processes. As an initial framework contribution, HERI should be understood as a foundation for continued refinement, independent evaluation and future validation. The framework is expected to evolve alongside advances in robotics, artificial intelligence, eldercare practice and international governance expectations. By translating ethical principles into an operational readiness structure, HERI seeks to contribute to the responsible, accountable and human-centred deployment of embodied AI humanoid robotics within eldercare environments.

Declarations

Prior Public Version Disclosure

This company-led technical preprint / technical disclosure version is made available for transparency, version control and public reference. It has not undergone independent peer review. If a substantially revised manuscript is later submitted to a peer-reviewed journal, this public version and DOI/URL should be disclosed to the target journal in accordance with its editorial policy.

Author Contributions

Conceptualization, Jian Zhang and Chunqiu Yan; methodology, Jian Zhang, Xin Zhao, Weijie Tan, Tian Shen and Hon Hsiang Ong; framework development, Jian Zhang, Xin Zhao, Tian Shen, Weijie Tan, Wentao Zhao, Qiang Huang, Wenqun Guo and Hon Hsiang Ong; institutional workflow interpretation and documentation review, Wentao Zhao, Qiang Huang, Wenqun Guo and Weijie Tan; robotics deployment context and technical interpretation, Tingting Shen, Tian Shen and Chunqiu Yan; evidence coordination and supplementary documentation, Xin Zhao, Tian Shen, Tingting Shen, Chunqiu Yan and Hon Hsiang Ong; writing—original draft preparation, Xin Zhao, Tian Shen, Weijie Tan, Hon Hsiang Ong and Jian Zhang; writing—review and editing, all authors; supervision, Jian Zhang. All authors have read and approved this technical report / technical disclosure version.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable. This technical report presents an ethical-readiness assessment framework and framework-development study. It does not report a clinical trial, clinical intervention, medical treatment evaluation or human-subject biomedical research. The evidence categories discussed in this report are anonymised, aggregated or governance-level records used for framework demonstration and evidence-boundary analysis. Any future empirical validation involving human participants, identifiable personal data, resident-level records, clinical evaluation or interventional deployment should obtain appropriate ethics, consent, institutional, data-protection and regulatory review.

Informed Consent Statement

Not applicable. No personally identifiable resident data, identifiable medical records, biometric identification data or individual clinical decision data are disclosed in this manuscript.

Data Availability Statement

The supporting evidence categories are anonymised, aggregated or retained internally under access-controlled conditions. Identifiable source records, signatures, photographs, resident-related data, staff-identifiable materials and institution-specific operational records are not publicly disclosed. Anonymised summaries and evidence-control materials may be made available by the corresponding author upon reasonable request, subject to institutional approval, confidentiality requirements and applicable data-protection obligations.

Conflicts of Interest

Several authors are affiliated with AJJ Healthcare Management Pte. Ltd. and Hangzhou Huaxi Intelligent Technology Co., Ltd., which are involved in healthcare technology, robotics-related deployment activities or related governance documentation. The RR-Ethics™ / HERI framework is proposed as an academic ethical-readiness assessment framework and should not be interpreted as regulatory approval, clinical validation, medical-device certification, commercial product certification, investment recommendation or universal deployment clearance. The authors declare that this technical report has been prepared for academic analysis, framework development and technical disclosure purposes, and that the interpretation and application of HERI remain subject to independent review, local regulation, institutional governance requirements and future validation.

Declaration of AI-Assisted Language and Editorial Support

AI-assisted tools were used during manuscript preparation to support English language editing, grammar refinement, formatting consistency, table organisation and structural clarity. The underlying research framework, field observations, operational evidence logic, validation design, model calculations, references, interpretation and conclusions were reviewed, verified and approved by the authors. No AI-assisted tool was used to generate original field data, fabricate source records, create unsupported references, replace institutional confirmation, perform independent research judgment or assume authorship responsibility.

Supplementary Materials

No supplementary files are publicly released with this company-led technical preprint / technical disclosure version. The HERI framework is supported by anonymised and institutionally controlled evidence materials retained by the research team for framework-development and evidence-control purposes. These materials are not published as supplementary files on this website and do not disclose identifiable resident data, individual medical records, biometric identifiers, facial images, voice recordings or individual clinical decision data. The evidence materials are used only to support the development of the HERI framework, pioneer-developer evidence mapping and future research refinement. They should not be interpreted as independent third-party validation, regulatory certification, clinical validation, medical-device approval, market assurance or final deployment approval.

References

[1] United Nations. World Population Ageing 2023: Challenges and opportunities of population ageing. United Nations; 2023.

[2] World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: World Health Organization; 2021.

[3] Vandemeulebroucke, T.; Dierckx de Casterlé, B.; Gastmans, C. The Use of Care Robots in Aged Care: A Systematic Review of Argument-Based Ethics Literature. Archives of Gerontology and Geriatrics 2018, 74, 15–25.

https://doi.org/10.1016/j.archger.2017.08.014

[4] Hung L, Zhao Y, Alfares H, Shafiekhani P. Ethical considerations in the use of social robots for supporting mental health and wellbeing in older adults in long-term care. Front Robot AI. 2025;12:1560214. doi:10.3389/frobt.2025.1560214.

[5] Leineweber M, Keusgen CV, Bubeck M, Haltaufderheide J, Ranisch R, Klingler C. Ethical aspects of the use of social robots in caring for older people: A systematic qualitative review. Med Health Care Philos. 2026;29:209-224. doi:10.1007/s11019-025-10313-3.

[6] Felber, N.A.; Pageau, F.; McLean, A.; Wangmo, T. The Concept of Social Dignity as a Yardstick to Delimit Ethical Use of Robotic Assistance in the Care of Older Persons. Medicine, Health Care and Philosophy 2022, 25, 99–110.

https://doi.org/10.1007/s11019-021-10054-z

[7] Tóth, Z.; Caruana, R.; Gruber, T.; Loebbecke, C. The Dawn of the AI Robots: Towards a New Framework of AI Robot Accountability. Journal of Business Ethics 2022, 178, 895–916. https://doi.org/10.1007/s10551-022-05050-z.

[8] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg, MD: NIST; 2023. doi:10.6028/NIST.AI.100-1.